Speech Recognition Front End Without Information Loss

Matthew Ager, Zoran Cvetković *Senior Member, IEEE,* and Peter Sollich

***Abstract*—Speech representation and modelling in high-dimensional spaces of acoustic waveforms, or a linear transformation thereof, is investigated with the aim of improving the robustness of automatic speech recognition to additive noise. The motivation behind this approach is twofold: (i) the information in acoustic waveforms that is usually removed in the process of extracting low-dimensional features might aid robust recognition by virtue of structured redundancy analogous to channel coding, (ii) linear feature domains allow for exact noise adaptation, as opposed to representations that involve non-linear processing which makes noise adaptation challenging. Thus, we develop a generative framework for phoneme modelling in high-dimensional linear feature domains, and use it in phoneme classification and recognition tasks. Results show that classification and recognition in this framework perform better than analogous PLP and MFCC classifiers below 18dB SNR. A combination of the high-dimensional and MFCC features at the likelihood level performs uniformly better than either of the individual representations across all noise levels.**

***Index Terms*—phoneme classification, speech recognition, robustness, noise.**

I. Introduction

A major problem faced by state-of-the-art automatic speech recognition (ASR) systems is a lack of robustness, manifested as a substantial performance degradation due to common environmental distortions, due to a discrepancy between training and run-time conditions, or due to spontaneous conversational pronunciation [2], [3]. It was long believed that context and language modelling would provide ASR with the level of robustness inherent to human speech recognition, hence substantial research efforts have been invested in these higher levels of speech recognition systems. At the same time, the importance of robust recognition of isolated phonemes, syllables and nonsense words has not been fully investigated, while it is well known that humans attain a major portion of their robustness in speech recognition early on in the process, before and independently of context effects [4], [5]. Moreover, for language and context modelling to work optimally, elementary speech units need to be recognised sufficiently accurately. In recognising syllables or isolated words, the human auditory system performs above chance level already at -18 dB SNR (signal-to-noise ratio) and significantly above it at -9 dB SNR [4], [5]. Recent more detailed studies show that human speech recognition remains unaffected by noise down to -2 dB SNR [6], [7]. No automatic speech classifier is able to achieve performance close to that of the human auditory system in recognising such isolated words or phonemes under severe noise conditions [8]. ASR systems deliver top performance owing to sophisticated language models, combined with hidden Markov model (HMM) advances [9]. However, the problem of robustness of ASR, or the lack thereof, still persists, and the underlying concern of robustly recognising isolated phonetic units remains an important unresolved issues. Hence, in this study we explore a novel approach to representing and modelling speech and investigate its effectiveness in the context of phoneme classification and recognition.

This work was done while M. Ager was with the Department of Mathematics, King's College London. P. Sollich is with the Department of Mathematics and Z. Cvetković is with the Department of Informatics, King's College London, Strand, London WC2R 2LS, UK.

This project was supported by EPSRC Grant EP/D053005/1.

This work was presented in part at ISIT 2011 [1].

While there are many reasons for the lack of robustness in ASR, one of the major factors could be the excessive nonlinear compression that takes place in the front-end of ASR systems. As the first step in all speech recognition algorithms, consecutive speech segments are represented by low-dimensional feature vectors, most commonly *cepstral* features such as Mel-Frequency Cepstral Coefficients (MFCC) [10] or Perceptual Linear Prediction (PLP) coefficients [11]. Using low-dimensional features was unavoidable when initially introduced in the 1970s [12], as it removes non-lexical variability irrelevant to recognition, and enables learning of statistical models from limited data and using very limited computational resources. However, this paradigm, which resulted in a major performance boost a few decades ago, might be a bottleneck towards achieving robustness nowadays, when massive amounts of training data are available and computers are orders of magnitude more powerful.

In the process of discarding components of the speech signal that are considered unnecessary for recognition, some information that makes speech such a robust message representation might be lost. It is commonly believed that speech representations that are used for compression also provide good feature vectors for speech recognition. The rationale is that, since speech can be reconstructed from its compressed form to sound like natural speech that the human auditory system can recognise quite reliably, then no relevant information is lost due to the compression. Speech production/recognition is, however, analogous to a channel coding/decoding problem, while speech compression is a source coding problem, and these two are fundamentally different. In particular, speech production embeds redundancy in speech waveforms in a highly structured manner, and distributions of different phonetic units can withstand a significant amount of additive noise and distortion before they overlap significantly. Speech compression and standard ASR front-ends, on the other hand, remove most of this redundancy in a manner that is optimal from a source coding perspective, and represent speech in a space of a relatively low dimension where different speech units, although separated, may not be sufficiently far apart from each other; they may then overlap considerably already at lower noise levels than in the original domain of acoustic waveforms. Human speech perception studies have shown explicitly that the information reduction that takes place in the conventional front-ends leads to a severe degradation in human speech recognition performance and, furthermore, that in noisy environments there is a high correlation between human and machine errors in recognition of speech with distortions introduced by typical ASR front-end processing [13], [14], [15], [16]. Hence, in this paper we study models of speech in the high-dimensional domain of acoustic waveforms or some linear transform thereof. An additional benefit from considering uncompressed waveforms is that modelling of noisy data, given models in quiet, is straightforward. This is in contrast with cepstral representations where, due to the nonlinearities involved, distributions change considerably both with noise type and level. This then makes efficient adaption of speech models to different noise conditions very challenging.

Linear representations have been considered previously by other authors. Sheikhzadeh and Deng [17] apply hidden filter models directly to acoustic waveforms, avoiding artificial frame boundaries and therefore allowing better modelling of short duration events. They consider consonant-vowel classification and illustrate the importance of power normalisation in the waveform domain, although a full implementation of the method and tests on benchmark tasks remain to be explored. Poritz [18], and later Ephraim and Roberts [19], consider modelling speech explicitly as a time series using autoregressive hidden Markov models. Their work inspired more recent advances in that direction by Mesot and Barber [20] who develop a switching linear dynamical systems (SLDS) framework. The SLDS approach exhibited significantly better performance at recognising spoken digits in additive Gaussian noise when compared to standard hidden Markov models (HMMs) used in combination with cepstral features; however, it is computationally very expensive even when approximate inference techniques are used. Turner and Sahani proposed using modulation cascade processes to model natural sounds simultaneously on many time-scales [21], but the application of this approach to ASR remains to be explored. In this paper we do not directly use the time series interpretation and impose no temporal constraints on the models. Instead, we investigate the effectiveness of the acoustic waveform front-end for robust phoneme classification and recognition using Gaussian mixture models (GMMs), as those models are commonly used in conjunction with HMMs in practical systems.

To assess the merits of representing speech without information loss and non-linear transformation, without the potentially interfering effects of segmentation, in Section II we first develop Gaussian mixture models for fixed-length segments of speech and use them for phoneme classification in the presence of additive noise. Next in Section III we investigate the effect of the segment duration on classification error. In the same section we compare the new high-dimensional representation with PLP features in a phoneme classification task that uses information from entire phonemes. Finally in Section IV we consider phoneme recognition from continuous speech in the standard HMM-GMM framework. In all scenarios that we investigate we find that PLP and MFCC features excel in low-noise conditions while high-dimensional linear representation achieve better results at higher noise levels, starting already at around 18dB SNR. We then demonstrate that a combination of high-dimensional linear and cepstral features achieves better results than either of the individual representations across all noise levels. Recently Power Normalised Cepstral Coefficients (PNCC) features [22] were proposed for robust speech recognition, but we became aware of these features at the time of the submission of this paper, hence these features are not considered in this study.

II. Models of Fixed-Length Segments of Speech

Generative classifiers use probability density estimates, $p(x)$, learned for each class of the training data. The predicted class of a test point, x, is determined as the class k with the greatest likelihood evaluated at x. Typically the log-likelihood, $\mathcal{L}(x) = \log(p(x))$, is used for the calculation. A test point x is thus assigned to one of K classes using the following function:

$$\mathcal{A}^L(x) = \arg \max_{k=1,\dots,K} \mathcal{L}^{(k)}(x) + \log(\pi_k) \ . \qquad (1)$$

The inclusion above of π_k, the prior probability of class k, means that we are effectively maximising the log-posterior probability of class k given x. In this section we build probabilistic models of fixed-length segments of acoustic waveforms of phoneme classes. Each waveform segment x is thus a vector in $\mathbb{R}^d$, where d is the number of time samples in the segment.

A. Exploratory Data Analysis

Towards constructing probabilistic models of high-dimensional speech representations in the acoustic waveform domain, it is of interest to investigate possible lower dimensional structures in the phoneme classes. Supposing that such structure exists and can be characterised then it could be used to find better representations for speech, and to construct more accurate probabilistic models. We thus initially deployed data-driven methods for learning possible low-dimensional manifolds, as explored in [23], [24], including locally linear embedding [25] and Isomap [26]. Many speech representations, typically some variant of MFCC or PLP, reduce the dimension of speech signals using non-linear processing. Those methods do not directly incorporate information about the structure of the phoneme class distributions, but model the properties of speech perception. The involved non-linear processing, however, makes exact noise adaptation very challenging (see Section II-D). Instead one would aim to find non-linear low-dimensional structures in the phoneme distributions, and exploit this information to build better models that remain defined in the original high-dimensional space. This could include Gaussian process latent variable models [27] (GP-LVM), which require as input an estimate of the dimension of the non-linear feature space. However, while we found that intrinsic dimension estimates suggest the existence of low dimensional nonlinear structures in the phoneme distributions, our investigation also showed that these structures still had sufficiently many dimensions to make it impractical to sample them adequately for ASR purposes [1]. Hence we turn to more generic density models. In particular we will construct generative models in the high-dimensional space which do not attempt to exploit any submanifold structure directly. We will see that approximations are still required, again due to the sparseness of the data and because of computational constraints.

