

(21) Application No **0111902.3**
 (22) Date of Filing **16.05.2001**
 (30) Priority Data
 (31) **09576429** (32) **22.05.2000** (33) **US**

(71) Applicant(s)
International Business Machines Corporation
(Incorporated in USA - New York)
New Orchard Road, Armonk, NY 10504,
United States of America

(72) Inventor(s)
Sankar Basu
Charles A Micchelli
Peder Olsen

(74) Agent and/or Address for Service
R J Burt
IBM United Kingdom Limited, Intellectual Property
Department, Hursley Park, WINCHESTER, Hampshire,
SO21 2JN, United Kingdom

(51) INT CL⁷
G10L 15/02 // G10L 15:14

(52) UK CL (Edition S)
G4R RHB RPW R1F

(56) Documents Cited
US 5764853 A

(58) Field of Search
 UK CL (Edition S) **G4R RHB RPP RPW RRL**
 INT CL⁷ **G10L 15/00 15/02 15/06 15/08 15/14**
 Online: **WPI, EPODOC, JAPIO, INSPEC**

(54) Abstract Title
A maximum entropy and maximum likelihood criteria for feature selection from multivariate data

(57) Improvements in speech recognition systems are achieved by considering projections of the high dimensional data on lower dimensional subspaces, subsequently by estimating the univariate probability densities via known univariate techniques, and then by reconstructing the density in the original higher dimensional space from the collection of univariate densities so obtained. The reconstructed density is by no means unique unless further restrictions on the estimated density are imposed. The variety of choices of candidate univariate densities as well as the choices of subspaces on which to project the data including their number further add to this non-uniqueness. Probability density functions are then considered that maximize certain optimality criterion as a solution to this problem. Specifically, those probability density functions that either maximize the entropy functional, or alternatively, the likelihood associated with the data are considered.




