

Maximum Likelihood and Minimum Classification Error Factor Analysis for Automatic Speech Recognition

Lawrence K. Saul and Mazin G. Rahim, *Senior Member, IEEE*

Abstract—Hidden Markov models (HMM's) for automatic speech recognition rely on high-dimensional feature vectors to summarize the short-time properties of speech. Correlations between features can arise when the speech signal is nonstationary or corrupted by noise. We investigate how to model these correlations using factor analysis, a statistical method for dimensionality reduction. Factor analysis uses a small number of parameters to model the covariance structure of high dimensional data. These parameters can be chosen in two ways: 1) to maximize the likelihood of observed speech signals, or 2) to minimize the number of classification errors. We derive an expectation-maximization (EM) algorithm for maximum likelihood estimation and a gradient descent algorithm for improved class discrimination. Speech recognizers are evaluated on two tasks, one small-sized vocabulary (connected alpha-digits) and one medium-sized vocabulary (New Jersey town names). We find that modeling feature correlations by factor analysis leads to significantly increased likelihoods and word accuracies. Moreover, the rate of improvement with model size often exceeds that observed in conventional HMM's.

Index Terms—Automatic speech recognition, discriminative training factor analysis, expectation-maximization (EM) algorithm.

I. INTRODUCTION

HIDDEN Markov models (HMM's) for automatic speech recognition [21] rely on high dimensional feature vectors to summarize the short-time, acoustic properties of speech. Though front-ends vary from recognizer to recognizer, the spectral information in each frame of speech is typically codified in a feature vector with thirty or more dimensions. In most systems, these vectors are conditionally modeled by mixtures of Gaussian probability density functions (PDF's). In this case, the correlations between different features are represented in two ways [16]: implicitly by the use of two or more mixture components, and explicitly by the nondiagonal elements in each covariance matrix. Naturally, these strategies for modeling correlations—implicit versus explicit—involve tradeoffs in accuracy, speed, and memory. This paper examines these tradeoffs using the statistical method of factor analysis.

The present work is motivated by the following observation. Currently, many HMM-based recognizers do not include any explicit modeling of correlations; that is to say—conditioned

on the hidden states, acoustic features are modeled by mixtures of Gaussian PDF's with *diagonal* covariance matrices. The reasons for this practice are well known. The use of full covariance matrices imposes a heavy computational burden, making it difficult to achieve real-time recognition. Moreover, one rarely has enough data to (reliably) estimate full covariance matrices. Some of these disadvantages can be overcome by parameter-tying [2], e.g., sharing the covariance matrices across different states or models. But parameter-tying has its own drawbacks: it considerably complicates the training and decoding procedures, and it requires some artistry to design the tied HMM's.

Unconstrained and diagonal covariance matrices clearly represent two extreme choices for the hidden Markov modeling of speech. The statistical method of factor analysis [7], [24] represents a compromise between these two extremes. The idea behind factor analysis is to map systematic variations of the data into a lower dimensional subspace. This enables one to represent, in a very compact way, the covariance matrices for high dimensional data. These matrices are expressed in terms of a small number of parameters that model the most significant correlations without incurring much overhead in time or memory. For this reason, we shall often refer to factor analysis as a strategy for dimensionality reduction [18]. However, it should be emphasized that factor analysis does not merely project a high dimensional feature vector into a low dimensional one.

In this paper, we investigate the combined use of mixture models and factor analysis in HMM's for automatic speech recognition. Applying factor analysis at the state and mixture component level [9], [11] results in a powerful form of dimensionality reduction, one tailored to the local properties of speech. This approach effectively combines two popular strategies in probabilistic modeling—clustering and dimensionality reduction. The combination of these two methods gives rise to a *local linear model*—local, because each mixture component models different facets of the data, and linear, because the dimensionality reduction is geared to Gaussian mixture components. Many variants on local linear models have been proposed in the machine learning literature. Successful applications include handwritten digit recognition [11], data compression [14], nonlinear image interpolation [3], and articulatory modeling of electropalatographic data [4].

Factor analysis uses a small number of parameters to model the covariance structure of high dimensional data. For automatic speech recognition, these parameters can be chosen in two ways: 1) to maximize the likelihood of observed speech signals, or 2) to minimize the number of classification errors [1], [12], [13],

Manuscript received November 6, 1997; revised March 5, 1999. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Joseph Picone.

The authors are with AT&T Labs—Research, Florham Park, NJ 07932 USA (e-mail: lsaul@research.att.com; mazin@research.att.com).

Publisher Item Identifier S 1063-6676(00)01708-9.

[17]. In this paper, we derive both an expectation-maximization (EM) algorithm for maximum likelihood estimation and a gradient descent algorithm for improved class discrimination. To the best of our knowledge, neither form of factor analysis has been previously applied to hidden Markov models for automatic speech recognition.

Briefly, the organization of this paper is as follows. In Section II, we review the method of factor analysis and describe what makes it attractive for large problems in speech recognition. In Sections III and IV, we present the learning algorithms for maximum likelihood (ML) and minimum classification error (MCE) factor analysis. In Section V, we give results on two tasks in automatic speech recognition: connected alpha-digits and New Jersey town names. Finally, in Section VI, we present our conclusions as well as ideas for future research.

II. DIMENSIONALITY REDUCTION

Factor analysis [7], [24] is a *linear* method for the dimensionality reduction of Gaussian random variables. It is closely related to principal components analysis (PCA) [6] in which one extracts the subspace defined by the largest eigenvectors of the covariance matrix. In traditional PCA, the data vectors are projected into this subspace, resulting in a simple form of dimensionality reduction. Though straightforward, PCA has two important disadvantages: 1) it does not define a proper density model outside the subspace of principal components; and 2) componentwise variations outside this subspace are modeled uniformly, even when the data does not warrant such an assumption. These disadvantages are serious drawbacks for HMM-based speech recognition. The first undermines the probabilistic formulation of HMM's, while the second precludes the use of composite data vectors (e.g., cepstra plus delta-cepstra) that combine features at different scales. Factor analysis does not suffer from either of these drawbacks, though it does contain traditional PCA (and also probabilistic versions of PCA [23], [27]) as a special limiting case.

Having described factor analysis as a linear method, we should emphasize that its application is not limited to strictly Gaussian PDF's. The combined use of mixture densities and factor analysis—resulting in a *nonlinear* form of dimensionality reduction—was first applied by Hinton *et al.* [11] to the modeling of handwritten digits. Essentially, a mixture of factor analyzers patches together the linear subspaces defined by its individual components to parameterize a globally nonlinear (but low-dimensional) manifold.

In this section, we first review the method of factor analysis for multivariate Gaussian PDF's. We then extend the method to mixture models and continuous density HMM's. The problem of parameter estimation is considered separately in Sections III and IV.

A. Multivariate Gaussian PDF

Let $\mathbf{x} \in \mathcal{R}^D$ denote a Gaussian random variable with mean $\boldsymbol{\mu}$. If the number of dimensions, D , is very large, it may be prohibitively expensive to estimate, store, multiply, or invert a full covariance matrix. The idea behind factor analysis is to find a subspace of much lower dimension, $f \ll D$, that captures most

of the variations in \mathbf{x} . To this end, let $\mathbf{z} \in \mathcal{R}^f$ denote a Gaussian random variable with zero mean and identity covariance matrix

$$P(\mathbf{z}) = (2\pi)^{-f/2} \exp \left\{ -\frac{1}{2} \mathbf{z}^T \mathbf{z} \right\}. \quad (1)$$

We now imagine that the variable \mathbf{x} is generated by a random process in which \mathbf{z} is a latent (or hidden) variable; the elements of \mathbf{z} are known as the *factors*. Let Λ denote an arbitrary $D \times f$ matrix, and let Ψ denote a *diagonal*, positive-definite $D \times D$ matrix. We imagine that \mathbf{x} is generated by sampling \mathbf{z} from (3), computing the D -dimensional vector $\hat{\boldsymbol{\mu}} + \Lambda \mathbf{z}$, then adding independent Gaussian noise with variances Ψ_{ii} to each component of this vector. The matrix Λ is known as the *factor loading* matrix. The relation between \mathbf{x} and \mathbf{z} is captured by the conditional distribution

$$P(\mathbf{x}|\mathbf{z}) = \frac{|\Psi|^{-1/2}}{(2\pi)^{D/2}} \cdot \exp \left\{ -\frac{1}{2} [\mathbf{x} - \boldsymbol{\mu} - \Lambda \mathbf{z}]^T \Psi^{-1} [\mathbf{x} - \boldsymbol{\mu} - \Lambda \mathbf{z}] \right\}. \quad (2)$$

