

Channel Factors Compensation in Model and Feature Domain for Speaker Recognition

Original

Channel Factors Compensation in Model and Feature Domain for Speaker Recognition / Vair, C; Colibro, D; Castaldo, Fabio; Dalmaso, E; Laface, Pietro. - I:(2006). (Intervento presentato al convegno IEEE Conference ODYSSEY 2006 tenutosi a Portorico nel Jun 2006).

Availability:

This version is available at: 11583/1413123 since:

Publisher:

IEEE

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Channel Factors Compensation in Model and Feature Domain for Speaker Recognition

Claudio Vair, Daniele Colibro*, Fabio Castaldo^, Emanuele Dalmasso^, Pietro Laface^*

Loquendo, Torino, Italy*

{Claudio.Vair, Daniele.Colibro}@loquendo.com

Politecnico di Torino, Italy^

{Fabio.Castaldo, Emanuele.Dalmasso, Pietro.Laface}@polito.it

Abstract

The variability of the channel and environment is one of the most important factors affecting the performance of text-independent speaker verification systems.

The best techniques for channel compensation are model based. Most of them have been proposed for Gaussian Mixture Models, while in the feature domain typically blind channel compensation is performed.

The aim of this work is to explore techniques that allow more accurate channel compensation in the domain of the features. Compensating the features rather than the models has the advantage that the transformed parameters can be used with models of different nature and complexity, and also for different tasks.

In this paper we evaluate the effects of the compensation of the channel variability obtained by means of the channel factors approach. In particular, we compare channel variability modeling in the usual Gaussian Mixture model domain, and our proposed feature domain compensation technique. We show that the two approaches lead to similar results on the NIST 2005 Speaker Recognition Evaluation data.

Moreover, the quality of the transformed features is also assessed in the Support Vector Machines framework for speaker recognition on the same data, and in preliminary experiments on Language Identification.

1. Introduction

It is well known that one of the main causes of relevant performance degradations in automatic speech recognition systems is the acoustic mismatch that occurs between training and test environments. In speaker recognition, errors are due not only to the similarity among voiceprints of different speakers, but also to the intrinsic variability of different utterances of the same speaker. Moreover, performance is heavily affected when a model, trained in a set of conditions, is used to test speaker data collected from different microphones, channels, and environments. In this paper we will refer to all these mismatching conditions as intersession variability or simply as channel variability.

Several proposals have been made to contrast these effects by means of feature transformations [1,2]. Since some feature based transformations, such as feature warping [1], do not rely on a specific model, they can be used as an additional front-end processing step for any recognition system that takes advantage of this compensation technique. However, this blind feature normalization does not exploit a priori knowledge of the condition as in [2], or other information that

can be obtained by a more detailed analysis of the variations of the speaker parameters in the acoustic space.

Feature mapping [2] uses the a priori information of a set of models trained in known conditions to map the feature vectors toward a channel independent feature space. The drawback of this approach is that it requires labeled training data that identify the conditions that one wants to compensate.

Thus, model-based techniques have been recently proposed that are able to compensate speaker and channel variations without requiring the explicit identification and labeling of different conditions. These techniques share a common background: modeling the variability of speaker utterances constraining them to a low dimensional space. This approach has proved to be effective for speaker adaptation both in speech [3] and speaker recognition [4,5], and for channel compensation [6-9]. All these methods are generative and use MAP adapted Gaussian Mixture Models (GMM) [10] for modeling the speakers.

Discriminative methods for speaker recognition are attractive because they are trained to minimize the errors. Discriminative models based on Support Vector Machines (SVMs) are typically smaller than the generative ones, trained with the same amount of data. SVM based systems require less computational resources both in training and testing. SVM have been shown to produce comparable results with respect to the state of the art GMM based systems and to improve the performance of a speaker recognition system that combines the scores produced by the generative and discriminative modeling approaches [11,12].