B. Gaussian Mixture Models

Without assuming any additional prior knowledge about the phoneme distributions, we use Gaussian mixture models (GMMs) to model phoneme densities. For the case of c mixture components this function has the form:

$$p(x) = \sum_{i=1}^{c} \frac{w_i}{(2\pi)^{\frac{d}{2}} |\mathbf{\Sigma}_i|^{\frac{1}{2}}} \exp\left[-\frac{1}{2}(x-\mu_i)^{\mathrm{T}} \mathbf{\Sigma}_i^{-1}(x-\mu_i)\right] \quad (2)$$

where w_i, μ_i and $\mathbf{\Sigma}_i$ are the weight, mean vector and covariance matrix of the i^{th} mixture component respectively. We additionally impose a zero mean constraint for models as a waveform x will be perceived the same as $-x$. With this constraint, the corresponding models represent all information about the phoneme distributions via the covariance matrices and component weights.

Reliable estimation of full covariance matrices is challenging even in the case of standard low-dimensional features when the number of components in the mixture becomes large. The problem is even more pronounced in the case of high-dimensional waveform features, where at 16kHz already with 70ms segments the dimension of the feature space becomes $d = 1020$. We therefore considered using density estimates derived from mixtures of probabilistic principal component analysis (MPPCA) [28], [29]. This method produces a Gaussian mixture model where the covariance matrix of each component is regularised by replacement with a rank-q approximation, where $q < d$. However, even PPCA requires an excessive number of parameters compared to the typical amount of available data [29]. Finally GMMs with diagonal covariance matrices were investigated, a common modelling approximation when training data is sparse, used also in state-of-the-art ASR systems for modelling distributions of cepstral features. This modelling approach achieved lowest classification error, so in all experiments reported in this paper GMMs with diagonal covariance matrices are used.

Diagonal covariance matrices will be a good approximation provided the data is presented in a basis where correlations between features are weak. For the acoustic waveform representation, this is clearly not the case on account of the strong temporal correlations in speech waveforms. The density model used for the phoneme classes thus becomes:

$$p(x) = \sum_{i=1}^{c} \frac{w_i}{(2\pi)^{\frac{d}{2}} |\mathbf{D}_i|^{\frac{1}{2}}} \exp\left[-\frac{1}{2}(x-\mu_i)^{\mathrm{T}} \mathbf{C}^{\mathrm{T}} \mathbf{D}_i^{-1} \mathbf{C}(x-\mu_i)\right] \quad (3)$$

where $\mathbf{C}$ is an orthogonal transformation selected to decorrelate the data at least approximately, while w_i, μ_i and $\mathbf{D}_i$ are the weight, mean vector and diagonal covariance matrix of the i^{th} mixture component, respectively. We systematically investigated candidate decorrelating bases derived from PCA, wavelet transforms and DCTs. Although the optimal basis for decorrelation on the training set is indeed formed by the phoneme-specific principal components, we found that the lowest test error is achieved with a DCT basis. Preliminary experiments with acoustic waveforms showed that, instead of performing a DCT on an entire phoneme segment, it is advantageous to separate DCTs into non-overlapping sub-segments. We systematically investigated different block lengths, and found that best classification results were obtained with 10 ms blocks, mirroring (except for the lack of overlaps) the frame decomposition of MFCC and PLP. For a sampling rate of 16 kHz, the transformation matrix $\mathbf{C}$ is then block diagonal consisting of 160×160 DCT blocks.

While the DCT in the density modelling in (3) was motivated by the need to decorrelate the data, and thus make the approximation of covariance matrices by diagonal ones more accurate, one may view the end result as a representation of speech waveforms via some form of short-time spectra, analogous to what is done towards extracting cepstral features. There is however a fundamental difference between the block-DCT transform in (3) and the short-time magnitude spectra used to derive PLP and MFCC features: the former is an orthonormal transform, that is, just a rotation of the coordinate system that preserves all the information, whereas the latter is a non-linear transform that incurs a dramatic information loss. In particular, while the discrete Fourier transform is also an orthogonal transform, retaining only its magnitude – as PLP and MFCC do – is equivalent to mapping at each discrete frequency a whole circle in $\mathbb{R}^2$ to just the value of its radius.

GMMs as given in (3) are also used for the MFCC and PLP features, except that component means are not constrained to be zero, and $\mathbf{C}$ is chosen to be the identity matrix, since these features already involve some form of DCT and are approximately decorrelated.

C. Model Average

In general, more variability of the training data can be captured with an increased number of mixture components; however, if too many components are used over-fitting will occur. The best compromise is usually located by cross validation using the classification error on a development set. The result is a single value for the number of components required. However, we observed that the optimal number of components decreases with SNR, hence we use an alternative approach and take the model average over the number of components, effectively a mixture of mixtures [30]; this gave consistent classification improvements over individual mixtures across all noise levels. Thus, we start from a selection of models parameterised by the number of components, c, which takes values in $\mathcal{C} = \{1, 2, 4, 8, 16, 32, 64, 128\}$ or subsets of it. The entries in this set are uniformly distributed on a log scale to give a good range of model complexity without including too many of the complex models. We compute the model average log-likelihood $\mathcal{M}(x)$ as:

$$\mathcal{M}(x) = \log\big(\sum_{c\in\mathcal{C}} u_c \exp(\mathcal{L}_c(x))\big) \tag{4}$$

with the model weights $u_c = \frac{1}{|\mathcal{C}|}$ and $\mathcal{L}_c(x)$ being the log-likelihood of x given the c-component model. Alternatively the mixture weights u_c could be determined from the posterior densities of the models on a development set to give a class dependent weighting, i.e.

$$u_c = \frac{\sum_{x\in\mathcal{D}} \exp(\mathcal{L}_c(x))}{\sum_{d\in\mathcal{C}} \sum_{x\in\mathcal{D}} \exp(\mathcal{L}_d(x))} \tag{5}$$

where $\mathcal{D}$ is a development set. Preliminary experiments suggested that using those posterior weights only gives a slight improvement over $u_c = \frac{1}{|\mathcal{C}|}$. We therefore adopt those uniform weights for all results shown in this paper.

D. Noise adaptation

The primary concern of this paper is to investigate the performance of the trained classifiers in the presence of additive noise. Throughout this study, when noise is present in testing, it is assumed that it can be modelled by a Gaussian distribution and that the covariance matrix is known exactly or as an estimate. This assumption may not be valid for real world scenarios but is a good approximation, providing the noise is stationary at the phoneme level. In practice the noise variance would have to be estimated from the input signal and there are many methods available that provide good estimates of the SNR [31], [32]. Two noise types are studied: white Gaussian noise, where the variance is known exactly, and pink noise extracted from NOISEX-92 [33] data base, which is not Gaussian and therefore tests the robustness of the system in the case where the Gaussian noise assumption does not hold. The distribution of the noise is then estimated via a Gaussian model that is later used for noise adaptation. Generative classification is particularly suited for achieving robustness as the estimated density models can capture the distribution of the noise corrupted phonemes. As the noise is additive in the acoustic waveform domain, signal and noise models can be specified separately and then combined exactly by convolution.

We consider the case where the SNR is set at sentence-level. All sentences will therefore be normalised to have unit energy per sample in quiet and noisy conditions. Different phonemes within these sentences can have higher or lower energies, as reflected in the density models by covariance $\mathbf{D}$ with trace above or below d, where d is the dimension of the feature vectors. The relative energy of each phoneme class is thus implicitly used during classification.

The adaptation to noise of phonetic classes in the acoustic waveform domain is performed by replacing each covariance matrix $\mathbf{D}$ with $\tilde{\mathbf{D}}(\sigma^2)$:

$$\tilde{\mathbf{D}}(\sigma^2) = \frac{\mathbf{D} + \sigma^2\mathbf{N}}{1+\sigma^2} \tag{6}$$

where $\mathbf{N}$ is the covariance matrix of the noise transformed by $\mathbf{C}$, normalised to have trace d, and σ^2 is the noise variance. For white noise, $\mathbf{N}$ is the identity matrix, otherwise it is estimated empirically from noise samples. In general a full covariance matrix will be required to specify the noise structure. However, with a suitable choice of $\mathbf{C}$ the resulting $\mathbf{N}$ will be close to diagonal, and indeed when $\mathbf{C}$ is a segmented DCT we find this to be true in our experiments with pink noise. To avoid the significant computational overheads of introducing non-diagonal matrices, we therefore retain only the diagonal elements of $\mathbf{N}$.

The normalisation by $1+\sigma^2$ arises as follows. Considering that sentences are normalised to unit energy per sample, a vector x containing d samples from the sentence has squared norm d. A vector n of d noise samples, of variance σ^2, will have average squared norm $E(\|n\|^2) = d\sigma^2$. Because the noise is assumed Gaussian, it can be shown that the fluctuations of $\|n\|^2$ away from its average are small, of relative order $1/\sqrt{d}$, or order $\sqrt{d}$ overall [34]. The cross-term in the squared norm of the noisy sentence vector,

$$\|y\|^2 = \|x+n\|^2 = \|x\|^2 + 2x^T n + \|n\|^2$$

is easily shown to be also of order $O(\sqrt{d})$ [34]. Thus altogether, $\|y\|^2 = d + d\sigma^2 + O(\sqrt{d})$, so that the energy per sample is

$$\frac{\|y\|^2}{d} = 1 + \sigma^2 + O\left(\frac{1}{\sqrt{d}}\right) \ .$$

For large, d as in our case, normalising y to unit energy per sample is therefore equivalent to rescaling by $1/\sqrt{1+\sigma^2}$. When this rescaling is applied to the noisy phoneme models, it gives precisely the normalisation factor in (6).

