## THE CORAL+ ALGORITHM FOR UNSUPERVISED DOMAIN ADAPTATION OF PLDA

Kong Aik Lee, Qiongqiong Wang, Takafumi Koshinaka

Biometrics Research Laboratories, NEC Corporation, Japan

k-lee@ax.jp.nec.com, q-wang@ah.jp.nec.com, koshinak@ap.jp.nec.com

## ABSTRACT

State-of-the-art speaker recognition systems comprise an xvector (or i-vector) speaker embedding front-end followed by a probabilistic linear discriminant analysis (PLDA) backend. The effectiveness of these components relies on the availability of a large collection of labeled training data. In practice, it is common that the domains (e.g., language, demographic) in which the system is deployed differ from that we trained the system. To close the gap due to the domain mismatch, we propose an unsupervised PLDA adaptation algorithm to learn from a small amount of unlabeled in-domain data. The proposed method was inspired by a prior work on featurebased domain adaptation technique known as the correlation alignment (CORAL). We refer to the model-based adaptation technique proposed in this paper as CORAL+. The efficacy of the proposed technique is experimentally validated on the recent NIST 2016 and 2018 Speaker Recognition Evaluation (SRE'16, SRE'18) datasets.

*Index Terms*— Speaker recognition, domain adaptation, unsupervised, discriminant analysis

## 1. INTRODUCTION

Speaker recognition is the task of recognizing a person from his/her voice given a small amount of speech utterance from the speaker [1]. Recent progresses have shown successful application of deep neural network to derive deep *speaker embeddings* from speech utterances [2, 3]. Analogous to word embeddings [4, 5], a speaker embedding is a fixedlength continuous-value vector that provides a succinct characterization of speakers voice rendered in a speech utterance. Similar to the classical i-vectors [6], deep speaker embeddings live in a simpler Euclidean space where distance could be measured easily, compared to the much complex input patterns. Techniques like within-class covariance normalization (WCCN) [7], linear discriminant analysis (LDA) [8], probabilistic LDA (PLDA) [9, 10, 11] can be applied.

Systems comprising x-vector speaker embedding (and ivector) followed by PLDA have shown state-of-the-art performances on speaker verification task [12]. Training an xvector PLDA system typically requires over hundred hours of training data with speaker labels, and with the requirement that the training set must contain multiple recordings of a speaker under different settings (recording devices, transmission channels, noise, reverberation etc.). These knowledge sources contribute to the robustness of the system against such nuisance factors. The challenging problem of domain mismatch arises when a speaker recognition system is used in a different domain (e.g., different languages, demographic etc.) than that of the training data. Its performance degrades considerably.

It is impractical to re-train the system for each and every domain as the effort at collecting large labelled data sets is expensive and time consuming. A more viable solution is to adapt the already trained model using a smaller, and possibly unlabeled, set of in-domain data. Domain adaptation could be accomplished at different stages of the x-vector (or i-vector) PLDA pipeline. PLDA adaptation is preferable in practice since the same feature extraction and speaker embedding front-end could be used while domain adapted PLDA backbends are used to cater for the condition in each specific deployment.

PLDA adaptation involves the adaptation of its mean vector <sup>1</sup> and covariance matrices. In the case of unsupervised adaptation (i.e., no labels are given), the major challenge is how the adaptation could be performed on the within and between class covariance matrices given that only the total covariance matrix could be estimated directly from the indomain data. In this paper, we show that this could be accomplished by applying similar principle as in the feature-based correlation alignment (CORAL) [14] from which a pseudoin-domain within and between class covariance matrices could be computed. We further improve the robustness by introducing additional adaptation parameter and regularization to the adaptation equation. The proposed unsupervised adaptation method is referred to as CORAL+.