The marginal distribution for \mathbf{x} is found by integrating out the hidden variable \mathbf{z} , or $P(\mathbf{x}) = \int d\mathbf{z} P(\mathbf{x}|\mathbf{z}) P(\mathbf{z})$. The calculation is straightforward because both $P(\mathbf{z})$ and $P(\mathbf{x}|\mathbf{z})$ are Gaussian

$$P(\mathbf{x}) = \frac{|\Psi + \Lambda \Lambda^T|^{-1/2}}{(2\pi)^{D/2}} \cdot \exp \left\{ -\frac{1}{2} [\mathbf{x} - \boldsymbol{\mu}]^T (\Psi + \Lambda \Lambda^T)^{-1} [\mathbf{x} - \boldsymbol{\mu}] \right\}. \quad (3)$$

From (3), we see that \mathbf{x} is normally distributed with mean $\boldsymbol{\mu}$ and covariance matrix $\Psi + \Lambda \Lambda^T$. It follows that when the diagonal elements of Ψ are small, most of the variation in \mathbf{x} occurs in the subspace $S(\Lambda)$ spanned by the columns of Λ . The variances Ψ_{ii} measure the typical size of componentwise fluctuations outside this subspace. In general, these variances may be quite different from one another. Hence, unlike PCA, (3) can model a distribution in which two vectors with the same projection onto $S(\Lambda)$ have quite different likelihoods.

Covariance matrices of the form $\Psi + \Lambda \Lambda^T$ have a number of useful properties. Most importantly, they are expressed in terms of a small number of parameters, namely the $D(f+1)$ nonzero elements of Λ and Ψ . If $f \ll D$, then storing Λ and Ψ requires much less memory than storing a full covariance matrix. Likewise, estimating Λ and Ψ also requires much less data than estimating a full covariance matrix. Covariance matrices of this form can be efficiently inverted using the matrix inversion lemma [20]

$$(\Psi + \Lambda \Lambda^T)^{-1} = \Psi^{-1} - \Psi^{-1} \Lambda (I + \Lambda^T \Psi^{-1} \Lambda)^{-1} \Lambda^T \Psi^{-1} \quad (4)$$

where I is the $f \times f$ identity matrix. This decomposition also allows one to compute the probability $P(\mathbf{x})$ with only $O(fD)$ multiplies, as opposed to the $O(D^2)$ multiplies that are normally required when the covariance matrix is nondiagonal. We will refer to covariance matrices of the form $\Psi + \Lambda \Lambda^T$ as *factored* covariance matrices. As a theoretical aside, note that if we allow the number of factors to equal the number of dimensions (i.e., $f = D$), then we recover the ability to parameterize an arbitrary covariance matrix, M . This is done by choosing the columns of the factor loading matrix, Λ , to be parallel to the eigenvectors of M .

Viewing factor analysis as a latent variable model, we can also compute the posterior distribution, $P(\mathbf{z}|\mathbf{x})$. The statistics of this distribution are derived in Appendix A. Computing the posterior statistics of the hidden variable \mathbf{z} is an important step in the EM algorithm for ML factor analysis. We will return to the problem of parameter estimation in Section III.

B. Hidden Markov Models

Nearly all speech recognizers rely on Gaussian PDF's to model the short-term properties of speech. These PDF's are associated with the state emission densities of continuous density HMM's. Often, the dimensionality of acoustic feature vectors is too large for these PDF's to employ full covariance matrices. Instead, correlations are modeled implicitly by mixtures of Gaussians with diagonal covariance matrices. Intuitively, one might expect the mixture components to model discrete types of variability, such as the speaker's gender or local dialect. However, mixture models are not geared to modeling continuous types of variability, as might arise from coarticulation effects or background noise. Clearly, both types of variability—discrete and continuous—are important for building accurate models of speech. Factor analysis provides a way to model continuous variability without incurring the overhead associated with full covariance matrices.

Consider a continuous density HMM whose feature vectors, conditioned on the hidden states, are modeled by mixtures of Gaussian PDF's. A mixture of factor analyzers [9] is a mixture of Gaussian PDF's, each having the form of (3). We can use these models as output distributions for each hidden state s

$$P(\mathbf{x}|s) = \sum_c P(c|s)P(\mathbf{x}|s, c) \quad (5)$$

$$P(\mathbf{x}|s, c) = \frac{|\Psi_{sc} + \Lambda_{sc}\Lambda_{sc}^T|^{-1/2}}{(2\pi)^{D/2}} \cdot \exp\left\{-\frac{1}{2}[\mathbf{x} - \boldsymbol{\mu}_{sc}]^T(\Psi_{sc} + \Lambda_{sc}\Lambda_{sc}^T)^{-1}[\mathbf{x} - \boldsymbol{\mu}_{sc}]\right\}. \quad (6)$$

The mixture components in (5) and (6) are indexed by the subscript c . All of the recursive algorithms for computing statistics in HMM's can be applied to these models. The main difference—or advantage—is that here one can compute $P(\mathbf{x}|s)$ with fewer operations than is normally required for full covariance matrices. This is done by exploiting the special structure of factored covariance matrices, as shown by (4). We will refer to these models as FA-HMM's and to HMM's with diagonal covariance matrices as DG-HMM's.

Clearly, an important consideration when applying factor analysis to speech recognition is the choice of acoustic features. In our experiments, we used a thirty-nine dimensional feature vector consisting of the first twelve cepstral coefficients (with first and second derivatives) and the normalized log-energy (with first and second derivatives). There are known to be correlations [16] between these features, especially between the different types of coefficients (e.g., cepstrum and delta-cepstrum). While these correlations have motivated our use of factor analysis, it is worth emphasizing that the method applies to arbitrary feature vectors. Indeed, whatever features are used to summarize the short-time properties of speech, one expects

correlations to arise from coarticulation effects, background noise, speaker idiosyncrasies, etc.

III. MAXIMUM LIKELIHOOD FACTOR ANALYSIS

The simplest learning criterion for continuous density HMM's is to maximize the likelihood of observed speech signals. The expectation-maximization (EM) algorithm [24] is a general iterative procedure for estimating the parameters of latent variable models. In this section, we begin by reviewing the EM algorithm for maximum likelihood factor analysis [24]. This is done for the multivariate Gaussian PDF of Section II-A. We then consider the extension to mixture models and continuous density HMM's. The EM algorithm for mixtures of factor analyzers was first derived by Ghahramani *et al.* [9].

A. Multivariate Gaussian PDF

Consider the multivariate Gaussian PDF from (3). Appealing to its representation as a latent variable model, we may derive an EM algorithm for obtaining maximum likelihood (ML) estimates of the parameters Λ , Ψ , and $\boldsymbol{\mu}$. For the parameter $\boldsymbol{\mu}$, an iterative procedure is not generally required, since one can simply set $\boldsymbol{\mu}$ equal to the sample mean (which is the ML estimate). However, by deriving the EM algorithm here in full generality, we will considerably simplify the extension (Section III-B) to mixture models and continuous density HMM's. In this respect, our derivation of ML factor analysis will differ slightly from standard treatments.