There is no exact method for combining models of the training data with noise models in the case of MFCC and PLP features, as these representations involve non-linear transforms of the waveforms. Cepstral mean and variance normalisation (CMVN) [35] is an approach commonly used in practice to compensate noise corrupted features. This method requires estimates of the mean and variance of the features, usually calculated sentence-wise on the test data or with a moving average over a similar time window. We take this to be a realistic baseline. Alternatively the required statistics can be estimated from a training set that has been corrupted by the same type and level of noise as used in testing. Clearly both approaches have merit. For example, sentence level CMVN requires no direct knowledge of the test conditions, and can remove speaker specific variation from the data. The estimates will be less accurate and as a consequence it is difficult to standardise all components in long feature vectors obtained by concatenating frames; instead, we considered standardising frame by frame. Using a noisy training set for CMVN requires that the test conditions are known so that data can be either collected or generated for training under the same conditions. The feature means and variances can be obtained accurately, and in particular we can standardise longer feature vectors. However, as the same standardisation is used for all sentences, any variation due to individual speakers will persist. We found that standardisation on the noisy training set gives lower error rates both in quiet conditions and in noise [34], hence all results for CMVN given below use this method.

At this exploratory stage, we study also the matched-condition scenario, where training and testing noise conditions are the same and a separate classifier is trained for each noise condition. While in practice it would be impossible to have a distinct classifier for every noise condition, matched conditions are nevertheless useful in our exploratory classification experiments: because training data comes directly from the desired noisy speech distribution, then assuming enough data is available to estimate class densities accurately this approach provides the optimal baseline for all noise adaptation methods [36], [37].

E. Classification Results

For all experiments reported in this paper realisations of phonemes were extracted from the SI and SX sentences of the TIMIT database. The training set consists of 3,696 sentences sampled at 16 kHz. Noisy data is generated by applying additive noise at nine SNRs. Recall that the SNRs were set at the sentence level, therefore the local SNR of the individual phonemes may differ significantly from the set value, causing mismatch in the classifiers. In total ten testing and training conditions were run: -18dB to 30dB in 6dB increments and quiet (Q). Following the extraction of the phonemes there are a total of 140,225 phoneme realisations. The glottal closures are removed and the remaining classes are then combined into 48 groups in accordance with [38], [39]. Even after this combination some of the resulting groups have too few realisations. To ensure that the training procedure is stable, the smallest groups with fewer than 1,500 realisations were increased in size by the addition of temporally shifted versions of the data, *i.e.* if x is an example in one of the small training classes then the phoneme segments extracted from positions shifted by $k = -100, -75, -50, \ldots, 75, 100$ samples were also included for training. For the purposes of calculating error rates, some very similar phoneme groups are further regarded as identical, resulting in 39 groups of effectively distinguishable phonemes [38]. PLP and MFCC features are obtained in the standard manner from frames of width 25 ms, with a shift of 10 ms between neighbouring frames. Standard implementations [40] of MFCC and PLP with default parameter values are used to produce a 13-dimensional feature vector from each time frame. The inclusion of dynamic Δ+$\Delta\Delta$ features [41] increases the dimension to 39.

For the MFCC and PLP representations, we consider in this experiment the five frames closest to the centre of each phoneme, covering 65 ms, and concatenate their feature vectors. Results are shown for the representations with and those without Δ+$\Delta\Delta$, giving feature vector dimensions of $5 \times 39 = 195$ and $5 \times 13 = 65$, respectively. The acoustic waveform representation is obtained by dividing each sentence into a sequence of 10ms non-overlapping frames, and then taking the seven frames (70ms) closest to the centre of each phoneme, resulting in a 1120-dimensional feature vector. Each frame is individually processed using the 160-point DCT. We present results for white and pink noise and will see that the approximation using diagonal covariances $\mathbf{D}$ in the DCT basis is sufficient to give good performance. The impact of the number of frames included in the MFCC, PLP and acoustic waveform representations is investigated in the next section.

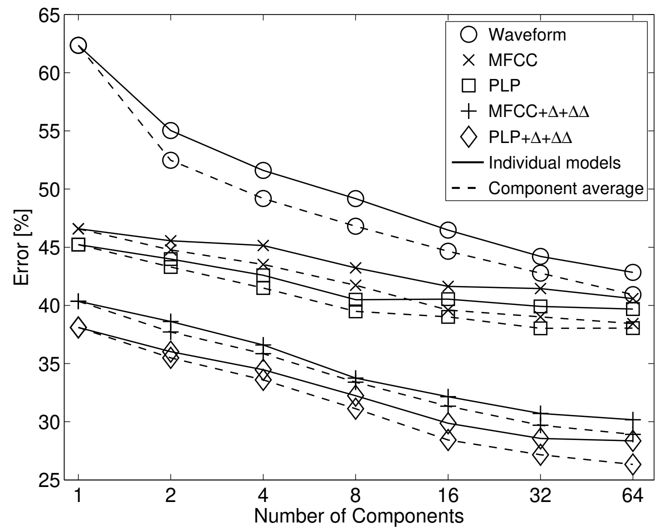

Fig. 1. *Model averaging for acoustic waveforms, MFCC and PLP models, all trained and tested in quiet conditions. Solid: GMMs with number of components shown; dashed: average over models up to number of components shown. The model average reduces the error rate in all cases.*

We comment briefly on the results for individual mixtures, i.e. with a fixed number of components. Gaussian mixture models were trained with up to 64 components for all representations. Typically performance on quiet data improved with the number of components, although this has significant cost for both training and testing. The optimal number of components for MFCC and PLP models in quiet conditions was 64, the maximum considered here. However, in the presence of noise the lowest error rates were obtained with few components; typically there was no improvement beyond four components. As explained in Section II-C, rather than working with models with fixed numbers of components, we average over models in all the results reported below. Figure 1 shows that the improvement obtained by this in quiet conditions is approximately 2% for both acoustic waveforms and PLP with a small improvement seen for MFCC also. The model average similarly improved results in noise.

One set of key results comparing the error rates in white Gaussian noise for phoneme classification in the three domains is shown in Figure 2. The MFCC and PLP classifiers are adapted to noise using CMVN. This method is comparable with the adapted waveform models as it only relies on the models trained in quiet conditions. The curve for acoustic waveforms is for models trained in quiet conditions and then adapted to the appropriate noise level using (6). Comparing waveforms first to MFCC and PLP without Δ+$\Delta\Delta$, we see that in quiet conditions the PLP representation gives the lowest error. The error rates for MFCC and PLP are significantly worse in the presence of noise, however, with acoustic waveforms giving an absolute reduction in error at 0dB SNR of 40.6% and 41.9% compared to MFCC and PLP, respectively. Curves are also shown for MFCC+Δ+$\Delta\Delta$ and PLP+Δ+$\Delta\Delta$, although these representations include information about underlying speech signals over 145 ms observation windows, *i.e.* significantly longer than the 70 ms observation windows used to form the acoustic waveform representation. Again the same trend holds; performance is good in quiet conditions but quickly deteriorates as the SNR decreases. The crossover point is between 24 dB and 30 dB SNR for both cepstral representations. The chance-level error rate of 93.5% can be seen below 0 dB SNR for the MFCC and PLP representations without deltas and below -6 dB SNR when deltas are included, whereas the acoustic waveform classifier performs significantly better than chance with an error of 76.7% even at -18 dB SNR. The dashed curves in Figure 2 represent the error rates obtained for classifiers trained in matched conditions. The results show that the waveform classifier compares favourably to MFCC and PLP below 24dB SNR when no deltas are appended. This results is in agreement with our working hypothesis that high-dimensional acoustic waveform representations may provide better separation of phoneme classes. Including Δ+$\Delta\Delta$ does reduce the error rates significantly and the crossover then occurs between 0 dB and 6 dB SNR. It is these observations that mainly motivate our further model developments below: clearly we should aim to include information similar to deltas in the waveform representation.

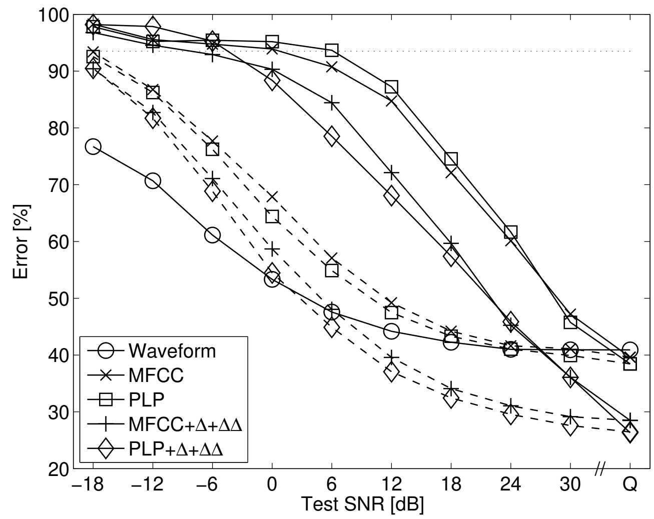

Fig. 2. *Results of phoneme classification using fixed-length segments in white Gaussian noise. All classifiers use the model average of mixtures up to 64 components. Dotted line indicates chance level at 93.5%. When the SNR is less that 24dB, acoustic waveforms are the significantly better representation, with an error rate below chance even at -18dB SNR. Dashed curves show results of matched training for corresponding MFCC and PLP representations.*

The same experiment was repeated using pink noise extracted from the NOISEX-92 database [33]. The results for both noise types were similar for the waveforms classifiers. For PLP+Δ+$\Delta\Delta$, adapted to noise using CMVN, there is a larger difference between the two noise types, with pink noise leading to lower errors. Nevertheless, the better performance is achieved by acoustic waveforms below a crossover point between 18 dB and 24 dB SNR [34]. Results for GMM classification on the TIMIT benchmark in quiet conditions have previously been reported in [39], [42] with errors of 25.9% and 26.3% respectively. To ensure that our baseline is valid we compared our experiment in quiet conditions for PLP+Δ+$\Delta\Delta$ and obtained a comparable error rate of 26.3% as indicated in the bottom right corner of Figure 2.