In [13] an original approach to channel compensation in the model space of SVM has been proposed. It evaluates the projection of the expanded vectors in a subspace that remove the dimensions that carry information not related to the speaker but only to the channel and the environment. This approach is interesting because it does not suffer the main drawback of the feature mapping technique, which relies on the knowledge and labeling of a discrete combination of conditions referring to the handset type, the transmission channel, the environment, and so forth.

In this work we mainly refer to [9] for intersession compensation in the model domain. We present our modifications to this method comparing the obtained results on the NIST 2005 Speaker Recognition Evaluation data (SRE-05) [14] and showing that our approach leads to similar results with a reduced computation cost.

The main objective of this work, however, has been to find a solution allowing compensating the observation features rather than the Gaussian means.

Compensating features rather than models has the advantage that the transformed parameters can be used as observation vectors for classifiers of different nature and complexity, and also for different tasks such as language or speech recognition.

The paper is organized as follows: the model based channel factors adaptation approach and our modifications are described in Section 2, together with our proposed channel factors feature adaptation technique. Section 3 summarizes the parameters of our baseline GMM systems. The experimental results, including the use of the compensated feature with a SVM classifier, are presented in Section 4. Some concluding remarks are given in Section 5.

2. Channel factors adaptation

Gaussian Mixture Models (GMMs) used in combination with Maximum A Posteriori (MAP) adaptation [10] represent the core technology of most of the state-of-the-art text-independent speaker recognition systems. In these systems the speaker models are derived from a common GMM root model, the so called Universal Background Model (UBM), by means of MAP adaptation. Usually, only mean vector adaptation is performed during model training. A speaker is, thus, represented by the set of the adapted mean vectors of all the Gaussians of the UBM.

A supervector that includes all the speaker specific parameters can be obtained simply appending the adapted mean value of all the Gaussians in a single stream. The same can be done for the UBM, obtaining the UBM supervector. The speaker model can be seen as a point in a high dimensional space, whose coordinates are the supervectors parameters. When some kind of mismatch, like the use of different microphones or communication channels, affects the input speech, all the speaker supervector parameters are possibly modified.

The idea behind the methods proposed in this paper is that the distortions in the large supervector space can be summarized by a small number of parameters in a lower dimensional subspace: the *channel factors* [15].

2.1. Model-domain adaptation

Channel factors adaptation for an utterance i and a supervector k is performed, in the supervector model space, as follows:

$$\boldsymbol{\mu}^{(i,k)} = \boldsymbol{\mu}^{(k)} + \mathbf{U}\mathbf{x}^{(i,k)} \quad (1)$$

where $\boldsymbol{\mu}^{(i,k)}$ and $\boldsymbol{\mu}^{(k)}$ are the adapted and the original supervector of GMM k respectively. \mathbf{U} is a low rank matrix projecting the channel factors subspace in the supervector domain. The N -dimensional vector $\mathbf{x}^{(i,k)}$ holds the channel factors for the current utterance i and GMM k .

The approach that we use is similar to the formulation in [9] with the difference that we do not perform channel compensation during training but we apply (1) only at testing time. The $\boldsymbol{\mu}^{(k)}$ supervectors are obtained by the classical MAP speaker adaptation, without any additional computation. The verification score is obtained computing the log-likelihood ratio of the test utterance using compensated speaker and UBM means. This gives us good performance improvements even without normalization of raw scores.

Since the vector $\mathbf{x}^{(i,k)}$ should account for the distortions produced in the supervector space by the intersession variability, we would expect that $\mathbf{x}^{(i,k)}$ depends on the utterance i , but only weakly on the speaker model k .