To begin, let $\{\mathbf{x}_n\}_{n=1}^N$ denote a sample of N data points. The EM procedure is a two-step iterative procedure for maximizing the log-likelihood, $\sum_n \ln P(\mathbf{x}_n)$, with $P(\mathbf{x}_n)$ given by (3). The E-step of this procedure is to compute the Q -function

$$Q(\tilde{\boldsymbol{\mu}}, \tilde{\Lambda}, \tilde{\Psi}; \boldsymbol{\mu}, \Lambda, \Psi) = \sum_n \int d\mathbf{z} P(\mathbf{z}|\mathbf{x}_n, \boldsymbol{\mu}, \Lambda, \Psi) \ln P(\mathbf{z}, \mathbf{x}_n|\tilde{\boldsymbol{\mu}}, \tilde{\Lambda}, \tilde{\Psi}). \quad (7)$$

The right hand side of (7) depends on $\boldsymbol{\mu}$, Λ and Ψ through the statistics of the posterior distribution, $P(\mathbf{z}|\mathbf{x}_n, \boldsymbol{\mu}, \Lambda, \Psi)$. Some useful properties of this distribution are summarized in Appendix A. Of particular importance are the posterior statistics

$$E[\mathbf{z}|\mathbf{x}_n] = [I + \Lambda^T \Psi^{-1} \Lambda]^{-1} \Lambda^T \Psi^{-1} (\mathbf{x}_n - \boldsymbol{\mu}) \quad (8)$$

$$E[\boldsymbol{\delta} \mathbf{z} \boldsymbol{\delta}^T | \mathbf{x}_n] = [I + \Lambda^T \Psi^{-1} \Lambda]^{-1} \quad (9)$$

where the shorthand $\boldsymbol{\delta} \mathbf{z}$ in (9) denotes the variation about the conditional mean, or $\boldsymbol{\delta} \mathbf{z} = \mathbf{z} - E[\mathbf{z}|\mathbf{x}_n]$. Note that the conditional covariance, $E[\boldsymbol{\delta} \mathbf{z} \boldsymbol{\delta}^T | \mathbf{x}_n]$, does not depend on the particular value of \mathbf{x}_n ; as a practical matter, this means that it does not have to be recomputed for each data point.

The M-step of the EM algorithm is to maximize the right hand side of (7) with respect to $\tilde{\boldsymbol{\mu}}$, $\tilde{\Psi}$, and $\tilde{\Lambda}$. To this end, let us define the new vector quantities

$$\Delta \mathbf{x}_n = \mathbf{x}_n - \frac{1}{N} \sum_k \mathbf{x}_k \quad (10)$$

$$\Delta \mathbf{z}_n = E[\mathbf{z}|\mathbf{x}_n] - \frac{1}{N} \sum_k E[\mathbf{z}|\mathbf{x}_k]. \quad (11)$$

These quantities measure the variation of each data point and its expected factors about the sample means. Note that $\Delta \mathbf{x}_n$ and $\Delta \mathbf{z}_n$ are measures of sample variance, as opposed to posterior averages such as $E[\delta \mathbf{z} \delta \mathbf{z}^T | \mathbf{x}_n]$, which measure the uncertainty in $P(\mathbf{z} | \mathbf{x}_n)$ induced by a single data point. The definitions of $\Delta \mathbf{x}_n$ and $\Delta \mathbf{z}_n$ are useful insofar as they allow us to write the EM updates in an especially compact form. In Appendix B, we show that maximizing (7) leads to the iterative updates

$$\tilde{\Lambda} \leftarrow \left(\sum_n (\Delta \mathbf{x}_n)(\Delta \mathbf{z}_n)^T \right) \cdot \left(\sum_n [E[\delta \mathbf{z} \delta \mathbf{z}^T | \mathbf{x}_n] + (\Delta \mathbf{z}_n)(\Delta \mathbf{z}_n)^T] \right)^{-1} \quad (12)$$

$$\tilde{\boldsymbol{\mu}} \leftarrow \frac{1}{N} \sum_n (\mathbf{x}_n - \tilde{\Lambda} E[\mathbf{z} | \mathbf{x}_n]) \quad (13)$$

$$\tilde{\Psi}_{ii} \leftarrow \frac{1}{N} \sum_n [(\Delta \mathbf{x}_n - \tilde{\Lambda} \Delta \mathbf{z}_n)_i^2 + (\tilde{\Lambda} E[\delta \mathbf{z} \delta \mathbf{z}^T | \mathbf{x}_n] \tilde{\Lambda}^T)_{ii}] \quad (14)$$

where $\tilde{\Psi}_{ii}$ are the diagonal elements of $\tilde{\Psi}$. These updates are guaranteed to converge monotonically to a (possibly local) extremum of the log-likelihood. Note that it is important to perform the updates in the order shown, since (for example) the $\tilde{\boldsymbol{\mu}}$ -update depends on the re-estimated value of $\tilde{\Lambda}$. Naturally, for $\Lambda = 0$ the updates in (13) and (14) reduce to the ML estimates for independent Gaussian random variables.

The reestimation equations are supported by various intuitions. In (12), for example, the factor loading matrix is reestimated by finding the least-squares solution that projects the variation in \mathbf{x} (about its mean) into a lower dimensional subspace. In (13), the mean is used to account for the average offset between each observation and its reconstruction by the factors. Finally, in (14), the variances are used to account for two types of variation in \mathbf{x} : those outside the dimensionality-reduced subspace (the first term on the right hand side), and those within the dimensionality-reduced subspace (the second term). All these intuitions extend to the more elaborate models considered in the next section.

B. Hidden Markov Models

One can readily integrate the EM algorithm for factor analysis into the ML training of continuous density HMM's. The parameter updates for FA-HMM's have a similar structure to the Baum-Welch updates for conventional HMM's (i.e., those whose Gaussian PDF's employ diagonal or full covariance matrices). In fact, the main novelties—which arise from the special form of factored covariance matrices—were introduced in the previous section.

With that in mind, let us reconsider the HMM whose observations are conditionally modeled by (5) and (6). Suppose that $\mathbf{x} = \{\mathbf{x}_n\}$ represents a sequence of training data. The forward-backward procedure [21] enables one to compute the posterior probability, $\gamma_n^{sc} = P(s_n = s, c_n = c | \mathbf{x})$, that the HMM used state s and mixture component c to generate \mathbf{x}_n . Let $N_{sc} = \sum_n \gamma_n^{sc}$ denote the total probability mass (or *effective*

number of feature vectors) attributed to state s and mixture component c . The parameter updates for FA-HMM's may be viewed as a specialization of (12)–(14) to each hidden state and mixture component, in which the observations \mathbf{x}_n are weighted by their posterior probabilities, γ_n^{sc} . Analogous to (10) and (11), we define for each state and mixture component the vector quantities

$$\Delta \mathbf{x}_n^{sc} = \mathbf{x}_n - \frac{1}{N_{sc}} \sum_k \gamma_k^{sc} \mathbf{x}_k \quad (15)$$

$$\Delta \mathbf{z}_n^{sc} = E_{sc}[\mathbf{z} | \mathbf{x}_n] - \frac{1}{N_{sc}} \sum_k \gamma_k^{sc} E_{sc}[\mathbf{z} | \mathbf{x}_k] \quad (16)$$

where $E_{sc}[\mathbf{z} | \mathbf{x}_n]$ denotes an expectation with respect to the posterior distribution, $P(\mathbf{z} | \mathbf{x}_n, s_n = s, c_n = c)$. In terms of these vectors, the EM updates are given by

$$\tilde{\Lambda}_{sc} \leftarrow \left(\sum_n \gamma_n^{sc} (\Delta \mathbf{x}_n^{sc})(\Delta \mathbf{z}_n^{sc})^T \right) \cdot \left(\sum_n \gamma_n^{sc} [E_{sc}[\delta \mathbf{z} \delta \mathbf{z}^T | \mathbf{x}_n] + (\Delta \mathbf{z}_n^{sc})(\Delta \mathbf{z}_n^{sc})^T] \right)^{-1} \quad (17)$$

$$\tilde{\boldsymbol{\mu}}_{sc} \leftarrow \frac{1}{N_{sc}} \sum_n \gamma_n^{sc} (\mathbf{x}_n - \tilde{\Lambda}_{sc} E_{sc}[\mathbf{z} | \mathbf{x}_n]) \quad (18)$$

$$[\tilde{\Psi}_{sc}]_{ii} \leftarrow \frac{1}{N_{sc}} \sum_n \gamma_n^{sc} [(\Delta \mathbf{x}_n^{sc} - \tilde{\Lambda}_{sc} \Delta \mathbf{z}_n^{sc})_i^2 + (\tilde{\Lambda}_{sc} E_{sc}[\delta \mathbf{z} \delta \mathbf{z}^T | \mathbf{x}_n] \tilde{\Lambda}_{sc}^T)_{ii}]. \quad (19)$$

The reader should note the similarities in structure between (17)–(19) and (12)–(14). The FA-HMM updates have essentially the same form as (12)–(14), except that now each observation \mathbf{x}_n is weighted by the posterior probability, γ_n^{sc} .