Following these encouraging results we seek to explore the effect of optimising the number of frames and the inclusion of information from the entire phoneme. The expectation is that including more frames in the concatenation for acoustic waveforms will have a similar effect to adding Δ+$\Delta\Delta$ for MFCC and PLP features. A direct analogue of deltas is unlikely to be useful for waveforms: MFCC and PLP are based on log magnitude spectra that change little during stationary phonemes, so that local averaging or differencing is meaningful. For waveforms, where we effectively retain not just Fourier component amplitudes but also phases, these phases combine essentially randomly during averaging or differencing, rendering the resulting delta-like features useless.

III. Segment Duration, Variable Duration Phoneme Mapping and Classifier Combination

A. Segment Duration

Ideally all relevant information should be retained by our phoneme representation, but as it is difficult to determine exactly which information is relevant we initially chose to take f consecutive frames closest to the centre of each phoneme and concatenate them. Whilst the precise number of frames required for accurate classification could in principle be inferred from the statistics of the phoneme segment durations, those durations not only vary significantly between classes but also the standard deviation within each class is considerable. Since no single value of f will be optimal for all phoneme classes, we consider the sum of the mixture log-likelihoods $\mathcal{M}_f$, as defined in (4), but now indexed by the number of frames used. The sum is taken over the set $\mathcal{F}$ which contains the values of f with the lowest corresponding error rates, for example $\mathcal{F} = \{7, 9, 11, 13, 15\}$ for PLP, giving:

$$\mathcal{R}(\bar{x}) = \sum_{f \in \mathcal{F}} \mathcal{M}_f(x^f) \tag{7}$$

where $\bar{x} = \{x^f | f \in \mathcal{F}\}$, with x^f being the vector with f frames. Note that we are adding the log-likelihoods for different f, which amounts to assuming independence between the different x^f in $\bar{x}$. Clearly this is an imperfect model, as *e.g.* all components of x^7 are also contained in x^{11} and so are fully correlated, but our experiments show that it is useful in practice. We also implemented the alternative of concatenating the x^f into one longer feature vector and then training a joint model on this, but the potential benefits of accounting for correlations are far outweighed by the disadvantages of having to fit density models in higher dimensional spaces. Consistent with the independence assumption in (7), in noise we adapt the models $\mathcal{M}_f$ separately and then combine them as above. The same applies to the further combinations discussed next.

B. Sector sum

Accounting for the information from the entire phoneme using the standard HMM-GMM framework is considered in the next section in the context of phoneme recognition. Here we are interested in assessing relative merits of waveform and cepstral representations in robust phoneme classification independently of segmentation errors and interfering effects of HMM assumptions. For that purpose, one can consider averaging of frames across the entire phoneme [42] or other methods for mapping variable phoneme duration to fixed-length representation, as proposed in [43], [44]. Towards GMM models which will be used in the HMM-GMM framework in the next section, here we extend the centre-only frame concatenation to use information from the entire phoneme by taking f frames with centres closest to each of the time instants A, B, C, D and E that are distributed along the duration of the phoneme as shown in Figure 3. In this manner the representation consists of five sequences of f frames per phoneme. Those sets of frames are then concatenated to give five vectors x_A, x_B, x_C, x_D and x_E. We train five models on those sectors and then combine the information they provide by taking the sum of the corresponding log-likelihoods:

$$\mathcal{S}(\hat{x}) = \sum_{s\in\{A,B,C,D,E\}} \mathcal{M}_s(x_s) \tag{8}$$

where $\hat{x} = \{x_A, x_B, x_C, x_D, x_E\}$ and $\mathcal{M}_s$ denotes the model for sector s, using some fixed number of frames f. Both improvements can be combined by taking the sum of the f-averaged log-likelihoods, $\mathcal{R}_s(\bar{x}_s)$, over the five sectors s:

$$\mathcal{T}(\hat{\bar{x}}) = \sum_{s\in\{A,B,C,D,E\}} \mathcal{R}_s(\bar{x}_s) \tag{9}$$

where $\bar{x}_s = \{x_s^f | f \in \mathcal{F}\}$ with x_s^f being the vector with f frames centred on sector s, and $\hat{\bar{x}}$ gathers all $\bar{x}_s$. Given the functions derived above, the class of a test point can be predicted using one of the following:

$$\mathcal{A}_f^M(x) = \arg\max_{k=1,\ldots,K} \mathcal{M}_f^{(k)}(x) + \log(\pi_k) \tag{10}$$

$$\mathcal{A}^R(\bar{x}) = \arg\max_{k=1,\ldots,K} \mathcal{R}^{(k)}(\bar{x}) + \log(\pi_k) \tag{11}$$

$$\mathcal{A}_f^S(\hat{x}) = \arg\max_{k=1,\ldots,K} \mathcal{S}_f^{(k)}(\hat{x}) + \log(\pi_k) \tag{12}$$

$$\mathcal{A}^T(\hat{\bar{x}}) = \arg\max_{k=1,\ldots,K} \mathcal{T}^{(k)}(\hat{\bar{x}}) + \log(\pi_k) \tag{13}$$

where π_k is the prior probability of predicting class k.

C. Classification Results

Figure 4 shows the impact of the number of frames concatenated from each sector and the impact of summing log-likelihoods over the five sectors on the classification error, focusing on quiet conditions. Since PLP features gave slightly lower error than MFCC features in the experiments using centre-only representation, in this section we show results for PLP features only. We see that the best results for acoustic waveform classifiers are achieved around 9 frames, and around 15 frames for PLP without deltas. The PLP+Δ+$\Delta\Delta$ features are less sensitive to the number of frames with little difference in error from 1 to 15 frames. It is worth noting that if we consider the best results obtained for PLP without deltas, 22.4% using 15 frames, with the best for PLP+Δ+$\Delta\Delta$, 19.6% with 9 frames, then the performance gap of 2.8% is much smaller than if we were to compare error rates where both classifiers used the same number of frames. The error rates obtained using the f-average over the five best values of f are 32.1%, 21.4% and 18.5% for acoustic waveforms, PLP and PLP+Δ+$\Delta\Delta$, respectively. Table I shows the absolute percentage error reduction for each of the four classifiers (10)–(13) in quiet conditions, compared to the GMM with the single best number of mixture components and number of frames f. The relative benefits of the f-average and the sector sum are clear. The sector sum gives the bigger improvements on its own in all cases compared to only the f-average, but the combination of the two methods is better still throughout. The same qualitative trend holds true in noise.

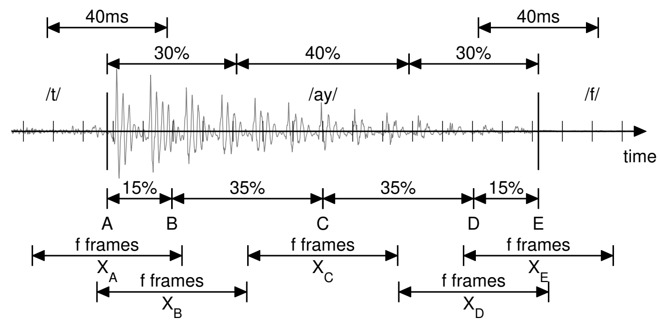

Fig. 3. *Comparison of phoneme representations. Top: Division described in [42] resulting in five sectors, three covering the duration of the phoneme and two of 40ms duration around the transitions. Bottom: f frames closest to the five points A, B, C, D and E (which correspond to the centres of the regions above) are selected to map the phoneme segment to five feature vectors x_A, x_B, x_C, x_D and x_E.*

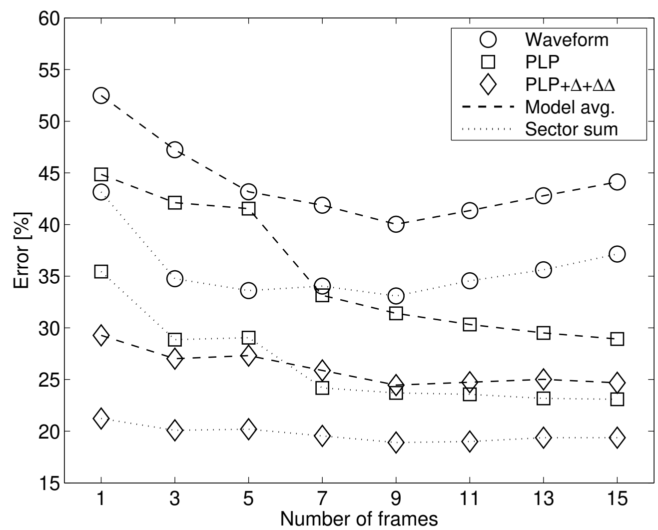

Fig. 4. *Error rates of the different representations in quiet conditions, as a function of f, the number of frames considered. Dashed: prediction (10) using only the central sector. Dotted: prediction (11) using the sum over all five sectors, leading to a clear improvement in all cases.*

Figure 5 compares the performance of the final classifiers, including both the f-average and the sector sum, on data corrupted by pink noise. The averaging over the number of frames is done for the five values of f achieving lowest errors, as shown in Figure 4. The solid curves give the results for the acoustic waveform classifier adapted to noise using (6), and for the PLP classifier with and without Δ+$\Delta\Delta$ trained in quiet conditions and adapted to noise by CMVN. The errors are generally significantly lower than in Figure 2, showing the benefits of f-averaging and the sector sum. PLP+Δ+$\Delta\Delta$ remains the best representation for very low noise, but waveforms give lower errors below crossover points around 18dB SNR. As before, they also perform better than chance down to -18dB SNR. The dashed lines in Figure 5 show for comparison the performance of PLP classifiers trained in matched conditions. The matched conditions PLP+Δ+$\Delta\Delta$ classifier has the best performance for all SNR. However, in high noise the adapted acoustic waveform classifier is significantly closer to matched PLP+Δ+$\Delta\Delta$ than PLP+Δ+$\Delta\Delta$ with CMVN. As explained, the CMVN and matched curves for PLP provide the extremes between which we would expect a PLP classifier to perform if model adaption analogous to that used with the acoustic waveforms was possible, or some other method to improve robustness was employed. In the next section, we explore the performance of MFCC features in combination with the ETSI advanced front-end (AFE) [45] and vector Taylor series [46] noise compensation techniques in the context of phoneme recognition.