## 2. DOMAIN ADAPTATION OF PLDA

This section presents a brief description of *probabilistic linear discriminant analysis* (PLDA) widely used in state-of-theart speaker recognition system. We then draw attention to the

<sup>&</sup>lt;sup>1</sup>Mean shift due to domain mismatch could be solved by centralizing the datasets to a common origin [13].

domain mismatch issue and how the *correlation alignment* (CORAL) [14, 15] technique deals with it via feature transformation.

#### 2.1. Probabilistic LDA

Let the vector  $\phi$  be a speaker embedding (e.g., x-vector, i-vector, etc.). We assume that the vector  $\phi$  is generated from a linear Gaussian model [8], as follows [9, 16]

$$p(\phi|\mathbf{h}, \mathbf{x}) = \mathcal{N}(\phi|\mu, \mathbf{F}\mathbf{h} + \mathbf{G}\mathbf{x} + \boldsymbol{\Sigma})$$
(1)

The vector  $\mu$  represents the global mean, while **F** and **G** are the speaker and channel loading matrices, and the diagonal matrix  $\Sigma$  models the residual variances. The variables **h** and **x** are the latent speaker and channel variables, respectively. A PLDA model is essentially a Gaussian distribution in the speaker embedding space. This could be seen more clearly in the form of the marginal density:

$$p(\phi) = \mathcal{N}(\phi | \mu, \mathbf{\Phi}_{\rm b} + \mathbf{\Phi}_{\rm w}) \tag{2}$$

The main idea here is to account for the speaker and channel variability with a between-class and a within-class covariance matrices

$$\begin{aligned} \mathbf{\Phi}_{\mathrm{b}} &= \mathbf{F}\mathbf{F}^{\mathsf{T}} \\ \mathbf{\Phi}_{\mathrm{w}} &= \mathbf{G}\mathbf{G}^{\mathsf{T}} + \boldsymbol{\Sigma} \end{aligned} \tag{3}$$

respectively. We refer the readers to [9, 10, 16] for details on the model training procedure.

In a speaker verification task, the PLDA model serves as a backend classifier. For a given pair of enrolment and test utterances, i.e, their speaker embeddings  $\phi_1$  and  $\phi_2$ , we compute the log-likelihood ratio score

$$l(\phi_1, \phi_2) = \frac{p(\phi_1, \phi_2)}{p(\phi_1) p(\phi_2)}$$
(4)

corresponding to the hypothesis test whether the two belong to the same or different speaker. The denominator is evaluated by substituting  $\phi_1$  and  $\phi_2$  in turn to (2). The numerator is computed using

$$p(\phi_1, \phi_2) = \mathcal{N}\left( \begin{bmatrix} \phi_1 \\ \phi_2 \end{bmatrix} \middle| \begin{bmatrix} \mu \\ \mu \end{bmatrix}, \begin{bmatrix} \mathbf{C} & \mathbf{\Phi}_{\mathrm{b}} \\ \mathbf{\Phi}_{\mathrm{b}} & \mathbf{C} \end{bmatrix} \right)$$
(5)

where  $\mathbf{C} = \boldsymbol{\Phi}_{\mathrm{b}} + \boldsymbol{\Phi}_{\mathrm{w}}$  is the total covariance matrix. The assumption is that the unseen data follow the same distribution as given by the within and between classes covariance matrices derived from the training set (i.e., the dataset we used to train the PLDA). Problem arises when the training set was drawn from a domain (out-of-domain) different from that of the enrollment and test utterances (in-domain).

### 2.2. Correlation Alignment

*Correlation alignment* (CORAL) [14] aims to align the second-order statistics, i.e., covariance matrices, of the out-of-domain (OOD) features to match the in-domain (InD)

features. No class (i.e., speaker) label is used and therefore it belongs to the class of unsupervised adaptation techniques. The algorithm consists of two steps, namely, whitening followed by re-coloring. Let  $C_o$  and  $C_I$  be the covariance matrices of the OOD and InD data, respectively. Denote  $\phi$ as a OOD vector, domain adaptation is performed by first whitening and then re-coloring, as follows

where

$$\mathbf{C}_{\mathrm{o}}^{-\frac{1}{2}} = \mathbf{Q}_{\mathrm{o}} \boldsymbol{\Lambda}_{\mathrm{o}}^{-\frac{1}{2}} \mathbf{Q}_{\mathrm{o}}^{\mathsf{T}}$$