As usual, in the case where multiple sequences are available as training data, the sums over n should be interpreted as sums over sequences as well [21]. Finally, we note that the updates for mixture weights and transition matrices in FA-HMM's are identical to those in conventional HMM's [21]. (This should be clear since FA-HMM's form a subset of HMM's whose Gaussian PDF's employ full covariance matrices.)

IV. MINIMUM CLASSIFICATION ERROR FACTOR ANALYSIS

Ultimately, we evaluate a speech recognizer by its accuracy, not by the likelihood scores of its component HMM's. Hence, maximizing the likelihood—while arguably the simplest criterion for parameter estimation—is not always the most desirable one. In fact, because continuous density HMM's represent only approximately correct models of speech, maximum likelihood estimates are not guaranteed to produce accurate recognizers, even in the limit of infinite training data. Moreover, empirically it is well known that maximizing likelihoods does not always translate into minimizing error rates. We will see an example of this—for connected alpha-digits—in Section V-A.

For all these reasons, many researchers have proposed alternative criteria for parameter estimation that more directly relate to the empirical error rate [1], [12], [13], [17]. In this paper, we consider the minimum classification error (MCE)

criterion as described by Juang *et al.* [12]. For each utterance, the MCE criterion relates the event of a recognition error to the log-likelihood ratio between correct and competing hypotheses, where each hypothesis represents one possible labeling of hidden states. In this section, we show that this type of discriminative training can be applied to FA-HMM's. To keep the paper self-contained, we begin by reviewing the basic framework for MCE parameter estimation. We then consider a simple example—how to train a classifier based on competing Gaussian PDF's. This example is analyzed for the special case of factored covariance matrices. Finally, in the last part of the section, we extend the method to continuous density HMM's.

A. Discriminative Training

Let $\{\mathbf{x}_n, y_n\}_{n=1}^N$ denote a labeled sample of data points, where the label y_n indicates that \mathbf{x}_n belongs to one of \mathcal{Y} classes. If we have statistical models for the class-conditional distributions, $P(\mathbf{x}|y)$, then we can construct the maximum *a posteriori* (MAP) classifier that labels each point by its most probable assignment

$$\mathcal{C}(\mathbf{x}) = \arg \max_y [\ln P(\mathbf{x}|y)P(y)]. \quad (20)$$

For simplicity, let us assume that the class labels have *a priori* equal probabilities, or $P(y) = 1/\mathcal{Y}$, so that the rightmost term in (20) can be ignored. The empirical error rate on the training data, \hat{J} , is then given by

$$\hat{J} = \frac{1}{N} \sum_n \Theta \left(\max_{y \neq y_n} \ln \left[\frac{P(\mathbf{x}_n|y)}{P(\mathbf{x}_n|y_n)} \right] \right) \quad (21)$$

where $\Theta(\cdot)$ denotes the unit threshold function, and the max operation ranges over all labels except the correct one. Note that (21) computes the classification error for each data point in two steps. First, the max operation is used to find the most likely *incorrect* hypothesis. Second, the threshold function is used to register an error if this hypothesis is *more* likely than the correct one.

Suppose now that the statistical models, $P(\mathbf{x}|y)$, are expressed smoothly in terms of parameters, Φ , at our disposal. In what follows, we will assume this dependence implicitly, rather than writing out $P(\mathbf{x}|y; \Phi)$ wherever we encounter probabilities. The goal of discriminative training is to find the parameter values that minimize the empirical error rate. Note, however, that (21) does not define a differentiable function of the log-likelihoods: both the step function and the max operation give rise to singular gradients. This makes it difficult to directly minimize the empirical error rate. Our strategy, following earlier work [12], is to construct a smooth objective function that mimics the behavior of (21). This is done in two steps. First, we replace the max operation by a softmax: $\max_k z_k \approx \ln \sum_k e^{z_k}$. Second, we replace the threshold function by a sigmoid function: $\Theta(z) \approx [1 + e^{-z}]^{-1}$. In particular, for the n th data point, let

$$d(\mathbf{x}_n) = \ln \sum_{y \neq y_n} P(\mathbf{x}|y) - \ln P(\mathbf{x}_n|y_n) \quad (22)$$

denote the log-likelihood-ratio between the softmax-smoothed competing hypothesis¹ and the correct one. The function $d(\mathbf{x}_n)$

provides a smooth measure of the misclassification of \mathbf{x}_n . The MCE loss function is obtained by substituting this smoothed log-likelihood ratio into a sigmoid function

$$J = \frac{1}{N} \sum_n \frac{1}{1 + \exp(-d(\mathbf{x}_n))}. \quad (23)$$

Minor variations [12] on this loss function can also be considered by generalizing the softmax and sigmoid functions. The important result is that (23) is a smooth function of the log-likelihoods, $\ln P(\mathbf{x}_n|y)$.

The goal of discriminative training is to minimize the MCE loss function with respect to the parameters Φ . We can update these parameters by gradient descent, in which at each iteration

$$\Phi \leftarrow \Phi - \eta(\Phi) \frac{\partial J}{\partial \Phi} \quad (24)$$

where $\eta(\Phi)$ is a positive learning rate, or more generally, a positive-definite matrix. Note that in the MCE loss function, (23), all the dependence on Φ enters through the log-likelihoods, $\ln P(\mathbf{x}|y)$. We can therefore use the chain rule

$$\frac{\partial J}{\partial \Phi} = \sum_y \frac{\partial J}{\partial \ln P(\mathbf{x}|y)} \frac{\partial \ln P(\mathbf{x}|y)}{\partial \Phi} \quad (25)$$

to decompose the partial derivatives in (24). Equations (22)–(25) provide the general framework for MCE parameter estimation. The specific form of the update equations necessarily depends on the parameterization of the class-conditional distributions, $P(\mathbf{x}|y)$. For a given parameterization, however, the only requirement is to compute the gradients with respect to Φ . In the rest of this section, we examine two cases of special interest.

B. Multivariate Gaussian PDF's

Suppose that the class-conditional distributions are multivariate Gaussian PDF's with factored covariance matrices. Discriminative training requires us to compute the gradients of $\ln P(\mathbf{x}|y)$, where for a particular class

$$P(\mathbf{x}|y) = \frac{|\Psi + \Lambda\Lambda^T|^{-1/2}}{(2\pi)^{D/2}} \cdot \exp \left\{ -\frac{1}{2} [\mathbf{x} - \boldsymbol{\mu}]^T (\Psi + \Lambda\Lambda^T)^{-1} [\mathbf{x} - \boldsymbol{\mu}] \right\}. \quad (26)$$

Let $M = \Psi + \Lambda\Lambda^T$ denote the covariance matrix in (26). In Appendix C, we show that the gradients of $\ln P(\mathbf{x}|y)$ are given by

$$\frac{\partial}{\partial \boldsymbol{\mu}_i} [\ln P(\mathbf{x}|y)] = -[M^{-1}(\mathbf{x} - \boldsymbol{\mu})]_i \quad (27)$$

$$\begin{aligned} \frac{\partial}{\partial \Lambda_{ij}} [\ln P(\mathbf{x}|y)] &= [M^{-1}\Lambda]_{ij} - [M^{-1}(\mathbf{x} - \boldsymbol{\mu})]_i [(\mathbf{x} - \boldsymbol{\mu})^T M^{-1}\Lambda]_j \quad (28) \end{aligned}$$

$$\frac{\partial}{\partial \Psi_{ii}} [\ln P(\mathbf{x}|y)] = \frac{1}{2} \{M_{ii}^{-1} - [M^{-1}(\mathbf{x} - \mathbf{u})]_i^2\} \quad (29)$$

¹If the number of classes, \mathcal{Y} , is very large, one can restrict the sum over y in (22) to the N -best alternatives, where $N \ll \mathcal{Y}$.

where $[\cdot]_i$ denotes the i th vector element. The inverse covariance matrix M^{-1} that appears in (27)–(29) can be computed efficiently using the lemma in (4). Naturally, these gradients also reduce to the correct form for diagonal covariance matrices, when no factors are present.

