Large-Vocabulary Speaker-Independent Continuous Speech Recognition with Semi-Continuous Hidden Markov Models

X.D. Huang, H.W. Hon, and K.F. Lee*

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
USA

ABSTRACT

A semi-continuous hidden Markov model based on the multiple vector quantization codebooks is used here for large-vocabulary speaker-independent continuous speech recognition. In the techniques employed here, the semi-continuous output probability density function for each codebook is represented by a combination of the corresponding discrete output probabilities of the hidden Markov model and the continuous Gaussian density functions of each individual codebook. Parameters of vector quantization codebook and hidden Markov model are mutually optimized to achieve an optimal model/codebook combination under a unified probabilistic framework. Another advantages of this approach is the enhanced robustness of the semi-continuous output probability by the combination of multiple codewords and multiple codebooks. For a 1000-word speaker-independent continuous speech recognition using a word-pair grammar, the recognition error rate of the semi-continuous hidden Markov model was reduced by more than 29% and 41% in comparison to the discrete and continuous mixture hidden Markov model respectively.

1. INTRODUCTION

In the discrete hidden Markov model (HMM), vector quantization (VQ) produces the closet codebword from the codebook for each acoustic observation. This mapping from continuous acoustic space to quantized discrete space may cause serious quantization errors for subsequent hidden Markov modeling. To reduce VQ errors, various smoothing techniques have been proposed for VQ and subsequent hidden Markov modeling [9,12]. A distinctive technique is multiple VQ codebook hidden Markov modeling, which has been shown to offer improved speech recognition accuracy [5,12]. In the multiple VQ codebook approach, VQ distortion can be significantly minimized by partitioning the parameters into separate codebooks. Another disadvantage of the discrete HMM is that the VQ codebook and the discrete HMM are separately modeled, which may not be an optimal combination for pattern classification [8]. The discrete HMM uses the discrete output probability distributions to model various acoustic events, which are inherently superior to the continuous mixture HMM with mixture of a small number of probability density functions since the discrete distributions could model events with any shapes provided enough training data exist.

On the other hand, the continuous mixture HMM models the acoustic observation directly using estimated continuous probability density functions without VQ, and has been shown to improve the recognition accuracy in comparison to the discrete HMM [15]. For speaker-independent speech recognition, mixture of a large number of probability density functions [14,16] or a large number of states in single-mixture case [4] are generally required to model characteristics of different speakers. However, mixture of a large number of probability density functions will considerably increase not only the computational complexity, but also the number of free parameters that can be reliablely estimated. In addition, the continuous mixture HMM has to be used with care as continuous probability density functions make more assumption than the discrete HMM, especially when the diagonal covariance Gaussian probability density is used for simplicity [15]. To obtain a better recognition accuracy, acoustic parameters must be well chosen according to the assumption of the continuous probability density functions used.

The semi-continuous hidden Markov model (SCHMM) has been proposed to extend the discrete HMM by replacing discrete output probability distributions with a combination of the original discrete output probability distributions and continuous probability density functions of a Gaussian codebook [6]. In the SCHMM, each VQ codeword is regarded as a Gaussian probability density. Intuitively, from the discrete HMM point of view, the SCHMM tries to smooth the discrete output probabilities with multiple codeword candidates in VQ procedure. From the continuous mixture HMM point of view, the SCHMM ties all the continuous output probability densities across each individual HMM to form a shared Gaussian codebook, i.e. a mixture of Gaussian probability densities. With the SCHMM, the codebook and HMM can be jointly re-estimated to achieve an optimal codebook/model combination in sense of maximum likelihood criterion. Such a tying can also substantially reduce the number of free parameters and computational complexity in comparison to the continuous mixture HMM, while maintain reasonablelv modeling power of a mixture of a large number of probability density functions. The SCHMM has shown to offer improved recognition accuracy in several speech recognition experiments [6,8,14,2].

