

The IBM Speaker Recognition System: Recent Advances and Error Analysis

Syed Omid Sadjadi¹, Jason W. Pelecanos¹, Sriram Ganapathy²

¹IBM Research, Yorktown Heights, NY, USA

²Electrical Eng. Dept., Indian Institute of Science, Bangalore, India

sadjadi@us.ibm.com

Abstract

We present the recent advances along with an error analysis of the IBM speaker recognition system for conversational speech. Some of the key advancements that contribute to our system include: a nearest-neighbor discriminant analysis (NDA) approach (as opposed to LDA) for intersession variability compensation in the i-vector space, the application of speaker and channel-adapted features derived from an automatic speech recognition (ASR) system for speaker recognition, and the use of a DNN acoustic model with a very large number of output units ($\sim 10k$ senones) to compute the frame-level soft alignments required in the i-vector estimation process. We evaluate these techniques on the NIST 2010 SRE extended core conditions (C1–C9), as well as the *10sec–10sec* condition. To our knowledge, results achieved by our system represent the best performances published to date on these conditions. For example, on the extended *tel-tel* condition (C5) the system achieves an EER of 0.59%. To garner further understanding of the remaining errors (on C5), we examine the recordings associated with the low scoring target trials, where various issues are identified for the problematic recordings/trials. Interestingly, it is observed that correcting the pathological recordings not only improves the scores for the target trials but also for the non-target trials.

Index Terms: deep neural networks, discriminant analysis, fMLLR, i-vector, nearest neighbor, speaker recognition

1. Introduction

In recent years, the research trend in the speaker recognition field has evolved from joint factor analysis (JFA) based methods, which attempt to model the speaker and channel subspaces separately [1], towards the i-vector approach that models both speaker and channel variabilities in a single low-dimensional (e.g., a few hundred) space termed the total variability subspace [2]. State-of-the-art i-vector based speaker recognition systems employ universal background models (UBM), which are based on either unsupervised GMMs [3] or supervised ASR acoustic models (e.g., GMM-HMM or DNN) [4, 5, 6], to generate frame-level soft alignments required in the i-vector estimation process. DNN Bottleneck and Tandem features have also been explored for speaker recognition [7, 8], and more recently successfully used in some state-of-the-art systems [9, 10]. The i-vectors are typically post-processed through a linear discriminant analysis (LDA) [11] stage to generate dimensionality reduced and channel-compensated features which can then be efficiently modeled and scored with various backends such as a probabilistic LDA (PLDA) [12, 13].

In this paper, we report on the latest advancements made in the IBM i-vector speaker recognition system [14] for conversational speech. Particularly, we describe the key compo-

nents that contribute significantly to our system performance. These components include: 1) a nearest-neighbor based discriminant analysis (NDA) approach [15] for channel compensation in i-vector space, which, unlike the commonly used Fisher LDA, is non-parametric and typically of full rank, 2) speaker- and channel-adapted features derived from feature-space maximum likelihood linear regression (fMLLR) transforms [16, 17], which are used both to train/evaluate the DNN and to compute the sufficient Baum-Welch statistics for i-vector extraction, and 3) a DNN acoustic model with a large number of output units ($\sim 10k$ senones) to compute the soft alignments (i.e., the posteriors). To quantify the contribution of these components, we evaluate our system in the context of speaker verification experiments using speech material from the NIST 2010 speaker recognition evaluation (SRE) which includes 9 extended core tasks as well as a *10sec–10sec* condition. Motivated by the relatively low speaker recognition error rates achieved by our system (e.g., 0.59% EER on C5 in SRE 2010), we also conduct an error analysis of low scoring target trials to gain insights regarding the nature of the issues associated with the remaining system errors on C5.

2. System Overview

In the following subsections, we briefly describe the major components of our speaker recognition system. A schematic block diagram of the system is depicted in Fig. 1.

