DEPARTMENT OF MATHEMATICS
UNIVERSITY OF HOUSTON HOUSTON, TEXAS

AN ITERATIVE PROCEDURE FOR OBTAINING MAXIMUM-LIKELIHOOD ESTIMATES OF THE PARAMETERS FOR A MIXTURE OF NORMAL DISTRIBUTIONS, II.

BY B. CHARLES PETERS, JR
AND HOMER F. WALKER

FEBRUARY 1976 REPORT #51

PREPARED FOR
EARTH OBSERVATION DIVISION, JSC
UNDER
CONTRACT NAS-9-12777

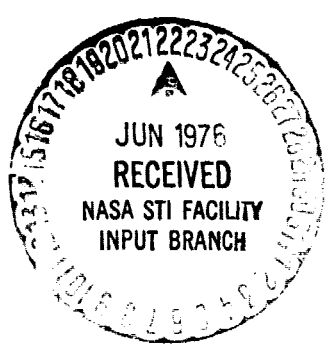

3801 CULLEN BLVD.
HOUSTON, TEXAS 77004

DEPARTMENT OF MATHEMATICS
UNIVERSITY OF HOUSTON HOUSTON, TEXAS

AN ITERATIVE PROCEDURE FOR OBTAINING MAXIMUM-LIKELIHOOD ESTIMATES OF THE PARAMETERS FOR A MIXTURE OF NORMAL DISTRIBUTIONS, II.

BY B. CHARLES PETERS, JR
AND HOMER F. WALKER

FEBRUARY 1976 REPORT #51

PREPARED FOR
EARTH OBSERVATION DIVISION, JSC
UNDER
CONTRACT NAS-9-12777

3801 CULLEN BLVD.
HOUSTON, TEXAS 77004

An Iterative Procedure for Obtaining Maximum-Likelihood Estimates of the Parameters for a Mixture of Normal Distributions, II.

by

B. Charles Peters, Jr.

NASA/National Research Council Research Associate

Earth Observations Division, Johnson Space Center

and

Homer F. Walker

Department of Mathematics, University of Houston

Houston, Texas

February, 1976

Report #51

ABSTRACT

AN ITERATIVE PROCEDURE FOR OBTAINING MAXIMUM-LIKELIHOOD ESTIMATES OF THE PARAMETERS FOR A MIXTURE OF NORMAL DISTRIBUTIONS

This paper addresses the problem of obtaining numerically maximum-likelihood estimates of the parameters for a mixture of normal distributions. In recent literature, a certain successive-approximations procedure, based on the likelihood equations, was shown empirically to be effective in numerically approximating such maximum-likelihood estimates; however, the reliability of this procedure was not established theoretically. Here, we introduce a general iterative prodedure, of the generalized steepest-ascent (deflected-gradient) type, which is just the procedure known in the literature when the step-size is taken to be 1. We show that, with probability 1 as the sample size grows large, this procedure converges locally to the strongly consistent maximum-likelihood estimate whenever the step-size lies between 0 and 2. We also show that the step-size which yields optimal local convergence rates for large samples is determined in a sense by the "separation" of the component normal densities and is bounded below by a number between 1 and 2.

An Iterative Procedure for Obtaining Maximum-Likelihood Estimates of the Parameters for a Mixture of Normal Distributions

by

B. Charles Peters, Jr.
NASA/National Research Council Research Associate
Earth Observations Division, Johnson Space Center

and

*Homer F. Walker
Department of Mathematics, University of Houston
Houston, Texas

1. Introduction.

Let x be an n-dimensional random variable whose density function p is a convex combination of normal densities, i.e.,

$$p(x) = \sum_{i=1}^{m} \alpha_i^o p_i(x) \quad \text{for } x \in \mathbb{R}^n,$$

where

$$\alpha_i^o > 0, \quad \sum_{i=1}^{m} \alpha_i^o = 1,$$

*This author was supported in part by NASA under Contract JSC-NAS-9-12777.

and

$$p_i(x) = \frac{1}{(2\pi)^{n/2}|\Sigma_i^o|^{1/2}} e^{-1/2(x-\mu_i^o)^T\Sigma_i^{o\,-1}(x-\mu_i^o)}.$$

If $\{x_k\}_{k=1,\ldots,N} \subseteq \mathbb{R}^n$ is an independent sample of observations on x, then a maximum-likelihood estimate of the parameters $\{\alpha_i^o,\mu_i^o,\Sigma_i^o\}_{i=1,\ldots,m}$ is a choice of parameters $\{\alpha_i,\mu_i,\Sigma_i\}_{i=1,\ldots,m}$ which locally maximizes the log-likelihood function

$$L = \sum_{k=1}^{N} \log p(x_k),$$

in which p is evaluated with the true parameters $\{\alpha_i^o,\mu_i^o,\Sigma_i^o\}_{i=1,\ldots,m}$ replaced by the estimate $\{\alpha_i,\mu_i,\Sigma_i\}_{i=1,\ldots,m}$. (In the following, it is usually clear from the context which parameters are used in evaluating the density functions p_i and p. Therefore, these parameters are explicitly pointed out only when some ambiguity exists.) We admit local maxima of L as maximum-likelihood estimates in order to avoid difficulties presented by the fact that L has no global maximum. It is observed below that this creates no problems when one is concerned with consistent maximum-likelihood estimates.

