

A Study on Universal Background Model Training in Speaker Verification

Taufiq Hasan and John H. L. Hansen, *Fellow, IEEE*

Abstract—State-of-the-art Gaussian mixture model (GMM)-based speaker recognition/verification systems utilize a universal background model (UBM), which typically requires extensive resources, especially if multiple channel and microphone categories are considered. In this study, a systematic analysis of speaker verification system performance is considered for which the UBM data is selected and purposefully altered in different ways, including variation in the amount of data, sub-sampling structure of the feature frames, and variation in the number of speakers. An objective measure is formulated from the UBM covariance matrix which is found to be highly correlated with system performance when the data amount was varied while keeping the UBM data set constant, and increasing the number of UBM speakers while keeping the data amount constant. The advantages of feature sub-sampling for improving UBM training speed is also discussed, and a novel and effective phonetic distance-based frame selection method is developed. The sub-sampling methods presented are shown to retain baseline equal error rate (EER) system performance using only 1% of the original UBM data, resulting in a drastic reduction in UBM training computation time. This, in theory, dispels the myth of “There’s no data like more data” for the purpose of UBM construction. With respect to the UBM speakers, the effect of systematically controlling the number of training (UBM) speakers versus overall system performance is analyzed. It is shown experimentally that increasing the inter-speaker variability in the UBM data while maintaining the overall total data size constant gradually improves system performance. Finally, two alternative speaker selection methods based on different speaker diversity measures are presented. Using the proposed schemes, it is shown that by selecting a diverse set of UBM speakers, the baseline system performance can be retained using less than 30% of the original UBM speakers.

Index Terms—Acoustic modeling, intelligent speaker selection, speaker recognition, speaker verification, universal background model (UBM).

Manuscript received February 16, 2010; revised August 14, 2010; accepted December 08, 2010. Date of publication February 14, 2011; date of current version July 15, 2011. This project was supported by the Air Force Research Laboratory (AFRL) through a subcontract to RADAC, Inc. under FA8750-09-C-0067 (Approved for public release. Distribution unlimited), and partially by the University of Texas at Dallas from the Distinguished University Chair in Telecommunications Engineering held by J. Hansen. A probe investigation on data selection for UBM training was presented in [1], with some of these findings motivating further investigation in Section IV. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Nestor Becerra Yoma.

The authors are with the Center for Robust Speech Systems (CRSS), Erik Jonsson School of Engineering and Computer Science, University of Texas at Dallas, Richardson, TX 75252 USA (e-mail: taufiq.hasan@utdallas.edu; john.hansen@utdallas.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TASL.2010.2102753

I. INTRODUCTION

IN recent years, Gaussian mixture model (GMM)-based approaches in text independent speaker identification systems have received considerable attention. Although, many distinct algorithms have been developed in this area, the use of GMMs for modeling acoustic features have become almost exclusive. The most fundamental GMM-based speaker recognition methods include the classical maximum *a posteriori* (MAP) adaptation of universal background model parameters [2] (GMM-UBM) and support vector machine (SVM) modeling of GMM super-vectors (GMM-SVM) [3]. Both of these approaches improve when using additional normalization schemes, such as factor analysis [4], eigenvoice [5], or nuisance attribute projection (NAP) [6]. Many groups employ these schemes successfully as individual subsystems in the recent 2008 National Institute of Science and Technology (NIST) speaker recognition (SRE) evaluations [7]–[10].

An important mutual element of these subsystems is the UBM. It is essentially a very large GMM trained to represent the speaker-independent distribution of the speech features [2] for all speakers in general, and is employed as the expected alternative speaker model during the verification task. It is also employed in open-set speaker recognition systems as well. In the two primary GMM-based systems (GMM-UBM, GMM-SVM), all speaker models are dependent on the UBM, making it a key element. However, despite its importance, focused research on UBM training has not yet been conducted in the literature [2]. The general trend is to use as many speakers and speech comprising a wide range of speech/channel conditions, without much thought regarding performance tradeoff. Though, recent research has focused on other aspects of the UBM, such as [11] where an adaptive individual background model training is considered, or [12], where the application of speaker normalization techniques on the UBM data is investigated, the basic and fundamental questions regarding the construction of a UBM and its implications in system performance is still an open unaddressed question. In this paper, we give an in depth consideration of the UBM training process and attempt to gain insight on how system performance is related to specific UBM composition.

There are a number of different parameters involved in the UBM training process. It is possible to classify these parameters into two broad categories as, 1) algorithm parameters and 2) data parameters. The algorithm parameters are the variations in the training process which include the number of mixtures, method of training, number of iterations, method of initialization, etc. The data parameters include different ways of defining the subset of available training data. These parameters consider

the corpus, the amount of data, number of speakers in the data, amount of data per speaker, method of selecting speakers, ways of using the feature vectors, data balancing according to channel, microphone, language, or other variability, and so on. Since the only available objective criteria that can measure the quality of a UBM is overall/final system performance, finding a better UBM becomes a challenge since it will generally rely on a trial and error-based endeavor, making it very impractical to vary all the mentioned parameters and find the optimal combination that gives the best performance. Thus, in order to limit the scope of this research, we only focus on a limited set of the data parameters, and attempt to analyze their effect on system performance in order to answer some fundamental questions concerning UBM training.

The first question that arises concerning the UBM data is the required amount. A common assumption in UBM training is that the more data used, the better the system performance. UBMs with 512, 1024, 2048 or more mixtures are sought after, with the assumption that they represent the definitive world speaker acoustic space. Research groups involved in the NIST SRE typically use 5-min utterances from all NIST 2004–2005 data along with the Switchboard Cellular I and II data [10]. However, there is no concrete evidence that using the maximum amount of data would guarantee the best overall performance. According to [2], as long as the development speaker population is kept the same, a small amount of data is sufficient for reasonable system performance. This suggests that, the degree of inter-speaker variation in the data is more important than the absolute amount of data per speaker. In this paper, we systematically analyze this aspect of UBM training, determine a measure of data variability from the UBM parameters, and relate this to system performance.

The issue of using a subset of features from a given UBM utterance arises as an inevitable consequence of using a reduced amount of data. The simplest methods for sub-sampling feature frames would include decimation and random feature selection, which have already been utilized to improve the CPU training time of a GMM [13]. Clearly, these feature sub-sampling methods do not consider the actual acoustic content of the features, and select features rather blindly. Here, we consider an adaptive phoneme dependent feature sub-sampling scheme for effectively capturing the subtle nuances of features in each utterance using a very small amount of data. The goal is to maximize the data variability that can be captured from a given UBM utterance using a minimal number of features. This is motivated by considering the feature selection issue at the phoneme level. Since inter-speaker variability in the UBM data has a higher contribution to system performance [2], intra-speaker phoneme variation should be less relevant for the UBM. When a long duration utterance is used for a speaker, some phones will occur more frequently and with greater duration, and therefore would contribute to probability density function (pdf) components in the UBM that represent the intra-speaker distribution of that phoneme, causing an imbalance. Thus, reducing the development data by means of proper selection of the training feature vectors will obviously improve computation speed, with a possible improvement in overall system performance as well.

The inter-speaker variability of the UBM data is directly related to the number of speakers present in the data-set. In our

study, we systematically control the number of speakers in the UBM data in an attempt to gain valuable insight on an overall system performance relation with inter-speaker variability. Considering the idea that speaker diversity is a key factor for UBM data selection, approaches of diverse speaker selection is also investigated.