FIG. 1

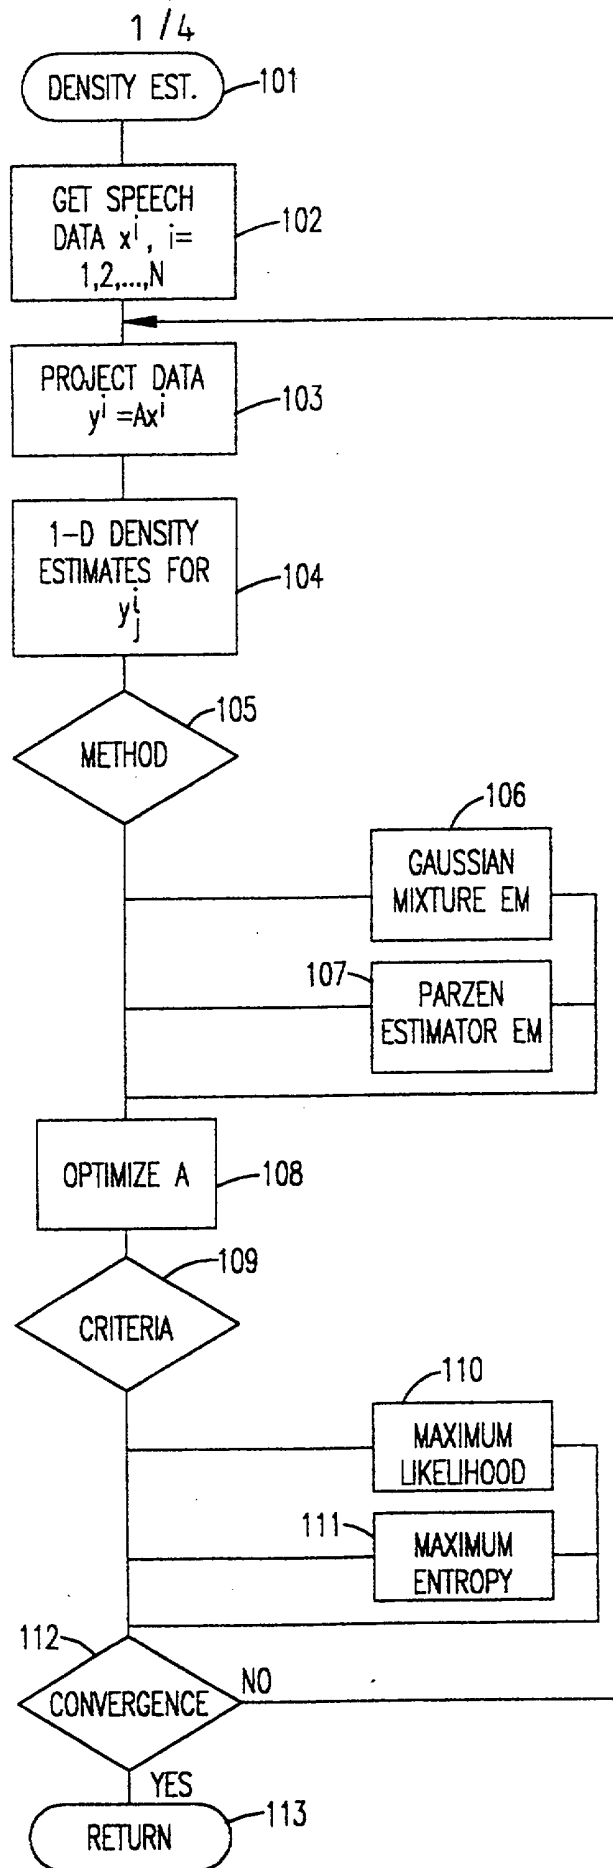


FIG. 1




FIG. 2




FIG. 3

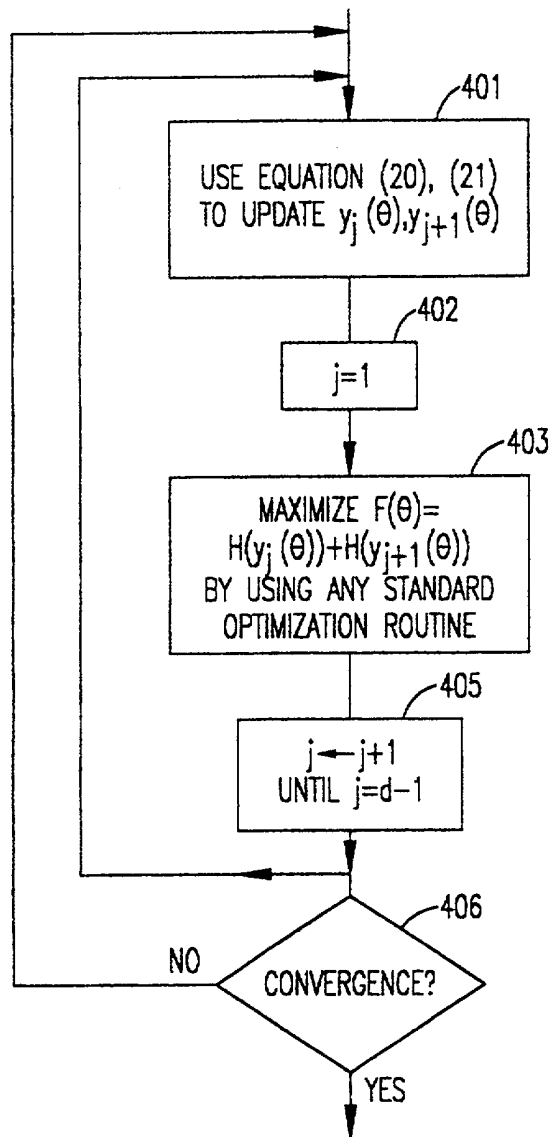


FIG. 4

A MAXIMUM ENTROPY AND MAXIMUM
LIKELIHOOD CRITERIA FOR FEATURE
SELECTION FROM MULTIVARIATE DATA

DESCRIPTION

5

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to speech recognition systems
and, more particularly, to the reconstructing of high dimensional density
feature sets from lower dimensional projections and the use of maximum
10 entropy and maximum likelihood criterion to optimize the directions of
projection to improve the performance of such systems.

Background Description

The basic problem of speech recognition is the identification of
sequences of words from a prescribed vocabulary corresponding to spoken
15 utterances. Each word in the vocabulary can be thought of as a sequence of
basic sounds, consisting of an alphabet of about fifty sounds (called
"phonemes"). Therefore, the goal of speech recognition is to model the speech
signal for each basic sound in such a way that it is possible to identify them by
observing the sampled speech signal. Due to inherent redundancies in the
20 acoustic signal a parsimonious representation is achieved by extracting feature
vectors periodically (every 10 ms). The feature vectors should be related to the

local spectral content of the speech signal, and should preserve the information necessary to enable discrimination between different phones. Furthermore, the acoustic features associated with a phoneme depend on a multitude of circumstances surrounding its formation (realization).

5 An important step in the speech recognition process is to isolate important features of the waveform over small time intervals (typically 25 ms). These features are represented by vector $x \in \mathbb{R}^d$ (where d usually is 39) which are then identified with context dependent sounds. Strings of such basic sounds are then converted into words using a dictionary of acoustic
10 representations of words. In an ideal situation the feature vectors generated by the speech waveform would be converted into a string of phonemes corresponding to the spoken utterance.

 A problem associated with this process is to identify a phoneme label for an individual acoustic vector x . Training data is provided for the purpose
15 of classifying a given acoustic vector. A standard approach for classification in speech recognition is to generate initial "prototypes" by K-means clustering and then to refine them by using the EM algorithm based on mixture models of gaussian densities. See, for example, Frederick Jelenik, *Statistical Methods for Speech Recognition*, MIT Press (1997). Moreover, in the decoding stage of
20 speech recognition the output probability density functions are most commonly assumed to be a mixture of Gaussian density functions.

 Density estimation of high dimensional data arises in speech recognition via classification of the training data. Specifically, acoustic vectors for a given sound are viewed as a random variable whose density is estimated
25 from the data. Consequently, the training stage requires that densities be found for all basic sounds. From this information we can assign to any acoustic vector the phoneme label corresponding to the highest likelihood obtained from these probability densities. This information is the basis of the

translation of acoustic vectors into text.

Speech data is characteristically represented by high dimensional vectors and each basic sound has several thousand data vectors to model it (typically, 3000 to 5000 for each of approximately 3500 basic sounds). Purely Gaussian densities have been known to be inadequate for this purpose due to the heavy tailed distributions observed by speech feature vectors. As an intended remedy to this problem, practically all speech recognition systems attempt modeling by using a mixture model with Gaussian densities for mixture components. Variants of the standard K-means clustering algorithm are used for this purpose. The classical version of the K-means algorithm can also be viewed as an special case of the EM algorithm (cf., David W. Scott, *Multivariate Density Estimation*, Wiley Interscience (1992)) for mixtures of Gaussians with variances tending to zero. Attempts to model the phonetic units in speech with non-Gaussian mixture densities are described by S. Basu and C. A. Micchelli, in "Parametric Density Estimation for the Classification of Acoustic Feature Vectors in Speech Recognition, *Nonlinear Modeling: Advanced Black-Box Techniques*, Eds. J. A. K. Suykens and J. Vandewalle, pp. 87–118, Kluwer Academic Publishers, Boston (1998).