Clearly, L is a differentiable function of the parameters to be estimated. Equating to zero the partial derivatives of L with respect to these parameters, one obtains, after a straightforward calculation, the following necessary condition for a maximum-likelihood estimate:

$$\text{(1.a)} \qquad \alpha_i = \frac{\alpha_i}{N} \sum_{k=1}^{N} \frac{p_i(x_k)}{p(x_k)}$$

$$\text{(1.b)} \qquad \mu_i = \left\{\frac{1}{N} \sum_{k=1}^{N} x_k \frac{p_i(x_k)}{p(x_k)}\right\} \Big/ \left\{\frac{1}{N} \sum_{k=1}^{N} \frac{p_i(x_k)}{p(x_k)}\right\} \qquad \Bigg\} \; i=1,\ldots,m$$

$$\text{(1.c)} \qquad \Sigma_i = \left\{\frac{1}{N} \sum_{k=1}^{N} (x_k-\mu_i)(x_k-\mu_i)^T \frac{p_i(x_k)}{p(x_k)}\right\} \Big/ \left\{\frac{1}{N} \sum_{k=1}^{N} \frac{p_i(x_k)}{p(x_k)}\right\}$$

These are known as the <u>likelihood equations</u>. A number of authors have investigated solutions of the likelihood equations and consistency of maximum-likelihood estimates. (See, for example, Cramer [2], Huzurbazar [7], Wald [12], Chanda [1], and the discussion in Zachs [14].) We observe that, loosely speaking, there is a unique solution of the likelihood equations which tends with probability 1 to the true parameters as the sample size N approaches infinity. Furthermore, this solution is a maximum-likelihood estimate, indeed, the unique strongly consistent maximum-likelihood estimate. (Strictly speaking, given any sufficiently small neighborhood of the true parameters, there is, with probability 1 as N approaches infinity, a unique solution of the likelihood equations in that neighborhood, and this solution is a maximum-likelihood estimate. For completeness, we present a brief proof of this result in Appendix 1.) This note is addressed to the problem of determining this strongly consistent maximum-likelihood estimate by successive approximations.

The likelihood equations, as written, suggest the following iterative procedure for obtaining a solution: Beginning with some set of starting

values, obtain successive approximations to a solution by inserting the preceding approximations in the expressions on the right-hand sides of (1.a), (1.b), and (1.c). This scheme is attractive for its relative ease of implementation, and we discuss below the findings of several authors concerning its use in obtaining maximum-likelihood estimates. For a discussion of other methods of determining maximum-likelihood estimates, see Kale [8] and Wolfe [13] as well as the authors given below.

Empirical studies of Day [3], Duda and Hart [4], and Hasselblad [5] suggest that this scheme is convergent and that convergence is particularly fast when the component normal densities in p are "widely separated" in a certain sense. Unfortunately, the likelihood equations have many solutions in general, and the iterates may converge to solutions, including "singular solutions" (see [4]), which are not the strongly consistent maximum-likelihood estimate if care is not taken in the choice of starting values. No theoretical evidence of convergence is given in [3], [4], or [5]. Peters and Coberly [10] have proved that, if all of the parameters μ_i and Σ_i are held fixed, then the iterative procedure suggested by the equation (1.a) alone converges locally to a maximum-likelihood estimate of the parameters α_i, $i=1,\ldots,m$. (An iterative procedure is said to converge locally to a limit if the iterates converge to that limit whenever the starting values are sufficiently near that limit.) They also report on numerical studies in which the computational feasibility of this procedure is demonstrated.

In the following, we present a general iterative procedure for determing the strongly consistent maximum-likelihood estimate, of which the above procedure is a special case. Indeed, our procedure is a generalized

steepest-ascent (deflected-gradient) method, and the above procedure is obtained when the step-size is taken to be 1. We show that, with probability 1 as the sample size grows large, this procedure converges locally to the strongly consistent maximum-likelihood estimate whenever the step-size is between 0 and 2. Furthermore, the value of the step-size which yields optimal local convergence rates is bounded from below by a number which always lies between 1 and 2. In fact, this optimal step-size lies near 1 if the component populations are "widely separated" in a certain sense and cannot be much smaller than 2 if two or more of the component populations have nearly identical means and covariance matrices. We also prove that, if the covariance matrices Σ_i are held fixed, then the restricted iterative procedure for the parameters α_i and μ_i has these local convergence properties with probability 1 whenever the sample size is at least $m(n+1)$. We conclude by comparing this procedure to other numerical methods for determining maximum-likelihood estimates.