To verify this hypothesis we run several tests estimating the channel factors of \mathbf{x} using the UBM, i.e. dropping the dependence on the GMM k . This is equivalent to set:

$$\mathbf{x}^{(i,k)} = \mathbf{x}^{(i)} \quad \forall k \quad (2)$$

and to apply the normalization:

$$\boldsymbol{\mu}^{(i,k)} = \boldsymbol{\mu}^{(k)} + \mathbf{U}\mathbf{x}^{(i)} \quad (3)$$

for all the models k that must be scored against utterance i . As reported in Session 4.3.1, the obtained results were almost equivalent to the ones obtained with the speaker-model dependent estimation of (1), but with relevant saving of computation time, in particular when T-Norm score normalization [16] is applied.

2.1.1. Training of the channel factors subspace

The channel factors subspace, modeled by the low rank matrix \mathbf{U} , is assumed to represent the distortion due to the intersession variability. This distortion can be estimated analyzing how the models of the same speaker are affected, when trained with utterances collected from different channels or conditions. Thus, a database has been set up including a large number of speakers, each one with multiple recordings collected from different calls and channels.

An EM training algorithm has been used to compute the \mathbf{U} matrix [17]. The number of columns N of the matrix \mathbf{U} defines the channel subspace dimension and it is typically less than 50.

2.1.2. Estimation of the channel factors parameters

To perform channel adaptation through equation (1) or (3), the channel factors vector \mathbf{x} must be estimated for each test utterance

A maximum likelihood solution to this problem has been proposed in [3] for speaker adaptation. For speaker verification, a technique called Probabilistic Subspace Adaptation (PSA), which uses MAP estimation of \mathbf{x} has been presented in [5].

In our experiments, we perform a single iteration of the PSA estimation, obtaining one vector $\mathbf{x}^{(i,k)}$ for each tuple {test utterance i , model k } in equation (1), or a single vector $\mathbf{x}^{(i)}$ for a test utterance i in equation (3).

2.2. Feature-domain adaptation

Channel factors adaptation in the model domain has proved to improve the performance of GMM speaker recognition systems. However, it is not readily applicable to GMMs with a different number of Gaussians, to other types of classifiers, like SVM or Artificial Neural Networks (ANNs), or to other tasks requiring more complex models, for example HMMs for speech recognition.

The feature domain method that we propose allows exploiting the benefits of the channel factors adaptation, mapping the compensation supervector on the acoustic features.

We rely on the hypotheses that led to equation (3): we assume that the acoustic space distortion, characterized by the vector $\mathbf{x}^{(i)}$, can be estimated using the UBM rather than the speaker dependent model GMM k . Neglecting, for the sake of conciseness, the model index k , we rewrite (3) for each Gaussian component m of the supervector as:

$$\boldsymbol{\mu}_m^{(i)} = \boldsymbol{\mu}_m + \mathbf{U}_m \mathbf{x}^{(i)} \quad \forall m \quad (4)$$

where $\boldsymbol{\mu}_m^{(i)}$, $\boldsymbol{\mu}_m$ and \mathbf{U}_m refers to the m -th Gaussian of the GMM. The number of rows of the mean vectors and of the subspace matrix \mathbf{U}_m , is equal to the dimension of the input feature vector.

The adaptation of the feature vector at time frame t , $\mathbf{O}(t)$, is obtained by subtracting to the observation feature a weighted sum of the channel compensation offset values:

$$\hat{\mathbf{O}}^{(i)}(t) = \mathbf{O}^{(i)}(t) - \sum_m \gamma_m(t) \mathbf{U}_m \mathbf{x}^{(i)} \quad (5)$$

where $\gamma_m(t)$ is the Gaussian occupation probability, and $\mathbf{U}_m \mathbf{x}^{(i)}$ is the channel compensation offset related to the m -th Gaussian of the UBM model. In the actual implementation, the right side summation of (5) is limited, for the sake of efficiency, to the first best contributors only.

Equation (5) allows obtaining adapted feature vectors suitable as front-end parameters to any further classification process. The quality of the transformed features has been verified using them for different classifiers and tasks.

3. Systems description

Two GMM systems have been trained and tested in this work: a Phonetic GMM (PGMM), and a classical GMM.