This paper is organized as follows. In Section II we discuss the definition of the UBM and how data parameters should be set for an ideal UBM. Section III describes our baseline system. In Section IV, we analyze the effect of changing the amount of data in the UBM and in Section V, we consider feature sub-sampling approaches. Next, in Section VI we analyze the effect of changing the number of unique speakers in the UBM data and in Section VII we discuss different ways of sub-sampling speakers. Finally, in Section VIII we draw conclusions. It should be noted that UBM training is analogous to baking a cake, there are many ingredients that can be adjusted, but the final outcome will still have the same basic structure. It is expected that improved UBM construction may not have a significant impact in performance metrics such as equal error rate (EER) or minDCF. However, from a scientific perspective, formulating a more effective UBM using less training data should reduce computational requirements and produce a more balanced UBM that represents the overall speaker acoustic space.

II. THE IDEAL UBM

As noted earlier, the UBM is a speaker-independent Gaussian mixture model (GMM) trained with speech acoustic features from a large set of speakers to represent the general, speaker-independent distribution of features. The concept of a UBM becomes important because of the likelihood ratio test performed in the verification task. Given an observation O , and a hypothesized speaker S , the task of speaker verification can be stated as a hypothesis test between

$$\begin{aligned} H_0 : O \text{ is from speaker } S \\ H_1 : O \text{ is not from speaker } S. \end{aligned} \quad (1)$$

In general, the hypothesis H_0 and H_1 are represented by a speaker-dependent model λ_S and the background model $\lambda_{\bar{S}}$. Thus, for the observed feature vectors \mathbf{X} , the likelihood ratio test is performed by evaluating

$$\Lambda(\mathbf{X}) = \frac{p(\mathbf{X}|\lambda_S)}{p(\mathbf{X}|\lambda_{\bar{S}})} \begin{cases} \geq \theta, & \text{accept } H_0 \\ < \theta, & \text{reject } H_0. \end{cases} \quad (2)$$

Thus, ideally, the background model $\lambda_{\bar{S}}$ should be a model that represents the entire space of all possible alternatives to the hypothesized speaker S , which leads to a speaker specific background model. This approach has been adopted by many researchers in the past [11], [14]–[16]. However, creating a speaker specific background model for each enrolled speaker can be computationally expensive, especially for a large number of speakers, which is typically the case in the NIST SRE evaluations [7]. Also, it may not always be necessary to represent all outside speakers, only those that may attempt to enter the speaker verification system as imposters. Thus, most modern speaker verification systems use a single speaker-independent background model, (i.e., the UBM) for modeling the alternative

hypothesis in the likelihood ratio test. Generally, the UBM is trained using a large amount of data coming from a variety of different speakers and channel/microphone conditions, so that the model contains at least some aspect of the variabilities that could be encountered in the unknown test data.

A. Data Balancing

Inspired by the general guideline of UBM data as presented in [2], the data requirements for an ideal UBM can be specified as follows. Let the following variables be

$$S = \{s_i\}, 1 \leq i \leq N_s \quad (3)$$

$$M = \{m_i\}, 1 \leq i \leq N_m \quad (4)$$

and

$$C = \{c_i\}, 1 \leq i \leq N_c \quad (5)$$

denote the set of all available speakers, microphone types, and transmission channel types, respectively, in a single gender database. Here N_s , N_m , and N_c denote the number of speakers, microphones and transmission channels, respectively. Let $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_N\}$ denote the set of all available features in the database. Next, define the following data sets

$$\begin{aligned} \mathbf{X}_{s_i} &= \{\mathbf{x}_j | \mathbf{x}_j \text{ belongs to speaker } s_i\} \\ \mathbf{X}_{m_i} &= \{\mathbf{x}_j | \mathbf{x}_j \text{ was recorded with mic. } m_i\} \\ \mathbf{X}_{c_i} &= \{\mathbf{x}_j | \mathbf{x}_j \text{ comes from channel } c_i\}. \end{aligned} \quad (6)$$

Obviously, considering the total number of speakers, microphones, and channels, the union of each reflects the total data available:

$$\mathbf{X} = \left(\bigcup_{i=1}^{N_s} \mathbf{X}_{s_i} \right) = \left(\bigcup_{i=1}^{N_m} \mathbf{X}_{m_i} \right) = \left(\bigcup_{i=1}^{N_c} \mathbf{X}_{c_i} \right). \quad (7)$$

If $\mathbf{X}_I \subset \mathbf{X}$ denote the set of features that should be used for the ideal UBM, it should contain features from all these variabilities in equal proportions with respect to the test data. If, the prior probabilities of the occurrence of a speaker, microphone type, or channel condition in the test data is known, the feature set \mathbf{X}_I should fulfill the following constraints:

$$n(\mathbf{X}_{s_i} \cap \mathbf{X}_I) = \mathcal{R}(p(s_i)n(\mathbf{X}_I)), \quad \forall s_i \in S \quad (8)$$

$$n(\mathbf{X}_{m_i} \cap \mathbf{X}_I) = \mathcal{R}(p(m_i)n(\mathbf{X}_I)), \quad \forall m_i \in M \quad (9)$$

$$n(\mathbf{X}_{c_i} \cap \mathbf{X}_I) = \mathcal{R}(p(c_i)n(\mathbf{X}_I)), \quad \forall c_i \in C \quad (10)$$

where $n(\cdot)$, $\mathcal{R}(\cdot)$ and $p(\cdot)$ indicate the number of elements in a set, the round-off operation and the prior probability of an element in the test data, respectively. Typically, there is no prior knowledge about the test data condition distribution, leaving the system designer with the only option of assuming these prior probabilities are equal. This is known as balancing of the UBM data as discussed in [2]. It should be noted that (8) assumes all the speakers to be considerably diverse in nature otherwise similar speakers (i.e., the cohorts) may introduce unbalance in the data. This issue is further discussed in Section VII.

B. Data Amount

It is clear that the set \mathbf{X} of all available data features should be large enough to represent all variabilities faithfully. For example, if there is only 5 min of data for a cordless phone in the entire 10-h UBM data set, trying to balance the microphone type would require us to use only 5 min of data from each microphone type, which may lead to insufficient amounts of UBM data. However, what defines a data amount to be sufficient or insufficient is rather vague and probably system dependent. Generally, for a given data set, the amount of data that appropriately represents the variability of the entire corpus should be sufficient for the UBM. Since human speech features can only occupy a limited region in the feature space due to physiology constraints, it is expected that the variability of the data would be saturated when the data amount becomes very large, assuming other conditions (e.g., channel, microphone, or language variability) are kept the same. Let $\vartheta(\cdot)$ represent a function that can measure the variability of the UBM data, then as the size of the set \mathbf{X}_I increases, $\vartheta(\mathbf{X}_I)$ should approach some constant value. Mathematically, this can be written as follows:

$$\lim_{n(\mathbf{X}_I) \rightarrow \infty} \vartheta(\mathbf{X}_I) = \sigma. \quad (11)$$

Having defined the scope of the data characteristics, it is now possible to consider alternative UBM training schemes. A baseline system scenario is first considered in the next section.

III. BASELINE SYSTEM DESCRIPTION

A. System Overview

Since the objective of this research is focused only on UBM training, a fairly standard GMM-UBM [2] baseline system is employed without any mismatch compensation and score normalization. A prime reason for not using expanded or enhanced GMM-based systems (i.e., joint factor analysis (JFA) or Eigenchannel [4]), is because they require time consuming training of the Eigenchannel and Eigenvoice matrices each time the UBM is retrained. Given the number and extent of the experiments required in this study, it was decided that using such enhancements would be impractical. Future studies could further explore the impact in UBM construction with other system processing tasks.