TABLE I
Absolute reduction in percentage error for each of the classifiers (10)–(13) in quiet conditions.

Model	Waveform	PLP	PLP+Δ+$\Delta\Delta$
Model average ($\mathcal{A}^M$)	1.6	2.8	4.4
f-average ($\mathcal{A}^R$)	5.6	6.0	6.3
Sector sum ($\mathcal{A}^S$)	6.7	8.4	8.7
f-average + sector sum ($\mathcal{A}^T$)	9.9	10.0	10.4

D. Combination of PLP and Acoustic Waveform Classifiers

We see from the results shown so far that, as in the preliminary experiments, PLP performs best in quiet conditions with acoustic waveforms being more robust to additive noise. To gain the benefits of both representations, we explore merging them via a convex combination of the corresponding log-likelihoods, parameterised by a coefficient α:

$$\mathcal{T}_\alpha(x) = (1-\alpha)\mathcal{T}_{\rm plp}(x) + \alpha\mathcal{T}_{\rm wave}(x) \qquad (14)$$

where $\mathcal{T}_{\rm plp}(x)$ and $\mathcal{T}_{\rm wave}(x)$ are the log-likelihoods of a point x corresponding to PLP and waveform features, respectively. $\mathcal{T}_\alpha(x)$ is then used in place of $\mathcal{T}(x)$ in (13) to predict the class. We would expect optimal α to be almost zero for high SNRs and close to one for low SNRs. In order to achieve the desired improvement in accuracy, we fit an appropriate combination function $\alpha(\sigma^2)$. To that end, a suitable range of possible values of α was identified at each noise level from the condition that the error rate is no more than 2% above the error for the best α. This range is broad, so the particular form of the fitted combination function is not critical [47]. We choose the following sigmoid function with two parameters σ_0^2 and β:

$$\alpha(\sigma^2) = \frac{1}{1+e^{\beta(\sigma_0^2-\sigma^2)}} \; . \qquad (15)$$

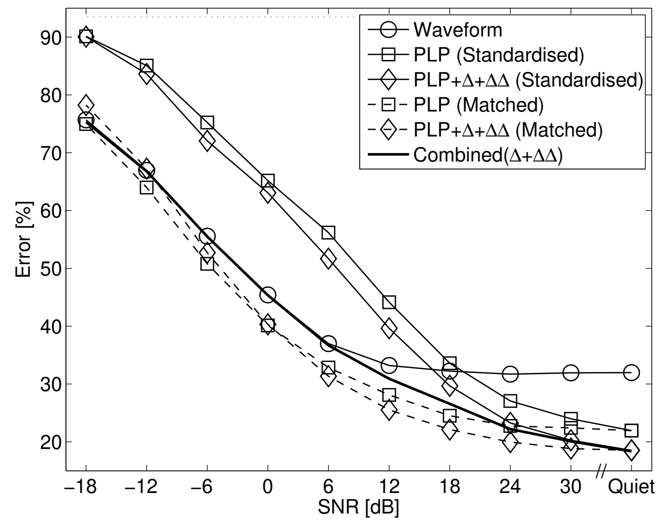

Fig. 5. *Performance of the classifiers in pink noise. Error curves shown are obtained using the sector sum and averaging over the number of frames for five values of f achieving lowest classification errors according to the results shown in Figure 4. Bold line: combined waveform and PLP+Δ+$\Delta\Delta$ classifier, with the latter adapted to noise by feature standarisation using CMVN.*

A fit through the numerically determined suitable ranges of α then gives $\sigma_0^2 = 17$ dB and $\beta = 0.3$. Note that this approach is equivalent to a multistream model, where each sector and value of f is an independent stream. The error of the combined classifier using models trained in quiet conditions is shown as the bold curve in Figure 5. In quiet conditions the combined classifier is slightly more accurate (18.4%) than PLP+Δ+$\Delta\Delta$ alone, corresponding to a small value of $\alpha = 0.003$. When noise is present the combined classifier is at least as accurate as the acoustic waveform classifier, and significantly better around 18 dB SNR. The combined classifier does improve upon PLP+Δ+$\Delta\Delta$ classifiers trained in matched conditions at very low SNR and narrows the performance gap to the order of no more than 9% throughout, rather than 22% when comparing to PLP+Δ+$\Delta\Delta$ adapted by CMVN.

IV. Phoneme Recognition from Continuous Speech

We now consider extending the classification results from the previous sections to the task of phoneme recognition from continuous speech using hidden Markov models. As GMMs are used for classification as well as for the emission density models of the HMMs, our developments so far can be transferred to recognition. The only exception is the sector sum, which was only intended to mimic the states of HMMs. The model average and frame average, on the other hand, remain suitable for HMMs. Most importantly, the noise adaptation for classifiers in the acoustic waveform domain given by equation (6) can be directly applied to provide a good model of continuous noisy speech, while the transition matrices of the HMM will remain unaltered in noise based on the assumption that the noise and speech are independent.

A. HMM Training

In this study we used the Hidden Markov Model Toolkit (HTK) [48] to model and recognise the phonemes. This is a flexible implementation that allows comparing results obtained using the standard representations to those obtained with the acoustic waveforms. However, a number of modifications to the standard training procedure are required to make the system work properly with acoustic waveform models.

The standard training process implements a flat-start procedure, which begins with single component GMMs as emission models of HMMs. The first stage initialises all the prototype emission models to the global statistics of the input frames from all sentences and all phonemes. Next the process updates the emission models for each state on a class-wise basis, where the frames from the training set belonging to each class are used to estimate the density of all states for that class. The final stage aims to maximise the training objective function by re-estimating all parameters of the HMM including the emission models and state transition matrices. These models are then further improved by splitting each emission model component to double the number of components. This process replaces each component by two new components, where the means of each component have been moved apart along the axis of maximum variance in the parent component. Following a split, a number of re-estimation iterations are performed. The splitting and re-estimation stages are repeated until each emission model consists of 128 components. Beyond this number of components, the error rates increase due to overfitting.

When using HTK with acoustic waveforms, the training objective needs to be selected first. There is a choice of maximum likelihood estimation, maximum mutual information, minimum phone error or large margin methods [39]. In order to provide the best adaptation to noise, maximum likelihood was used as this gives models that should best reflect the true distribution of the phonemes. The other methods may provide better performance in quiet or matched conditions, but it is not clear how optimising the phoneme class models in such manner will affect the noise adaptation, since models trained in this way will not be true generative models of the observed speech and so the noise adaptation of (6) may fail to produce an accurate model of the noisy speech.

Additionally, HTK prunes potential alignments with likelihoods that are too low relative to the best sequence. If no alignment for a particular sentence can be found with sufficiently high likelihood, then the sentence is excluded from the training dataset as this exclusion will reduce the training time by ignoring training data and alignments that have very low likelihood given the current model parameters. Due to the higher dimensionality of the feature vectors for the acoustic waveform representation, the likelihood of potential alignments can cover a much wider range than it does in the case of standard cepstral representations. When the default pruning parameter value – indicating the allowable tolerance on likelihood values – is used to train acoustic waveform models, the majority of the data is rejected. Consequently, the pruning parameter was increased significantly. A small fraction of data was still rejected, but the training process was faster than if the pruning option was disabled.

The training initialisation also needs to be adjusted to ensure convergence towards adequate models for acoustic waveforms. Initially the standard training method with the flat start was used, but the results obtained using this training procedure were poor for acoustic waveforms. We traced the problem to the component splitting stage, where the means of the split Gaussian components are moved apart. For acoustic waveform distributions, however, zero mean GMMs provide better models, since a waveform x and $-x$ are identically perceived. Following this, a passive method was considered to enforce the zero mean constraint, by including negative instances of the training data. However, the model components still had means that differed significantly from zero. To overcome this problem, the means were instead constrained to zero after splitting and then an update was applied to all parameters except for the mean vectors. This provided a modest improvement, but it appears that the splitting process is the major issue in training acoustic waveform models as the mean vectors will be displaced. Ultimately the best results are obtained when the GMMs trained on regions x_B, x_C and x_D shown in Figure 3 are used to initialise the emission models for the three states of each phoneme. This initialisation is repeated for each number of components and therefore avoids any splitting of components as the required number of components is specified during the training of the GMMs. Following the HMM initialisation, the variances, transition matrices and component weights are re-estimated using the standard Baum-Welch method.

While model averaging according to (4) is still a valid concept, it was not used in the phoneme recognition experiments reported in the next subsection, because in the HMM framework the likelihood of the data must be evaluated on all models simultaneously, and this would increase the computational load significantly. We would expect reductions in error rates if we did implement model averaging for HMMs, but the improvement would be similar for all representations as observed for the GMM results and would not lead to different conclusions about comparisons of the representations.

As we saw in the previous section, the number of frames, f, used for the acoustic waveform representation was also critical for classification. Again a range of values for f were investigated to find those that are most suitable for recognition. This was necessary as the optimal feature duration of acoustic waveform representations is not known, however the parameters used for the standard representations and the results of the classification experiments of the previous sections provided an initial range for investigation.