3.1. Phonetic GMM system

We describe first the speaker recognition system that has been used to produce the results submitted to the NIST SRE-2005 evaluation [14].

The system decodes the speaker utterance, both in enrolment and in verification, producing phonetic labeled segments. The decoder is a hybrid HMM-ANN model trained to recognize 11 language independent phone classes. Each phone class is modeled by a three state left-to-right automaton with self-loops. The ANN is a Multilayer Perceptron that estimates the posterior probability of each phone class state, given an acoustic feature vector. The ANN has been trained using 20 hours of speech of 10 different languages using corpora not specifically collected for speaker recognition evaluations.

The UBM and the voiceprints consist of a set of phonetic GMMs: each phone class state has associated its GMM. The maximum number of (diagonal covariance) Gaussians per mixture per state is 64, for a total of approximately 2000 Gaussians.

A gender – and nearly language – independent UBM has been trained on the same data that were used for training the ANN model.

In enrolment, the labels and the boundaries of the phonetic segments are used for MAP adaptation of the parameters of the phone class-dependent GMMs.

In recognition, the phonetically labeled audio segments are scored against their corresponding GMMs. Thus, the likelihood of a given observation vector is computed selecting

the GMM corresponding to the phone class decoded at that time frame. The PGMM classifier is currently the core technology of the Loquendo Free Speech Identification (LFSI) system.

The system uses 19 Mel Frequency Cepstral Coefficients (MFCC). Feature warping to a Gaussian distribution is then performed, for each static parameter stream, on a 3 sec sliding window excluding silence frames [1]. 36 parameters per frame are obtained discarding the C0 cepstral parameters, and computing the usual delta parameters on a symmetric 5 frame window.

3.2. GMM system

A classical GMM have been used in this work for the development of the channel factors compensation approach. This system can be considered a special case of the Phonetic GMM system where the number of classes is just one. The procedures for the adaptation the speaker models, and for scoring the test utterances are the same. The GMM system is characterized by a reduced set of mixtures (512), and features (the first 12 cepstral parameters and their deltas). Again, the gender independent UBM has been trained using 20 hours of speech of 10 different languages, rather than databases specifically collected for speaker recognition.

4. Experiments

4.1. Testing protocol

The system was assessed using a common protocol in terms of evaluation data set, score normalization procedure and performance measure.

4.1.1. Evaluation data set

The speaker recognition methods described in this paper were evaluated on the NIST 2005 Speaker Recognition Evaluation data (SRE05) [14]. All tests are related to the core test condition, as defined by NIST, including all trials in the enrollment and verification lists (2771 true speaker and 28472 impostor trials). The core test condition of the NIST SRE05 is related to four wire conversations, lasting approximately five minutes. Each side of a conversation is recorded on a separate channel and the silence is not removed from the recordings.

4.1.2. Score Normalization

The evaluation has been carried out with and without score normalization. First the raw score are speaker-normalized by means of Z-norm. The Z-norm parameters for each speaker model have been evaluated using a subset of speaker samples included in the NIST SRE04 database [14]. Separate statistics have been collected for the female and male speakers, using 2 audio samples of 80 speakers for each gender.

Test dependent normalization is performed using T-norm [16]. A fixed set of impostor models have been selected among the voiceprints enrolled with data belonging to the SRE04 evaluation. The T-norm parameters for each test sample were estimated using the Z-normalized scores of the impostor voiceprints. We refer to the Z-Norm followed by T-Norm as ZT-Norm.




Figure 1: DET plots for GMM and PGMM baseline systems, with and without score normalization

System	EER	DCF
GMM raw	13.8	0.548
PGMM raw	10.8	0.468
GMM ZTNorm	10.7	0.404
PGMM ZTNorm	9.2	0.343

Table 1: EER and minimum DCF for GMM and PGMM systems, with and without score normalization

4.1.3. Performance measures

The performance of the systems proposed in this paper was evaluated in terms of Detection Error Tradeoff curves (DET), Equal Error Rate (EER) and minimum normalized Detection Cost Function (DCF) (as defined by NIST [14]).