The parameter updates for class-conditional Gaussian PDF's are computed from (24) and (25). In practice, one chooses a positive-definite matrix, $\eta(\Phi)$, to ensure that each parameter update has roughly the same magnitude effect on the MCE loss function. This is achieved by scaling the $((\partial/\partial\mu), (\partial/\partial\Lambda))$ gradients by the variance, Ψ . Also, as is standard practice [12], the variance parameters are updated in the log domain to enforce the constraint of nonnegativity.

C. Hidden Markov Models

In this section, we consider how discriminative training can be applied to FA-HMM's for automatic speech recognition. Suppose that FA-HMM's are being used to model subword units, such as context-dependent phones. Training data in this setting consists of variable-length utterances and their segmentations into strings of phones or sub-phones. In particular, the segmentations label each frame of speech by the state of a particular HMM. We can naturally view these labels as *targets* for the Viterbi decoding of the recognizer. Learning to match these targets is a form of MCE parameter estimation, where the class-conditional distributions of (21) are the state-emission densities of HMM's, and the hypotheses in (22) range over the possible segmentations of each recognized utterance. Here we consider the special case where the feature vectors, conditioned on the hidden states, are modeled by mixtures of Gaussian PDF's with factored covariance matrices.

For this type of discriminative training, we must compute the gradients of the state-emission densities with respect to their parameters. In FA-HMM's, these distributions are given by (5) and (6). The gradients for mixtures of factor analyzers are a straightforward extension of our previous results. In particular, consider a segmented utterance in which the feature vector \mathbf{x} is assigned to state s of an FA-HMM. Also, let $P(c|\mathbf{x}, s)$ denote the posterior probability that within state s , the feature vector \mathbf{x} is attributed to mixture component c . Then the required gradients are

$$\begin{aligned} & \left(\frac{\partial}{\partial \mu_{sc}}, \frac{\partial}{\partial \Lambda_{sc}}, \frac{\partial}{\partial \Psi_{sc}} \right) [\ln P(\mathbf{x}|s)] \\ &= P(c|\mathbf{x}, s) \left(\frac{\partial}{\partial \mu_{sc}}, \frac{\partial}{\partial \Lambda_{sc}}, \frac{\partial}{\partial \Psi_{sc}} \right) [\ln P(\mathbf{x}|s, c)]. \end{aligned} \quad (30)$$

The gradients on the right hand side of these equations are given by (27)–(29) from the previous section. Besides these parameters, one may also adapt the mixture weights, $P(c|s)$, by gradient descent. In practice, one adapts the transformed parameters w_{sc} , where $P(c|s) = e^{w_{sc}} / \sum_k e^{w_{sk}}$; this is done to enforce the constraints $P(c|s) \geq 0$ and $\sum_c P(c|s) = 1$. In this case, the necessary gradients are given by

$$\frac{\partial}{\partial w_{sc}} [\ln P(\mathbf{x}|s)] = P(c|\mathbf{x}, s) - P(c|s). \quad (31)$$

Given these gradients, the parameter updates are again computed from (24) and (25). For each state and mixture compo-

nent, one considers a different positive-definite matrix, $\eta(\Phi)$, that multiplies the gradients in the learning rule. In particular, this matrix is used to scale the $((\partial/\partial\mu_{sc}), (\partial/\partial\Lambda_{sc}))$ gradients by the variances, Ψ_{sc} . In the experiments that follow, we computed the MCE loss function for each utterance, and updated the parameters on an utterance-by-utterance basis. In addition, the softmax in (22) was approximated by considering only the four best alternatives to the correct segmentation. In these respects and others, we followed the procedure outlined in earlier work [12].

V. EXPERIMENTS

Continuous density HMM's were evaluated on two tasks in automatic speech recognition, one small-sized vocabulary (36 words) and one medium-sized vocabulary (1219 words). The same front end was applied to both tasks. For signal processing, waveforms were pre-emphasized and blocked into 30 ms frames at every 10 ms interval. For feature extraction, frames were Hamming windowed, autocorrelated, and processed by LPC cepstral analysis to produce a vector of twelve lifted cepstral coefficients. The feature vector was then augmented by its normalized log energy value, as well as temporal derivatives of first and second order. Overall, each frame of speech was described by thirty nine features. These input features were fed directly to the HMM's described below.

A. Connected Alpha-Digits

As a small vocabulary task, we considered the recognition of alphanumeric strings (e.g., N Z 3 V J 4 E 3 U 2). Easily confused letters such as B/P, C/Z, and M/N make this a challenging problem in speech recognition. The training and test data [25] for this task were recorded over a telephone network and consisted of 14 622 and 7255 utterances, respectively. The speech recognizers were built using 285 left-to-right HMM's, each of which modeled a context-dependent sub-word unit. Nonzero transition probabilities were fixed to uniform values. Testing was done with a free grammar network (i.e., no grammar constraints).

We trained continuous density HMM's with both diagonal and factored covariance matrices. Our first experiments considered the maximum likelihood (ML) training procedure described in Section III-B. We ran several experiments, varying both the number of mixture components and the number of factors allocated to each state in the HMM's. The goal was to determine the best model of acoustic feature correlations.

Table I summarizes the results of these experiments. The columns from left to right show the number of mixture components (C), the number of factors (f), the word error rates (including insertion, deletion, and substitution errors) on the test set, the average log-likelihood per frame of speech on the test set, and the CPU time to recognize twenty test utterances (on an SGI R4000). Not surprisingly, given the large amount of training data, both word accuracies and likelihood scores increase with the number of modeling parameters. Nevertheless, the table shows that adding factors leads to significant improvements in performance, at roughly the same rate as adding mixture components.

TABLE I
RESULTS FOR DIFFERENT ML
RECOGNIZERS ON CONNECTED ALPHADIGITS. THE COLUMNS INDICATE THE
NUMBER OF MIXTURE COMPONENTS (C), THE NUMBER OF FACTORS (f), THE
WORD ERROR RATES AND AVERAGE LOG-LIKELIHOOD SCORES ON THE TEST
SET, AND THE CPU TIME TO RECOGNIZE 20 UTTERANCES

C	f	word error (%)	log-likelihood	CPU time (sec)
1	0	16.2	32.9	25
1	1	14.6	34.2	30
1	2	13.7	34.9	30
1	3	13.0	35.3	38
1	4	12.5	35.8	39
2	0	13.4	34.0	30
2	1	12.0	35.1	44
2	2	11.4	35.8	48
2	3	10.9	36.2	61
2	4	10.8	36.6	67
4	0	11.5	34.9	46
4	1	10.4	35.9	80
4	2	10.1	36.5	93
4	3	10.0	36.9	132
4	4	9.8	37.3	153
8	0	10.2	35.6	93
8	1	9.7	36.5	179
8	2	9.6	37.0	226
16	0	9.5	36.2	222

Perhaps the most interesting comparisons are between models with the same number of parameters. The overall number of parameters is proportional to $C(f + 2)$. A useful rule of thumb is that FA-HMM's with two factors have the same number of parameters as DG-HMM's with twice as many mixture components. The left graph in Fig. 1 shows a plot of the average log-likelihood versus the number of parameters, or $C(f + 2)$; the stars and circles in this plot indicate recognizers with DG-HMM's and FA-HMM's, respectively. One sees quite clearly from this plot that given a *fixed* number of parameters, models with factored covariance matrices tend to have significantly higher likelihoods. The right graph in Fig. 1 shows a similar plot of the word error rates versus the number of parameters. Here one does not see much difference; presumably, because HMM's represent only approximately correct models of speech, higher likelihoods do not necessarily translate into lower error rates. We will return to this point later.