There exists a large literature on estimating probability densities associated with univariate data. However, corresponding methods for estimating multivariate probability densities prove to be problematic due to various reasons. See again David W. Scott, *Multivariate Density Estimation*, Wiley Interscience (1992). This is especially true in the context of very high dimensions. Crucial among these difficulties is the fact that the data appears to be increasingly sparse with the increase in the number of dimensions.

SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide improvements in speech recognition systems.

5 Recognizing that acoustic vectors are important instances of high dimensional statistical data for which feature recognition can enhance the performance of density estimation in classification and training, we approach the problem by considering projections of the high dimensional data on lower dimensional subspaces, say a single dimensional subspace, subsequently by estimating the univariate probability densities via known univariate techniques, and then by reconstructing the density in the original higher dimensional space from the collection of univariate densities so obtained. In some sense the approach is reminiscent of function reconstruction from projections (e.g., in computerized tomography). The reconstructed density is by no means unique unless further restrictions on the estimated density are imposed. The variety of choices of candidate univariate densities as well as the choices of subspaces on which to project the data including their number further add to this non-uniqueness. One can then consider probability density functions that maximize certain optimality criterion as a solution to this problem. For the purpose of the present invention, we consider those probability density functions that either maximize the entropy functional, or alternatively, the likelihood associated with the data.

BRIEF DESCRIPTION OF THE DRAWINGS

25 The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

Figure 1 is a flow diagram illustrating the process of estimating densities of high dimensional speech data;

Figure 2 is a flow diagram showing in more detail the process of block 106 in Figure 1;

5 Figure 3 is a flow diagram showing in more detail the process of block 107 in Figure 1; and

Figure 4 is a flow diagram showing in more detail the process of block 108 in Figure 1.

10 DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT OF THE INVENTION

Maximum Entropy from Marginals

15 In this section we compute the maximum of the entropy of a probability density function on \mathbf{R}^d with prescribed marginals in d linearly independent (column) vectors, $y^1, y^2, \dots, y^d \in \mathbf{R}^d$. Let Y be the $d \times d$ matrix whose columns are these vectors. Thus, we are given univariate probability density functions, p_1, p_2, \dots, p_d and consider all multivariate probability density functions $P : \mathbf{R}^d \rightarrow \mathbf{R}$ such that for all measurable function f on \mathbf{R} with $f \cdot p_1, f \cdot p_2, \dots, f \cdot p_d \in C(\mathbf{R})$ (the space of absolutely integrable functions on \mathbf{R}) we have that

$$20 \quad \int_{\mathbf{R}^d} f(y_i^T x) P(x) dx = \int_{\mathbf{R}} f(t) p_i(t) dt, \quad i = 1, 2, \dots, d. \quad (1)$$

We denote the class of all densities P satisfying the above equations by $C(p)$. Recall that the entropy of P is given by

$$H(P) = - \int_{\mathbb{R}^d} P(x) \log P(x) dx. \quad (2)$$

5 See I. M. Cover and J. A. Thomas, *Elements of Information Theory*, John Wiley & Sons, Inc. (1991). Our first observation computes the maximum of $H(P)$ for all $P \in C(p)$.

Theorem 1: *Given any probability density functions p_1, p_2, \dots, p_d on \mathbb{R} such that for $i = 1, 2, \dots, d, p_i \log p_i \in L(\mathbb{R})$, we have that*

$$\max\{H(P) : P \in C(p)\} = -\frac{1}{2} \log \det G + \sum_{i=1}^d H(p_i), \quad (3)$$

10 where $G = Y^T Y$ is the Gram matrix

$$G(Y) = [y_i^T y_j] \quad i, j = 1, \dots, d \quad (4)$$

Moreover,

$$P^*(x) = \sqrt{\det G(Y)} \prod_{i=1}^d p_i(y_i^T x), \quad x \in \mathbb{R}^d \quad (5)$$

is the unique probability density function in $C(p)$ at which the maximum in

equation (3) is achieved.

In the next section we shall use this observation for maximum entropy based methods for feature selection of data.

Feature Selection Using Maximum Entropy

5 Our goal in this section is to use Theorem 1 as a means to isolate desirable features of multivariate data. Thus, our point of view here is to take observed data $x = (x^1, x^2, \dots, x^N)^T \in \mathbf{R}^N$ considered to be samples of some unknown probability function P . We project this data onto d directions, estimate the probability density function of the *univariate* data, find the
10 maximum entropy of all multivariate probability density functions in \mathbf{R}^d with marginals in these directions prescribed by the estimated univariate probability density functions and then *maximize* the entropy further over the choice of these directions. This strategy for optimal feature selection depends on the method used to estimate *univariate* data.

15 Thus, for univariate data $x = (x_1, x_2, \dots, x_N)^T \in \mathbf{R}^N$, we suppose $p(\cdot; x)$ is a probability density function on \mathbf{R} , which is a good estimator for the random variable from which x is drawn. For a given $y \in \mathbf{R}^d$ of norm one, we form the vector of univariate data $X^T y \in \mathbf{R}^N$ where X is the $d \times N$ matrix whose columns are the multivariate data x^1, x^2, \dots, x^N and use $p(\cdot; X^T y)$ as an estimate
20 for the marginal probability density function of P in the direction y . Given d directions y^1, y^2, \dots, y^d following Theorem 1, our maximum entropy estimator for the unknown probability density function P is

$$P^*(x) = \sqrt{\det G(Y)} \prod_{i=1}^d p(y_i^T x; X^T y_i), \quad x \in \mathbf{R}^d, \quad x \in \mathbf{R}^d, \quad (6)$$

and

$$\mathcal{F}(Y) := H(P^*) = \frac{1}{2} \log \det G(Y) + \sum_{i=1}^d \mathcal{H}(y_i), \quad (7)$$