 $\boldsymbol{\phi}^{'} = \mathbf{C}_{\mathrm{I}}^{\frac{1}{2}} \mathbf{C}_{\mathrm{o}}^{-\frac{1}{2}} \boldsymbol{\phi}$ 

(6)

whitens the input vector, and

$$\mathbf{C}_{\mathrm{I}}^{\frac{1}{2}} = \mathbf{Q}_{\mathrm{I}} \boldsymbol{\Lambda}_{\mathrm{I}}^{\frac{1}{2}} \mathbf{Q}_{\mathrm{I}}^{\mathsf{T}}$$

does the re-coloring. Here, **Q** and **A** are the eigenvectors and eigenvalues pertaining to the covariance matrices<sup>2</sup>. Such simpler and "frustratingly easy" approach [15] has shown to outperform a more complicated non-linear transformation reported in [18]. In [15], CORAL is performed on the OOD x-vectors (or i-vectors) embeddings, and the transformed vectors (pseudo in-domain) are used to re-train the PLDA. Note that speaker labels of the OOD training data remain the same.

#### 3. THE CORAL+ ALGORITM

CORAL is a feature-based domain adaptation technique [14]. We propose integrating CORAL to PLDA leading to a modelbased domain adaptation.

#### 3.1. Domain adaptation

It is commonly known that a linear transformation on a normally distributed vector leads to an equivalent transformation on the mean vector and covariance matrix of its density function. Let  $\mathbf{A} = \mathbf{C}_{\mathrm{I}}^{1/2} \mathbf{C}_{\mathrm{o}}^{-1/2}$  be the transformation matrix and  $\phi' = \mathbf{A}^{\mathrm{T}} \phi$  the transformed vector. The covariance matrix of the pseudo in-domain vector  $\phi'$  is given by

$$\mathbf{C}_{\mathrm{o}}^{'} = \mathbf{A}^{\mathsf{T}} \mathbf{C}_{\mathrm{o}} \mathbf{A} = \mathbf{A}^{\mathsf{T}} \boldsymbol{\Phi}_{\mathrm{w,o}} \mathbf{A} + \mathbf{A}^{\mathsf{T}} \boldsymbol{\Phi}_{\mathrm{b,o}} \mathbf{A}$$
(7)

Here, we have considered a PLDA trained on OOD data with a total covariance matrix  $C_o = \Phi_{w,o} + \Phi_{b,o}$  given by the sum of within and between class covariance matrices, as noted in Section 2.1. The above equation shows that training a PLDA on the transformed vectors  $\phi'$ , as proposed in [15], is equivalent to transforming the within-class, between-class, and total covariance matrices of a PLDA trained on OOD data.

<sup>&</sup>lt;sup>2</sup>The whitening and re-coloring procedures are better known as the zerophase component analysis (ZCA) transformation [17]. As opposed to principal component analysis (PCA) and Cholesky whitening (and re-coloring), ZCA preserves the maximal similarity of the transformed feature to the original space.



# **Fig. 1**. The effects of regularization. Elements with negative variances are removed automatically.

#### 3.2. Model-level adaptation

Instead of replacing the covariance matrices in an OOD PLDA with pseudo in-domain matrices, model-level adaptation allows us to consider their interpolation

$$\begin{aligned} \mathbf{\Phi}_{\mathrm{b}}^{+} = & (1-\beta)\mathbf{\Phi}_{\mathrm{b,o}} + \beta \mathbf{A}^{\mathsf{T}}\mathbf{\Phi}_{\mathrm{b,o}} \mathbf{A} \\ \mathbf{\Phi}_{\mathrm{w}}^{+} = & (1-\gamma)\mathbf{\Phi}_{\mathrm{w,o}} + \gamma \mathbf{A}^{\mathsf{T}}\mathbf{\Phi}_{\mathrm{w,o}} \mathbf{A} \end{aligned}$$