2. The general iterative procedure.

In order to minimize notational difficulties, we introduce several vector spaces and give useful representations of their elements. For each i, $1 \le i \le m$, α_i, μ_i, and Σ_i are elements of the vector spaces $\mathbb{R}^1$, $\mathbb{R}^n$, and the set of all real, symmetric $n \times n$ matrices, respectively. We denote by $\mathcal{A}, \mathcal{M}$, and $\mathcal{S}$ the respective m-fold direct sums of these spaces with themselves, and we represent elements of $\mathcal{A}, \mathcal{M}$, and $\mathcal{S}$ as columns

$$\bar{\alpha} = \begin{pmatrix} \alpha_1 \\ \cdot \\ \cdot \\ \alpha_m \end{pmatrix} \in \mathcal{A}, \quad \bar{\mu} = \begin{pmatrix} \mu_1 \\ \cdot \\ \cdot \\ \mu_m \end{pmatrix} \in \mathcal{M}, \quad \bar{\Sigma} = \begin{pmatrix} \Sigma_1 \\ \cdot \\ \cdot \\ \Sigma_m \end{pmatrix} \in \mathcal{S}.$$

It will be convenient to adopt the following notational equivalence for elements of the direct sum $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$:

$$\Theta = \begin{pmatrix} \bar{\alpha} \\ \bar{\mu} \\ \bar{\Sigma} \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_m \\ \mu_1 \\ \vdots \\ \mu_m \\ \Sigma_1 \\ \vdots \\ \Sigma_m \end{pmatrix}.$$

If, for $i = 1, \ldots, m$ and $\Theta \in \mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$, we denote

$$A_i(\Theta) = \frac{\alpha_i}{N} \sum_{k=1}^{N} \frac{p_i(x_k)}{p(x_k)},$$

$$M_i(\Theta) = \left\{\frac{1}{N} \sum_{k=1}^{N} x_k \frac{p_i(x_k)}{p(x_k)}\right\} \Big/ \left\{\frac{1}{N} \sum_{k=1}^{N} \frac{p_i(x_k)}{p(x_k)}\right\},$$

$$S_i(\Theta) = \left\{\frac{1}{N} \sum_{k=1}^{N} (x_k-\mu_i)(x_k-\mu_i)^T \frac{p_i(x_k)}{p(x_k)}\right\} \Big/ \left\{\frac{1}{N} \sum_{k=1}^{N} \frac{p_i(x_k)}{p(x_k)}\right\},$$

then the likelihood equations can be written as

$$\Theta = \begin{pmatrix} A(\Theta) \\ M(\Theta) \\ S(\Theta) \end{pmatrix} \tag{2}$$

where

$$A(\Theta) = \begin{pmatrix} A_1(\Theta) \\ \vdots \\ A_m(\Theta) \end{pmatrix}, \quad M(\Theta) = \begin{pmatrix} M_1(\Theta) \\ \vdots \\ M_m(\Theta) \end{pmatrix}, \quad S(\Theta) = \begin{pmatrix} S_1(\Theta) \\ \vdots \\ S_m(\Theta) \end{pmatrix}.$$

One can write (2) equivalently as

$$\Theta = \Phi_\epsilon(\Theta) \equiv (1-\epsilon)\Theta + \epsilon \begin{pmatrix} A(\Theta) \\ M(\Theta) \\ S(\Theta) \end{pmatrix} \tag{3}$$

for any value of ϵ. Of course, (3) becomes (2) when $\epsilon = 1$.

The following iterative procedure is suggested by (3) for obtaining a solution of the likelihood equations: Beginning with some starting value $\Theta^{(1)}$,

define successive iterates iductively by

$$\Theta^{(k+1)} = \Phi_\epsilon(\Theta^{(k)}) \tag{4}$$

for $k = 1,2,3,\ldots$. This is the general iterative procedure with which this note is concerned. Clearly, this procedure becomes the procedure given in the introduction when $\epsilon = 1$.