In this study, the SCHMM is applied to Sphinx, a speaker-independent continuous speech recognition system. Sphinx uses multiple VQ codebooks for each acoustic observation [12]. To apply the SCHMM to Sphinx, the SCHMM algorithm must be modified to accommodate multiple codebooks and multiple codewords combination. For the SCHMM re-estimation algorithm, the modified unified re-estimation algorithm for multiple VQ codebooks and hidden Markov models are proposed in this paper. The applicability of the SCHMM to speaker-independent continuous speech is explored based on 200 generalized triphone models [12]. In the 1000-word speaker-independent continuous speech recognition task using word-pair grammar, the error rate was reduced by more than 29% and 41% in comparison to the corresponding discrete HMM and continuous mixture HMM respectively.

2. SEMI-CONTINUOUS HIDDEN MARKOV MODELS

2.1. Discrete HMMs and Continuous HMMs

An N-state Markov chain with state transition matrix $A = [a_{ij}]$, $i,j = 1, 2, \ldots, N$, where a_{ij} denotes the transition probability from state i to state j; and a discrete output probability distribution, $b_j(O_k)$, or continuous output probability density function $b_j(\mathbf{x})$ associated with each state j of the unobservable Markov chain is considered here. Here O_k represents discrete observation symbols (usually VQ indices), and $\mathbf{x}$ represents continuous observations (usually speech frame vectors) of K-dimensional random vectors.

With the discrete HMM, there are L discrete output symbols from a L-level VQ, and the output probability is modeled with discrete probability distributions of these discrete symbols. Let O be the observed sequence, $O = O_{k_1} O_{k_2} \cdots O_{k_T}$ observed over T samples. Here O_{k_t} denotes the VQ codeword k_t observed at time t. The

* and University of Edinburgh, 80 South Bridge, Edinburgh EH1 1HN, UK

observation probability of such an observed sequence, $Pr(\mathbf{O}|\lambda)$, can be expressed as:

$$Pr(\mathbf{O}|\lambda) = \sum_{S} \pi_{s_0} \prod_{t=1}^{T} a_{s_{t-1}s_t} b_{s_t}(O_{k_t}) \tag{1}$$

where S is a particular state sequence, $S \in (s_0 s_1 \cdots s_t)$, $s_t \in (1, 2, \ldots, N)$, and the summation is taken over all of the possible state sequences, S, of the given model λ, which is represented by (π, A, B), where π is the initial state probability vector, A is the state transition matrix, and B is the output probability distribution matrix. In the discrete HMM, classification of O_{k_t} from $\mathbf{x}_t$ in the VQ may not be accurate.

If the observation to be decoded is not vector quantized, then the probability density function, $f(\mathbf{X}|\lambda)$, of producing an observation of continuous vector sequences given the model λ, would be computed, instead of the probability of generating a discrete observation symbol, $Pr(\mathbf{O}|\lambda)$. Here X is a sequence of continuous acoustic vectors $\mathbf{x}$, $\mathbf{X} = \mathbf{x}_1 \mathbf{x}_2 \cdots \mathbf{x}_T$. The principal advantage of using the continuous HMM is the ability to directly model speech parameters without involving VQ. However, the continuous HMM requires considerably longer training and recognition times, especially when a mixture of several Gaussian probability density components is used. In the continuous Gaussian (M-component) mixture HMM [11], the output probability density of state j, $b_j(\mathbf{x})$, can be represented as

$$b_j(\mathbf{x}) = \sum_{k=1}^{M} c_{jk} N(\mathbf{x}, \mu_{jk}, \sum_{jk}) \tag{2}$$

where $N(\mathbf{x}, \mu, \sum)$ denotes a multi-dimensional Gaussian density function of mean vector μ and covariance matrix $\sum$, and c_{jk} is a weighting coefficient for the kth Gaussian component. With such a mixture, any arbitrary distribution can be approximately modeled, provided the mixture is large enough.