B. *Recognition results*

Figure 6 shows the phoneme recognition error rates for HMMs using acoustic waveforms, for the four training-initialisation approaches detailed in Section IV-A. This investigation considers the performance of HMMs, when 1, 2, 4, 8, 16, 32 and 64 component GMMs are used as emission density models. Evidently, the standard training procedure is ineffective for acoustic waveforms. Notice that the inclusion of negative instances of the training data fails to improve the outcome. In both cases, direct inspection of the model parameters revealed that the models contained components with mean vectors that are significantly different from zero. To overcome this issue, the next method forces the means to zero after each stage of component splitting. This gives much better performance when few components are used. However, the improvement is minor when the models consist of 64 components. The final training method, denoted as *GMM init.* in the legend of Figure 6, initialises the emission models using the GMMs trained in Section III. In this case the error rates are significantly lower than all of the other methods for every number of components. This initialisation method reduces the error rate, ranging from 54.3% with a single component to 34.4% with 64 components.

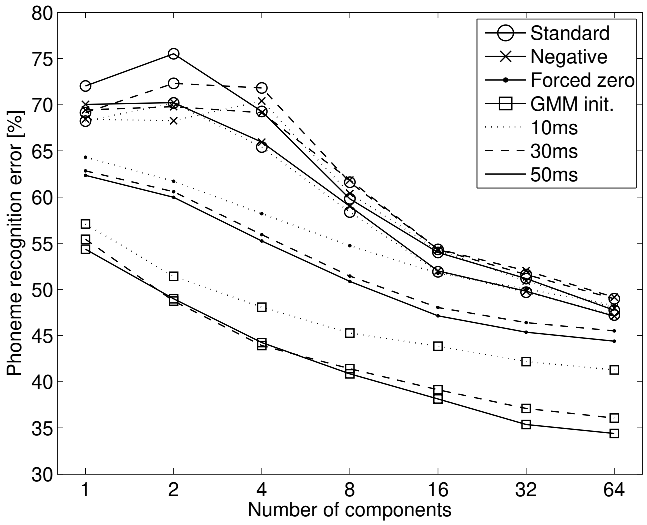

Fig. 6. *Comparison of the four training approaches for the acoustic waveform representation in quiet conditions: default training procedure with component splitting (Standard), default with negative instances added to the training data (Negative), default with emission model mean vectors forced to zero after splitting (Forced zero), and emission models initialised using the zero-mean sector models from Section III (GMM init.). Three curves are shown for each of the methods, corresponding to different number of frames in the concatenation.*

The model training process for acoustic waveforms was initially considered for $f = 1$, $f = 3$ and $f = 5$ frames, corresponding to feature vectors covering 10 ms, 30 ms and 50 ms in duration, respectively, extracted every 10 ms. The duration of 10 ms corresponds to no overlap of adjacent feature vectors, while 30 ms is closest to the standard cepstral representation. The effect of the number of frames is also illustrated in Figure 6. It can be observed that the results improve when going from $f = 1$ to $f = 3$ and $f = 5$, with no indication that $f = 5$ is the optimal value. This suggests that it is necessary to investigate the acoustic waveform representation using more than five frames to find the optimum. In the previous section representations up to $f = 15$ were considered for classification, and this range was sufficient to find the optimal value of f. Figure 7 shows the results for the phoneme recognition task over the range from $f = 1$ to $f = 13$. Each curve represents one fixed global SNR of additive pink noise. There is a slight improvement of around 1.1% that results from increasing the acoustic waveform duration from 50 ms to 90 ms. Beyond this duration, the error rates increase, which could be due to the representation becoming less localised.

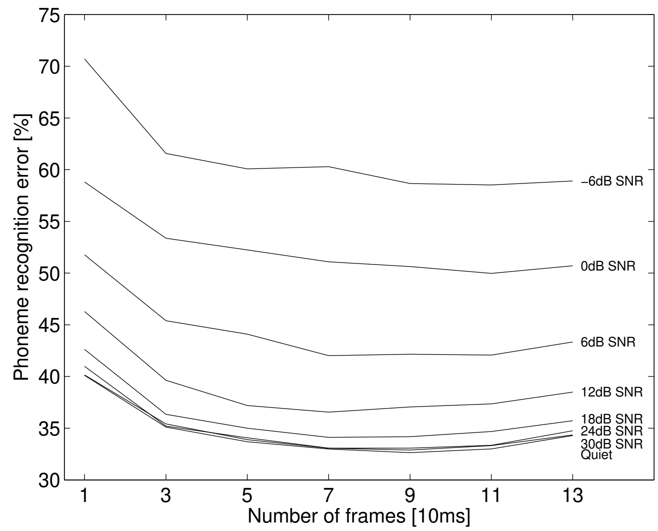

Fig. 7. *The effect of representation duration, denoted by number of 10 ms frames concatenated. Curves are shown for quiet conditions, and in additive pink noise for seven values of SNR. Overall the best performance is obtained using 9 frames, a 90ms duration representation.*

A comparison between acoustic waveforms and cepstral baselines is shown in Figure 8. The models for acoustic waveforms are adapted to additive noise according to (6), while in the case of cepstral features, we now consider standard noise compensation techniques: the ETSI advanced front-end [45] and vector Taylor series [46], in conjunction with MFCC features to obtain noise-compensated cepstral representations MFCC-AFE and MFCC-VTS, respectively. The compensated features are then standardised using sentence-wise cepstral mean and variance normalisation (CMVN) as that significantly improved performance [34]. A single frame is used for MFCC and its variants, with Δ+$\Delta\Delta$ computed in the usual manner. This corresponds to a time interval of 25 ms for the frame, 105 ms for Δ and 145 ms for Δ+$\Delta\Delta$. It was assumed that these methods provide optimal baselines for comparison as they are commonly used and have been tuned for the task. Figure 8 shows recognition results obtained with the acoustic waveform representation for 10 ms, 50 ms, 90 ms and 130 ms. Considering the comparison to the cepstral baselines without Δ+$\Delta\Delta$, the acoustic waveform representation with 10 ms duration gives similar error rates to both cepstral baselines in quiet conditions and better performance down to 0 dB SNR. Note that a similar observation can be made about classification results obtained with fixed-length segments of acoustic waveforms and cepstral features without Δ+$\Delta\Delta$ information, as shown in Figure 2. Results shown in Figure 8 demonstrate that in phoneme recognition too the addition of Δ+$\Delta\Delta$ features provides a significant performance gain for cepstral representations. We can observe again that the performance of cepstral representations with Δ+$\Delta\Delta$ information is significantly better than that achieved with acoustic waveforms at low SNRs, but that the 90 ms acoustic waveform representation improves further in noisy conditions, giving lower error rates than MFCC-AFE+Δ+$\Delta\Delta$ and MFCC-VTS+Δ+$\Delta\Delta$ between -6 dB and 12 dB SNR.

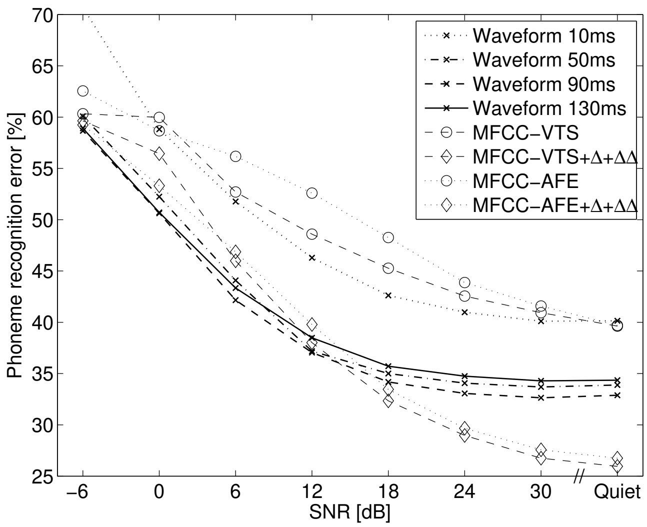

Fig. 8. *Phoneme recognition error rates for the acoustic waveform representations of different durations. Tests are carried out in additive pink noise. Noise compensated MFCC features using vector Taylor series (MFCC-VTS) and ETSI advanced front-end (MFCC-AFE) methods are used as baselines.*

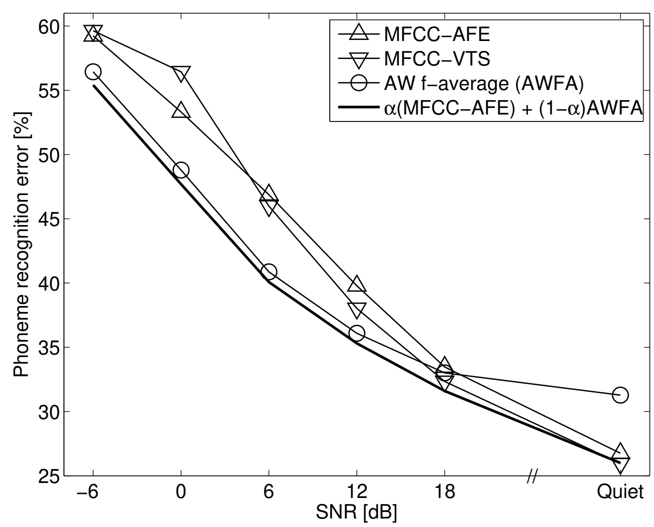

Fig. 9. *Comparison of phoneme recognition error rates in additive pink noise. Baseline comparisons are provided by MFCC-AFE and MFCC-VTS. The f-average (AWFA) is taken for the range 50ms - 130ms. Uniformly lower error rates are achieved by the convex combination of MFCC-AFE and AWFA.*

We also re-investigate in the HMM context the effect of model averaging by combining models corresponding to different number of frames used to represent acoustic waveforms. Having established that the lowest error rates are achieved using $f = 9$ frames, the models corresponding to $f \in \mathcal{F} = \{5, 7, 9, 11, 13\}$ are averaged by a multi-stream process [48] and used for recognition. The resulting acoustic waveform f-average (AWFA) model achieved better performance than the individual models with a fixed value of f. Finally, we re-investigate in the recognition context combining acoustic waveform and MFCC-AFE HMMs according to (14). Note again that the larger f-values for waveforms effectively incorporate information that is provided for MFCCs via $\Delta + \Delta\Delta$, which is why we do not consider similar f-frame concatenations for MFCCs. Phoneme recognition results obtained with the frame averaging and representation combination in additive pink noise are shown in Figure 9. The acoustic waveform representation with averaging over the number of frames (AWFA) gives lower error rates than the cepstral representations for SNR below a cross-over point around 18 dB. Error rates achieved by the combined representation are uniformly lower than any of the individual representations considered. As a consequence, the combined HMM model achieves significant reductions in error rates compared to cepstral-only models with standard noise compensation techniques.