In the next section, we show that if $0 < \epsilon < 2$, then, with probability 1 as N approaches infinity, this procedure converges locally to the strongly consistent maximum-likelihood estimate. This is done by showing that, with probability 1 as N approaches infinity, the operator Φ_ϵ is <u>locally contractive</u> (in a suitable vector norm) near that estimate, provided $0 < \epsilon < 2$. In saying that Φ_ϵ is locally contractive near a point $\Theta \in \mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$, we mean that there is a vector norm $||\ ||$ on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$ and a number λ, $0 \leq \lambda < 1$, such that

$$\left\| \Phi_\epsilon(\Theta') - \Theta \right\| \leq \lambda \left\| \Theta' - \Theta \right\| \tag{5}$$

whenever Θ' lies sufficiently near Θ.

3. <u>The local contractibility and convergence results</u>.

We now establish the following

<u>THEOREM</u>. With probability 1 as N approaches infinity, Φ_ϵ is a locally contractive operator (in some norm on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$) near the strongly consistent maximum-likelihood estimate whenever $0 < \epsilon < 2$.

Our main result, given by the following corollary, is an immediate consequence of this theorem.

COROLLARY. With probability 1 as N approaches infinity, the iterative procedure (4) converges locally to the strongly consistent maximum-likelihood estimate whenever $0 < \epsilon < 2$.

Throughout the remainder of this paper, the symbol "∇" denotes the Fréchet derivative of a vector-valued function of a vector variable. When ambiguity exists, the specific vector variable of differentiation appears as a subscript of this symbol. For questions concerning the definition and properties of Fréchet derivatives, see Luenberger [10].

Proof of the theorem: Let

$$\Theta = \begin{pmatrix} \bar{\alpha} \\ \bar{\mu} \\ \bar{\Sigma} \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_m \\ \mu_1 \\ \vdots \\ \mu_m \\ \Sigma_1 \\ \vdots \\ \Sigma_m \end{pmatrix}$$

be the strongly consistent maximum-likelihood estimate. We assume that $\alpha_i \neq 0$, $i = 1,\ldots,m$. (As N tends to infinity, the probability is 1 that this is the case.) It must be shown that, with probability 1 as N approaches infinity, an inequality of the form (5) holds whenever ϵ is a sufficiently small positive number.

For any norm on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$, one can write

$$\Phi_\epsilon(\Theta') - \Theta = \nabla\Phi_\epsilon(\Theta)\,[\Theta' - \Theta] + O(\|\Theta' - \Theta\|^2)$$

Consequently, the theorem will be proved if it can be shown that, for $0 < \epsilon < 2$, $\nabla\Phi_\epsilon(\Theta)$ converges with probability 1 to an operator which has operator norm less than 1 with respect to a suitable vector norm on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$.

One can write $\nabla\Phi_\epsilon$ as $(1-\epsilon)I$ plus a matrix of Frechet derivatives:

$$\nabla\Phi_\epsilon = (1-\epsilon)I + \epsilon \begin{pmatrix} \nabla_{\vec{\alpha}}A & \nabla_{\vec{\mu}}A & \nabla_{\vec{\Sigma}}A \\ \nabla_{\vec{\alpha}}M & \nabla_{\vec{\mu}}M & \nabla_{\vec{\Sigma}}M \\ \nabla_{\vec{\alpha}}S & \nabla_{\vec{\mu}}S & \nabla_{\vec{\Sigma}}S \end{pmatrix}$$

This is consistent with our representation of elements of $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$ as columns.

The entries of the above matrix can themselves be represented as matrices of Frechet derivatives. For $i = 1,\ldots,m$, we introduce inner products $\langle x,y\rangle_i' = x^T(\alpha_i\Sigma_i^{-1})y$ on $\mathbb{R}^n$ and $\langle A,B\rangle_i'' = \mathrm{tr}\{A(\frac{\alpha_i}{2}\Sigma_i^{-1})B^T\}$ on the space of real, symmetric $n\times n$ matrices, and we define $\beta_i(x) = \frac{p_i(x)}{p(x)}$, $\gamma_i(x) = (x-\mu_i)$, and $\delta_i(x) = [\Sigma_i^{-1}(x-\mu_i)(x-\mu_i)^T - I]$.

After a straightforward but extremely tedious calculation, one obtains with the aid of equations (1) that

$$\nabla_{\bar{\alpha}} A(\Theta) = I - (\operatorname{diag} \alpha_i)\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k) \\ \vdots \\ \beta_m(x_k) \end{pmatrix} \begin{pmatrix} \beta_1(x_k) \\ \vdots \\ \beta_m(x_k) \end{pmatrix}^T \}$$

$$\nabla_{\bar{\mu}} A(\Theta) = -(\operatorname{diag} \alpha_i)\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k) \\ \vdots \\ \beta_m(x_k) \end{pmatrix} \begin{pmatrix} <\beta_1(x_k)\gamma_1(x_k), \cdot >'_1 \\ \vdots \\ <\beta_m(x_k)\gamma_m(x_k), \cdot >'_m \end{pmatrix}^T \}$$