2.2. Semi-Continuous Hidden Markov Models

In the discrete HMM, the discrete probability distributions are sufficiently powerful to model any random events with a reasonable number of parameters. The major problem with the discrete output probability is that the VQ operation partitions the acoustic space into separate regions according to some distortion measure. This introduces errors as the partition operations may destroy the original signal structure. An improvement is to model the VQ codebook as a family of Gaussian density functions such that the distributions are overlaped, rather than disjointed. Each codeword of the codebook can then be represented by one of the Gaussian probability density functions and may be used together with others to model the acoustic event. The use of a parametric family of finite mixture densities (a mixture density VQ) can then be closely combined with the HMM methodology. From the continuous mixture HMM point of view, the output probability in the continuous mixture HMM is shared among the Gaussian probability density functions of the VQ. This can reduce the number of free parameters to be estimated as well as the computational complexity. From the discrete HMM point of view, the partition of the VQ is unnecessary, and is replaced by the mixture density modeling with overlap, which can effectively minimize the VQ errors. The procedure, known as the EM algorithm [3], is a specialization, to the mixture density context, of a general algorithm for obtaining maximum likelihood estimates. This has been defined earlier by Baum [1] in a similar way and has been widely used in HMM-based speech recognition methods. Thus, the VQ problems and HMM modeling problems can be unified under the same probabilistic framework to obtain an optimized VQ/HMM combination, which forms the foundation of the SCHMM.

Provided that each codeword of the VQ codebook is represented by a Gaussian density function, for a given state s_t of HMM, the probability density function that s_t produces a vector $\mathbf{x}$ can then be written as:

$$b_{s_t}(\mathbf{x}) = f(\mathbf{x}|s_t) = \sum_{j=1}^{L} f(\mathbf{x}|O_{j_t}, s_t) Pr(O_{j_t}|s_t) \tag{3}$$

where L denotes the VQ codebook level. For the sake of simplicity, the output probability density function conditioned on the codewords can be assumed to be independent of the Markov states s_t, (3) can then be written as:

$$f(\mathbf{x}|s_t) = \sum_{j=1}^{L} f(\mathbf{x}|O_{j_t}) Pr(O_{j_t}|s_t) = \sum_{j=1}^{L} f(\mathbf{x}|O_{j_t}) b_{s_t}(O_{j_t}) \tag{4}$$

This equation is the key to the semi-continuous hidden Markov modeling. Given the VQ codebook index O_{j_t}, the probability density function $f(\mathbf{x}|O_{j_t})$ can be estimated with the EM algorithm [17], or maximum likelihood clustering. It can also be obtained from the HMM parameter estimation directly as explained later. Using (4) to represent the semi-continuous output probability density, it is possible to combine the codebook distortion characteristics with the parameters of the discrete HMM under a unified probabilistic framework. Here, each discrete output probability is weighted by the continuous conditional Gaussian probability density function derived from VQ. If these continuous VQ density functions are considered as the continuous output probability density function in the continuous mixture HMM, this also resembles the L-mixture continuous HMM with all the continuous output probability density functions shared with each other in the VQ codebook. Here the discrete output probability in state i, $b_i(O_{j_t})$, becomes the weighting coefficients for the mixture components.

In implementation of the SCHMM [8], Eq. (4) can be replaced by finding M most significant values of $f(\mathbf{x}|O_j)$ (with M be one to six, the algorithm converges well in practice) over all possible codebook indices O_j, which can be easily obtained in the VQ procedure. This can significantly reduce the amount of computational load for subsequent output probability computation since M is of lower order than L. Experimental results show this to perform well in speech recognition [8], and result in an L-mixture continuous HMM with a computational complexity significantly lower than the continuous mixture HMM.

2.3. Re-estimation formulas for the SCHMM

If the $b_i(O_{j_t})$ are considered as the weighting coefficients of different mixture output probability density functions in the continuous mixture HMM, the re-estimation algorithm for the weighting coefficients can be extended to re-estimate $b_i(O_{j_t})$ of the SCHMM [11]. The re-estimation formulations can be more readily computed by defining a forward partial probability, $\alpha_t(i)$, and a backward partial probability, $\beta_t(i)$ for any time t and state i as:

$$\begin{aligned} \alpha_t(i) &= Pr(\mathbf{x}_1, \mathbf{x}_2, \cdots \mathbf{x}_t, s_t = i|\lambda) \\ \beta_t(i) &= Pr(\mathbf{x}_{t+1}, \mathbf{x}_{t+2}, \cdots \mathbf{x}_T|s_t = i, \lambda) \end{aligned} \tag{5 a}$$