where  $\{\beta, \gamma\}$  are the adaptation parameters constrained to lie between zero and one. Notice that the first term on the righthand-side of the equations is the OOD between/within covariance matrix while the second term is the pseudo-in-domain covariance matrix. For clarity, we further simplify the adaptation equations, as follows

$$\Phi_{\rm b}^{+} = \Phi_{\rm b,o} + \beta \left( \mathbf{A}^{\mathsf{T}} \Phi_{\rm b,o} \mathbf{A} - \Phi_{\rm b,o} \right)$$

$$\Phi_{\rm w}^{+} = \Phi_{\rm w,o} + \gamma \left( \mathbf{A}^{\mathsf{T}} \Phi_{\rm w,o} \mathbf{A} - \Phi_{\rm w,o} \right)$$

$$(8)$$

The second term on the right-hand-side of the equations represents the new information seen in the in-domain data to be added to the PLDA model.

#### 3.3. Regularized adaptation

The central idea of domain adaptation is to propagate the uncertainty seen in the in-domain data to the PLDA model. The adaptation equations in (8), do not guarantee that the variances, and therefore the uncertainty, increase. In this section, we achieve this goal in the transform space where both the OOD and pseudo-in-domain matrices are simultaneously diagonalized.

Let **B** be an orthogonal matrix such that  $\mathbf{B}^{\mathsf{T}} \Phi \mathbf{B} = \mathbf{I}$  and  $\mathbf{B}^{\mathsf{T}} (\mathbf{A}^{\mathsf{T}} \Phi \mathbf{A}) \mathbf{B} = \mathbf{E}$ , where **E** is a diagonal matrix. This procedure is referred to as *simultaneous diagonalization*. The transformation matrix **B** is obtained by performing twice the *eigenvalue decomposition* (EVD) on the matrix  $\Phi$  and then  $\mathbf{A}^{\mathsf{T}} \Phi \mathbf{A}$  after the first transformation has been applied. The procedure is illustrated in Algorithm 1.

By applying the simultaneous diagonalization on (8), the following adaptation could be obtained:

$$\boldsymbol{\Phi}_{b}^{+} = \boldsymbol{\Phi}_{b,o} + \beta \mathbf{B}_{b}^{-\mathsf{T}} \left( \mathbf{E}_{b} - \mathbf{I} \right) \mathbf{B}_{b}^{-1}$$

$$\boldsymbol{\Phi}_{w}^{+} = \boldsymbol{\Phi}_{w,o} + \gamma \mathbf{B}_{w}^{-\mathsf{T}} \left( \mathbf{E}_{w} - \mathbf{I} \right) \mathbf{B}_{w}^{-1}$$