V. Conclusion and Discussion

In this paper we have explored the potential of improving the robustness of ASR to additive noise by posing the problem directly in the space of acoustic waveforms of speech, or some linear transform thereof which does not incur any information loss. In order to assess the separation of phonetic classes in the acoustic waveform domain and the domains of standard cepstral features, we first considered classification of phonemes using observations of a given fixed duration. A remarkable result of that experiment is that phoneme classification in the acoustic waveform domain significantly outperforms classification in the cepstral domains already at very low noise levels, and that even at an SNR as low as -18 dB it achieves accuracy distinctly above that of random guessing (see Figure 2). A significant performance gain in classification using cepstral features is achieved via Δ+$\Delta\Delta$ features, which reflect temporal dynamics of speech. While finding analogous features in the acoustic waveform domain is a nontrivial open problem, considerations in Section III demonstrated that a performance gain in the acoustic waveform domain can be achieved by extending the length of the basic observation unit and thus including the information on the temporal dynamics implicitly. Next, acoustic waveforms were considered in the context of phoneme recognition from continuous speech. To allow a direct comparison with the cepstral features, and compatibility with existing technologies, the recognition was implemented in the standard HMM-GMM framework. Note however that [49]: *"The HMM and frame-based cepstra have co-evolved as ASR system components and hence are very much tuned to each other. The use of HMMs as the core acoustic modeling technology might obscure the gains from new features, especially those from long time scales; this may be one reason why progress with novel techniques has been so difficult. In particular, systems incorporating new signal processing methods in the front end are at a disadvantage when tested using standard HMMs."* Still phoneme recognition using acoustic waveforms achieved lower error rates than cepstral features for noise levels below 18 dB SNR.

Recognition and classification results presented in this proof-of-concept study demonstrate that representing speech using high dimensional linear features, which incur no loss of information, is a promising direction towards achieving the long sought-after robustness of ASR systems. Significant improvements can potentially be made by more sophisticated modelling of speech in the acoustic waveform domain. The choice of decorrelating basis is another area that opens up scope for further work. In these directions methods such as MLLT [50] and semi-tied covariance matrices [51] may lead to improved results. We conducted some preliminary experiments with semi-tied covariance matrices in the context of phoneme recognition. This improved recognition performance in quiet, but suffered from increased error levels in noise. However, more extensive investigations may prove fruitful. Even within the models introduced in this study, given the wide range of possibilities for choice of parameters, many of these have not been extensively tuned. This includes the weighting of the model averaging over the number of components in GMMs, and the averaging over the number of frames f, either in a class-independent manner, or tuned even further to vary with the phoneme or phonetic group being evaluated; a uniform weight was assumed throughout (averaging over the number of components was not explored at all in the recognition experiment), but it is likely that the information content used for discrimination is not distributed in exactly this way. Then models could potentially also be developed to explicitly model correlations between feature vectors obtained for different number of frames f.

Towards other classes of models, Gaussian scale mixture models (GSMMs) are a subset of GMMs that seem particularly suitable for modelling the distributions of the speech signal. This is motivated by the observation in our experiments that the covariance matrices of some of the components were approximately scaled copies of each other. Densities of the acoustic waveform representation can thus be modelled using mixtures of scaled components where again the means are constrained to zero and the covariance matrices are constrained to be diagonal. The probability density function of the GSMM is thus given by:

$$p(x) = \sum_{i=1}^{c}\sum_{j=1}^{S} w_i v_j \mathcal{N}(x; 0, s_j^2 \mathbf{C}^{-1}\mathbf{D}_i(\mathbf{C}^{\mathrm{T}})^{-1})$$

where the symbols are as in (3), with the addition of S scales and scale weights, s_j and v_j, respectively. Here, the range of scales corresponds to a model of the amplitude distribution for a given phoneme. The eight scale model with 128 components provided an error rate that improved by 4.2% in quiet, compared to the standard 128-component GMM when used for phoneme classification as considered in Section II [34]. Further reductions in the error rates may be expected if more scales are used, given that some of the scale distributions are not adequately modelled with only eight scales. Although these GSMM models achieved lower error rates than the corresponding GMMs with the same number of components, computational constraints prevented us from exploring this direction further. Note that GSMM models are an example of a Richter distribution, which have been previously used to model heavy tailed distributions [52].

We expect that the results can be further improved, and that acoustic waveform features may by effective if incorporated into techniques considered by other authors, in particular, committee classifiers and the use of a hierarchy to reduce broad phoneme class confusions [53], [54]. Indeed, in the context of phoneme recognition, much lower confusion rates between broad phoneme classes are achieved using acoustic waveforms than cepstral features, as illustrated by results shown in Figure 10. Finally considering the success of deep neural networks (DNN) in ASR tasks [55], it is imperative to explore DNN modelling of acoustic waveforms of speech. While DNNs have exhibited robustness to modest differences between training and test settings, their performance is poor if the mismatch is more severe [56]. A recent study by Mitra *et al.* [57] shows, however, that the choice of features provides scope for improving the robustness of DNNs. Hence it is of interest to explore possible gains achievable by using DNNs in conjunction with high dimensional linear representations as considered in this study.

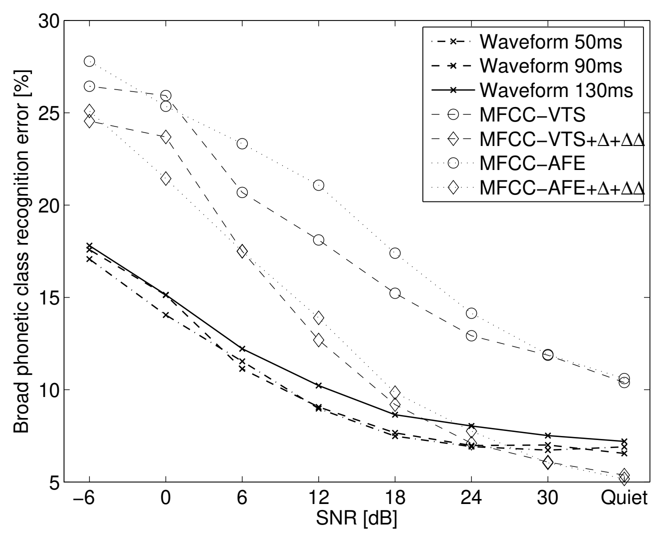

Fig. 10. *Broad phoneme class (vowels, nasals, strong fricatives, weak fricatives, stops and silence) recognition error rates.*

ACKNOWLEDGMENT

Zoran Cvetkovic would like to thank Jont Allen, Bishnu Atal and Herve Bourlard for their encouragement, support and inspiration, and Malcolm Slaney for his valuable suggestions about the presentation of this paper.

REFERENCES

[1] M. Ager, Z. Cvetković, and P. Sollich, “Combined waveform-cepstral representation for robust speech recognition,” in *Proc. ISIT*, July 2011, pp. 125–128.

[2] J. M. Baker, L. Deng, J. Glass, S. Khudanpur, C. Lee, N. Morgan, and D. O’Shaugnessy, “Research developments and directions in speech recognition and understanding, part 1,” *IEEE Signal Process. Mag.*, vol. 26, pp. 75–80, May 2009.

[3] ——, “Updated minds report on speech recognition and understanding, part 2,,” *IEEE Signal Process. Mag.*, vol. 26, pp. 78–85, July 2009.

[4] G. Miller, A. Heise, and W. Lichten, “The Intelligibility of Speech as a Function of the Context of the Test Material,” *J. Exp. Psychol.*, vol. 41, pp. 329–335, 1951.

[5] G. Miller and P. Nicely, “An Analysis of Perceptual Confusions among some English Consonants,” *J. Acoust. Soc. Am.*, vol. 27, pp. 338–352, 1955.

[6] R. Singh and J. B. Allen, “The influence of stop consonants’ perceptual features on the articulation index model,” *J. Acoust. Soc. Am.*, vol. 131, no. 4, pp. 3051 – 3068, April 2012.

[7] J. C. Toscano and J. B. Allen, "Across and within consonant errors for isolated syllables in noise," *J. Speech, Language, Hearing Res.*, vol. 57, pp. 2293–2307, December 2014.

[8] J. Sroka and L. Braida, "Human and Machine Consonant Recognition," *Speech Comm.*, vol. 45, no. 4, pp. 401–423, 2005.

[9] M. J. F. Gales and S. J. Young, "The application of hidden markov models in speech recognition," *Foundations and Trends in Signal Processing*, vol. 1, no. 3, pp. 195–304, 2007.