It is worth noting that the above experiments used a fixed number of factors per mixture component. In fact, because the variability of speech is highly context-dependent, it makes sense to vary the number of factors, even across states within the same HMM. A simple heuristic is to adjust the number of factors depending on the amount of training data for each state (as determined by an initial segmentation of the training utterances). We found that this heuristic led to more pronounced differences in likelihood scores and error rates. In particular, substantial improvements were observed for three recognizers whose HMM's employed an *average* of two factors per mixture component. Table II summarizes these results. The reader will notice that these recognizers are extremely competitive in all aspects of performance—accuracy, memory, and speed—with the baseline recognizers in Table I. The likelihoods and error rates of these recognizers are indicated by the dashed lines in Fig. 1.




Fig. 1. Plots of (a) log-likelihood scores and (b) word error rates on the test set versus the number of parameters. The stars, connected by solid lines, indicate DG-HMM's; the circles indicate FA-HMM's. The dashed lines connect the recognizers in Table II.

TABLE II
RESULTS FOR ML RECOGNIZERS ON CONNECTED ALPHADIGITS
WITH VARIABLE NUMBERS OF FACTORS; f DENOTES THE
AVERAGE NUMBER OF FACTORS PER MIXTURE COMPONENT

C	f	word error (%)	log-likelihood	CPU time (sec)
1	2	12.3	35.4	32
2	2	10.5	36.3	53
4	2	9.6	37.0	108

All the recognizers in Tables I and II were trained by maximum likelihood (ML) estimation. The plots in Fig. 1, however, show that increased log-likelihoods in FA-HMM's do not translate directly into lower error rates. This observation motivated us to train a number of recognizers using the minimum classification error (MCE) criterion described in Section IV. This was done for the three DG-HMM recognizers in Table I and the three FA-HMM recognizers in Table II. In each case, the parameters of the MCE recognizer were initialized by those of its ML counterpart. Training consisted of five iterations of gradient descent, where each iteration involved a pass through all the utterances in the




Fig. 2. Word accuracies of ML, MCE, and MCE+SBR recognizers on connected alphas digits shown in black, gray, and white. Each subplot compares recognizers with the same overall number of parameters, but with different types of covariance matrices (FA versus DG).

training set. The test set error rates of these recognizers, before and after discriminative training, are shown in Fig. 2. The figure shows that all the recognizers had significantly lower error rates as a result of MCE training. At the same time, the recognizers using factor analysis retained their edge in accuracy. This shows that MCE training and factor analysis are complementary methods for improving overall performance. Fig. 2 also shows the results of a further training procedure—hierarchical signal bias removal (SBR) [22]—designed to reduce the acoustic mismatch between training and testing environments. This method led to further improvements in recognition accuracy.

B. New Jersey Town Names

As a medium vocabulary task, we considered the recognition of New Jersey town names (e.g., Cranbury). Telephone speech from two different collections [25] was used in this experiment. The training data consisted of 12 100 short phrases, each 2–4 words long, spoken in the seven major dialects of American English. This collection of phrases was carefully designed to obtain a fairly even coverage over triphone units. The test data consisted of 2426 isolated utterances of New Jersey town names, collected from nearly 100 speakers. The number of town names was 1219. All the speech data was digitized at the caller's local switch and transmitted in this form to the recognizers. The recognizers were built using 43 left-to-right HMM's, each of which modeled a context-independent English phone. Phones were modeled by three-state HMM's; the only exception to this rule was that silence and background noise were modeled

TABLE III
RESULTS ON NEW JERSEY TOWN NAMES. THE COLUMNS INDICATE THE NUMBER OF MIXTURE COMPONENTS, THE NUMBER OF FACTORS, AND THE ERROR RATES FOR ML AND MCE RECOGNIZERS

C	f	ML error (%)	MCE error (%)
4	0	19.0	17.4
4	1	17.6	15.0
8	0	16.8	14.5
8	1	15.0	12.5
16	0	14.6	12.6
16	1	13.8	12.8
32	0	13.6	12.0
32	1	12.7	11.7
64	0	12.6	11.7

by a single state. All nonzero transition probabilities were fixed to uniform values.

Again, we trained continuous density HMM's with both diagonal and factored covariance matrices. Since in the alphas digit experiments, the largest gains (per parameter) occurred with small numbers of factors, here we only considered FA-HMM's with exactly one factor per mixture component. The results of ML and MCE training for several recognizers are shown in Table III. Again we see that parameters devoted to the explicit modeling of correlations lead to sizable reductions in the error rate. Also, both the DG-HMM and FA-HMM recognizers are substantially improved by discriminative training. Fig. 3 shows plots of the classification error rate versus the number of modeling parameters [i.e., $C(f+2)$] for the ML and MCE rec-




Fig. 3. Plots of town name error rates versus the number of parameters for ML and MCE recognizers. The solid and dashed lines indicate recognizers with DG-HMM's and FA-HMM's, respectively.

ognizers. The rate of improvement in performance (as a function of model size) is roughly comparable for DG-HMM's and FA-HMM's.

VI. SUMMARY

In this paper, we have studied the combined use of mixture densities and factor analysis for speech recognition. This was done in the framework of hidden Markov modeling, where acoustic features are conditionally modeled by mixtures of Gaussian PDF's. On two tasks—connected alpha-digits and New Jersey town names—we have shown that mixture densities and factor analysis are complementary means of modeling acoustic correlations. Moreover, when used together, they can lead to smaller, faster, and more accurate recognizers than either method on its own. For example, comparing the last lines of Tables I and II reveals a factor of two improvement in speed and memory with hardly any loss in accuracy.

Both mixture models and factor analysis may be understood as latent variable methods—the former for clustering, the latter for dimensionality reduction. Just as mixture densities have

proved indispensable in automatic speech recognition, we believe that factor analysis can make a similarly important contribution. Both methods have a sound probabilistic framework that enables them to be fully exploited in statistical pattern recognition.

It is worth comparing factor analysis to other approaches [5], [10], [26], [28] proposed for dimensionality reduction and feature space analysis in automatic speech recognition. We believe that factor analysis has two distinguishing features. First, unlike many strategies for dimensionality reduction, factor analysis does not merely project a high dimensional feature vector into a low dimensional one; variations outside the reduced-dimensionality subspace are also modeled by the variance parameters, Ψ . Second, ML parameter estimation in factor analysis is supported by an efficient, batch-update EM algorithm; this is of considerable value even when the parameter estimates are ultimately refined by gradient-based discriminative training.

In this paper, we used factor analysis to model correlations between cepstra, delta-cepstra, and delta-delta-cepstra. It is worth emphasizing, however, that the method applies to arbitrary features. Indeed, the ability to model correlations efficiently should enable researchers to consider other features besides cepstra. While cepstra have the advantage of being only weakly correlated, it may be that other features (e.g., narrow-band statistics) actually convey more information about the speech signal.

We believe that factor analysis has many potential applications to automatic speech recognition. For example, in segmental HMM's [19], where feature vectors encode multiple (successive) frames of speech, factored covariance matrices provide a natural way to model correlations over time. In model adaptation [8], [15], [29], where parameters are reestimated to match new testing conditions, they provide a more robust alternative to adapting full covariance matrices. These are only two examples; in principle, factor analysis can be applied wherever one introduces Gaussian PDF's.

Factor analysis should be evaluated in light of other methods for modeling correlations. One popular method is to employ full covariance matrices through some form of parameter-tying. While we have presented factor analysis as an alternative to this approach, more generally it should be viewed as a complement. Certainly, the clever tying of factor loading matrices across units, states, and/or mixture components would lead to further improvements in overall performance (as a function of model size). Thus, whatever gains have been achieved by tying full covariance matrices, one would expect *additional* gains to be achieved from tying factored ones.

Overall performance in automatic speech recognition is measured in terms of speed, memory, and accuracy. Factor analysis can help in all three respects. Compared to full covariance matrices, factored ones are easier to manipulate and more robust to overfitting. In general, factor analysis represents a useful compromise between diagonal and full covariance matrices, one that promises improvements over both extremes.