$$(9)$$

**Algorithm 1:** The CORAL+ algorithm for unsupervised adaptation of PLDA.

Input Output	Out-of-domain PLDA matrices $\{ \Phi_{w,o}, \Phi_{b,o} \}$ In-domain data $X_{InD}$ Adaptation hyper-parameters $\{ \gamma, \beta \}$ Adapted covariance matrices $\{ \Phi_{w}^+, \Phi_b^+ \}$			
Estimate en <b>C</b> I = C	npirical covariance matrix from in-domain data $X_{ m InD}$ ov $(X_{ m InD})$			
Compute ou C <sub>o</sub> = 4	t-of-domain covariance matrix $\Phi_{w,o} + \; \Phi_{b,o}$			
for each $oldsymbol{\Phi}~$ in $ig\{oldsymbol{\Phi}_{ ext{w,o}},oldsymbol{\Phi}_{ ext{b,o}}ig\}$ do				
Com	pute Pseudo In-domain covariance matrix $\mathbf{S} = \mathbf{C}_{I}^{1/2} \mathbf{C}_{o}^{-1/2} \boldsymbol{\Phi} \mathbf{C}_{o}^{-1/2} \mathbf{C}_{I}^{1/2}$			
Find and !	$\begin{array}{l} \{B, E\} \text{ via simultaneous diagonalization of } \Phi\\ S\\ \{Q, \Lambda\} \leftarrow \text{Evd}(\Phi)\\ \{P, E\} \leftarrow \text{Evd}(\Lambda^{-1/2}Q^TSQ\Lambda^{-1/2})\\ B = Q\Lambda^{-1/2}P \end{array}$			
Regularized adaptation of PLDA, $\alpha \in \{\gamma, \beta\}$ $\Phi^+ = \Phi + \alpha \mathbf{B}^{-T} \max(0, \mathbf{E} - I) \mathbf{B}^{-1}$				
Notation	EvD(.) returns a matrix of eigenvectors and the corresponding eigenvalues in a diagonal matrix.			

As before, the between and within class covariance matrices are adapted separately. Notice that the term  $(\mathbf{E} - \mathbf{I})$  will ends up with negative variances if any diagonal elements of  $\mathbf{E}$  is less than one. We propose the following regularized adaptation:

$$\boldsymbol{\Phi}_{\mathrm{b}}^{+} = \boldsymbol{\Phi}_{\mathrm{b},\mathrm{o}} + \beta \mathbf{B}_{\mathrm{b}}^{-\mathsf{T}} max \left(\mathbf{E}_{\mathrm{b}} - \mathbf{I}\right) \mathbf{B}_{\mathrm{b}}^{-1}$$

$$\boldsymbol{\Phi}_{\mathrm{w}}^{+} = \boldsymbol{\Phi}_{\mathrm{w},\mathrm{o}} + \gamma \mathbf{B}_{\mathrm{w}}^{-\mathsf{T}} max \left(\mathbf{E}_{\mathrm{w}} - \mathbf{I}\right) \mathbf{B}_{\mathrm{w}}^{-1}$$

$$(10)$$

The max(.) operator ensures that the variance increases. We refer to the regularized adaptation in (10) as the CORAL+ algorithm, while (9) corresponds to the CORAL+ algorithm without regularization. Algorithm 1 summarizes the CORAL+ algorithm. Figure 1 shows a plot of the diagonal elements of the term ( $\mathbf{E}_{\rm b}$  - I) in (10). Those entries with negative variances were removed automatically by the max(.)operator. It ensures that the uncertainty increases (or stays the same) in the adaptation process.

It is worth noticing that, one could recover the subspace matrices  $\{\mathbf{F}, \mathbf{G}\}$  via EVD. Nevertheless, this is not generally required as scores could be computed by plugging in the adapted covariance matrices  $\Phi_{\rm b}^+$ ,  $\Phi_{\rm w}^+$  and  $\mathbf{C}^+ = \Phi_{\rm b}^+ + \Phi_{\rm w}^+$  into (2) and (5).

#### 4. EXPERIMENT

Experiments were conducted on the the recent SRE'16 and SRE'18 datasets. The performance was measured in terms of *equal error rate* (EER) and *minimum detection cost* (Min-Cost) [19, 20]. The latest SREs organized by NIST have been

focusing on domain mismatch as one of the technical challenges. In both SRE'16 and SRE'18, the training set consists primarily English speech corpora collected over multiple years in the North America. This dataset encompasses Switchboard, Fisher, and the MIXER corpora used in SREs 04 - 06, 08, 10, and 12. The enrollment and test segments are in Tagalog and Cantonese for SRE'16, and Tunisian Arabic for SRE'18. Domain adaptation was performed using the unlabeled subsets provided for the evaluation.