The intermediate probabilities, $\chi_t(i,j,k)$, $\gamma_t(i,j)$, $\gamma_t(i)$, $\zeta_t(i,j)$, and $\zeta_t(j)$ can be defined as follows for efficient re-estimation of the model parameters:

$$\begin{aligned} \chi_t(i,j,k) &= Pr(s_t = i, s_{t+1} = j, O_{k_{t+1}}|\mathbf{X}, \lambda) \\ \gamma_t(i,j) &= Pr(s_t = i, s_{t+1} = j|\mathbf{X}, \lambda) \\ \gamma_t(i) &= Pr(s_t = i|\mathbf{X}, \lambda) \\ \zeta_t(i,k) &= Pr(s_t = i, O_{k_t}|\mathbf{X}, \lambda) \\ \zeta_t(k) &= Pr(O_{k_t}|\mathbf{X}, \lambda) \end{aligned} \tag{6}$$

All these intermediate probabilities can be represented by $\chi_t(\cdot)$. Using Eq. (5) and (6), the re-estimation equations for π_i, a_{ij}, and $b_i(O_j)$ can be written as:

$$\pi_i = \gamma_1(i), \quad 1 \le i \le N; \tag{7}$$

$$\bar{a}_{ij} = \frac{\sum_{t=1}^{T} \gamma_t(i,j)}{\sum_{t=1}^{T} \gamma_t(i)}, \quad 1 \le i, j \le N; \tag{8}$$

$$\bar{b}_i(O_j) = \frac{\sum_{t=1}^{T} \zeta_t(i,j)}{\sum_{t=1}^{T} \gamma_t(i)}, \quad 1 \le i \le N; \ 1 \le j \le L. \tag{9}$$

The means and covariances of the Gaussian probability density functions can also be re-estimated to update the VQ codebook separately with Eq. (5) and (6). The feedback from the HMM esti-

mation results to the VQ codebook implies that the VQ codebook is optimized based on the HMM likelihood maximization rather than minimizing the total distortion errors from the set of training data. Although re-estimation of means and covariances of different models will involve inter-dependencies, the different density functions which are re-estimated are strongly correlated. To re-estimate the parameters of the VQ codebook, i.e. the means, μ_j, and covariance matrices, Σ_j, of the codebook index j, it is not difficult to extend the continuous mixture HMM re-estimation algorithm with modified Q function. In general, it can be written as:

$$\bar{\mu}_j = \frac{\sum_c [\sum_{t=1}^{T} \zeta_t(j) \mathbf{x}_t]}{\sum_c [\sum_{t=1}^{T} \zeta_t(j)]}, \quad 1 \le j \le L; \tag{10}$$

and

$$\bar{\Sigma}_j = \frac{\sum_c [\sum_{t=1}^{T} \zeta_t(j) (\mathbf{x}_t - \bar{\mu}_j)(\mathbf{x}_t - \bar{\mu}_j)']}{\sum_c [\sum_{t=1}^{T} \zeta_t(j)]}, \quad 1 \le j \le L. \tag{11}$$

where c denotes the HMM used; and expressions in [ ] are variables of model c. In Eq. (10) and (11), the re-estimation for the means and covariance matrices in the output probability density function of the SCHMM are tied up with all the HMM models, which is similar to the approach with tied transition probability inside the model [10]. From Eq. (10) and (11), it can be observed that they are merely a special form of EM algorithm for the parameter estimation of mixture density functions [17], which are closely welded into the HMM re-estimation equations.

When multiple codebooks are used, each codebook represents a set of different speech parameters. One way to combine these multiple output observations is to assume that they are independent, and the output probability is computed as the product of the output probability of each codebook. It has been shown that performance using multiple codebook can be substationally improved [13]. In the semi-continuous HMM, the semi-continuous output probability of multiple codebooks can also be computed as the product of the semi-continuous output probability for each codebook as Eq. (4), which consists of L-mixture continuous density functions. In other word, the semi-continuous output probability could be modified as:

$$b_c(\mathbf{x}) = \prod_{c} \sum_{j=1}^{L} f(\mathbf{x} | O_{j_t}^c) b_{j_t}^c(O_{j_t}^c) \tag{12}$$

where c denotes the codebook used. The re-estimation algorithm for the multiple codebook based HMM could be extended if Eq. (6.a) is computed for each codeword of each codebook c with the combination of the rest codebook probability [7].