2.1. DNN i-vector extraction

The i-vector representation is based on the total variability modeling concept which assumes that speaker- and channel-dependent variabilities reside in the same low-dimensional subspace [2]. In order to learn the bases for the total variability subspace, one needs to first compute the Baum-Welch statistics which are defined as,

$$N_g(s) = \sum_t \gamma_{tg}(s), \quad (1)$$

$$\mathbf{F}_g(s) = \sum_t \gamma_{tg}(s) \mathbf{O}_t(s), \quad (2)$$

where $N_g(s)$ and $\mathbf{F}_g(s)$ denote the zeroth- and first-order statistics for speech session s , respectively, with $\gamma_{tg}(s)$ being the posterior probability of the mixture component g given the observation vector $\mathbf{O}_t(s)$ at time frame t .

The observation vector $\mathbf{O}_t(s)$ can be either the conventional raw acoustic features such as MFCCs or their speaker- and channel-adapted versions which is computed through a per recording fMLLR transform [17, 16] typically obtained with a GMM-HMM system. Note from Fig. 1 that the same fMLLR transformed features can be used to train/evaluate the DNN




Figure 1: Block diagram of the IBM speaker recognition system with fMLLR speaker- and channel-adapted features, DNN posterior based i-vectors, and NDA dimensionality reduction.

as well as compute the sufficient Baum-Welch statistics for i-vector extraction.

Traditionally, the frame-level soft alignments, $\gamma_{tg}(s)$, in (1) and (2) are computed with a GMM acoustic model trained in an unsupervised fashion (i.e., with no phonetic labels). However, in [4], a supervised GMM-HMM acoustic model (derived from a speech recognition system) was utilized to estimate the GMM-UBM hyperparameters for speaker recognition, assuming that class-conditional distributions for the various phonetic classes are Gaussian. More recently, inspired by the success of DNN acoustic models in automatic speech recognition (ASR) field, [5] proposed the use of DNN senone (context-dependent triphones) posteriors for computing the soft alignments, $\gamma_{tg}(s)$, which resulted in remarkable reductions in speaker recognition error rates. Motivated by these results, in this paper, we explore DNN i-vectors extracted with a very large number of senones for speaker recognition, and compare their effectiveness against GMM i-vectors on this task.

2.2. Nearest-neighbor discriminant analysis (NDA)

As noted previously, state-of-the-art speaker recognition systems employ LDA for intersession variability compensation in the i-vector space. However, there are some limitations associated with the parametric LDA where the underlying distribution of classes is assumed to be Gaussian and unimodal. Nevertheless, it is well known in the speaker recognition community that the actual distribution of i-vectors may not necessarily be Gaussian [18], particularly in the presence of noise and channel distortions [15, 19]. In addition, for the NIST SRE scenarios, speech recordings come from various sources (sometimes out-of-domain), therefore unimodality of the distributions cannot be guaranteed.

In order to alleviate some of the limitations identified for LDA, a nonparametric nearest-neighbor based discriminant analysis techniques was proposed in [20], and recently evaluated for both speaker and language recognition tasks on high-frequency (HF) radio channel degraded data [15, 19], where it compared favorably to LDA. In NDA, the expected values that represent the global information about each class are replaced with local sample averages computed based on the k -NN of individual samples. More specifically, in the NDA approach, the between-class scatter matrix is defined as,

$$\tilde{\mathbf{S}}_b = \sum_{i=1}^C \sum_{\substack{j=1 \\ j \neq i}}^C \sum_{l=1}^{N_i} w_l^{ij} (\mathbf{x}_l^i - \mathcal{M}_l^{ij}) (\mathbf{x}_l^i - \mathcal{M}_l^{ij})^T, \quad (3)$$

where \mathbf{x}_l^i denotes the l^{th} sample from class i , and \mathcal{M}_l^{ij} is the local mean of k -NN samples for \mathbf{x}_l^i from class j . Here, C and N_i denote the number of classes and the number of samples in class i , respectively. The weighting function w_l^{ij} is introduced in (3) to deemphasize the local gradients that are large in magnitude to mitigate their influence on the scatter matrix. The weight pa-

rameter, w_l^{ij} , is larger for samples near the classification boundary, while it drops off to approximately 0 for samples that are far from the boundary. In this study, the within-class scatter matrix, \mathbf{S}_w , is computed similarly as in LDA. The NDA transform is then formed by calculating the eigenvectors of $\mathbf{S}_w^{-1} \tilde{\mathbf{S}}_b$.