For the front-end, 39-dimensional MFCC features (MFCC+ Δ + $\Delta\Delta$) are extracted using a 25-ms analysis window with 10-ms shift. Next, feature warping [17] based on applying a 3-s sliding window, is performed. To remove silence frames, a phone recognizer based voice activity detector (VAD) is utilized. For baseline UBM training, 1024 mixtures are used. UBM training is performed using the maximum-likelihood (ML) criterion with HTK [18] tools, performing 15 iterations per mixture split. Here, only male speaker trials are considered. For modeling, gender dependent UBMs are adapted to each enrollment speaker-dependent model using classical MAP adaptation [2] with one iteration and a relevance factor of 19. In scoring, a standard 20-best expected log likelihood ratio scoring is employed. The 5-min tel-tel condition trials [7] of the NIST 2008 SRE are used for all evaluations.

B. UBM Database

Here, two different database sets are employed for constructing the UBM: 1) 2019 male utterances from the NIST SRE 2004 1-s (5-min training) data, consisting of 126 unique speakers; 2) 5685 male utterances from NIST SRE 2004 and NIST SRE 2006 (that include channel and microphone labels), consisting of 392 unique speakers. For these two data set cases, the resulting baseline system EER was 11.43% and 11.41%, respectively.

C. Computational Resources

The speaker ID system was implemented on a high-performance Rocks computing cluster running the CentOS Linux distribution. The cluster comprises 18 HP Intel Quad-Core Xeon 2.33-GHz CPU's, yielding 72 CPU cores. A total of 126-GB RAM is available internally on the system. A 4-TB external RAID disk array is attached to the cluster by means of a storage area network (SAN). The array is connected with the cluster nodes through a 1-Gbit Ethernet switch. For calculation of CPU times for UBM training, the time required for each parallel process in the 72 CPUs were separately calculated and accumulated. Thus, the CPU times reported in this work represent a close estimation of the training time that would be required if only a single CPU was available.

IV. UBM DATA: WHAT IS A SUFFICIENT AMOUNT

As discussed in Section II, the UBM data should contain enough variability in a sufficient amount to represent each of the different channel and microphone conditions. However, using an enormously large database may not be necessary for the best performance, as long as the required variability is maintained. In this section, system performance variation is examined by increasing the total amount of data for UBM training, while keeping the speaker population the same.

An experiment was performed using the UBM database-1 as described in Section III. It should be noted that in this UBM data set, a single speaker may occur multiple times in different channel/mic conditions. This database is used mainly because of the following reasons:

- the NIST SRE 04 1-s corpus contains sufficient variability;
- other research groups have shown success in UBM training using this set alone ([8], [9], [19]);
- the actual number of speakers in this experiment is not a primary concern;
- each utterance is 5 min in duration, making it convenient to extract equal amounts of data from each utterance uniformly.

The total duration of the data is 168.25 h. Next the UBM is trained using only the first n feature frames from each utterance, and evaluated the GMM-UBM system (described in Section III) for the male trials only. For different values of n , the equivalent total data amount in hours, the EER and CPU times required for training are obtained.

A simple formula is proposed to measure the variability of the data. Since a variance measure of a 39-dimensional feature vector may not provide an accurate measure of variability of the

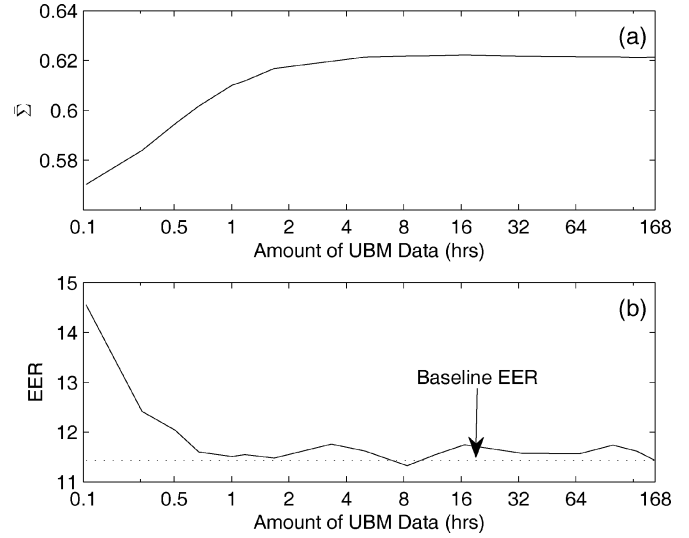


Fig. 1. (a) Variation of UBM average weighted variance ($\bar{\Sigma}$) with total amount of UBM data (h), (b) variation of system EER with total amount of UBM data (h). Feature frames are selected uniformly from each utterance.

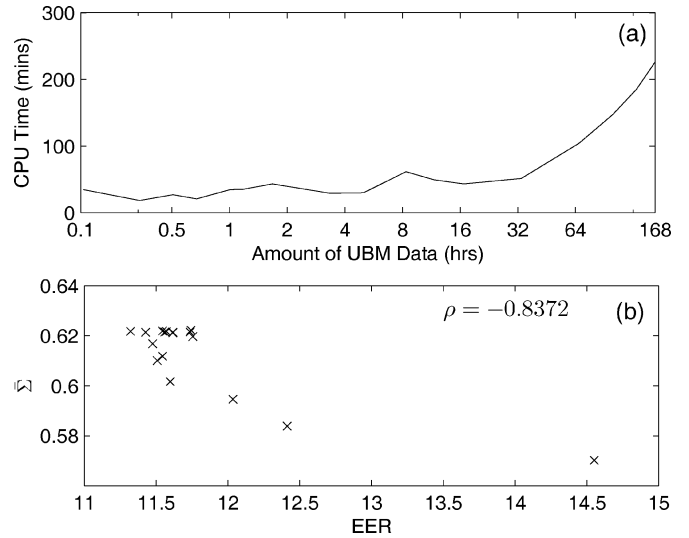


Fig. 2. (a) UBM training CPU time variation with changing amount of UBM data (h). (b) Scatter plot showing the correlation between EER and average weighted variance ($\bar{\Sigma}$). All data points from Fig. 1(a) and (b) are used to generate this scatter plot.

data, an analysis of the UBM covariance matrix is performed. Here, a parameter, average weighted variance (AWV) is defined which is computed from the UBM diagonal covariance matrices as follows. Let the UBM be expressed as

$$f(\mathbf{x}|\lambda_{\text{UBM}}) = \sum_{i=1}^M w_i \frac{1}{(2\pi)^{K/2} |\Sigma_i|^{1/2}} e^{-(1/2)(\mathbf{x}-\boldsymbol{\mu}_i)^T \Sigma_i^{-1} (\mathbf{x}-\boldsymbol{\mu}_i)} \quad (12)$$

where w_i , $\boldsymbol{\mu}_i$, Σ_i , K , and M denote the weights, mean vectors, covariance matrices, feature dimension, and number of Gaussian mixtures, respectively. Assuming a diagonal covariance matrix Σ_i , let

$$\Sigma_i = \begin{bmatrix} \sigma_{i,1}^2 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \sigma_{i,K}^2 \end{bmatrix}. \quad (13)$$