3. EXPERIMENTAL EVALUATION

3.1. Analysis Conditions

For both training and evaluation, the standard Sphinx front-end consists of 12th order bilinear transformed LPC cepstrum [12]. The complete database consists of 4358 training sentences from 105 speakers (june-train) and 300 test sentences from 12 speakers.

The vocabulary of the Resource Management database is 991 words. There is also an *official* word-pair recognition grammar, which is just a list of allowable word pairs without probabilities for the purpose of reducing the recognition perplexity to about 60.

3.2. Experimental Results Using Bilinear Transformed Cepstrum

Discrete HMMs and continuous mixture HMMs based on 200 generalized triphones are first experimented as benchmarks. The discrete HMM is the same as Sphinx except only 200 generalized triphones are used [12].

In the continuous mixture HMM implemented here, the cepstrum, difference cepstrum, normalized energy, and difference energy are packed into one vector. This is similar to the one codebook implementation of the discrete HMM [12]. Each continuous output probability consists of 4 diagonal Gaussian probability density function as Eq. (2). To obtain reliable initial models for the continuous mixture HMM, the Viterbi alignment with the discrete HMM is used to phonetically segment and label training speech. These labeled segments are then clustered by using the k-means clustering algorithm to obtain initial means and diagonal covariances. The forward-backward algorithm is used iteratively for the monophone models, which are then used as initial models for the generalized triphone models. Though continuous mixture HMM was reported to significantly better the performance of the discrete HMM [15], for the experiments conducted here, it is *significantly* worse than the discrete HMM. Why is this paradox? One explanation is that multiple codebooks are used in the discrete HMM, therefore the VQ errors for the discrete HMM are not so serious here. Another reason may be that the diagonal covariance assumption is not appropriate for the bilinear transformed LPC cepstrum since many coefficients are strongly correlated after the transformation. Indeed, observation of average covariance matrix for the bilinear transformed LPC cepstrum shows that values of off-diagonal components are generally quite large.

For the semi-continuous model, multiple codebooks are used instead of packing different feature parameters into one vector. The initial model for the SCHMM comes directly from the discrete HMM with the VQ variance obtained from k-means clustering for each codeword. In computing the semi-continuous output probability density function, only the M (1, 4 here) most significant codewords are used for subsequent processing. Under the same analysis condition, the percent correct (correct word percentage) and word accuracy (percent correct - percent insertion) results of the discrete HMM, the continuous mixture HMM, and the SCHMM are shown in Table 1.

Table 1
Average recognition accuracy

types	*percent correct (word accuracy)*
Discrete HMM	89.5% (88.0%)
Continuous Mixture HMM	84.2% (81.3%)
SCHMM - top1	87.2% (84.0%)
SCHMM - top4	90.6% (89.1%)

From Table 1, it can be observed that the SCHMM with top 4 codewords works better than both the discrete and continuous mixture HMM. The SCHMM with top 1 codeword works actually worse than the discrete HMM, which indicates that diagonal Gaussian assumption may be inappropriate here. Though bilinear transformed cepstral coefficients could not be well modeled by the diagonal Gaussian assumption (which was proven by the poor performance of the continuous mixture HMM and the SCHMM with Gaussian assumption (which was proven by the poor performance of the continuous mixture HMM and the SCHMM with top 1 codeword), the SCHMM with top 4 codewords works modestly better than the discrete HMM. The improvement may primarily come from smoothing effect of the SCHMM, i.e. the robustness of multiple codewords and multiple codebooks in the semi-continuous output probability representation, albeit 200 generalized triphone models are relatively well trained in comparison to standard Sphinx version [12], where 1000 generalized triphone models are used.