Three important observations can be made from a careful examination of the nonparametric between-class scatter matrix in (3). First, notice that the mean vector, \mathcal{M}_l^{ij} , is calculated locally (as opposed to globally in LDA), thereby resulting in more robust transforms with the NDA, particularly for scenarios where unimodality of the class conditional distributions cannot be guaranteed. Second, because all the samples are taken into account for the calculation of the nonparametric between-class scatter matrix (as opposed to only the class centroids in LDA), $\tilde{\mathbf{S}}_b$, and hence the NDA projection, is generally of full rank. Finally, compared to LDA, NDA is more effective in preserving the complex structure (i.e., local and boundary structure) within and across different classes because LDA only uses the global gradient obtained with the centroids of the classes to measure the between-class scatter. On the other hand, NDA uses the local gradients that are emphasized along the boundary through the weighting function, w_l^{ij} .

3. Experiments

This section provides a description of our experimental setup including speech data, the ASR system configuration, and the speaker recognition (SR) system configuration.

3.1. Data

We conduct the core of our speaker recognition experiments using conversational telephone and microphone (phone call and interview) speech material extracted from datasets released through the LDC for the NIST 2004-2010 SRE [21, 22], as well as Switchboard Cellular (SWBCELL) Parts I and II and Switchboard2 (SWB2) Phase II and Phase III corpora. These datasets contain speech spoken in U.S. English (the non English portion was filtered out) from a large number of male and female speakers with multiple sessions per speaker. The NIST SRE 2010 data is held out for evaluations, while the remaining data are used to train the system hyper-parameters (i.e., the i-vector extractor, LDA/NDA, and PLDA). In our experiments, we consider all 9 extended core tasks (C1–C9) along with the *10sec-10sec* condition (C10s) in the NIST SRE 2010 that involve telephone and microphone trials from both male and female speakers (consult [23] for a more detailed description of the tasks).

3.2. DNN system configuration

A DNN model, with 7 fully connected hidden layers with 2048 units per layer except for the bottleneck layer that has 512 units, is discriminatively trained using the standard error

back-propagation and cross-entropy objective function to estimate posterior probabilities of 10,000 senones (HMM triphone states). The training is accomplished using the IBM Attila toolkit [24] on 600 hours of conversational telephone speech (CTS) data from the Fisher corpus [25] with a 9-frame context of 40-dimensional speaker-adapted feature vectors obtained through per recording fMLLR transforms [16, 17]. The fMLLR transforms are generated for each recording with decoding alignments obtained from a GMM-HMM acoustic model (see [26, 27] for more details).

3.3. SR system configuration

For speech parameterization (other than the fMLLR based features), we extract 13-dimensional MFCCs (including c_0) from 25 ms frames every 10 ms using a 24-channel mel filterbank. The first and second temporal cepstral derivatives are also computed over a 5-frame window and appended to the static features to capture the dynamic pattern of speech over time. This results in 39-dimensional feature vectors. For non-speech frame dropping, we employ an unsupervised speech activity detector (SAD) based on voicing energy features [28]. After dropping non-speech, short-time cepstral mean subtraction (CMS) is applied.

In this paper, a 500-dimensional total variability subspace is learned and used to extract i-vectors from the recordings. To learn the i-vector extractor, out of a total of 60,178 recordings available from 1884 male and 2601 female speakers, we select 48,325 recordings from NIST SRE 2004-2008, SWB-CELL, and SWB2 corpora. The zeroth and first order Baum-Welch statistics are computed for each recording using soft alignments obtained from either a gender-independent 2048-component GMM-UBM with diagonal covariance matrices, or the DNN acoustic model with 2k, 4k, and 10k senones. The GMM-UBM is trained using 21,207 recordings selected from the NIST SRE 2004-2006, SWBCELL, and SWB2 corpora.