[10] S. Davis and P. Mermelstein, "Comparison of Parametric Representations for Monosyllabic Word Recognition in Continuously Spoken Sentences," *IEEE Trans. Acoustics, Speech Signal Process.*, vol. 28, no. 4, pp. 357–366, 1980.

[11] H. Hermansky, "Perceptual Linear Predictive (PLP) Analysis of Speech," *J. Acoust. Soc. Am.*, vol. 87, no. 4, pp. 1738–1752, 1990.

[12] B. S. Atal and S. L. Hanauer, "Speech analysis and synthesis by linear prediction of the speech wave," *J. Acoust. Soc. Am.*, vol. 50, pp. 637–655, 1971.

[13] S. D. Peters, P. Stubley, and J. Valin, "On the Limits of Speech Recognition in Noise," *Proc. ICASSP*, pp. 365–368, 1999.

[14] B. Meyer, M. Wächter, T. Brand, and B. Kollmeier, "Phoneme Confusions in Human and Automatic Speech Recognition," 2007, pp. 2740–2743.

[15] K. K. Paliwal and L. D. Alsteris, "On the Usefulness of STFT Phase Spectrum in Human Listening Tests," *Speech Comm.*, vol. 45, no. 2, pp. 153–170, 2005.

[16] L. D. Alsteris and K. K. Paliwal, "Further Intelligibility Results from Human Listening Tests using the Short-Time Phase Spectrum," *Speech Comm.*, vol. 48, no. 6, pp. 727–736, 2006.

[17] H. Sheikhzadeh and L. Deng, "Waveform-Based Speech Recognition Using Hidden Filter Models: Parameter Selection and Sensitivity to Power Normalization," *IEEE Trans. Speech Audio Process.*, vol. 2, no. 1, pp. 80 –89, Jan 1994.

[18] A. Poritz, "Linear Predictive Hidden Markov Models and the Speech Signal," in *Proc. ICASSP*, 1982.

[19] Y. Ephraim and W. Roberts, "Revisting Autoregressive Hidden Markov Modeling of Speech Signals," vol. 12, no. 2, pp. 166–169, Feb. 2005.

[20] B. Mesot and D. Barber, "Switching Linear Dynamical Systems for Noise Robust Speech Recognition," *IEEE Trans. Audio, Speech Language Process.*, vol. 15, no. 6, pp. 1850–1858, 2007.

[21] R. E. Turner and M. Sahani, "Modeling Natural Sounds with Modulation Cascade Processes," in *Advances in Neural Information Processing Systems*, vol. 20, 2008.

[22] C. Kim and R. M. Stern, "Power-normalized cepstral coefficients (pncc) for robust speech recognition," in *Proc. ICASSP*, March 2012, pp. 4101–4014.

[23] N. Vasiloglou and A. G. D. Anderson, "Learning the Intrinsic Dimensions of the TIMIT Speech Database with Maximum Variance Unfolding," in *IEEE Workshops on DSP and SPE*, 2009.

[24] S. Zhang and Z. Zhao, "Dimensionality Reduction-based Phoneme Recognition," in *Proc. ICSP*, October 2008, pp. 667–670.

[25] S. Roweis and L. Saul, "Nonlinear Dimensionality Reduction by Locally Linear Embedding," *Science*, vol. 290, pp. 2323–2326, 2000.

[26] L. Saul, K. Weinberger, J. Ham, F. Sha, and D. Lee, "Spectral Methods for Dimensionality Reduction," in *Semisupervised Learning*. Cambridge, MA: MIT Press, 2006.

[27] N. Lawrence, "Probabilistic Non-Linear Principal Component Analysis with Gaussian Process Latent Variable Models," *J. Mach. Learn. Res.*, vol. 6, pp. 1783 –1816, 2005.

[28] M. Tipping and C. Bishop, "Mixtures of Probabilistic Principal Component Analysers," *Neural Computation*, vol. 11, no. 2, pp. 443–482, 1999.

[29] J. Yousafzai, M. Ager, Z. Cvetković, and P. Sollich, "Discriminative and generative machine learning approaches towards robust phoneme classification," *Workshop on Information Theory and Applications, ITA*, 2008.

[30] S. Srivastava, M. Gupta, and B. Frigyik, "Bayesian Quadratic Discriminant Analysis," *J. Mach. Learn. Res.*, vol. 8, pp. 1287–1314, 2007.

[31] C. Kim and R. Stern, "Robust Signal-to-Noise Ratio Estimation Based on Waveform Amplitude Distribution Analysis," in *Proc. INTERSPEECH*, 2008.

[32] J. Tchorz and B. Kollmeier, "Estimation of the signal-to-noise ratio with amplitude modulation spectrograms," *Speech Comm.*, vol. 38, no. 1, pp. 1–17, 2002.

[33] A. Varga, H. Steeneken, M. Tomlinson, and D. Jones, "The NOISEX-92 study on the effect of additive noise on automatic speech recognition," DRA Speech Research Unit, Tech. Rep., 1992.

[34] M. P. Ager, "Phoneme classification and noise adaptation in linear feature domains," Ph.D. dissertation, King's College London, 2010.

[35] P. Jain and H. Hermansky, "Improved Mean and Variance Normalization for Robust Speech Recognition," in *Proc. ICASSP*, 2001.

[36] M. Gales and S. Young, "Robust continuous speech recognition using parallel model combination," *IEEE Trans. Speech Audio Process.*, vol. 4, pp. 352–359, 1996.

[37] R. Rose, "Environmental Robustness in Automatic Speech Recognition," *Robust2004 - ISCA and COST278 Workshop on Robustness in Conversational Interaction*, 2004.

[38] K.-F. Lee and H.-W. Hon, "Speaker-Independent Phone Recognition using Hidden Markov Models," *IEEE Trans. Acoustics, Speech Signal Process.*, vol. 37, no. 11, pp. 1641–1648, 1989.

[39] F. Sha and L. Saul, "Large Margin Gaussian Mixture Modeling for Phonetic Classification and Recognition," in *Proc. ICASSP*, 2006.

[40] D. Ellis, "PLP and RASTA (and MFCC, and inversion) in Matlab," 2005, online web resource, http://labrosa.ee.columbia.edu/matlab/rastamat/.

[41] S. Furui, "Speaker-Independent Isolated Word Recognition using Dynamic Features of Speech Spectrum," *IEEE Trans. Acoustics, Speech Signal Process.*, vol. 34, no. 1, pp. 52–59, 1986.

[42] P. Clarkson and P. Moreno, "On the Use of Support Vector Machines for Phonetic Classification," in *Proc. ICASSP*, vol. 2, 1999, pp. 585–588.

[43] M. J. F. Gales and C. Longworth, "Discriminative classifiers with generative kernels for noise robust asr," in *Proc. INTERSPEECH*, September 2008, pp. 1996–1999.

[44] M. Ostendorf and S. Roukos, "A stochastic segment model for phoneme-based continuous speech recognition," *IEEE Trans. Acoustics, Speech Signal Process.*, vol. 37, no. 12, pp. 1857–1869, December 1989.

[45] "ETSI document - ES 202 050 - STQ: DSR. Distributed speech recognition; Advance front-end feature extraction algorithm;Compression algorithms;," 2007.

[46] P. Moreno, B. Raj, and R. M. Stern, "A Vector Taylor Series Approach for Environment-Independent Speech Recognition," in *Proc. ICASSP*, vol. 1, 1996, pp. 733–736.

[47] M. Ager, Z. Cvetković, P. Sollich, and B. Yu, "Towards Robust Phoneme Classification: Augmentation of PLP Models with Acoustic Waveforms," in *Proc. EUSIPCO*, 2008.

[48] S. Young, G. Evermann, M. Gales, T. Hain, D. Kershaw, G. Moore, J. Odell, D. Ollason, D. Povey, V. Valtchev, and P. Woodland, *The HTK Book, version 3.4*. Cambridge, UK: Cambridge University, 2006.

[49] N. Morgan, Q. Zhu, A. Stolcke, K. Sonmez, S. Sivadas, T. Shinozaki, M. Ostendorf, P. Jain, H. Hermansky, D. Ellis, G. Diddington, B. Chen, O. Cetin, H. Bourlard, and M. Athineos, "Pushing the envelope – aside," *IEEE Signal Process. Mag.*, pp. 81 – 88, September 2005.

[50] M. J. F. Gales, "Maximum likelihood linear transformations for hmm-based speech recognition," *Computer Speech and Language*, vol. 12, pp. 75–98, 1998.

[51] ——, "Semi-tied covariance matrices for hidden markov models," *IEEE Trans. Speech Audio Process*, vol. 7, no. 3, pp. 272–281, May 1999.

[52] M. J. F. Gales and P. A. Olsen, "Tail distribution modelling using the richter and power exponential distributions," in *Proc. EUROSPEECH*, 1999, pp. 1507–1510.

[53] H. Chang and J. Glass, "Hierarchical Large-Margin Gaussian Mixture Models For Phonetic Classification," in *Proc. IEEE ASRU Workshop*, 2007, pp. 272–275.

[54] F. Pernkopf, T. Pham, and J. Bilmes, "Broad Phonetic Classification using Discriminative Bayesian Networks," *Speech Comm.*, vol. 51, no. 2, 2009.

[55] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. rahman Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Process. Mag.*, pp. 82 – 97, November 2012.

[56] D. Yu, M. Seltzer, J. Li, J.-T. Huang, and F. Seide, "Feature learning in deep neural networks - studies on speech recognition," in *Int. Conf. Learning Representations*, May 2013.

[57] V. Mitra, W. Wang, H. Franco, Y. Lei, C. Bartels, and M. Graciarena, "Evaluating robust features on deep neural networks for speech recognition in noisy and channel mismatched conditions," in *Proc. INTERSPEECH*, September 2014.