The enrollment utterances have a nominal duration of 60 seconds, while the test duration ranges from 10 to 60 seconds. We used x-vector speaker embedding, which has shown to be very effective for speaker verification task over short utterances. (Recent results show that i-vector is more effective for longer utterance of over 2 minutes). The x-vector extractor follows the same configuration and was trained using the same setup as the Kaldi recipe <sup>3</sup>. A slight difference here is that we used an attention model in the pooling layer and extended the data augmentation [21].

In our experiments, the dimension of the x-vector was 512. As commonly used in most state-of-the-art systems, LDA was used to reduce the dimensionality. We investigated the cases of 150- and 200-dimensional x-vector after LDA projection. CORAL [14] transformation was applied on the raw x-vectors before LDA. The transformed, and then projected x-vectors were used to train a PLDA for the CORAL PLDA baseline. It is worth noticing that the LDA projection matrix was computed from the raw x-vectors, from which the CORAL transformation was also derived. We find that this gives the best performance compared to that reported in [15].

The proposed CORAL+ is a model-based adaptation technique. Domain adaptation is achieved by adapting the parameters (i.e., covariance matrices) pertaining to the OOD PLDA as in (10) and Algorithm 1 using the unlabeled indomain dataset. The adaptation parameters were set empirically to 0.80 in the experiments. Tables 1 and 2 show the performance of the baseline PLDA model trained on the out-of-domain English dataset (OOD PLDA), the PLDA trained on the x-vectors which have been adapted using CORAL (CORAL PLDA), and the OOD PLDA adapted to in-domain with CORAL+ algorithm (CORAL+ PLDA). Also shown in the tables is the CORAL+ adaptation without regularization (w/o reg). This correspond to the use of (9) replacing (10) in Algorithm 1.

The results on both SRE'16 and SRE'18 show consistent improvement of CORAL+ PLDA compared to the OOD PLDA baseline. The relative improvement amounts to 36.6% and 22.35% reduction in EER, and 32.0% and 23.0% reduction in MinCost on SRE'16 and SRE'18, respectively, for LDA dimension of 200. Also shown in the tables is an unsupervised adaptation method implemented in Kaldi<sup>3</sup> (Kaldi PLDA). The proposed CORAL+ PLDA consistently outperforms this baseline on both SRE'16 and SRE'18 though the

**Table 1**. Performance comparison on SRE'16 (CMN). The dimension of x-vector after LDA is 150 and 200. Boldface denotes the best performance for each column.

	LDA 150		LDA 200	
	EER (%)	MinCost	EER (%)	MinCost
OOD PLDA	9.69	0.783	9.94	0.813
Kaldi PLDA	6.82	0.552	6.57	0.558
CORAL PLDA	6.50	0.539	6.31	0.543
CORAL+ PLDA	6.62	0.540	6.30	0.553
w/o reg	6.93	0.544	6.51	0.547

**Table 2.** Performance comparison on SRE'18 (CMN2). The dimension of x-vector after LDA is 150 and 200. Boldface denotes the best performance for each column.

	LDA 150		LDA 200	
	EER (%)	MinCost	EER (%)	MinCost
OOD PLDA	7.19	0.538	7.47	0.569
Kaldi PLDA	6.25	0.435	6.48	0.466
CORAL PLDA	6.22	0.449	6.42	0.482
CORAL+ PLDA	5.95	0.421	5.80	0.438
w/o reg	6.49	0.441	6.33	0.460

improvement over this baseline is more apparent on SRE'18. At LDA dimension of 200, the relative improvement amounts to 10.5% reduction in EER, and 6.0% reduction in MinCost on SRE'18.

Compared to the feature-based CORAL (CORAL PLDA), the benefit of CORAL+ (CORAL+ PLDA) is more apparent on SRE'18. We obtained a relative reduction of 9.7% in EER and 9.1% in MinCost at LDA dimension of 200. It is worth mentioning that SRE'16 has a unlabeled set with about the same size compared to that of SRE'18. Nevertheless, SRE'18 unlabeled set exhibits less variability (speaker and channel). This also explains the benefit of regularized adaptation on SRE'18 when a smaller and constrained unlabelled dataset is available for domain adaptation.