APPENDIX A POSTERIOR DISTRIBUTION

In this Appendix, we consider the properties of the posterior distribution, $P(z|\mathbf{x})$, for the multivariate Gaussian PDF men-

tioned at the end of Section II-A. The posterior distribution is computed from Bayes' rule as

$$P(\mathbf{z}|\mathbf{x}) = \frac{P(\mathbf{x}|\mathbf{z})P(\mathbf{z})}{P(\mathbf{x})}. \quad (32)$$

The three terms on the right hand side are given explicitly by (1)–(3). Because $P(\mathbf{z})$ and $P(\mathbf{z}|\mathbf{x})$ are Gaussian, it follows from the form of Bayes' rule that $P(\mathbf{z}|\mathbf{x})$ is also Gaussian. In particular, up to a normalization constant, the posterior distribution is given by

$$P(\mathbf{z}|\mathbf{x}) \sim \exp \left\{ -\frac{1}{2} \mathbf{z}^T (I + \Lambda^T \Psi^{-1} \Lambda) \mathbf{z} + (\mathbf{x} - \boldsymbol{\mu})^T \Psi^{-1} \Lambda \mathbf{z} \right\}. \quad (33)$$

The statistics of the hidden variable \mathbf{z} , conditioned on the observation \mathbf{x} , are determined by the linear and quadratic terms in the exponent of (33). In particular, the conditional mean and covariance are given by

$$\mathbb{E}[\mathbf{z}|\mathbf{x}] = [I + \Lambda^T \Psi^{-1} \Lambda]^{-1} \Lambda^T \Psi^{-1} (\mathbf{x} - \boldsymbol{\mu}) \quad (34)$$

$$\mathbb{E}[\boldsymbol{\delta} \mathbf{z} \boldsymbol{\delta}^T | \mathbf{x}] = [I + \Lambda^T \Psi^{-1} \Lambda]^{-1} \quad (35)$$

where $\mathbb{E}[\cdot|\mathbf{x}]$ denotes an average with respect to the posterior distribution, $P(\mathbf{z}|\mathbf{x})$, and $\boldsymbol{\delta} \mathbf{z} = \mathbf{z} - \mathbb{E}[\mathbf{z}|\mathbf{x}]$ denotes the variation about the conditional mean. Note that the right hand side of (35) does not depend on the value of \mathbf{x} . As a practical matter, this means that the conditional covariance matrix $\mathbb{E}[\boldsymbol{\delta} \mathbf{z} \boldsymbol{\delta}^T | \mathbf{x}]$ does not have to be recomputed for each data point.

APPENDIX B EM ALGORITHM

In this Appendix, we provide a more detailed derivation of the EM algorithm in Section III-A. We begin by evaluating the Q -function, (7). This is done by averaging $\ln P(\mathbf{z}, \mathbf{x}_n | \tilde{\boldsymbol{\mu}}, \tilde{\Lambda}, \tilde{\Psi})$ over the posterior distribution, $P(\mathbf{z}|\mathbf{x}_n, \boldsymbol{\mu}, \Lambda, \Psi)$, then taking the sum over all data points. Up to a constant term (which does not depend on $\tilde{\boldsymbol{\mu}}, \tilde{\Lambda}$, or $\tilde{\Psi}$), the Q -function is given by

$$\begin{aligned} Q(\tilde{\boldsymbol{\mu}}, \tilde{\Lambda}, \tilde{\Psi}; \boldsymbol{\mu}, \Lambda, \Psi) &= -\frac{1}{2} \sum_n \mathbb{E}[(\mathbf{x}_n - \tilde{\boldsymbol{\mu}} - \tilde{\Lambda} \mathbf{z})^T \tilde{\Psi}^{-1} (\mathbf{x}_n - \tilde{\boldsymbol{\mu}} - \tilde{\Lambda} \mathbf{z}) | \mathbf{x}_n] \\ &\quad - \frac{N}{2} \ln |\tilde{\Psi}| + \text{const} \end{aligned} \quad (36)$$

where $\mathbb{E}[\cdot|\mathbf{x}_n]$ denotes an expectation over the posterior distribution, $P(\mathbf{z}|\mathbf{x}_n, \boldsymbol{\mu}, \Lambda, \Psi)$. Let $\boldsymbol{\delta} \mathbf{z} = \mathbf{z} - \mathbb{E}[\mathbf{z}|\mathbf{x}_n]$ denote the variation about the conditional mean of the hidden variable. We can decompose the terms inside the expectation as

$$\begin{aligned} Q(\tilde{\boldsymbol{\mu}}, \tilde{\Lambda}, \tilde{\Psi}; \boldsymbol{\mu}, \Lambda, \Psi) &= -\frac{1}{2} \sum_n (\mathbf{x}_n - \tilde{\boldsymbol{\mu}} - \tilde{\Lambda} \mathbb{E}[\mathbf{z}|\mathbf{x}_n])^T \tilde{\Psi}^{-1} \\ &\quad \cdot (\mathbf{x}_n - \tilde{\boldsymbol{\mu}} - \tilde{\Lambda} \mathbb{E}[\mathbf{z}|\mathbf{x}_n]) \\ &\quad - \frac{1}{2} \sum_n \mathbb{E}[(\tilde{\Lambda} \boldsymbol{\delta} \mathbf{z})^T \tilde{\Psi}^{-1} (\tilde{\Lambda} \boldsymbol{\delta} \mathbf{z}) | \mathbf{x}_n] - \frac{N}{2} \ln |\tilde{\Psi}| + \text{const}. \end{aligned} \quad (37)$$

The M-step of the EM algorithm is to maximize this expression with respect to $\tilde{\boldsymbol{\mu}}, \tilde{\Lambda}$, and $\tilde{\Psi}$. Setting $\partial Q / \partial \tilde{\boldsymbol{\mu}} = 0$ gives at once

the reestimation formulae for the mean, (13). Setting $\partial Q / \partial \tilde{\Lambda} = 0$ gives

$$\begin{aligned} \tilde{\Lambda} \left(\sum_n \mathbb{E}[\boldsymbol{\delta} \mathbf{z} \boldsymbol{\delta}^T | \mathbf{x}_n] \right) &= \sum_n (\mathbf{x}_n - \tilde{\boldsymbol{\mu}} - \tilde{\Lambda} \mathbb{E}[\mathbf{z}|\mathbf{x}_n])^T \mathbb{E}[\mathbf{z}|\mathbf{x}_n]. \end{aligned} \quad (38)$$

It is straightforward to eliminate $\tilde{\boldsymbol{\mu}}$ from this equation using (13). This yields the reestimation formula for the factor loading matrix, (12). Finally, setting $\partial Q / \partial \tilde{\Psi}_{(ii)} = 0$ gives

$$\begin{aligned} \tilde{\Psi}_{ii} &= \frac{1}{N} \sum_n [(\mathbf{x}_n - \tilde{\boldsymbol{\mu}} - \tilde{\Lambda} \mathbb{E}[\mathbf{z}|\mathbf{x}_n])_i^2 \\ &\quad + (\tilde{\Lambda} \mathbb{E}[\boldsymbol{\delta} \mathbf{z} \boldsymbol{\delta}^T | \mathbf{x}_n] \tilde{\Lambda}^T)_{ii}]. \end{aligned} \quad (39)$$

Here again, eliminating $\tilde{\boldsymbol{\mu}}$ via (14) leads to the desired result—in this case, the reestimation formula for the variance, (14).

APPENDIX C GRADIENTS

In this Appendix, we derive the gradients required for discriminative training of FA-HMM's, namely (27)–(29). To this end, consider a multivariate Gaussian PDF with mean $\boldsymbol{\mu}$ and factored covariance matrix $\Psi + \Lambda \Lambda^T$. If we denote the covariance matrix by M , then the log-likelihood is given by

$$\ln P(\mathbf{x}) = \frac{1}{2} [\ln \det M - (\mathbf{x} - \boldsymbol{\mu})^T M^{-1} (\mathbf{x} - \boldsymbol{\mu})]. \quad (40)$$

For discriminative training, we must compute the gradients of $\ln P(\mathbf{x})$ with respect to the parameters $\boldsymbol{\mu}$, Ψ , and Λ . The first of these is straightforward and gives (27). For the others, we can use the chain rule—differentiating the right hand side of (40) with respect to M , then differentiating M with respect to Ψ and Λ . The first step is facilitated by the identities

$$\frac{\partial}{\partial M_{ij}} [\ln \det M] = M_{ji}^{-1} \quad (41)$$

$$\frac{\partial}{\partial M_{ij}} [M_{kl}^{-1}] = -M_{ki}^{-1} M_{jl}^{-1}. \quad (42)$$

Using these identities to differentiate (40), we obtain

$$\begin{aligned} \frac{\partial}{\partial M_{ij}} [\ln P(\mathbf{x})] &= \frac{1}{2} \{ M_{ji}^{-1} - [M^{-1} (\mathbf{x} - \boldsymbol{\mu})]_i [M^{-1} (\mathbf{x} - \boldsymbol{\mu})]_j \}. \end{aligned} \quad (43)$$

From the definition $M = \Psi + \Lambda \Lambda^T$, it is straightforward to compute the gradients $\partial M / \partial \Psi$ and $\partial M / \partial \Lambda$. Finally, inserting these gradients into the chain rule gives (28) and (29).