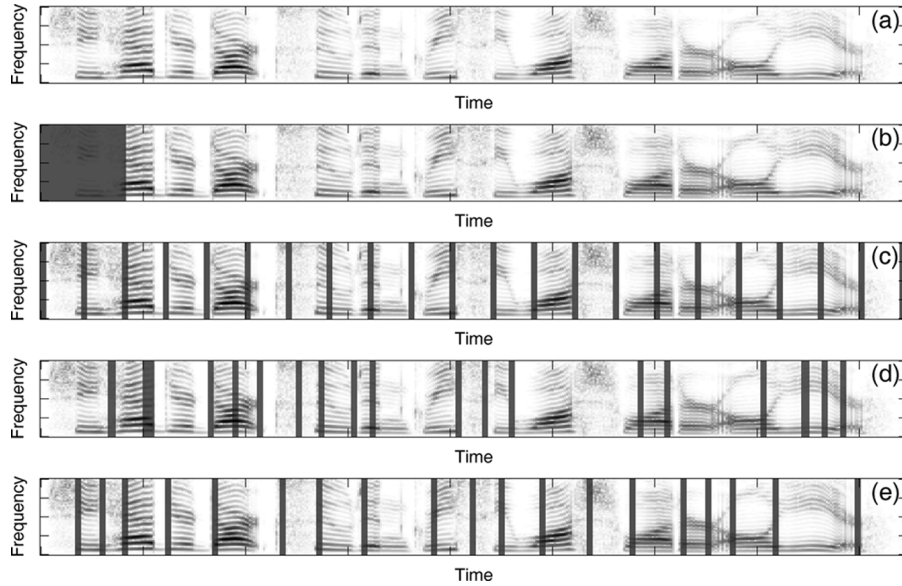


Fig. 3. Conceptual illustration of the feature selection schemes. (Selected frames are shown in dark.) (a) Original utterance spectrogram, (b) LFS, (c) UFS, (d) RFS, and (e) IFS.

Next, define the average weighted variance (AWV), $\bar{\Sigma}$, as

$$\bar{\Sigma} = \frac{1}{K} \sum_{i=1}^M \sum_{j=1}^K w_i \sigma_{i,j}^2. \quad (14)$$

From Fig. 1(b), it is clear that performance is comparable to the baseline system using only ~ 1.5 h of UBM data, which results in about ~ 2.7 s of data from each utterance. This is not very surprising since this ~ 1.5 h of data contains all the inter-speaker variabilities present in all the utterances in the original NIST SRE 04 corpus used. Very interestingly, from Fig. 1(a), a clear relation with $\bar{\Sigma}$ and system EER is observed. After using more than 1.5 h of data, the variance parameter $\bar{\Sigma}$, saturates, which indicates that the 2.7 s of data used from each utterance is actually sufficient to represent the variability it has to offer for the UBM, in this case. Obviously this duration measure cannot be completely generalized. The scatter plot in Fig. 2(b) shows the correlation between $\bar{\Sigma}$ and EER more clearly, with these two parameters having a correlation coefficient of -0.8372 . From Fig. 2(a), an exponential relation can be seen between UBM training CPU time and the total UBM data amount. In these experiments, it was observed that using about 1.5 h of data requires about 30 CPU minutes to train the UBM, while more than 200 CPU minutes are needed if all the data is used. Since CPU computation time can be considered linearly proportional to the complexity of training, this indicates that, contrary to popular belief, that “more training data is better,” the addition of five times the computational resources with more than 160 hours of training data actually has a negligible contribution to improving overall system performance.

Thus, the conclusion drawn here is that if the selected UBM data set is well chosen, (i.e., contains sufficient speech/speaker variability) using all the features of each utterance is not necessary. In the next section, we investigate how system

performance is affected based on alternative feature selection methods.

V. SUBSAMPLING OF FEATURE FRAMES

In the previous section, it was established that only 2.7 s of data from each 5 min utterance of the UBM data is sufficient for system performance equivalent to the baseline configuration. In this section, several alternative approaches for selecting this subset of development data for UBM training are considered for more effective representation of the feature space of each utterance. Time versus frequency spectrograms of these approaches are illustrated in Fig. 3(b)–(e) using a spectrogram of an original TIMIT utterance, shown in Fig. 3(a). The use of the first n feature vectors from each utterance, as done in the previous section, is termed as “leading feature selection” (LFS) and is depicted in Fig. 3(b). As noted earlier, subsampling the feature frames can also be done uniformly or randomly [13]. These methods are denoted as UFS (“uniform feature selection”) and RFS (“random feature selection”) [13], and illustrated in Fig. 3(c) and (d), respectively. Now, though the subsampling methods LFS, UFS and RFS would reduce computation time, they are overly simplified and completely data independent. These methods do not consider the specific distribution of the phonetic content over time. Thus, we propose a generic method termed as “intelligent” feature selection (IFS), which aims to select a diverse set of n training feature frames from input training utterances. This method assesses the similarity of successive frames using a phonetically motivated distance measure, with selection of a feature frame only if the corresponding dissimilarity is higher than some threshold. In Fig. 3(e), a conceptual IFS method that attempts to select a frame from the beginning of each distinct phoneme is illustrated.

Clearly, there can be variations in this approach if alternative distance criteria between features are used. Since one design criteria here is training speed, in this phase the Euclidean distance is used due to its simplicity.

A. Feature Selection Based on Euclidean Distance

In this section, an intelligent feature selection (IFS) scheme is proposed based on the simple Euclidean distance between features (IFS-EU). The aim here is to estimate similarity between successive feature frames using this distance measure, and select a feature frame only if the frame is sufficiently different from those previously selected. From an intuitive understanding of the Euclidean distance, it is noted that this distance measure is related to the smoothed log-spectral distance [20] when applied to cepstral feature vectors. The formulation is started by deriving the pdf of the distance function between feature vectors.

1) *PDF of Euclidean Distance Between Features:* Assume that the K -dimensional feature vectors of the development speaker data, originating from a specific phone, can be modeled by an independent, wide sense stationary (WSS), white Gaussian vector random sequence $\mathbf{X}[n]$ with a covariance function matrix $\mathbf{K}_{XX}[m, n]$ given by

$$\mathbf{K}_{XX}[m, n] = \text{diag}(\lambda_1 \dots \lambda_K) \delta[m - n] \quad (15)$$

where m, n denotes the feature indices, and λ_p ($p = 1, \dots, K$) are the variances of the individual cepstral coefficients. The Euclidean distance between the m th and n th feature vector will be

$$d(m, n) = \|\mathbf{X}[m] - \mathbf{X}[n]\|^{1/2}. \quad (16)$$

The feature vectors have a common mean, and thus the term inside the parenthesis in (16) will be a zero mean vector random sequence. Also, due to the independence assumption, $d(m, n)$ is independent of m and n . Thus,

$$d(m, n) = d = \|\mathbf{Z}\|^{1/2} \quad (17)$$

where \mathbf{Z} is a zero mean Gaussian random vector having a covariance matrix \mathbf{K}_{ZZ} and found to be

$$\mathbf{K}_{ZZ} = \text{diag}(2\lambda_1 \dots 2\lambda_K). \quad (18)$$

The factor of 2s are introduced because each element of \mathbf{Z} is constructed from the subtraction of two independent white Gaussian random variables having variances λ_p ($p = 1, \dots, K$). From (17), it is possible to write

$$d^2 = \sum_{i=1}^K Z_i^2 = \sum_{i=1}^K (2\lambda_i) W_i^2, \text{ where } W_i \sim \mathcal{N}(0, 1). \quad (19)$$

For simplification, assume that the effect of the individual λ_i values in (19) can be approximated using a lumped parameter $\bar{\lambda}$. Thus,

$$d^2 \approx 2\bar{\lambda} \sum_{i=1}^K W_i^2 = 2\bar{\lambda} Y \quad (20)$$

where $\bar{\lambda}$ is defined as the average variance given by

$$\bar{\lambda} = \frac{1}{K} \sum_{i=1}^K \lambda_i. \quad (21)$$

In (20), $Y = \sum_{i=1}^K W_i^2$ is a squared sum of zero mean independent Gaussian random variables, and thus will follow a chi-squared distribution given by

$$f_Y(y) = \frac{\left(\frac{1}{2}\right)^{K/2}}{\Gamma\left(\frac{K}{2}\right)} y^{(K/2-1)} e^{-y/2}. \quad (22)$$