3.3. Experimental Results Using Less Correlated Data

If the diagonal Gaussian covariance is used, each dimension in speech vector should be un-correlated. In practice, this can be partially satisfied by using less correlated feature as acoustic observation representation. One way to reduce correlation is principal component projection. In the implementation here, the projection matrix is computed by first pooling together the bilinear transformed cepstrum of the whole training sentences, and then computing the eigenvector of that pooled covariance matrix. Unfortunately, only insignificant improvements are obtained based on such a projection [7]. This is because the covariance for each codeword is quite different, and such a projection only makes *average* covariance diagonal, which is inadequate.

As bilinear transformed cepstral coefficients could not be well modeled by diagonal Gaussian probability density function, experiments without bilinear transformation are conducted. The 18th order cepstrum is used here for the SCHMM because of less correlated characteristics of the cepstrum. With 4358 training sentences (june-train), test results of 300 sentences (june-test) are listed in Table 2.

Table 2
Average accuracy of 18th order cepstrum

types	percent correct (word accuracy)
Discrete HMM	86.3% (83.8%)
SCHMM + top1	86.6% (85.5%)
SCHMM + top2	88.8% (87.6%)
SCHMM + top4	89.3% (88.5%)
SCHMM + top6	89.6% (88.6%)
SCHMM + top8	89.3% (88.2%)

Here, the recognition accuracy of the SCHMM is significantly improved in comparison with the discrete HMM, and error reduction is over 29%. Even the SCHMM with top one codeword is used, it is still better than the discrete HMM (85.5% vs. 83.8%). Use of multiple codewords (top4 and top6) in the semi-continuous output probability density function greatly improves the word accuracy (from 85.5% to 88.6%). Further increase of codewords used in the semi-continuous output probability density functions shows no improvement on word accuracy, but substantial growth of computational complexity. From Table 2, it can be seen that the SCHMM with top four codewords is adequate (88.5%). In contrast, when bilinear transformed data was used, the error reduction is less than 10% in comparison to the discrete HMM, and the SCHMM with top one codeword is actually slightly worse than the discrete HMM. This strongly indicates that appropriate feature is very important if continuous probability density function, especially diagonal covariance assumption, is used. If assumption is inappropriate, maximum likelihood estimation will only maximize the *wrong* assumption. Although more than 29% error reduction has been achieved for 18th order LPC analysis using diagonal covariance assumption, the last results with the discrete HMM (bilinear transformed cepstrum, 88.3%) and the SCHMM (18th order cepstrum, 88.6%) are about the same. This suggest that bilinear transformation is helpful for recognition, but have correlated coefficients, which is inappropriate to the diagonal Gaussian assumption. It can be expected that with the full covariance SCHMM and bilinear transformed cepstral data, better recognition accuracy can be obtained.

4. CONCLUSIONS

Semi-continuous hidden Markov models based on multiple vector quantization codebooks take the advantages of both the discrete HMM and continuous HMM. With the SCHMM, it is possible to model a mixture of a large number of probability density functions with a limited amount of training data and computational complexity. Robustness is enhanced by using multiple codewords and multiple codebooks for the semi-continuous output probability representation. In addition, the VQ codebook itself can be adjusted together with the HMM parameters in order to obtain the optimum maximum likelihood of the HMM. The applicability of the continuous mixture HMM or the SCHMM relies on appropriately chosen acoustic parameters and assumption of the continuous probability density function. Acoustic features must be well represented if diagonal covariance is applied to the Gaussian probability density function. This is strongly indicated by the experimental results based on the bilinear transformed cepstrum and cepstrum. With bilinear transformation, high frequency components are compressed in comparison to low frequency components [2,3]. Such a transformation converts the linear frequency axis into a mel-scale-like one. The discrete HMM can be substantially improved by bilinear transformation. However, bilinear transformation introduces strong correlations, which is inappropriate for the diagonal Gaussian assumption modeling. Using the cepstrum without bilinear transformation, the diagonal SCHMM can be substantially improved in comparison to the discrete HMM.