where we define

$$\mathcal{H}(y) = H(p(\cdot; X^T y)), \quad y \in \mathbb{R}^d \quad (8)$$

5 We desire to make the entropy equation (7) as large as possible by
 varying our choice of vectors y^1, y^2, \dots, y^d . This is not fruitful without
 additional constraints on the vectors y^1, y^2, \dots, y^d . Indeed, we can have
 independent vectors which are nearly dependent so that their Gram matrix is
 nearly singular and so the first term in equation (7) will become large. Our
 10 preference is to choose an *orthonormal basis* so that the corresponding Gram
 matrix is the identity matrix. We turn to an example of the above which is
 instructive. Specifically, we have in mind as our univariate estimate for the
 data vector $x = (x_1, x_2, \dots, x_N)^T \in \mathbb{R}^N$ the Gaussian density $G(\cdot; \mu(x), \sigma(x))$,
 where

$$15 \quad G(t; \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(t-\mu)^2}{2\sigma}}, \quad t \in \mathbb{R}, \quad (9)$$

with mean

$$\mu(x) = \frac{1}{N} \sum_{j=1}^N x_j \quad (10)$$

and variance

$$\sigma(x) = \frac{1}{N} \sum_{j=1}^N (x_j - \mu)^2. \quad (11)$$

5 Thus, we have that $H(G(\cdot; \mu, \sigma)) = \frac{1}{2} \log(2\pi e \sigma)$ and the function that we shall optimize is

$$\mathcal{F}(y) = -\frac{1}{2} \log \det G(Y) + \frac{d}{2} \log(2\pi e) + \frac{1}{2} \sum_{j=1}^d \log y_j^T V y_j, \quad (12)$$

where V is the $d \times d$ matrix $V = U U^T$ and U is the $d \times N$ matrix whose columns are the vectors

$$u_j = \frac{1}{\sqrt{N}} (x^j - \bar{x}), \quad j=1, 2, \dots, N \quad (13)$$

10 where

$$\bar{x} = \frac{1}{N} \sum_{i=1}^N x^i. \quad (14)$$

First, we consider the *minimum* of $\mathcal{F}_G(Y)$ over all $d \times d$ orthogonal matrices Y . To this end, we introduce the $d \times d$ matrix $W = Y^T V Y$, whose elements are given as

$$\omega_{ij} := \langle y^j, V y^j \rangle, \quad i, j = 1, 2, \dots, d.$$

5 Since $\det W = \det G \det V$, we conclude that

$$\mathcal{F}_G(Y) = \frac{d}{2} \log(2\pi e) + \frac{1}{2} \log \det V + \frac{1}{2} \log \left(\prod_{j=1}^d \omega_{jj} \right) - \frac{1}{2} \log \det W.$$

Recall from Hadamard's inequality that $\det W \leq \prod_{j=1}^d \omega_{jj}$ where equality occurs

if and only if $\omega_{ij} = 0$ for all $i \neq j$, $i, j = 1, 2, \dots, d$. Thus, we obtain that

$$\mathcal{F}_G(Y) \geq \frac{d}{2} \log(2\pi e) + \frac{1}{2} \log \det V.$$

10 Moreover, equality occurs if and only if

$$\langle y^i, V y^j \rangle = 0, \quad i \neq j, \quad i, j = 1, 2, \dots, d.$$

We suppose that v_1, v_2, \dots, v_d are orthonormal eigenvectors corresponding to eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_d$ with $0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d$ for V . We let

$$A = \text{diag}[\lambda_1, \lambda_2, \dots, \lambda_d]$$

15 and S be the orthogonal $d \times d$ matrix whose columns are v^1, v^2, \dots, v^d , i.e., $V = S^T A S$. We then have that

$$\min\{\mathcal{F}_G(Y) : Y\} = \mathcal{F}_G(S) = \frac{1}{2}d \log(2\pi e) + \frac{1}{2}\sum_{j=1}^d \log \lambda_j$$

Note that we have demonstrated that there exist *orthonormal* vectors at which the minimum is attained.

Let us now consider the problem of *maximizing* the function \mathcal{F}_G over all $d \times d$ matrices Y which are orthogonal. In this case, we have from equation (12) that

$$\mathcal{F}_G(Y) = \frac{d}{2} \log(2\pi e) + \frac{1}{2} \sum_{j=1}^d \log \langle y^j, \mathcal{V}y^j \rangle$$

and by the concavity of the log function obtain the inequality

$$\mathcal{F}_G(Y) \leq \frac{d}{2} \log \frac{2\pi e}{d} + \frac{d}{2} \log \sum_{j=1}^d d\lambda_j \quad (15)$$

where equality holds if, for all $j = 1, 2, \dots, d$,

$$\langle y^j, \mathcal{V}y^j \rangle = \frac{1}{d} \sum_{i=1}^d \lambda_i \quad (16)$$

In other words, the matrix W has *constant* diagonal elements. To ensure that the bound in inequality (15) is sharp, we need to establish the existence of an orthogonal matrix Y such that equation (16) is satisfied. Our first observation in this regard is the following fact. Suppose H is a $d \times d$ *Hadamard matrix*; that is, all the entries of H are either 1 or -1 and satisfies the condition that

$$HH^T = dI.$$

We set $J = d^{-1/2}H$ and $Y = SJ$ so that Y is orthonormal. Since $V = SAS^T$, we have, for $i = 1, 2, \dots, d$ that

$$\langle y^i, Vy^i \rangle = \sum_{j=1}^d J_{ij}^2 \lambda_j = \frac{1}{d} \sum_{j=1}^d \lambda_j$$

5 which is the desired equation (16). Hadamard matrices do not exist for all orders. When they do exist, the above remark tells us how to construct an orthonormal matrix from the eigenvalues of V which maximizes the function \mathcal{F}_G .