From (20), $d = \sqrt{2\bar{\lambda}Y}$. Using this transformation in (22), the pdf of d can be obtained as

$$f_D(d) = \frac{2^{1-K}}{\Gamma\left(\frac{K}{2}\right)} \frac{d^{K-1}}{\bar{\lambda}^{K/2}} \exp\left(-\frac{d^2}{4\bar{\lambda}}\right). \quad (23)$$

The mean and variance of this distribution can be found as

$$\mu_D = \frac{2\sqrt{\bar{\lambda}}\Gamma\left(\frac{1+K}{2}\right)}{\Gamma\left(\frac{K}{2}\right)} \text{ and} \quad (24)$$

$$\sigma_D^2 = 2K\bar{\lambda} - \mu_D^2, \text{ respectively.} \quad (25)$$

Note that the pdf of d will provide the likelihood of the distance between any two features in the data set. In other words, if the goal is to select feature vectors that are farther apart on average, a set of features should be selected in which each pair has a distance greater than μ_D .

2) *Calculation of Distance Threshold:* In this feature selection problem, the data is processed on a frame-by-frame basis. Assuming that the pdf parameters are known for the current frame, select the next frame if its distance from the current frame is greater than a threshold d_{th} . For a fixed value $\alpha \in [0, 1]$, define d_{th} as

$$P[d > d_{th}] = \int_{d_{th}}^{\infty} f_D(z) dz = \alpha. \quad (26)$$

The process is illustrated in Fig. 4(b) for $\alpha = 0.1$ and 0.2. This implies that a feature vector is selected only if its distance from the current feature is so high that the event is less probable than α , suggesting a high likelihood of a change in the phone represented by the feature. It is observed that the pdf $f_D(d)$ can be closely approximated by a Gaussian distribution having a mean μ_D and variance σ_D^2 . Fig. 4(a) compares the function $f_D(d)$, its Gaussian approximation $\hat{f}_D(d)$, and a histogram plot, estimated from 13-dimensional MFCC coefficients of a UBM utterance. Using the Gaussian approximation, it is possible to obtain from (26)

$$d_{th} = \mu_D + \sqrt{2}\sigma_D \text{erfc}^{-1}(2\alpha) \quad (27)$$

where erfc^{-1} is the inverse of the complementary error function (erfc). Here, $\text{erfc}()$ is defined as

$$\text{erfc}(x) = \sqrt{\frac{2}{\pi}} \int_0^x e^{-t^2} dt. \quad (28)$$

3) *Estimation of PDF Parameters:* Here a recursive method is employed for estimating the feature vector mean and variance

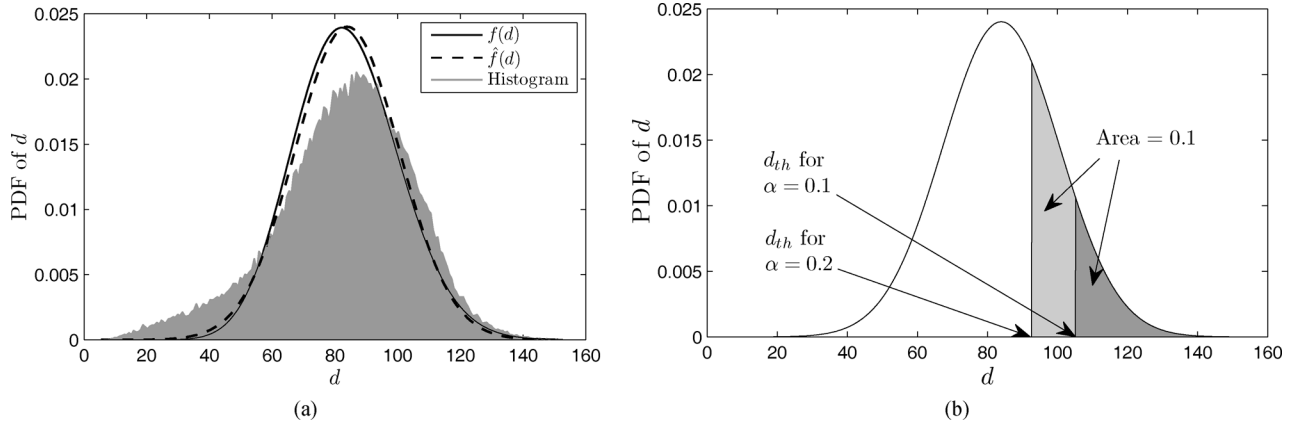


Fig. 4. (a) Comparison of the theoretical pdf, its Gaussian approximation and the actual pdf obtained from histogram. 13-dimensional MFCC coefficients were used and the parameter $\bar{\lambda}$ was calculated directly from the data. For this data, $\bar{\lambda} = 281.6836$, $\mu_D = 83.9506$, and $\sigma_D^2 = 276.07$. (b) A pdf of inter-feature Euclidean distance and the proposed distance threshold (shown for $\alpha = 0.1$ and 0.2).

similar to [21]. Denoting $\lambda[n]$ as the vector containing the diagonal elements of $\mathbf{K}_{\mathbf{X}\mathbf{X}}[0, 0]$, and $\mu_{\mathbf{X}}[n]$ as the mean vector of the n th frame, the equations

$$\begin{aligned} \hat{\mu}_{\mathbf{X}}[n] &= \beta_m \hat{\mu}_{\mathbf{X}}[n-1] + (1 - \beta_m) \mathbf{X}[n], \text{ and} \\ \hat{\lambda}_{\mathbf{X}}[n] &= \beta_v \hat{\lambda}_{\mathbf{X}}[n-1] + (1 - \beta_v) \|\mathbf{X}[n] - \hat{\mu}_{\mathbf{X}}[n]\|^2 \end{aligned} \quad (29)$$

are used, where $\beta_m, \beta_v \in [0, 1)$ are overall smoothing parameters.

4) *Implementation:* Let i denote the current frame index and set $j = i + 1$. For initialization ($i = 1$), $\mathbf{X}[i]$ is always selected, and $\hat{\mu}_{\mathbf{X}}[i]$ and $\hat{\lambda}_{\mathbf{X}}[i]$ are calculated from $\mathbf{X}[i]$ and $\mathbf{X}[j]$ as

$$\hat{\mu}_{\mathbf{X}}[i] = 0.5(\mathbf{X}[i] + \mathbf{X}[j]) \text{ and} \quad (30)$$

$$\hat{\lambda}_{\mathbf{X}}[i] = 0.5(\mathbf{X}[i] \cdot^2 + \mathbf{X}[j] \cdot^2) - \hat{\mu}_{\mathbf{X}}[i] \cdot^2 \quad (31)$$

where $(\cdot)^2$ denotes an element-wise square operation. Next, $\bar{\lambda}$ and d_{th} are calculated using (21) and (27). Now, j is iteratively incremented by 1 and $d(i, j)$ is calculated from (16). The values $\hat{\mu}_{\mathbf{X}}[i]$ and $\hat{\lambda}_{\mathbf{X}}[i]$ are updated in each step using (28) and (29), along with the threshold d_{th} . If $d(i, j) > d_{th}$ is found, $\mathbf{X}[j]$ is selected. Next, $i = j$ and $j = i + 1$ is set and the process is repeated until the desired number of features are selected. In our experiments, the settings used are $\alpha = 0.1$, $\beta_m = 0.8$, and $\beta_v = 0.6$.