ACKNOWLEDGMENT

The authors are grateful to A. Ljolje, AT&T Labs, and Z. Ghahramani, University College, London, U.K., for useful discussions. They also thank P. Modi, AT&T Labs, for providing an initial segmentation of the training utterances.

REFERENCES

- [1] L. Bahl, P. Brown, P. deSouza, and L. Mercer, "Maximum mutual information estimation of hidden Markov model parameters for speech recognition," in *Proc. ICASSP*, vol. 86, 1986, pp. 49–52.

- [2] J. Bellegarda and D. Nahamoo, "Tied mixture continuous parameter modeling for speech recognition," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 38, pp. 2033–2045, 1990.
- [3] C. Bregler and S. Omohundro, "Nonlinear image interpolation using manifold learning," in *Advances in Neural Information Processing Systems*, G. Tesauro, D. Touretzky, and T. Leen, Eds. Cambridge, MA: MIT Press, 1995, vol. 7, pp. 971–980.
- [4] M. Á. Carreira-Perpiñán and S. J. Renals, "Dimensionality reduction of electropalatographic data using latent variable models," *Speech Commun.*, vol. 26, pp. 259–282, 1998.
- [5] R. Chengalvarayan and L. Deng, "HMM-based speech recognition using state-dependent, discriminatively-derived transforms on Mel-warped DFT features," *IEEE Trans. Speech Audio Processing*, vol. 5, pp. 243–256, 1997.
- [6] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.
- [7] B. Everitt, *An Introduction to Latent Variable Models*. London, U.K.: Chapman and Hall, 1984.
- [8] J. Gauvain and C. Lee, "Maximum a posteriori estimation for multivariate Gaussian mixture observations of Markov chains," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 2, pp. 291–298, 1994.
- [9] Z. Ghahramani and G. Hinton, "The EM algorithm for mixtures of factor analyzers," Univ. Toronto, Toronto, Ont., Canada, Tech. Rep. CRG-TR-96-1, 1996.
- [10] H. Haeb-Umbach and H. Ney, "Linear discriminant analysis for improved large vocabulary continuous speech recognition," in *Proc. ICASSP '92*, 1992, pp. 13–16.
- [11] G. Hinton, P. Dayan, and M. Revow, "Modeling the manifolds of images of handwritten digits," *IEEE Trans. Neural Networks*, vol. 8, pp. 65–74, 1997.
- [12] B. H. Juang, W. Chou, and C. H. Lee, "Minimum classification error rate methods for speech recognition," *IEEE Trans. Speech Audio Processing*, vol. 5, pp. 266–277, 1997.
- [13] B. Juang and S. Katagiri, "Discriminative learning for minimum error classification," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 40, pp. 3043–3054, 1992.
- [14] N. Kambhatla and T. Leen, "Fast nonlinear dimension reduction," in *Advances in Neural Information Processing Systems*, J. Cowan, G. Tesauro, and J. Alspector, Eds. San Mateo, CA: Morgan Kaufman, 1994, vol. 6, pp. 152–159.
- [15] C. Leggetter and P. Woodland, "Maximum likelihood linear regression for speaker adaptation of continuous density hidden Markov models," *Comput., Speech, Lang.*, vol. 9, pp. 171–185, 1995.
- [16] A. Ljolje, "The importance of cepstral parameter correlations in speech recognition," *Comput., Speech, Lang.*, vol. 8, pp. 223–232, 1994.
- [17] A. Ljolje, Y. Ephraim, and L. Rabiner, "Estimation of hidden Markov model parameters by minimizing empirical error rate," in *Proc. ICASSP 1990*, pp. 709–712.
- [18] E. Oja, *Subspace Methods of Pattern Recognition*. Cambridge, U.K.: Research Studies, 1983.
- [19] M. Ostendorf, V. Digalakis, and O. Kimball, "From HMM's to segment models: A unified view of stochastic modeling for speech recognition," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 4, pp. 360–378, 1996.
- [20] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge, U.K.: Cambridge Univ. Press, 1992.
- [21] L. Rabiner and B. Juang, *Fundamentals of Speech Recognition*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [22] M. Rahim, B. Juang, W. Chou, and E. Buhrke, "Signal conditioning techniques for robust speech recognition," *IEEE Signal Processing Lett.*, vol. 3, pp. 107–109, 1996.
- [23] S. Roweis, "EM algorithms for PCA and SPCA," in *Advances in Neural Information Processing Systems*, M. Jordan, M. Kearns, and S.olla, Eds. Cambridge, MA: MIT Press, 1998, vol. 10, pp. 626–632.
- [24] D. Rubin and D. Thayer, "EM algorithms for factor analysis," *Psychometrika*, vol. 47, pp. 69–76, 1982.
- [25] R. Sachs, M. Tikjian, and E. Roskos, "United States English subword speech data," unpublished AT&T rep., 1994.
- [26] E. Schukat-Talamazzini, J. Hornegger, and H. Niemann, "Optimal linear feature transformations for semi-continuous hidden Markov models," in *Proc. ICASSP '95*, 1995, pp. 369–372.
- [27] M. E. Tipping and C. Bishop, "Mixtures of principal component analyzers," in *Proc. IEEE 5th Int. Conf. Artificial Neural Networks*, 1997.
- [28] H. Watanabe, T. Yamaguchi, and S. Katagiri, "Discriminative metric design for robust pattern recognition," *IEEE Trans. Signal Processing*, vol. 45, pp. 2655–2662, 1997.
- [29] G. Zavaliagos, R. Schwartz, and R. Makhoul, "Batch, incremental and instantaneous adaptation techniques for speech recognition," in *Proc. ICASSP 1995*, pp. 676–679.



Lawrence K. Saul received the B.A. degree in physics from Harvard University, Cambridge, MA, in 1990 and the Ph.D. degree in physics from the Massachusetts Institute of Technology, Cambridge, in 1994.

He is currently a Principal Technical Staff Member at AT&T Labs—Research, Florham Park, NJ, where he pursues research in speech and language processing, neural computation, and machine learning. Before joining AT&T, he was a Postdoctoral Fellow at the Center for Biological and Computational Learning, Massachusetts Institute of Technology. He is currently on the Editorial Board of the *Journal of Machine Learning*.

Dr. Saul served on the program committees for the 1997–1999 Conferences on Neural Information Processing Systems.



Mazin G. Rahim (S'86–M'91–SM'96) received the B.Eng. and Ph.D. degrees from the University of Liverpool, Liverpool, U.K., in 1987 and 1991, respectively.

He is currently a Division Manager at AT&T Labs—Research, Florham Park, NJ, where he is pursuing research in the areas of robustness, acoustic modeling, and utterance verification for automatic speech recognition. Prior to joining AT&T, he was a Research Professor with the Center for Computer Aids for Industrial Productivity, Rutgers University, New Brunswick, NJ, where he was engaged in research in neural networks for speech and speaker recognition. He has more than 50 publications in the area of speech processing and is the author of the book *Artificial Neural Networks for Speech Analysis/Synthesis* (London, U.K.: Chapman & Hall, 1994).

Dr. Rahim was an Associate Editor for the IEEE TRANSACTIONS ON SPEECH AND AUDIO PROCESSING from 1995 to 1999. He was a Chair of the 1999 Workshop on Automatic Speech Recognition and Understanding.