Feature Selection Using Maximum Entropy and Mixtures of Gaussians

10 Often one finds modeling univariate data by a single Gaussian to be inadequate. We, therefore, consider a mixture model of m Gaussians

$$p(t|\mu, \omega, \sigma) = \sum_{j=1}^m \omega_j G(t; \mu_j, \sigma_j), \quad t \in R, \quad (18)$$

where the mixture weights satisfy the constraints

$$\sum_{j=1}^m \omega_j = 1, \quad \omega_j > 0; \quad i=1, 2, \dots, m. \quad (19)$$

15 For any choice of parameters μ , ω and σ , the entropy of the probability density of equation (18) can easily be computed by a Monte-Carlo method using random samples drawn from this probability density function.

We turn to the problem of selecting the means and variances of the probability density function in equation (18) from a given but arbitrary data vector $(x_1, x_2, \dots, x_N)^T \in \mathbb{R}^N$. As before, we take an MLE (maximum likelihood estimation) approach. However, as is well known for $m > 2$ the likelihood function will not have a maximum. See, for example, V. N. Vapnik, *The Nature of Statistical Learning Theory*, Springer Verlag (1995). Instead, we cannot use it as an absolute criterion to select the means and variances. Therefore, we turn to the EM algorithm to iteratively set the means and variances. Here we fix a number of iterations in the EM update formulas appearing below independent of the data vector x . The first step is the initialization of the means and variances. This is done by clustering the data vector into m groups of (approximately) equal size starting with its first coordinate and proceeding sequentially through all its components. On each of the m clusters we use the formulas (10) and (11) for the mean and the variances. We then update them by the EM formulas, (see, for instance, C. M. Bishop, *Neural Networks for Pattern Recognition*, Clarendon Press, Oxford (1995), and A. P. Dempster, N. M. Laird and D. B. Baum, "Maximum likelihood from incomplete data via the EM algorithm", *Journal of Royal Statistical Soc., Ser. B.* vol. 39, pp. 1–38, 1977), which produces a sequence of MLE's whose likelihoods are nondecreasing. One step of the update formulas is given by

$$\hat{\omega}_j = \frac{1}{N} \sum_{i=1}^N P_{ij} \quad j=1, 2, \dots, m \quad (20)$$

$$\mu_j = \frac{\sum_{i=1}^N x_i P_{ij}}{\sum_{i=1}^N P_{ij}}, \quad j=1, 2, \dots, m \quad (21)$$

$$\sigma_j = \frac{\sum_{i=1}^N P_{ij} (x_i - \mu_j)^2}{\sum_{i=1}^N P_{ij}}, \quad j=1, 2, \dots, m \quad (22)$$

where for $j = 1, 2, \dots, m, i = 1, 2, \dots, N$, we define

$$P_{ij} = \frac{\omega_j G(x_i; \mu_j, \sigma_j)}{p(x_i; \mu, \omega, \sigma)}.$$

5 With these formulas, we used a *fixed* number of iterations to set the values ω ,
 μ and σ and denote by $p(\cdot; x)$ our univariate estimate for our data vector
 $x \in \mathbf{R}^N$ by mixture models of Gaussians.

The next step is to iteratively update the orthogonal matrix Y . Our
iteration is initialized per our previous discussion of the single Gaussian
10 model. Consequently, all our experiments are performed in d -dimensions
where d is chosen to admit a Hadamard matrix. We then modify the
eigenvalues of the matrix $V = UU^T$ corresponding to the column vectors
 w_1, w_2, \dots, w_N given equation (17) per our discussions of the single gaussian
model. Thus, for any vector $Y \in \mathbf{R}^d$ and data $x_1, x_2, \dots, x_N \in \mathbf{R}^d$ organized into
15 a $d \times N$ matrix X whose columns are $x_1, x_2, \dots, x_N \in \mathbf{R}^d$ we can evaluate the

function

$$\mathcal{H}(y) = H(p(\cdot; X^T y)).$$

Our next task is to iteratively update our orthonormal feature vectors y_1, y_2, \dots, y_d to enhance our entropy functional (8) which specializes to the form

5

$$\mathcal{F}(Y) = \sum_{j=1}^d \mathcal{H}(y_j).$$

To update the vectors we used two dimensional planar rotations. Thus, we choose two consecutive vectors y_j, y_{j+1} and form new vectors

$$y_l(\theta) = y_l, \quad l \neq j, j+1, \quad j = 1, 2, \dots, d, \quad (23)$$

10

$$y_j(\theta) = y_j \cos \theta - y_{j+1} \sin \theta \quad (24)$$

$$y_{j+1}(\theta) = y_j \sin \theta + y_{j+1} \cos \theta \quad (25)$$

$$\dots, \pi_m, \quad (26)$$

where $\theta \in [0, 2\pi]$. We then optimize the univariate function

$$F(\theta) = \mathcal{H}(y,(\theta)) + \mathcal{H}(y_{j+1}(\theta))$$

15

over $\theta \in [0, 2\pi]$ and replace y_1, y_2, \dots, y_d by $y_1(\hat{\theta}), y_2(\hat{\theta}), \dots, y_d(\hat{\theta})$ where $\hat{\theta}$ is a good choice for the maximum of the function F . If $j < d-1$, we increment j by one and repeat the procedure; otherwise, we set $j = 1$ and repeat the procedure, cycling several times through the columns of Y . To avoid

computing the derivatives of F we compute its values on an acceptably fine grid and choose its largest value.

Feature Selection Using Maximum Likelihood and Parzen Estimators

In the preceding section, we considered modeling our multivariable data by choosing our univariate marginals $p(\cdot; x)$, $x = (x_1, x_2, \dots, x_N)^T \in \mathbf{R}^N$ to be an MLE mixture model of Gaussians.