B. Performance of Subsampling Schemes

The EER performance along with the computation time required for UBM training using the set of presented approaches is shown in Table I. Baseline performance with 100% of the data used to train the UBM is 11.43%. It is clear that all four UBM training methods considered here, using a mere 1% of UBM data employed can provide performance equivalent to the baseline system with up to a 7 times reduction in CPU computation time. In addition, using the proposed feature selection scheme, denoted by IFS-EU, it is noted that a $\sim 0.4\%$ reduction in EER is achieved in comparison to the baseline system. This is because the selected features in the IFS-EU method are better able to represent the diverse speaker pool, while suppressing some of the fine model traits of the intra-speaker phoneme variability, which, it is believed, are less important for construction of an

TABLE I
COMPARISON OF DIFFERENT UBM TRAINING SCHEMES WITH RESPECT TO EER AND TRAINING CPU TIME

Method	% data	EER (%)	CPU Time h:mm
Baseline	100%	11.43	3:46
LFS	1%	11.48	0:24
UFS	1%	11.54	0:22
RFS	1%	11.41	0:18
IFS-EU	1%	10.99	0:27

effective UBM. However, we clarify that we do not claim that the 99% data not used for UBM training are definitely harmful for the system. We simply emphasize the fact that more data is not necessarily better, and sub-sampling this data properly can provide equivalent or even better performance than using all the data.

VI. UBM DATA: NUMBER OF SPEAKERS

In this section, the impact of changing the number of unique speakers in the UBM training data on system performance is considered. It was shown in the previous section, that if the data contains sufficient variability, a very small portion of data should be sufficient for training the UBM. Now, different speakers should possess different feature speech/physiology characteristics, indicating that an increase in the number of speakers in the UBM should lead to an increase in the variance of the data. Intuitively, this should be beneficial to overall system performance. An experiment is performed to validate this hypothesis.

In this experiment, the amount of data was kept fixed and the number of unique speakers was varied from 10 to 320 in an exponential manner. The UBM dataset-2 was used in this case because it has a larger number of speakers. For each UBM training run, the specified number of speakers are selected randomly from the pool and system performance is computed for the UBM trained with those speakers' features. Five independent experimental runs are performed and the average of those EER are calculated. The AWV values are also calculated for each UBM using (14). In Fig. 5(a) and (b), AWV ($\bar{\Sigma}$) and EER are plotted against the number of UBM speakers. As we have expected, the system performance is improved drastically, as the

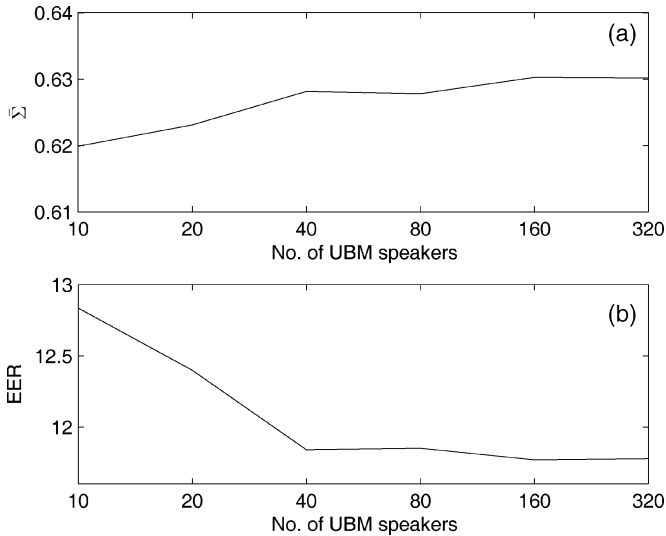


Fig. 5. Variation of (a) AWV ($\bar{\Sigma}$) and (b) system performance with the change of number of UBM speakers.

number of UBM speakers are increased, with an increase of the AWV. At some level, overall performance saturates. Thus, we justify that introducing a new speaker increases the variability of the UBM data (as reflected in the increase in AWV), which can benefit system performance. This in effect helps further justify the argument of [2] regarding the amount of training data.

VII. SELECTION OF UBM SPEAKERS

In this section, the goal is to investigate the issue of using a subset of all the available speakers in the UBM. It has been established that having a large number of dissimilar speakers in the UBM data aids in improving system performance. However, it is known that many speakers have similar acoustic properties, (i.e., the cohort speakers) that may again introduce an unbalance in the UBM data. This issue can be illustrated with a hypothetical feature space in Fig. 6. In this data-set, if equal amounts of data from all speakers are used to train the UBM, the similar speakers that are clustered together, (i.e., cohort groups 1, 2, and 3) will be emphasized in the UBM. This would result in a higher score from the UBM in the likelihood ratio test if the test speaker is from that cohort group. Now, it would be better if these speakers are spread out in the feature space as much as possible so that the entire acoustic space is uniformly covered (assuming a uniform open speaker test space). However, practically there are some problems in that scenario.

- The feature space is multidimensional, which means uniformly covering this space would require an infinite number of speakers.
- In reality, the speaker features are not very easily distinguishable as in the simplistic illustration in Fig. 6, rather they are highly overlapping.

Thus, the motivation here is to use a reduced number of speakers than all the available speakers according to some speaker divergence criteria, so that closely related speakers are not used in the UBM (i.e., if the speaker is already included in the training set, do not include an acoustically close neighbor as well).

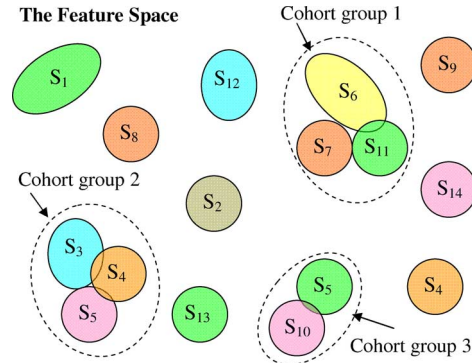


Fig. 6. Schematic diagram of the speaker space in the UBM.

A. KL Divergence-Based Speaker Selection (KL-D)

Here a UBM speaker selection method is developed using the Kullback–Leibler (KL) divergence between speaker models. For each UBM speaker s_i , $i \in (1, N_s)$, a GMM model Λ_i is trained. To calculate the similarity between GMMs, the symmetric KL divergence [22], [23], is used, given by

$$D_{KL}^s(\Lambda_i, \Lambda_j) = E_{\Lambda_i(\mathbf{X})} \left[\log \frac{\Lambda_i(\mathbf{X})}{\Lambda_j(\mathbf{X})} \right] E_{\Lambda_j(\mathbf{X})} \left[\log \frac{\Lambda_j(\mathbf{X})}{\Lambda_i(\mathbf{X})} \right] \quad (32)$$

where $\Lambda_i(\mathbf{X})$ and $\Lambda_j(\mathbf{X})$ are likelihoods of occurrence of the observation vector \mathbf{X} , given that it belongs to speaker model Λ_i and Λ_j , respectively. Next, compute the $N_s \times N_s$ divergence matrix, obtaining the KL score for each pair of speakers. To measure how diverse a speaker i is from the all other speakers, define a diversity factor D_i , given by

$$D_i^{(KL)} = \frac{1}{N_s} \sum_{j \in N_s, j \neq i} D_{KL}^s(\Lambda_i, \Lambda_j). \quad (33)$$

This relation means D_i is a measure of the average divergence of the model Λ_i from all other speaker models. Thus, after computing all D_i values, they are sorted according to their absolute value, and the top N_D most divergent speakers are selected for the UBM.