After extracting 500-dimensional i-vectors, we either use LDA or NDA for inter-session variability compensation by reducing the dimensionality to 250. In order to train the NDA, we employ a one-versus-rest strategy to compute the inter-speaker scatter matrix in (3). This provides flexibility on the number of nearest neighbors used for computing the local means. A cosine similarity metric (as opposed to Euclidean) is used to find the k -nearest neighbors for each sample. The dimensionality reduced i-vectors are then centered, whitened, and unit-length normalized. For scoring, a Gaussian PLDA model with a full covariance residual noise term [12, 13] is learned using the i-vectors extracted from all 60,178 speech segments (1884 male and 2601 female speakers) as noted previously. The Eigenvoice subspace in the PLDA model is assumed full-rank.

4. Results and Discussion

4.1. Experimental results

In this section, we summarize the results obtained with the experimental setup presented in Section 3. In the first experiment, we evaluated the effectiveness of NDA versus LDA for inter-session variability compensation in the i-vector space. The outcome of this experiment on C5 is presented in Table 1, in terms of the equal error rate (EER), minimum detection cost function with the NIST SRE 2008 [29] and 2010 [23] definitions (minDCF08 and minDCF10). It can be seen from the table that the systems with NDA consistently provide better speaker recognition performance across all three metrics. For

Table 1: Performance comparison of IBM speaker recognition systems with various configurations on C5, with 10k senones.

System	EER [%]	minDCF08	minDCF10
GMM-MFCC-LDA	2.11	0.113	0.440
GMM-MFCC-NDA	1.49	0.071	0.280
DNN-MFCC-LDA	0.91	0.045	0.172
DNN-MFCC-NDA	0.68	0.034	0.140
DNN-fMLLR-LDA	0.75	0.032	0.125
DNN-fMLLR-NDA	0.59	0.025	0.095

the GMM based system, a relative improvement of 30% in EER is achieved with NDA over LDA, while for the DNN based systems with MFCCs and fMLLR features relative improvements of 25% and 21% are obtained, respectively. As we discussed before, this is due to the nonparametric nature of the scatter matrices in NDA that makes no assumptions regarding the underlying class-conditional distributions (i.e., Gaussianity and unimodality). Another important observation that can be made from Table 1 is that, irrespective of the dimensionality reduction algorithm used, the systems with fMLLR features outperform the MFCC based systems. This is attributed to the ability of the fMLLR transforms in reducing the speaker and channel variabilities in the acoustic feature space. Finally, consistent with the results reported in recent studies [5, 6], the DNN based systems outperform the GMM based systems by significant margin (e.g., resulting in a relative improvement of more than 54% in terms of the EER with MFCCs and NDA). In [14], we also investigated the impact of the senone set size (2k, 4k, and 10k) on speaker recognition performance, where we observed that the larger the number of senones, the better the performance. It is worth noting that increasing the number of components in the unsupervised GMM acoustic model (with diagonal covariance matrices) for speaker recognition did not seem to result in much performance gains in the recent studies [5, 6].

In the next set of experiments, we investigated the impact of in-domain and out-of-domain training for LDA versus NDA. This was accomplished by first splitting the training data into in- and out-of-domain parts, and then retraining the LDA and NDA models on these parts. Similar to the setup used in the domain adaptation challenge (e.g., see [9, 30]), the in-domain part contains CTS data that come from the *Mixer* collection, i.e., SRE2004–2008 corpora, while the out-of-domain part contains data from the older SWB2 and SWBCELL corpora. Table 2 shows the results from these experiments on C5 which are obtained with i-vectors computed using 10k DNN senones and fMLLR features. It is observed that while NDA outperforms LDA with out-of-domain training, it offers minimal improvement (at least in terms of EER) when only in-domain training data is used. However, with pooled in-domain and out-of-domain data (which is no longer unimodal), a significant improvement in performance is obtained with NDA over LDA

Table 2: Comparison of LDA vs NDA for in-domain and out-of-domain training on C5, with fMLLR features and 10k senones.

System	Domain	EER [%]	minDCF08	minDCF10
DNN-LDA	out	1.80	0.085	0.256
	in	0.78	0.034	0.149
	in+out	0.75	0.032	0.125
DNN-NDA	out	1.55	0.071	0.226
	in	0.75	0.033	0.119
	in+out	0.59	0.025	0.095

Table 3: Performance comparison of IBM speaker recognition systems with various configurations on C1–C10s (excluding C5).