In this section, we study the possibility of using Parzen estimators for our univariate data and maximum likelihood as a feature selection criterion instead of maximum entropy. Starting with a bin width h and a univariate kernel K , the corresponding Parzen estimator for $x = (x_1, x_2, \dots, x_N)^T$ is given by

$$p(t, x) = \frac{1}{Nh} \sum_{j=1}^N K \left(\frac{t - x_j}{h} \right), \quad t \in \mathbf{R}.$$

In what follows, we choose the kernel K to be the standard Gaussian $K = G(\cdot; 0, 1)$. Given d orthonormal vectors y^1, y^2, \dots, y^d and corresponding marginals $p(\cdot; X^T y)$, the maximum entropy estimator given by Theorem 1 is

$$P^*(x) = \prod_{j=1}^d p(y_j^T x; X^T y_j), \quad x \in \mathbf{R}^d, \quad (26)$$

and the *log-likelihood* function of the data has the form

$$L(Y) := \sum_{i=1}^N \sum_{j=1}^d \log \left(\frac{1}{Nh} \sum_{k=1}^N K(y^j, g^{ik}) \right)$$

where

$$g^{ik} := \frac{x^i - x^k}{h}, \quad i, k=1, 2, \dots, N.$$

We resort to an EM strategy and develop an iterative update formula that
 5 increases $L(Y)$. To this end, we let the matrix \hat{Y} be our initial guess for the
 desired orthogonal matrix Y and seek another orthogonal matrix Y such that
 $L(Y) \geq L(\hat{Y})$. To find a suitable Y , we consider the equation

$$L(Y) - L(\hat{Y}) = \sum_{i=1}^N \sum_{j=1}^d \log \left(\sum_{k=1}^N c_{kj} a_k \right), \quad (28)$$

where

$$10 \quad a_k = \frac{K(\hat{y}^j, g^{ik})}{K(\hat{y}^j, g^{ik})} \quad (29)$$

and

$$c_{kj} = \frac{K(\hat{y}^j, g^{ik})}{\sum_{m=1}^N K(\hat{y}^j, g^{im})} \quad (30)$$

Substituting equations (30) and (29) into equation (28) and using the concavity of the log function as well as the form of the kernel K we get

$$\begin{aligned} L(Y) - L(\hat{Y}) &\geq \sum_{i=1}^N \sum_{j=1}^d \sum_{k=1}^N c_{kij} \log a_k \\ &= Q(Y) - Q(\hat{Y}) \end{aligned}$$

where

$$5 \quad Q(Y) := -\frac{1}{2} \sum_{j=1}^d \langle y^j, A_j y^j \rangle$$

and

$$A_j = \sum_{i=1}^N \sum_{k=1}^N c_{kij} (x^i - x^k)(x^i - x^k)^T.$$

The task is now to compute the maximum of $\sum_{j=1}^d \langle y^j, A_j y^j \rangle$ over all

10 orthonormal matrices Y . We approach this problem as in the preceding section by using planar rotations. In the case at hand, our angle θ can be *determined explicitly*. To this end, we consider the function

$$F(\theta) = \langle y^1(\theta), A_1 y^1(\theta) \rangle + \langle y^2(\theta), A_2 y^2(\theta) \rangle$$

which simplifies to

$$F(\theta) = \alpha \sin 2\theta + \beta \cos 2\theta + \gamma$$

where

$$\begin{aligned}
 B &= \frac{1}{2}(A_1 - A_2), \\
 C &= \frac{1}{2}(A_1 + A_2), \\
 \alpha &= -2\langle y^1, By^2 \rangle, \\
 \beta &= \langle y^1, By^1 \rangle - \langle y^2, By^2 \rangle, \text{ and} \\
 \gamma &= \langle y^1, Cy^1 \rangle - \langle y^2, Cy^2 \rangle
 \end{aligned}$$

and its minimum occurs at

$$\theta = k\pi \pm \frac{\psi}{2}$$

where k is an integer and

$$\cos\psi = \frac{\alpha}{\sqrt{\alpha^2 + \beta^2}}.$$

Feature Selection Using Maximum Likelihood Gaussian Mixture Estimators

In this section, we discuss optional feature selection for the data $x^1, x^2, \dots, x^N \in \mathbf{R}^d$, the log-likelihood function of interest to us here is

$$L(Y, \omega, \mu, \sigma) = \sum_{i=1}^N \sum_{j=1}^d \log \left(\sum_{t=1}^m \omega_{jt} G(y^j, x^i; \mu_{jt}, \sigma_{jt}) \right),$$

15 where

$$\omega_{ij} > 0, \quad \sum_{i=1}^m \omega_{ij} = 1, \quad j = 1, 2, \dots, d.$$

We describe a method to iteratively increase L . First, we fix Y and point out how to update ω , μ and σ to increase L .

We let the parameters $\hat{\omega}$, $\hat{\mu}$ and $\hat{\sigma}$ be the initial guess. We shall
 5 identify ω , μ and σ such that $L(Y, \omega, \mu, \sigma) \geq L(Y, \hat{\omega}, \hat{\mu}, \hat{\sigma})$. To this end, we consider the equation

$$L(Y, \omega, \mu, \sigma) - L(Y, \hat{\omega}, \hat{\mu}, \hat{\sigma}) = \sum_{i=1}^N \sum_{j=1}^d \log \left(\sum_{t=1}^m \alpha_{ijt} \beta_{ijt} \right), \quad (31)$$

where

$$\alpha_{ijt} = \frac{\hat{\omega}_{ij} G(y^j, x^i; \hat{\mu}_{ij}, \hat{\sigma}_{ij})}{\sum_{k=1}^m \hat{\omega}_{kj} G(y^j, x^i; \hat{\mu}_{kj}, \hat{\sigma}_{kj})}$$