B. Speaker Selection Using Prototype UBM (P-UBM)

In this method, to find the most divergent speakers, all the data is pooled and a prototype UBM model, Λ_0 is trained. Assuming this UBM holds a central position in the GMM space, an attempt to find speakers that are most divergent from this UBM is performed. The diversity factor for each speaker i is computed simply from the likelihood of occurrence of that speaker's features given the model Λ_0

$$D_i^{(P)} = \Lambda_0(\mathbf{X}_i). \quad (34)$$

In a similar way, the $D_i^{(P)}$ values are sorted and the top N_D divergent speakers from the prototype UBM is selected. It is noted that this is a simplistic method for speaker selection and does not guarantee that the selected speakers are diverse among themselves.

C. Results of Speaker Selection Methods

In these experiments, the UBM database-2 is used and the total amount of data is held fixed to 1.5 h. The system is evalu-

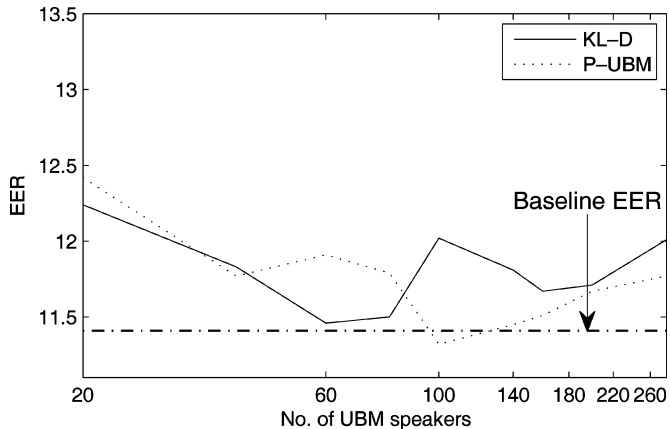


Fig. 7. System performance variation with the change of number of UBM speakers selected using different methods.

ated by varying the number of speakers, N_D from 20 to 300 in an exponential fashion while the proposed KL-D and P-UBM methods were used for selecting the best speakers. The exact number of frames were selected from each utterance using the LFS method (described in Section V) from the selected speakers' data such that the total amount of data equals 1.5 h. For all values of N_D , the EER values obtained are plotted in Fig. 7. It should be noted that 1.5 h of data used in this case is only 0.3% of the complete UBM database-2 (which contains 473.75 h of data). Thus, according to experiments in Section IV, this amount is not sufficient to retain the baseline system performance. However, this lower amount is still employed so that (a) it is more convenient to analyze the effect of the selected speakers data, and (b) enough data per speaker is available for the case of a lower number of selected speakers.

From Fig. 7 we observe that, though there are some fluctuations,¹ it can be seen that both methods perform very close to the baseline system using a significantly lower number of speakers (i.e., 60 and 100 for methods KL-D and P-UBM, respectively). Notice that this close to baseline performance is achieved with only 1.5 h of data, instead of 473.75 h of data. Interestingly, after a certain point, system performance actually degrades as the number of speakers is increased in the proposed methods, which we believe is due to the introduction of similar/redundant speakers that are creating an imbalance in the UBM data. As expected, the performance does not reach the baseline for a larger number of speakers for either method due to the use of 1.5 h of data which is not sufficient to retain the variability of the dataset. Now that we identified the regions in the plot where the proposed methods perform the best, we attempt to increase the data amount for the selected speakers for further performance improvement. For the KL-D approach using 60 speakers, we increased the amount of data from 1.5 h to 3 h and obtained an improvement in EER from 11.46% to 11.30%, which is better than the baseline EER. For the P-UBM method, using 100 speakers

¹We believe this fluctuation in EER is due to the fact that the GMMs are trained on multiple utterances of the same speaker from different channels. This creates a mild bias toward the dominating channel type in each GMM, resulting in slight channel dependent clustering in some cases. Future studies could explore UBM construction by removing channel effects from the UBM utterances using techniques like JFA, total variability features[24], IIFA [25], etc.

TABLE II
COMPARISON OF DIFFERENT SPEAKER SELECTION APPROACH FOR UBM TRAINING WITH RESPECT TO EER AND NO. OF SPEAKERS

Method	% data	Number of speakers	EER (%)
Baseline	100%	392	11.41
KL-D	0.31%	60	11.46
KL-D	0.62%	60	11.30
P-UBM	0.31%	100	11.32

and 1.5 h of data already provides a slight improvement over baseline system which uses the full UBM database-2. The results are summarized in Table II.

Note that there are 392 speakers in UBM database-2, which means less than 30% of the speakers were used in both proposed methods. Also, less than 1% of the total amount of data was used for training the UBM. Thus, we conclude that if a diverse set of speakers can be carefully selected, a much lower number of speaker data can provide performance equivalent to/better than the baseline system.

VIII. CONCLUSION

In this paper, an organized method is developed for determining the data to be selected for effective UBM training. Rigorous experiments were performed showing the relationship between data variance and overall speaker verification system performance. Four efficient sub-sampling schemes for feature frame selection were presented with potential benefits of reducing the computation time by up to seven-fold. A new intelligent feature frame sub-sampling algorithm is proposed which is experimentally shown to outperform the baseline system that uses all the available data. The implication of selectively using speaker data for UBM construction was analyzed and two effective speaker selection methods were proposed and evaluated. The results show that carefully selected reduced speech data size and speaker count are sufficient to achieve effective speaker verification performance.

REFERENCES

- [1] T. Hasan, Y. Lei, A. Chandrasekaran, and J. H. L. Hansen, "A novel feature sub-sampling method for efficient universal background model training in speaker verification," in *Proc. IEEE ICASSP'10*, Dallas, TX, Mar. 2010, pp. 4494–4497.
- [2] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker verification using adapted Gaussian mixture models," *Digital Signal Process.*, vol. 10, no. 1–3, pp. 19–41, 2000.
- [3] W. Campbell, D. Sturim, D. Reynolds, and A. Solomonoff, "SVM based speaker verification using a GMM supervector kernel and NAP variability compensation," in *Proc. ICASSP'06*, Toulouse, France, May 2006, vol. 1, pp. 97–100.
- [4] P. Kenny, P. Ouellet, N. Dehak, V. Gupta, and P. Dumouchel, "A study of interspeaker variability in speaker verification," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 16, no. 5, pp. 980–988, Jul. 2008.
- [5] R. Kuhn, J. Junqua, P. Nguyen, and N. Niedzielski, "Rapid speaker adaptation in eigenvoice space," *IEEE Trans. Speech Audio Process.*, vol. 8, no. 6, pp. 695–707, Nov. 2000.
- [6] A. Solomonoff, W. Campbell, and I. Boardman, "Advances in channel compensation for SVM speaker recognition," in *Proc. IEEE ICASSP'05*, 2005, vol. 1, pp. 629–632.
- [7] The NIST Year 2008 Speaker Recognition Evaluation Plan 2006 [Online]. Available: <http://www.nist.gov>, .