System	EER [%] (minDCF10)								
	C1	C2	C3	C4	C6	C7	C8	C9	C10s
GMM-MFCC-NDA	1.2 (0.22)	1.8 (0.32)	1.8 (0.30)	1.1 (0.25)	2.9 (0.55)	3.3 (0.59)	1.1 (0.25)	0.7 (0.11)	11.7 (0.99)
DNN-MFCC-NDA	0.8 (0.11)	1.3 (0.16)	0.9 (0.13)	0.6 (0.11)	1.5 (0.34)	1.1 (0.32)	0.6 (0.14)	0.4 (0.05)	8.3 (0.94)
DNN-fMLLR-NDA	0.9 (0.13)	1.3 (0.17)	0.9 (0.10)	0.7 (0.08)	1.0 (0.24)	1.4 (0.31)	0.4 (0.10)	0.3 (0.05)	12.3 (0.99)

(i.e., 21% relative in terms of EER). This improvement may be attributed to the robustness of NDA to multimodality in the data, as discussed in Section 2.2.

For completeness, we also evaluated the performance of our speaker recognition system on extended microphone and telephone tasks, under normal and high/low vocal effort, (C1–C9) as well as the *10sec–10sec* condition in the NIST SRE 2010. The results are provided in Table 3 for both the GMM and DNN based systems. It is clear that the DNN based systems, with either MFCCs or fMLLR features, perform significantly better than the GMM based system. Additionally, the DNN based system trained with raw MFCCs tend to perform better than the fMLLR based system on interview microphone conditions (C1–C2). We speculate that this is because the fMLLR transforms, which are obtained using GMM-HMMs trained only on telephony data, are unable to effectively reduce the variability due to channel mismatch on microphone recordings. For other tasks, overall, MFCC and fMLLR based systems perform similarly except on the C10s where, among other challenges, the VTL normalizations and the fMLLR transforms are also adversely impacted by the short duration of recordings (i.e., 10 s).

4.2. Error analysis

In consideration of the obtained results, for example 0.59% EER on the extended condition 5, we attempt to gain an understanding of some of the underlying issues in the system. In particular, we would like to examine the characteristics of the recordings contributing to some of the lowest scoring target trials. Towards this, we analyzed 265 recordings relating to the 200 lowest scoring target trials and attempted to characterize them with various potentially problematic properties. It was found that, while many of the recordings are audibly acceptable, there are 76 with co-channel speech (either through cross-channel feedback or from background competing speakers), 51 with background noise, music, breath, snuffle, and fidgeting sounds, 33 with very sparse speech activity, 8 with cocktail party noise, and 2 exhibiting signal clipping effects.

Our objective is to improve the target trial scores relating to the problematic files. In particular, we focused our attention on three categories: (i) cocktail party effect, (ii) general noise, and (iii) co-channel speaker. These categories were chosen since they may benefit from manually changing the SAD information to more appropriately reflect the speech portions of the speaker of interest (i.e., the target speaker).

Figure 2 shows the result of the manual modification of the SAD labels by plotting the trial scores (both target and non-target) that related to a selection of 18 modified SAD recordings. The x-axis represents the scores from trials with the original recordings while the y-axis shows the scores for the modified recordings. Here, there were three target trials that greatly improved in score. An interesting point is that not only did some of the target scores significantly improve, but many of the non-target scores also decreased. One view is that the presence of a co-channel speaker has the effect of bringing the result-

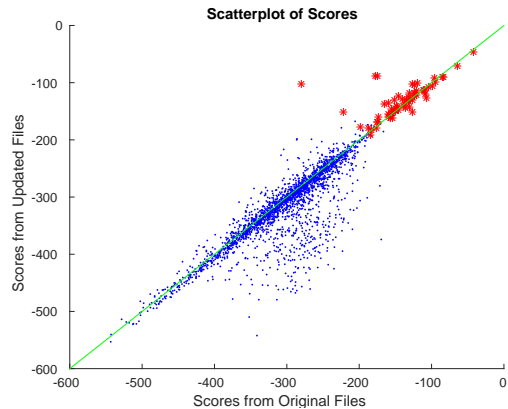


Figure 2: Scatter plot of the scores of repaired trials versus the original trials. The asterisks (*) represent the target trials while the dots (·) represent the non-target trials.

ing i-vector closer to the general population of speaker i-vectors (on average). This raises an important observation that not only should a speech detection component be utilized but also a component that only focuses on the speaker of interest (or the homogeneity of the class). The large movement in both target and non-target trials was basically described by 8 of the 18 recordings. These consisted of intermittent noise conditions, cocktail party noise, and co-channel (cross-channel) speech where the interference was significant or comparable to the level of the target speech. Interestingly, by *manually* modifying the SAD for 18 recordings the EER is reduced from 0.59% to 0.56%.