10 and

$$\beta_{ijt} = \frac{\omega_{ij} G(y^j, x^i; \mu_{ij}, \sigma_{ij})}{\hat{\omega}_{ij} G(y^j, x^i; \hat{\mu}_{ij}, \hat{\sigma}_{ij})}.$$

Note that, by definition, $\sum_{t=1}^m \alpha_{ijt} = 1$. Thus, using the concavity of the log

function in equation (29) we obtain

$$\begin{aligned} L(Y, \omega, \mu, \sigma) - L(Y, \hat{\omega}, \hat{\mu}, \hat{\sigma}) &\geq \sum_{i=1}^N \sum_{j=1}^d \sum_{t=1}^m \alpha_{ijt} \log \beta_{ijt} \\ &= Q(Y, \omega, \mu, \sigma) - Q(Y, \hat{\omega}, \hat{\mu}, \hat{\sigma}), \end{aligned}$$

where

$$Q(Y, \omega, \mu, \sigma) = \sum_{i=1}^N \sum_{j=1}^d \sum_{t=1}^m \alpha_{ijt} \log(\omega_{ij} G(y^j, x^i; \mu_{ij}, \sigma_{ij})).$$

- 5 We can *explicitly* maximize $Q(Y, \omega, \mu, \sigma)$ with respect to ω , μ , σ to get the equations

$$\omega_{ij} = \frac{1}{N} \sum_{i=1}^N \alpha_{ijt},$$

$$\mu_{ij} = \frac{\sum_{i=1}^N \alpha_{ijt} (y^j, x^i)}{\sum_{i=1}^N \alpha_{ijt}}$$

and

$$\sigma_{ij} = \frac{\sum_{i=1}^N \alpha_{ijt} (y^j, x^i) - \mu_{ij}^2}{\sum_{i=1}^N \alpha_{ijt}}$$

10

These formulas provide an effective means for updating the parameters ω , μ and σ . We will now consider an update formula for Y . To this end, we use the new parameters ω , μ and σ for this step and reformulate the log-likelihood function in terms of unique vectors μ^l , $l=1, 2, \dots, m$, defined by the equation

$$5 \quad \mu_{j,l} = (y^j, \mu^l)$$

$$\log L(Y, \omega, \mu, \sigma) = \sum_{i=1}^N \sum_{j=1}^d \log \left(\sum_{\ell=1}^m \omega_{j\ell} G(y^j, x^i; (y^j, \mu^\ell), \alpha_{j\ell}) \right).$$

Suppose that \hat{Y} is our initial choice. We seek an orthogonal matrix Y such that $L(Y, \omega, \mu, \sigma) \geq L(\hat{Y}, \omega, \mu, \sigma)$. Similar to the argument used before, we have that

$$\log L(Y, \omega, \mu, \sigma) - \log L(\hat{Y}, \omega, \mu, \sigma) \geq \sum_{i=1}^N \sum_{j=1}^d \sum_{\ell=1}^m \alpha_{j\ell} \log \beta_{j\ell}$$

10 where

$$\alpha_{j\ell} = \frac{\omega_{j\ell} G(\hat{y}^j, x^i; (\hat{y}^j, \mu^\ell), \sigma_{j\ell})}{\sum_{k=1}^m \omega_{jk} G(\hat{y}^j, x^i; (\hat{y}^j, \mu^k), \sigma_{jk})},$$

and

$$\beta_{j\ell} = \frac{\omega_{j\ell} G(y^j, x^i; (y^j, \mu^\ell), \sigma_{j\ell})}{\omega_{j\ell} G(\hat{y}^j, x^i; (\hat{y}^j, \mu^\ell), \sigma_{j\ell})}.$$

Hence, we conclude that $L(Y, \omega, \mu, \sigma) - L(\hat{Y}, \omega, \mu, \sigma) \geq Q(Y, \omega, \mu, \sigma) - Q(\hat{Y}, \omega, \mu, \sigma)$, where

$$\begin{aligned} Q(Y, \omega, \mu, \sigma) &= \sum_{i=1}^N \sum_{j=1}^d \sum_{t=1}^m \alpha_{ijt} \log (\omega_{ij} G(y^j, x^i); (y^j, \mu^j), \sigma_{ij}) \\ &= \frac{1}{2} \sum_{j=1}^d (y^j, A_j^j) + Q_C \end{aligned}$$

with

5

$$A_j = \sum_{i=1}^N \sum_{t=1}^m \frac{\alpha_{ijt}}{\sigma_{ij}} (x_i - \mu_j)(x_i - \mu_j)^T$$

and

$$Q_C = \sum_{i=1}^N \sum_{j=1}^d \sum_{t=1}^m \alpha_{ijt} \log \left(\frac{\omega_{ij}}{\sqrt{\sigma_{ij}}} \right).$$

Since Q_C is independent of Y , we can use the methods described above to increase the function Q .

10

Applications in Speech Recognition

15

From a practical standpoint, estimation of densities in speech data is accompanied by all the difficulties characteristic of high dimensional density estimation problems. Feature vectors of dimension fifty or more are typical in speech recognition systems, and consequently, the data can be considered to be highly sparse. In contrast, the literature on multivariate density estimation

puts high emphasis on “exploratory data analysis”, the goal of which is to glean insight about the densities via visualization of the data. This is not feasible for dimensions of the order of fifty or more, even when projections on lower dimensional spaces are considered.