4.2. GMM versus PGMM

Figure 1 shows DET plots of the GMM baseline systems. The figure includes the results of the standard (GMM) and the Phonetic (PGMM) Gaussian Mixture classifiers, obtained with and without ZT-Norm score normalization. The related scores, in terms of Equal Error Rate and minimum normalized DCF are given in Table 1.

The better accuracy of the PGMM can be explained by the more accurate modeling of the phonetic system compared with the standard GMM, in terms of acoustic features (36 vs. 24), number of Gaussians (~2000 vs 512) and different Gaussians selection method.

4.3. GMM with channel factors compensation

In this section we show the improvement obtained with the channel factors adaptation both in model and in feature domains, applied to GMM and PGMM systems. The channel subspace dimension has been set to 20 in these experiments.

4.3.1. UBM Channel factors

First, using the GMM system, the channel factors were computed on the UBM and kept fixed for all the speaker models verified against a given audio. Figure 2 shows the DET curves of the GMM system with and without compensation, applying the UBM channel factors both in model (MD) and in feature (FD) domain respectively.




Figure 2: DET plots for GMM system and UBM compensation, in model (MD) and feature (FD) domain

System - UBM Compensation	EER	DCF
GMM MD raw	9.48	0.348
GMM FD raw	9.16	0.357
GMM MD ZTNorm	8.07	0.280
GMM FD ZTNorm	6.80	0.241

Table 2: EER and minimum DCF for GMM system, and UBM channel factor compensation, in model (MD) and feature (FD) domain

The effectiveness of the channel factors compensation is significant both on the raw and ZT-normed scores, as can be observed looking at the EER and DCF rates reported in Table 2. Moreover, better performance is obtained by the feature domain UBM compensation. This can be probably ascribed to the fact that in feature domain the same adaptation is performed both in enrollment and in verification. In the model domain, instead, channel compensation was performed only in testing, while the models were trained using the conventional MAP adaptation, because no improvement was obtained including the channel factors compensation in training (see next subsection).

4.3.2. Speaker-Dependent Channel factors

Speaker dependent channel factors compensation was tested on GMM and PGMM systems in the model domain.

Using the GMM system, we compared the results obtained by means of standard MAP training and channel factor compensated MAP training similar to [6]. The DET plots are shown in Figure 3. Without score normalization, standard MAP training outperforms channel factors MAP, but using ZT-Norm the two techniques perform very similarly. The computation requirements of the channel factor compensated MAP don't seem to justify its use. We didn't obtain any improvement using more sophisticated training procedures even using 3 or 8 conversations sides for training.

Similar results were obtained using the PGMM system. Table 3 summarizes the scores of the GMM and PGMM systems with channel factor compensation.

The PGMM system with MAP training and channel factors compensation is only slightly better than the equivalent GMM system, loosing the gap appreciable on the uncompensated systems.




Figure 3: DET plots for GMM system and speaker dependent (SD) channel factors compensation, in model domain (MD). Training MAP (TrMAP) and channel factors compensated MAP (TrCFM)

Systems – SD compensation MD	EER	DCF
GMM TrMAP raw	8.72	0.333
GMM TrCFM raw	11.87	0.493
PGMM TrMAP raw	8.61	0.406
GMM TrMAP ZT-Norm	7.02	0.240
GMM TrCFM ZT-Norm	7.49	0.244
PGMM TrMAP ZT-Norm	6.62	0.229
GMM FD + PGMM ZT-Norm	5.94	0.202

Table 3: EER and minimum DCF for GMM and PGMM system with speaker dependent (SD), model domain compensation

The simple linear combination, with equal weights, of the GMM with the feature domain compensation (Table 2, last row) and the PGMM gives a significant improvement as shown in the last row of Table 4.