Future work will involve the automation of such manual processes (perhaps through the use of speaker diarization and audio enhancement (e.g. see [31]), as well as expand the types of issues addressed such as short-duration modeling [32] and signal quantization issues.

5. Conclusions

In this paper, we presented the recent improvements made in our state-of-the-art i-vector speaker recognition system. We investigated the impact of several key components of the system on performance using extended core tasks in the NIST 2010 SRE that involved both microphone and telephone trials. Some important observations made from our experiments were as follows: 1) the NDA was found to be consistently more effective than LDA for inter-session variability compensation in i-vector based speaker recognition, 2) the fMLLR based features provided better representation than raw MFCCs for matched data conditions (i.e., telephony trials), and 3) the DNN based UBM with large number of components (i.e., 10k senones) resulted in remarkable improvements in the performance of our system. To the best of our knowledge, the results presented in this paper represent the best performances reported to date on the extended core tasks in the NIST 2010 SRE.

6. References

- [1] P. Kenny, G. Boulianne, P. Ouellet, and P. Dumouchel, "Joint factor analysis versus eigenchannels in speaker recognition," *IEEE Trans. Audio Speech Lang. Process.*, vol. 15, no. 4, pp. 1435–1447, 2007.
- [2] N. Dehak, P. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Trans. Audio Speech Lang. Process.*, vol. 19, no. 4, pp. 788–798, 2011.
- [3] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker verification using adapted Gaussian mixture models," *Digital Signal Process.*, vol. 10, pp. 19–41, 2000.
- [4] M. K. Omar and J. Pelecanos, "Training universal background models for speaker recognition," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2010)*, Brno, Czech, June 2010, pp. 52–57.
- [5] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, "A novel scheme for speaker recognition using a phonetically-aware deep neural network," in *Proc. IEEE ICASSP*, Florence, Italy, May 2014, pp. 1695–1699.
- [6] D. Snyder, D. Garcia-Romero, and D. Povey, "Time delay deep neural network-based universal background models for speaker recognition," in *Proc. IEEE ASRU*, Scottsdale, AZ, December 2015, pp. 92–97.
- [7] L. P. Heck, Y. Konig, M. K. Sönmez, and M. Weintraub, "Robustness to telephone handset distortion in speaker recognition by discriminative feature design," *Speech Commun.*, vol. 31, no. 2-3, pp. 181–192, 2000.
- [8] S. Yaman, J. W. Pelecanos, and R. Sarikaya, "Bottleneck features for speaker recognition," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2012)*, Singapore, Singapore, June 2012, pp. 105–108.
- [9] F. Richardson, D. A. Reynolds, and N. Dehak, "A unified deep neural network for speaker and language recognition," in *Proc. INTERSPEECH*, Dresden, Germany, September 2015, pp. 1146–1150.
- [10] P. Matějka *et al.*, "Analysis of DNN approaches to speaker identification," in *Proc. IEEE ICASSP*, Shanghai, China, March 2016, pp. 5100–5104.
- [11] K. Fukunaga, *Introduction to Statistical Pattern Recognition*, 2nd ed. New York: Academic press, 1990.
- [12] S. J. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *Proc. IEEE ICCV*, Rio De Janeiro, October 2007, pp. 1–8.
- [13] D. Garcia-Romero and C. Y. Espy-Wilson, "Analysis of i-vector length normalization in speaker recognition systems," in *Proc. INTERSPEECH*, Florence, Italy, August 2011, pp. 249–252.