All experiments conducted here were based on only 200 generalized triphones as smoothing can play a more important role in those less-well-trained models, more improvement can be expected for 1000 generalized triphones (where the word accuracy for the discrete HMM is 91% with bilinear transformed data). In addition, removal of diagonal covariance assumption by use of full covariance can be expected to further improve recognition accuracy [1]. Regarding use of full covariance, the SCHMM has a distinctive advantage. Since Gaussian probability density functions are tied to the VQ codebook, by chosing M most significant codewords, computational complexity can be several order lower than the conventional continuous mixture HMM while maintaining the modeling power of large mixture components.

Experimental results have clearly demonstrated that the SCHMM offers improved recognition accuracy in comparison to both the discrete HMM and the continuous mixture HMM in speaker-independent continuous speech recognition. We conclude that the SCHMM is indeed a powerful technique for modeling non-stationary stochastic processes with multi-modal probabilistic functions of Markov chains.

ACKNOWLEDGEMENTS

We would like to thank Professor Mervyn Jack and Professor Raj Reddy for their help and insight shared in this research.

REFERENCES

[1] Baum, L.E. and *et al.*, "A maximization technique occuring in the statistical analysis of probabilistic functions of Markov chains", J. Ann. Math. Stat. Vol. 41, pp. 164-171, 1970.

[2] Bellegarda, J. and Nahamoo, D., "Tied mixture continuous parameter models for large vocabulary isolated speech recognition", ICASSP 89, Glasgow, Scotland, 1989.

[3] Dempster, A.P., Laird, N.M., and Rubin, D.B. "Maximum-likelihood from incomplete data via the EM algorithm", J. Royal Statist. Soc. Ser. B (methodological), Vol. 39, pp. 1-38, 1977.

[4] Doddington, G.R., "Phonetically sensitive discriminants for improved speech recognition", ICASSP 89, Glasgow, Scotland, 1989.

[5] Gupta, V.N., Lennig, M., and Mermelstein, P. "Integration of Acoustic Information in a Large Vocabulary Word Recognizer", ICASSP 87, pp. 697-700, 1987.

[6] Huang, X.D. and Jack, M.A., "Hidden Markov modelling of speech based on a semi-continuous model", Electronics Letters, Vol. 24, pp. 6-7, 1988.

[7] Huang, X.D., Hon, H.W., and Lee, K.F., "Multiple codebook semi-continuous hidden Markov models for speaker-independent continuous speech recognition", CMU Technical Report CMU-CS-89-136, C. Science, 1989.

[8] Huang, X.D. and Jack, M.A., "Semi-continuous hidden Markov models for speech recognition", Computer Speech and Language, Vol. 3, 1989.

[9] Jelinek, F., "Continuous speech recognition by statistical methods", Proceedings of IEEE, Vol. 64, pp. 532-556, 1976.

[10] Jelinek, F. and Mercer, R.L., "Interpolated estimation of Markov source parameters from sparse data", The workshop on pattern recognition in practice, Amsterdam, 1980.

[11] Juang, B.H., "Maximum-likelihood estimation for mixture multivariate stochastic observations of Markov chain", AT&T Technical Journal, Vol. 64, pp. 1235-1249, 1985.

[12] Lee, K.F., "Large-vocabulary speaker-independent continuous speech recognition: The SPHINX system", Ph.D. thesis, C. Science, CMU, 1988.

[13] Lee, K.F., Hon, H.W., and Reddy, R., "The SPHINX speech recognition system", ICASSP 89, Glasgow, Scotland, 1989.

[14] Paul, D., "The Lincoln continuous speech recognition system: recent developments and results", DARPA 1989 Feb. meeting, 1989.

[15] Rabiner, L.R. and *et al.*, "Recognition of isolated digits using hidden Markov models with continuous mixture densities", AT&T Technical Journal, Vol. 64, pp. 1211-1234, 1985.

[16] Rabiner, L.R., Wilpon, J.G., and Soong, F.K., "High Performance Connected Digit Recognition Using Hidden Markov Models", ICASSP 88, NY, 1988.

[17] Redner, R.A. and Walker, H.F., "Mixture densities, maximum likelihood and the EM algorithm", SIAM review, Vol. 26, pp. 195-239, 1984.

[18] Shikano, K., "Evaluation of LPC spectral matching measures for phonetic unit recognition", CMU Technical Report CMU-CS-86-108, C. Science, 1986.