$$\nabla_{\bar{\Sigma}} A(\Theta) = -(\operatorname{diag} \alpha_i)\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k) \\ \vdots \\ \beta_m(x_k) \end{pmatrix} \begin{pmatrix} <\beta_1(x_k)\delta_1(x_k), \cdot >''_1 \\ \vdots \\ <\beta_m(x_k)\delta_m(x_k), \cdot >''_m \end{pmatrix}^T \}$$

$$\nabla_{\bar{\alpha}} M(\Theta) = -\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k)\gamma_1(x_k) \\ \vdots \\ \beta_m(x_k)\gamma_m(x_k) \end{pmatrix} \begin{pmatrix} \beta_1(x_k) \\ \vdots \\ \beta_m(x_k) \end{pmatrix}^T \}$$

$$\nabla_{\bar{\mu}} M(\Theta) = I - \{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k)\gamma_1(x_k) \\ \vdots \\ \beta_m(x_k)\gamma_m(x_k) \end{pmatrix} \begin{pmatrix} <\beta_1(x_k)\gamma_1(x_k), \cdot >'_1 \\ \vdots \\ <\beta_m(x_k)\gamma_m(x_k), \cdot >'_m \end{pmatrix}^T \}$$

$$\nabla_{\bar{\Sigma}} M(\theta) = (\operatorname{diag} \frac{1}{\alpha_i N} \sum_{k=1}^{N} \beta_i(x_k)\gamma_i(x_k) <\delta_i(x_k), \cdot >''_i) -$$

$$-\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k)\gamma_1(x_k) \\ \vdots \\ \beta_m(x_k)\gamma_m(x_k) \end{pmatrix} \begin{pmatrix} <\beta_1(x_k)\gamma_1(x_k), \cdot >''_1 \\ \vdots \\ <\beta_m(x_k)\gamma_m(x_k), \cdot >''_m \end{pmatrix}^T \}$$

$$\nabla_{\bar{\alpha}} S(\Theta) = -(\operatorname{diag} \Sigma_i)\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k)\delta_1(x_k) \\ \vdots \\ \beta_m(x_k)\delta_m(x_k) \end{pmatrix} \begin{pmatrix} \beta_1(x_k) \\ \vdots \\ \beta_m(x_k) \end{pmatrix}^T \}$$

$$\nabla_{\bar{\mu}} S(\Theta) = (\operatorname{diag} \Sigma_i \frac{1}{\alpha_i N} \sum_{k=1}^{N} \beta_i(x_k)\delta_i(x_k) \langle \gamma_i(x_k), \cdot \rangle_i')-$$

$$-(\operatorname{diag} \Sigma_i)\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k)\delta_1(x_k) \\ \vdots \\ \beta_m(x_k)\delta_m(x_k) \end{pmatrix} \begin{pmatrix} \langle \beta_1(x_k)\gamma_1(x_k), \cdot \rangle_1' \\ \vdots \\ \langle \beta_m(x_k)\gamma_m(x_k), \cdot \rangle_m' \end{pmatrix}^T \}$$

$$\nabla_{\bar{\Sigma}} S(\Theta) = (\operatorname{diag} \Sigma_i \frac{1}{\alpha_i N} \sum_{k=1}^{N} \beta_i(x_k)\delta_i(x_k) \langle \delta_i(x_k), \cdot \rangle_i'') -$$

$$- (\operatorname{diag} \Sigma_i)\{\frac{1}{N} \sum_{k=1}^{N} \begin{pmatrix} \beta_1(x_k)\delta_1(x_k) \\ \vdots \\ \beta_m(x_k)\delta_m(x_k) \end{pmatrix} \begin{pmatrix} \langle \beta_1(x_k)\delta_1(x_k), \cdot \rangle_1'' \\ \vdots \\ \langle \beta_m(x_k)\delta_m(x_k), \cdot \rangle_m'' \end{pmatrix}^T \}$$

The inner products $\langle \cdot, \cdot \rangle_i'$ and $\langle \cdot, \cdot \rangle_i''$, together with scalar multiplication on $\mathbb{R}^1$, induce an inner product $\langle \cdot, \cdot \rangle$ on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$. Setting

$$V(x) = \begin{pmatrix} \beta_1(x) \\ \vdots \\ \beta_m(x) \\ \beta_1(x)\gamma_1(x) \\ \vdots \\ \beta_m(x)\gamma_m(x) \\ \beta_1(x)\delta_1(x) \\ \vdots \\ \beta_m(x)\delta_m(x) \end{pmatrix} \in \mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S},$$

$$B_{23} = (\text{diag } \frac{1}{\alpha_i N} \sum_{k=1}^{N} \beta_i(x_k)\gamma_i(x_k) \langle \delta_i(x_k), \cdot \rangle_i''),$$