5 While the use of linear discriminant analysis to reduce the dimension of the feature set is classically known (see, e.g., Frederick Jelenik, *Statistical Methods for Speech Recognition*, MIT Press (1997) and Nagendra Kumar, “Investigation of silicon auditory model and generalization of linear discriminant analysis for improved speech recognition”, Ph.D dissertation, Johns Hopkins University, 1997), the idea of reconstructing the high di-
10 dimensional density from lower dimensional projections as well as the use of maximum entropy and maximum likelihood criterion to optimize the directions of projection is novel for this purpose.

Referring now to Figure 1 of the drawings, there is shown a flow
15 diagram of the logic which implements the process of the invention. Density estimates are input at 101. In function block 102, speech data x^t , $t=1,2,\dots,N$, is retrieved before entering a processing loop. Function block 103 is the first block in the processing loop. Data is projected; i.e., $y^t=Ax^t$. Then, in function block 104, one-dimensional density estimates are computed for y_j^t . At this
20 point, processing method calls are made at block 105. The first method called is the Gaussian mixture EM method in function block 106, shown in more detail in Figure 2. The second method called is the Parzen estimator EM method in function block 107, shown in more detail in Figure 3. When these methods have completed, A is optimized in function block 108, shown in more
25 detail in Figure 4. Next, criteria method calls are made at block 109. The first method called is the maximum likelihood method in function block 110. The second method called is the maximum entropy method in function block 111.

In decision block 112, a determination is made as to whether convergence has been obtained and, if not, the process loops back to function block 103; otherwise, the process has completed and a return 113 is made to the main program.

5 Figure 2 shows in more detail the process Gaussian mixture EM method in function block 106. Cluster data is received at function block 201, and prototype initialization is performed in function block 202, in which mixture weights ω , means μ and variances σ are computed. The formulas for $\hat{\omega}_j$, $\hat{\mu}_j$ and $\hat{\sigma}_j$ are updated in function block 203 using formulas (20), (21) and (22). A determination is made in decision block 204 as to whether a convergence is detected. If not, the process loops back to function block 201 to iteratively repeat the process. When convergence is detected, a return is made to the main process.

10 Figure 3 shows in more detail the Parzen estimator EM method in function block 107. Cluster data is input at function block 301. An integer N , kernel type K and bin width h are chosen for equation (28) in function block 302. The index j is initialized to 1 in function block 303. Then in function block 304, $y_j(\theta)$ and $y_{j+1}(\theta)$ are chosen as in equations (23) and (24) with $\theta = \hat{\theta}$. The index j is incremented by 1 in function block 305, and then a determination is made in decision block 306 as to whether $j = d - 1$. If not, the process loops back to function block 303, but if so, a return is made to the main program.

15 Figure 4 shows in more detail the optimization of A in function block 108. The process is entered at function block 401 where equations (20) and (21) are used to update $y_j(\theta)$ and $y_{j+1}(\theta)$ using equations (23) and (24). The index j is initialized to 1 in function block 402. The equation $F(\theta) = \mathcal{J}(y_j(\theta)) + \mathcal{J}(y_{j+1}(\theta))$ is maximized in function block 403 using any

standard optimization routine. The index j is incremented by 1 in function block 404 and the process loops back to function block 401. This iteration continues until $j = d - 1$, at which point a determination is made in decision block 405 as to whether a convergence is detected. If not, the process loops
5 back to function block 401. When convergence is detected, the process makes a return to the main routine.

While the invention has been described in terms of a single preferred embodiment, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended
10 claims.

CLAIMS

Having thus described our invention, what we claim as new and desire to secure by Letters Patent is as follows:

- 1 1. A method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition where the number of directions is
3 the same as the dimension of sampled data comprising the steps of:
4 projecting the multivariate data on chosen directions;
5 reconstructing a univariate probability density from the projected data
6 for each choice of directions; and
7 using the univariate probability densities to reconstruct an unknown
8 multivariate density.

- 1 2. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 1, wherein the
3 univariate probability densities are reconstructed by a maximum entropy
4 criterion.

- 1 3. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 1, wherein the
3 univariate probability densities are reconstructed by a maximum likelihood
4 criterion.

- 1 4. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 1, further
3 comprising the step of choosing directions for projecting the multivariate data
4 using an interative numerical algorithm for estimating multivariate density and

5 computing directions for projections via rotational transformations when the
6 univariate densities are mixtures of Gaussians.

1 5. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 4, wherein the
3 univariate probability densities are reconstructed by a maximum entropy
4 criterion.

1 6. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 4, wherein the
3 univariate probability densities are reconstructed by a maximum likelihood
4 criterion.

1 7. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 1, further
3 comprising the step of choosing directions for projecting the multivariate data
4 using an iterative numerical algorithm for estimating multivariate density and
5 computing directions for projections via rotational transformations when the
6 univariate densities are Parzen density estimators.

1 8. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 7, wherein the
3 univariate probability densities are reconstructed by a maximum entropy
4 criterion.

1 9. The method of high dimensional density estimation and feature extraction
2 from multivariate data in speech recognition recited in claim 7, wherein the
3 univariate probability densities are reconstructed by a maximum likelihood

4 criterion.



Application No: GB 0111902.3
Claims searched: 1 to 9

Examiner: John Donaldson
Date of search: 19 June 2001

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
 UK Cl (Ed.S): G4R(RHB, RPP, RPW, RRL)
 Int Cl (Ed.7): G10L 15/00, 15/02, 15/06, 15/08, 15/14
 Other: Online: WPI, EPODOC, JAPIO, INSPEC

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	US 5764853 (WATARI), see abstract	-

X Document indicating lack of novelty or inventive step Y Document indicating lack of inventive step if combined with one or more other documents of same category. & Member of the same patent family	A Document indicating technological background and/or state of the art. P Document published on or after the declared priority date but before the filing date of this invention. E Patent document published on or after, but with priority date earlier than, the filing date of this application.
--	--