4.4. SVM channel compensation

Our work draws on the results of the generalized linear discriminant sequential (GLDS) kernel approach of [11]. However, since for computational reasons the autocorrelation matrix \mathbf{R} in [11] is usually approximated by its diagonal elements, it turns out that it is possible to feed a SVM that uses a linear inner-product kernel, with polynomial features where each component is properly normalized by its standard deviation.

The channel compensated features of (5) can be readily used as observation vectors for the SVM classifiers. In particular, the channel factors $\mathbf{x}^{(i)}$ are estimated for each test or training utterance i (including the ones related to the set of impostors).

Using $\mathbf{x}^{(i)}$, every frame of that utterance is channel compensated according to (5). A polynomial expansion of the third order is then performed, and the mean and variance of every component of all the expanded vectors are evaluated.




Figure 4: DET plots for SVM and GMM system and UBM compensation in feature domain

System	EER	DCF
1. SVM raw	9.41	0.369
2. SVM FD Comp. raw	8.79	0.318
3. SVM ZT-Norm	9.81	0.362
4. SVM FD Comp. ZT-Norm	8.65	0.299
5. 2. + PGMM (ZT-Norm)	6.08	0.214
6. 4. + PGMM (ZT-Norm)	6.18	0.211

Table 4: EER and minimum DCF for SVM system, and channel factor compensation in feature domain

The expanded vector of an utterance – variance-normalized – is the channel compensated pattern for the SVM classifiers.

The observation vectors for the SVM classifiers are the same 24 parameters of the GMM system, and their expansion up to the third order polynomial.

The gender independent impostor set required to train these discriminative models includes the utterances of 1619 speakers obtained from the train splits of the NIST SRE-2000 and SRE-2004 databases.

Figure 4 shows the DET plots of the SVM system with and without feature compensation and score normalization. EER and minimum DCF rates are shown in Table 4. The raw scores, with and without channel compensation, are better than the corresponding scores of the GMM systems. However, the GMM system outperforms SVMs using ZT-Norm. Score normalization does not give appreciable performance improvements to the SVM system.

Again, linear combination of the SVM and PGMM systems improves appreciably the performances, even when raw SVMs scores are used.

Although less precise than the GMM system using the same parameters, the advantage of using SVM classifiers is not only their reduced computational cost both in training and testing, but also their ability to produce scores that tend to be intrinsically normalized. This happens because each speaker model is trained against the same set of impostors, and both the speaker and impostor utterances are channel compensated.

4.5. Language identification

To verify the quality of the channel compensated features in a completely different task, we performed an experiment on language identification comparing the performance of a gender independent classifier, based on SVMs, using three sets of basic features: the 24 MFCC features, their channel compensated counterparts, and the shifted-delta parameters proposed in [18]. Again, the vectors were subjected to a polynomial expansion of the third order and the SVMs trained using a linear kernel.

Basic features	ERR %
1. 12 MFCC+delta	18.16
2. Channel compensated 12 MFCC+delta	9.80
3. 49 shifted-delta	7.99
4. 2. and 3. fused	5.67

Table 5: Language Identification ERRs

From the OGI 22 Languages database, 8 languages were selected among the ones appearing also in the OGI Multilanguage Telephone Speech: English, German, Hindi, Italian, Korean, Mandarin, Spanish, and Tamil.

For each language, the conversations were equally split into a train and test list. The impostor set for a given language was composed of the set of conversations of all the remaining languages. Segments of 30 seconds have been used for testing.

The results, in terms of EER percentage, are shown in Tab. 6. Comparing the first and second rows we see that the feature domain channel factors compensation halves the Equal Error Rate. It is worth noting that the features were compensated using the same transformation matrix U computed for the speaker recognition experiments. This result not only shows that the channel compensation approach in feature space can be applied to other tasks, but also that the channels subspace is fairly task and language independent.

5. Conclusions

We have shown that the feature adaptation approach proposed in this paper has the same benefits of the channel factors model domain adaptation. Moreover it can be used with other types of classifiers, like SVM or ANNs and for other tasks.