$$B_{32} = (\text{diag } \Sigma_i \frac{1}{\alpha_i N} \sum_{k=1}^{N} \beta_i(x_k)\delta_i(x_k) \langle \gamma_i(x_k), \cdot \rangle_i'),$$

$$B_{33} = (\text{diag } \Sigma_i \frac{1}{\alpha_i N} \sum_{k=1}^{N} \beta_i(x_k)\delta_i(x_k) \langle \delta_i(x_k), \cdot \rangle_i''),$$

one obtains

$$\nabla\Phi_\epsilon(\Theta) = \begin{pmatrix} I & 0 & 0 \\ 0 & I & \epsilon B_{23} \\ 0 & \epsilon B_{32} & (1-\epsilon)I + \epsilon B_{33} \end{pmatrix} -$$

$$- \epsilon \begin{pmatrix} (\text{diag } \alpha_i) & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (\text{diag } \Sigma_i) \end{pmatrix} \{\frac{1}{N} \sum_{k=1}^{N} V(x_k) \langle V(x_k), \cdot \rangle\}.$$

Denoting the vector of true parameters by Θ, one verifies without difficulty that $\nabla\Phi_\epsilon$ is of the form

$$\nabla\Phi_\epsilon(\Theta) = \frac{1}{N} \sum_{k=1}^{N} F(x_k, \Theta),$$

where the operator $F(x,\Theta)$ not only has finite expectation (in norm) at Θ but also has a Fréchet derivative with respect to Θ for which the following holds: If $|| \; ||$ is any operator norm on $\nabla_\Theta F$, then there exists a real-valued function f on $\mathbb{R}^n$ such that

$$\int_{\mathbb{R}^n} f(x)p(x)dx < \infty$$

at Θ° and such that $||\nabla_\Theta F(x,\Theta)|| \leq f(x)$ for all $x \in \mathbb{R}^n$ and all Θ in a sufficiently small neighborhood of Θ°. Since the solution Θ of the likelihood equations is strongly consistent, it follows from the Strong Law of Large Numbers (see Loéve [9]) that $\nabla\Phi_\epsilon(\Theta)$ converges with probability 1 to $E(\nabla\Phi_\epsilon(\Theta^\circ))$ as N approaches infinity.

To complete the proof of the theorem, it must be shown that $E(\nabla\Phi_\epsilon(\Theta^\circ))$ has operator norm less than 1 with respect to some vector norm on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$ whenever $0 < \epsilon < 2$. A straightforward calculation yields

$$E(\nabla\Phi_\epsilon(\Theta^\circ)) = \begin{pmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix} - \epsilon \begin{pmatrix} (\text{diag } \alpha_i^\circ) & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (\text{diag } \Sigma_i^\circ) \end{pmatrix} \{ \int_{\mathbb{R}^n} V(x) \langle V(x), \cdot \rangle p(x)dx \},$$

where $\alpha_i^\circ, \mu_i^\circ$, and Σ_i°, $i = 1,\ldots,m$, are the components of Θ°. Thus $E(\nabla\Phi_\epsilon(\Theta^\circ))$ is an operator on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$ of the form $I - \epsilon QR$, where

$$Q = \begin{pmatrix} (\text{diag } \alpha_i^\circ) & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (\text{diag } \Sigma_i^\circ) \end{pmatrix}$$

and

$$R = \int_{\mathbb{R}^n} V(x) \langle V(x), \cdot \rangle p(x)dx$$

are positive-definite and symmetric with respect to the inner product $\langle \cdot, \cdot \rangle$. Since QR is positive-definite and symmetric with respect to the inner product $\langle \cdot, Q^{-1} \cdot \rangle$ on $\mathcal{A} \oplus \mathcal{M} \oplus \mathcal{B}$, it suffices to show that

$$\langle W, RW \rangle = \langle W, Q^{-1}[QR]W \rangle \leq \langle W, Q^{-1}W \rangle$$

for all $W \in \mathcal{A} \oplus \mathcal{M} \oplus \mathcal{B}$. Indeed, it follows from this inequality that, with respect to the inner product $\langle \cdot, Q^{-1} \cdot \rangle$, the operator norm of QR is no greater than 1 and, hence, the operator norm of $E(\nabla \Phi_{\epsilon}(\Theta^{o}))$ is less than 1 whenever $0 < \epsilon < 2$.