- [8] H. Li, B. Ma, K.-A. Lee, H. Sun, D. Zhu, K. C. Sim, C. You, R. Tong, I. Karkkainen, C.-L. Huang, V. Pervouchine, W. Guo, Y. Li, L. Dai, M. Nosrathighods, T. Tharmarajah, J. Epps, E. Ambikairajah, E. S. Chng, T. Schultz, and Q. Jin, "The I4U system in NIST 2008 speaker recognition evaluation," in *Proc. IEEE ICASSP'09*, Apr. 2009, pp. 4201–4204.
- [9] W. Guo, Y. Long, Y. Li, L. Pan, E. Wang, and L. Dai, "IFLY system for the NIST 2008 speaker recognition evaluation," in *Proc. IEEE ICASSP'09*, Apr. 2009, pp. 4209–4212.
- [10] L. Burget, M. Fapso, V. Hubeika, O. Glembek, M. Karafiát, M. Kockmann, P. Matejka, P. Schwarz, and J. Cernocký, "BUT system description: NIST SRE 2008," in *Proc. NIST Speaker Recognition Evaluation Workshop*, 2008, pp. 1–4.
- [11] Y. Bar-Yosef and Y. Bistriz, "Adaptive individual background model for speaker verification," in *Proc. Interspeech'09*, 2009.
- [12] A. Sarkar, S. Umesh, and S. P. Rath, "Text-independent speaker identification using vocal tract length normalization for building universal background model," in *Proc. Interspeech'09*, 2009.
- [13] C. Barras, X. Zhu, J.-L. Gauvain, and L. Lamel, *The CLEAR'06 LIMSI Acoustic Speaker Identification System For CHIL Seminars*. Berlin/Heidelberg, Germany: Springer, 2007, vol. 4122, Lecture Notes in Computer Science.
- [14] D. A. Reynolds, "Speaker identification and verification using gaussian mixture speaker models," *Speech Commun.*, vol. 17, no. 1-2, pp. 91–108, Aug. 1995.
- [15] A. E. Rosenberg, J. DeLong, C. H. Lee, B. H. Juang, and F. K. Soong, "The use of cohort normalized scores for speaker verification," in *Proc. ISCA ICSLP'92*, 1992, pp. 599–602.
- [16] T. Matsui and S. Furui, "Similarity normalization method for speaker verification based on a posteriori probability," in *Proc. ISCA Workshop Autom. Speaker Recognition, Identification Verification*, 1994.
- [17] J. Pelecanos and S. Sridharan, "Feature warping for robust speaker verification," in *Proc. 2001: Speaker Odyssey*, 2001, pp. 213–218.
- [18] S. Young, *HTK Reference Manual*. Cambridge, U.K.: Cambridge Univ. Eng. Dept., 1993.
- [19] E. Dalmaso, F. Castaldo, P. Laface, D. Colibro, and C. Vair, "Loquendo – Politecnico di Torino's 2008 NIST speaker recognition evaluation system," in *Proc. IEEE ICASSP'09*, Apr. 2009, pp. 4213–4216.
- [20] S. Furui, *Digital Speech Processing, Synthesis, and Recognition*. New York: Marcel Dekker, 1989.
- [21] M. M. Bruce, Estimation of variance by a recursive equation. Technical note NASA, 1969.
- [22] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, 1951.
- [23] M. Ben and F. Bimbot, "D-map: A distance-normalized map estimation of speaker models for automatic speaker verification," in *Proc. IEEE ICASSP'03*, Apr. 2003, vol. 2, pp. II-69–II-72.
- [24] N. Dehak, P. Kenny, R. Dehak, P. Ouellet, and P. Dumouchel, "Front-end factor analysis for speaker verification," *IEEE Transaction Audio, Speech Language Processing.*, vol. 19, no. 4, pp. 788–798, May 2011.
- [25] Y. Lei and J. Hansen, "Factor analysis-based information integration for arabic dialect identification," in *Proc. IEEE ICASSP'10*, Dallas, TX, Mar. 2010, pp. 4337–4340.



Taufiq Hasan received the B.Sc. and M.Sc. degrees in electrical and electronic engineering from Bangladesh University of Engineering and Technology, Dhaka, in 2006 and 2008, respectively. He is currently pursuing the Ph.D. degree in the Erik Jonsson School of Engineering and Computer Science, University of Texas at Dallas (UTD), Richardson.

He is currently a Research Assistant in the Erik Jonsson School of Engineering and Computer Science, UTD. He was a Lecturer in the Electrical and Electronic Engineering Department, United International University, Dhaka, Bangladesh, from December 2006 to June 2008. He is a member of the Center for Robust Speech Systems (CRSS) at UTD. His research interests are robust speaker recognition in noise and channel mismatch, speech enhancement, and automatic video summarization.



John H. L. Hansen (S'81–M'82–SM'93–F'07) received the B.S.E.E. degree from the College of Engineering, Rutgers University, New Brunswick, NJ, in 1982 and the M.S. and Ph.D. degrees in electrical engineering from Georgia Institute of Technology, Atlanta, GA, in 1983 and 1988, respectively.

He joined the Erik Jonsson School of Engineering and Computer Science, University of Texas at Dallas (UTD), Richardson, in the fall of 2005, where he is Professor and Department Head of Electrical Engineering, and holds the Distinguished University

Chair in Telecommunications Engineering. He also holds a joint appointment as Professor in the School of Behavioral and Brain Sciences (Speech and Hearing). At UTD, he established the Center for Robust Speech Systems (CRSS) which is part of the Human Language Technology Research Institute. Previously, he served as Department Chairman and Professor in the Department of Speech, Language, and Hearing Sciences (SLHS), and Professor in the Department of Electrical and Computer Engineering, at the University of Colorado at Boulder (1998–2005), where he co-founded the Center for Spoken Language Research. In 1988, he established the Robust Speech Processing Laboratory (RSPL) and continues to direct research activities in CRSS at UTD. His research interests span the areas of digital speech processing, analysis and modeling of speech and speaker traits, speech enhancement, feature estimation in noise, robust speech recognition with emphasis on spoken document retrieval, and in-vehicle interactive systems for hands-free human-computer interaction. He has supervised 51 (22 Ph.D., 29 M.S./M.A.) thesis candidates. He is author/coauthor of 380 journal and conference papers and eight textbooks in the field of speech processing and language technology, coauthor of the textbook *Discrete-Time Processing of Speech Signals*, (IEEE Press, 2000), co-editor of *DSP for In-Vehicle and Mobile Systems* (Springer, 2004), *Advances for In-Vehicle and Mobile Systems: Challenges for International Standards* (Springer, 2006), *In-Vehicle Corpus and Signal Processing for Driver Behavior* (Springer, 2008), and lead author of the report "The Impact of Speech Under 'Stress' on Military Speech Technology," (NATO RTO-TR-10, 2000).

Prof. Hansen was named IEEE Fellow in 2007 for contributions in "Robust Speech Recognition in Stress and Noise," and is currently serving as Technical Chair of the IEEE Signal Processing Society Speech and Language Processing Technical Committee (2010–2013; member 2005–2008), and Educational Technical Committee (2005–2008; 2008–2010). Previously, he has served as Technical Advisor to U.S. Delegate for NATO (IST/TG-01), IEEE Signal Processing Society Distinguished Lecturer (2005/2006), Associate Editor for the IEEE TRANSACTIONS ON SPEECH AND AUDIO PROCESSING (1992–1999), Associate Editor for the IEEE SIGNAL PROCESSING LETTERS (1998–2000), Editorial Board Member for the IEEE Signal Processing Magazine (2001–2003). He has also served as a Guest Editor of the October 1994 special issue on Robust Speech Recognition for the IEEE TRANSACTIONS ON SPEECH AND AUDIO PROCESSING. He has served on the Speech Communications Technical Committee for the Acoustical Society of America (2000–2003), and is serving as a member of the ISCA (International Speech Communications Association) Advisory Council. In 2010, he was recognized as an ISCA Fellow, for contributions on "research for speech processing of signals under adverse conditions." He was recipient of The 2005 University of Colorado Teacher Recognition Award as voted on by the student body. He also organized and served as General Chair for ICSLP/Interspeech-2002: International Conference on Spoken Language Processing, September 16–20, 2002, and served as Co-Organizer and Technical Program Chair for IEEE ICASSP-2010, Dallas, TX.