Future research will be devoted to apply this technique decoupling the model used for feature compensation from the ones used for recognition, even within the GMM framework.

The system based on SVMs is attractive from an application point of view because the produced scores are fairly well stable when there are variations in the training and test conditions. This characteristic may avoid the burdensome task of the score normalizations.

6. References

[1] J. Pelecanos, and S. Sridharan, "Feature Warping for Robust Speaker Verification", *Proc. 2001: a Speaker Odyssey*, pp. 213-218, 2001.

[2] D. Reynolds, "Channel Robust Speaker Verification via Feature Mapping", in *Proc. ICASSP 2003*, pp. II-53-6, 2003.

[3] R. Kuhn J.C. Junqua, P. Nguyen, and N. Niedzielski, "Rapid Speaker Adaptation in Eigenvoice Space", *IEEE Trans. on Speech and Audio Processing*, Vol.8, No.6, Nov. 2000, pp. 695-707.

[4] O. Thyes, R. Kuhn, P. Nguyen, and J.C. Junqua, "Speaker Identification and Verification Using Eigenvoices", *Proc. ICSLP2000*, pp. 242-245, 2000.

[5] S. Lucey, and T. Chen, "Improved Speaker Verification Through Probabilistic Subspace Adaptation", *Proc. EUROSPEECH-2003*, pp. 2021-2024, 2003.

[6] P. Kenny, M. Mihoubi, and P. Dumouchel, "New MAP Estimators for Speaker Recognition", *Proc. EUROSPEECH-2003*, pp. 2964-2967, 2003.

[7] P. Kenny, G. Boulianne, and P. Dumouchel, "Eigenvoice Modeling with Sparse Training Data", *IEEE Trans. on Speech and Audio Processing*, Vol.13, No.3, May. 2005, pp. 345-354.

[8] N. Brümmner, "NIST SRE 2004 Evaluation Workshop", Toledo, Spain, 2004.

[9] R. Vogt, B. Baker and S. Sridharan, "Modelling Session Variability in Text-independent Speaker Verification", *Proc. INTERSPEECH-2005*, pp. 3117-3120, 2005.

[10] D. A. Reynolds, T. F. Quatieri and R. B. Dunn, "Speaker Verification Using Adapted Gaussian Mixture Models", *Digital Signal Processing*, Vol. 10, pp. 19-41, 2000.

[11] W. M. Campbell, "Generalized Linear Discriminant Sequence Kernels for Speaker Recognition", *Proc. ICASSP, 2002*, pp. I-164-167, 2002.

[12] S. Kajarekar, "Four Weightings and a Fusion: a Cepstral-SVM System for Speaker Recognition", *ASRU 2005*.

[13] A. Solomonoff, W.M. Campbell and I. Boardman, "Advances In Channel Compensation For SVM Speaker Recognition," *Proc. ICASSP 2005*, pp. I-629-632, 2005.

[14] National Institute of Standards and Technology, "NIST speech group website," <http://www.nist.gov/speech>, 2005.

[15] P. Kenny, P. Dumouchel, "Disentangling Speaker and Channel Effects in Speaker Verification" *Proc. ICASSP 2004*, pp. I-37-40., 2004.

[16] R. Auckenthaler, M. Carey and H. Lloyd-Thomas, "Score Normalization for Text-Independent Speaker Verification Systems", *Digital Signal Processing*, 10, pp. 42-54, 2000.

[17] M. E. Tipping and C. M. Bishop, "Mixtures of Probabilistic Principal Component Analysis", *Neural Computation*, Vol.11, No.2, pp. 443-482, 1999.

[18] P. A. Torres-Carrasquillo, E. Singer, M. A. Kohler, R. Greene, D. A. Reynolds, and J. R. Deller Jr., "Approaches to Language Identification using Gaussian Mixture Models and Shifted Delta Cepstral Features", *Proc. ICSLP 2002*, pp. 90-93, 2002.