#### 5. CONCLUSION

We have presented the CORAL+ algorithm for unsupervised adaptation of PLDA backend to deal with the domain mismatch issue in practical applications. Similar to the feature-based correlation alignment (CORAL) technique, the CORAL+ domain adaptation is accomplished by matching the out-of-domain statistics to that of the in-domain. We show that statistics matching could be directly applied on PLDA model. We further improve the robustness by introducing additional adaptation parameter and regularization to the adaptation equation. The proposed method shows significant improvement compared to the PLDA baseline. Results also show the benefit of model-based adaptation especially when the data available for adaptation is relatively small and constrained.

<sup>&</sup>lt;sup>3</sup>https://github.com/kaldi-asr/kaldi/tree/master/egs/sre16/v2

## 6. REFERENCES

- J. H. L. Hansen and T. Hasan, "How humans and machines recognize voices: a tutorial review," *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 74–99, 2015.
- [2] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for textindependent speaker verification," in *Proc. Interspeech*, 2017, pp. 999–1003.
- [3] E. Variani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *Proc. ICASSP*, 2014, pp. 4052–4056.
- [4] Y. Bengio, R. Ducharme, and P. Vincent, "A neural probabilistic language model," in *NIPS*, 2000.
- [5] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *NIPS*, 2013.
- [6] N. Dehak, P. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [7] A. O. Hatch, S. Kajarekar, and A Stolcke, "Withinclass covariance normalization for SVM-based speaker recognition," in *Proc. Interspeech*, 2006, pp. 1471– 1474.
- [8] C. Bishop, *Pattern recognition and machine learning*, Springer, New York, 2006.
- [9] S. J. D. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *Proc. ICCV*, 2007, pp. 1–8.
- [10] S Ioffe, "Probabilistic linear discriminant analysis," in *Proc. ECCV, Part IV, LNCS 3954*, 2006, pp. 531–542.
- [11] P. Kenny, "Bayesian speaker verification with heavytailed priors.," in *Proc. Odyssey: Speaker and Language Recognition Workshop*, 2010.
- [12] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *ICASSP*, 2018, pp. 5329– 5333.
- [13] K. A. Lee, V. Hautamaki, T. Kinnunen, et al., "The I4U mega fusion and collaboration for NIST speaker recognition evaluation 2016," in *Proc. Interspeech*, 2017, pp. 1328–1332.
- [14] B. Sun, J. Feng, and K. Saenko, "Return of frustratingly easy domain adaptation," in *Proc. AAAI*, vol. 6, 2016, p. 8.

- [15] J. Alam, G. Bhattacharya, and P. Kenny, "Speaker verification in mismatched conditions with frustratingly easy domain adaptation," in *Proc. Odyssey*, 2018, pp. 176– 180.
- [16] S. J. D. Prince, *Computer vision: models, learning, and inference*, Cambridge University Press, 2012.
- [17] A. Kessy, A. Lewin, and K. Strimmer, "Optimal whitening and decorrelation," *The American Statistician*, vol. 2018, pp. 1–6, 2018.
- [18] W. Lin, M.-W. Mak, L Li, and J.-T. Chien, "Reducing domain mismatch by maximum mean discrepancy based autoencoders," in *Proc. Odyssey*, 2018, pp. 162–167.
- [19] National Institute of Standards and Technology, "NIST 2016 Speaker Recognition Evaluation Plan," *NIST SRE*, 2016.
- [20] National Institute of Standards and Technology, "NIST 2018 Speaker Recognition Evaluation Plan," *NIST SRE*, 2018.
- [21] Koji Okabe, Takafumi Koshinaka, and Koichi Shinoda, "Attentive statistics pooling for deep speaker embedding," in *Proc. Interspeech*, 2018, pp. 2252–2256.