For

$$W = \begin{pmatrix} y_1 \\ \vdots \\ y_m \\ v_1 \\ \vdots \\ v_m \\ B_1 \\ \vdots \\ B_m \end{pmatrix} \in \mathcal{A} \oplus \mathcal{M} \oplus \mathcal{B},$$

one has

$$\langle W, RW \rangle = \int_{\mathbb{R}^n} [\sum_{i=1}^{m} y_i \beta_i(x) + \sum_{i=1}^{m} v_i^T (\alpha_i^o \Sigma_i^{o-1}) \beta_i(x) \gamma_i(x) + \sum_{i=1}^{m} \operatorname{tr}\{B_i (\frac{\alpha_i^o}{2} \Sigma_i^{o-1}) \beta_i(x) \delta_i(x)^T\}]^2 p(x) dx$$

$$= \int_{\mathbb{R}^n} [\sum_{i=1}^{m} (\alpha_i^{\sigma-1} y_i + v_i^T \Sigma_i^{\sigma-1} \gamma_i(x) + \mathrm{tr}\{B_i (\tfrac{1}{2}\Sigma_i^{\sigma-1}) \delta_i(x)^T\}) \alpha_i^{\sigma} \beta_i(x)]^2 p(x) dx$$

$$\leq \int_{\mathbb{R}^n} [\sum_{i=1}^{m} (\alpha_i^{\sigma-1} y_i + v_i^T \Sigma_i^{\sigma-1} \gamma_i(x) + \mathrm{tr}\{B_i (\tfrac{1}{2}\Sigma_i^{\sigma-1}) \delta_i(x)^T\})^2 \alpha_i^{\sigma} \beta_i(x)] p(x) dx$$

The inequality is a consequence of the following corollary of Schwarz's inequality: If $\eta_i \geq 0$ for $i = 1,\ldots,m$ and if $\sum_{i=1}^{m} \eta_i = 1$, then $|\sum_{i=1}^{m} \xi_i \eta_i|^2 \leq \sum_{i=1}^{m} \xi_i^2 \eta_i$ for all $\{\xi_i\}_{i=1,\ldots,m}$. If the squared expressions in the last sum above are written out in full, one sees that the integrals of the cross terms in these expressions vanish. Consequently

$$\langle W, RW \rangle \leq \int_{\mathbb{R}^n} \sum_{i=1}^{m} [\alpha_i^{\sigma-2} y_i^2 + (v_i^T \Sigma_i^{\sigma-1} \gamma_i(x))^2 + (\mathrm{tr}\{B_i (\tfrac{1}{2}\Sigma_i^{\sigma-1}) \delta_i(x)^T\})^2] \alpha_i^{\sigma} p_i(x) dx$$

Now

(6.a) $$\int_{\mathbb{R}^n} \alpha_i^{\sigma-1} y_i^2 p_i(x) dx = \alpha_i^{\sigma-1} y_i^2$$

(6.b) $$\int_{\mathbb{R}^n} (v_i^T \Sigma_i^{\sigma-1} \gamma_i(x))^2 \alpha_i^{\sigma} p_i(x) dx = \int_{\mathbb{R}^n} v_i^T \Sigma_i^{\sigma-1} (x-\mu_i^{\sigma})(x-\mu_i^{\sigma})^T \Sigma_i^{\sigma-1} v_i \alpha_i^{\sigma} p_i(x) dx$$

$$= \langle v_i, v_i \rangle_i'$$

(6.c) $$\int_{\mathbb{R}^n} (\mathrm{tr}\{B_i (\tfrac{1}{2}\Sigma_i^{\sigma-1}) \delta_i(x)^T\})^2 \alpha_i^{\sigma} p_i(x) dx = \langle B_i, \Sigma_i^{\sigma-1} B_i \rangle_i''.$$

(A proof of (6.c) follows below.) From (6.a), (6.b), and (6.c), one concludes that

$$\langle W,RW\rangle \le \sum_{i=1}^{m}\alpha_i^{o-1}y_i^2 + \sum_{i=1}^{m}\langle v_i,v_i\rangle_i' + \sum_{i=1}^{m}\langle B_i,\Sigma_i^{o-1}B_i\rangle_i'' = \langle W,Q^{-1}W\rangle.$$

This completes the proof of the theorem.

Proof of (6.c): Setting $y = \Sigma_i^{o-1/2}(x-\mu_i^o)$ and

$$C = \int_{\mathbb{R}^n} (\mathrm{tr}\{B_i(\tfrac{1}{2}\Sigma_i^{o-1})\delta_i(x)^T\})^2\alpha_i^o p_i(x)dx,$$

one verifies that

$$C = \frac{\alpha_i^o}{4}\int_{\mathbb{R}^n} (\mathrm{tr}\{B_i[\Sigma_i^{o-1/2}yy^T\Sigma_i^{o-1/2} - \Sigma_i^{o-1}]^T\})^2 p_o(y)dy,$$

where $p_o \sim N(0,I)$. Denoting $\Sigma_i^{o-1/2}B_i\Sigma_i^{o-1/2} = D = (d_{jk})$, one then derives