- [14] S. O. Sadjadi, S. Ganapathy, and J. Pelecanos, "The IBM 2016 speaker recognition system," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2016)*, Bilbao, Spain, June 2016.
- [15] S. O. Sadjadi, J. W. Pelecanos, and W. Zhu, "Nearest neighbor discriminant analysis for robust speaker recognition," in *Proc. INTERSPEECH*, Singapore, Singapore, September 2014, pp. 1860–1864.
- [16] V. Digalakis, D. Rtischev, and L. Neumeyer, "Speaker adaptation using constrained estimation of Gaussian mixtures," *IEEE Trans. Speech Audio Process.*, vol. 3, no. 5, pp. 357–366, September 1995.
- [17] M. J. F. Gales, "Maximum likelihood linear transformations for HMM-based speech recognition," *Comput. Speech Lang.*, vol. 12, no. 2, pp. 75–98, 1998.
- [18] P. Kenny, "Bayesian speaker verification with heavy tailed priors," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2010)*, Brno, Czech, June 2010.
- [19] S. O. Sadjadi, S. Ganapathy, and J. W. Pelecanos, "Nearest neighbor discriminant analysis for language recognition," in *Proc. IEEE ICASSP*, Brisbane, Australia, April 2015, pp. 4205–4209.
- [20] K. Fukunaga and J. Mantock, "Nonparametric discriminant analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 5, no. 6, pp. 671–678, 1983.
- [21] C. Cieri, L. Corson, D. Graff, and K. Walker, "Resources for new research directions in speaker recognition: The Mixer 3, 4 and 5 corpora," in *Proc. INTERSPEECH*, Antwerp, Belgium, August 2007, pp. 950–953.
- [22] L. Brandschain, D. Graff, C. Cieri, K. Walker, C. Caruso, and A. Neely, "Mixer 6," in *Proc. LREC*, Valletta, Malta, May 2010.
- [23] NIST, "The NIST Year 2010 Speaker Recognition Evaluation Plan," http://www.nist.gov/itl/iad/mig/upload/NIST.SRE10_evalplan-r6.pdf, 2010.
- [24] H. Soltau, G. Saon, and B. Kingsbury, "The IBM Attila speech recognition toolkit," in *Proc. IEEE SLT*, Berkeley, CA, December 2010, pp. 97–102.
- [25] C. Cieri, D. Miller, and K. Walker, "The Fisher corpus: A resource for the next generations of speech-to-text," in *Proc. LREC*, Lisbon, Portugal, May 2004.
- [26] S. Ganapathy, S. Thomas, D. Dimitriadis, and S. Rennie, "Investigating factor analysis features for deep neural networks in noisy speech recognition," in *Proc. INTERSPEECH*, Dresden, Germany, September 2015, pp. 1898–1902.
- [27] G. Saon, H. K. Kuo, S. Rennie, and M. Picheny, "The IBM 2015 English conversational telephone speech recognition system," in *Proc. INTERSPEECH*, Dresden, Germany, September 2015, pp. 3140–3144.
- [28] S. O. Sadjadi and J. H. L. Hansen, "Unsupervised speech activity detection using voicing measures and perceptual spectral flux," *IEEE Signal Process. Lett.*, vol. 20, no. 3, pp. 197–200, 2013.
- [29] NIST, "The NIST Year 2008 Speaker Recognition Evaluation Plan," http://www.itl.nist.gov/iad/tests/sre/2008/sre08_evalplan_release4.pdf, 2008.
- [30] D. Garcia-Romero and A. McCree, "Supervised domain adaptation for i-vector based speaker recognition," in *Proc. IEEE ICASSP*, Florence, Italy, May 2014, pp. 4047–4051.
- [31] O. Plocht, L. Burget, H. Aronowitz, and P. Matějka, "Audio enhancing with DNN autoencoder for speaker recognition," in *Proc. IEEE ICASSP*, Shanghai, China, March 2016, pp. 5090–5094.
- [32] P. Kenny, T. Stafylakis, P. Ouellet, M. J. Alam, and P. Dumouchel, "PLDA for speaker verification with utterances of arbitrary duration," in *Proc. IEEE ICASSP*, Vancouver, BC, May 2013, pp. 7649–7653.