$$C = \frac{\alpha_i^o}{4}\int_{\mathbb{R}^n} (\mathrm{tr}\{D[yy^T - I]\})^2 p_o(y)dy$$

$$= \frac{\alpha_i}{4}\int_{\mathbb{R}^n} [(\mathrm{tr}\{Dyy^T\})^2 - 2\mathrm{tr}\{D\}\mathrm{tr}\{Dyy^T\} + (\mathrm{tr}\{D\})^2]p_o(y)dy$$

$$= \frac{\alpha_i^o}{4}\{\sum_{j,k,p,q} d_{jk}d_{pq}\int_{\mathbb{R}^n} y_k y_j y_q y_p p_o(y)dy - 2(\mathrm{tr}\{D\})^2 + (\mathrm{tr}\{B\})^2\}$$

$$= \frac{\alpha_i^o}{4}\{\sum_k\sum_{p\neq k} d_{kk}d_{pp} + \sum_k\sum_{j\neq k} d_{jk}d_{jk} + \sum_k\sum_{j\neq k} d_{jk}d_{kj} + 3\sum_k d_{kk}^2 - (\mathrm{tr}\{D\})^2\}$$

$$= \frac{\alpha_i^o}{2}\mathrm{tr}\{D^2\} = \frac{\alpha_i^o}{2}\mathrm{tr}\{\Sigma_i^{o-1/2}B_i\Sigma_i^{o-1}B_i\Sigma_i^{o-1/2}\} = \mathrm{tr}\{B_i(\frac{\alpha_i}{2}\Sigma_i^{o-1})(\Sigma_i^{o-1}B_i)^T\}$$

$$= \langle B_i,\Sigma_i^{o-1}B_i\rangle_i''.$$

4. The optimal ϵ.

The results just obtained state that, with probability 1 as N approaches infinity, the iterative procedure (4) converges locally to the strongly consistent maximum-likelihood estimate Θ whenever $0 < \epsilon < 2$. In this section we observe that there exists a particular value of ϵ, referred to as "the optimal ϵ," which yields, with probability 1, the fastest asymptotic uniform rate of local convergence of (4) near Θ. We derive a lower bound between 1 and 2 on the optimal ϵ and relate it to the separation of the component populations in the mixture.

From the proof of the theorem, one sees that the optimal ϵ is that which minimizes the spectral radius of the operator $E(\nabla\Phi_\epsilon(\Theta^o))$ restricted to the space $\mathcal{E} \oplus \mathcal{M} \oplus \mathcal{S}$, where $\mathcal{E}$ is the subspace of $\mathcal{A}$ whose components sum to zero. Indeed, the restricted operator $E(\nabla\Phi_\epsilon(\Theta^o)) = I - \epsilon QR$ is symmetric on $\mathcal{E} \oplus \mathcal{M} \oplus \mathcal{S}$ with respect to the inner product $\langle \cdot, Q^{-1} \cdot \rangle$. Consequently, its operator norm with respect to this inner product is equal to its spectral radius and, hence, minimal. We observe that the restriction of QR to $\mathcal{E} \oplus \mathcal{M} \oplus \mathcal{S}$ is positive-definite and symmetric with respect to the inner product $\langle \cdot, Q^{-1} \cdot \rangle$. Letting ρ and τ denote, respectively, the largest and smallest eigenvalues of this restriction of QR, one verifies that the spectral radius of $E(\nabla\Phi_\epsilon(\Theta^o))$, restricted to $\mathcal{E} \oplus \mathcal{M} \oplus \mathcal{S}$, is minimized when $1 - \epsilon\tau = \epsilon\rho - 1$, i.e., when $\epsilon = \frac{2}{\rho+\tau}$.

It follows from the proof of the theorem that ρ is never greater than 1. Thus the optimal ϵ is bounded below by $2/1+\tau$, where τ lies between 0 and 1. In particular, this lower bound on the optimal ϵ lies between 1 and 2. We have been unable to determine ρ more precisely in general.

It should be noted that, if ρ is strictly less than $\frac{1}{2}$, then the optimal ϵ is actually greater than 2, even though the theorem just proved fails to guarantee the local convergence of (4) for such values of ϵ.

Suppose that the component populations in the mixture are "widely separated" in the sense that each pair (μ_i^o, Σ_i^o) differs greatly from every other such pair. Then, for $i,j = 1,\ldots,m$,

$$\frac{\alpha_i^o p_i(x)}{p(x)} \quad \frac{\alpha_j^o p_j(x)}{p(x)} \approx 0 \quad \text{for } x \in \mathbb{R}^n \text{ whenever } i \neq j.$$

One sees that $QR \approx I$ and, hence, ρ and τ must both lie near 1. Consequently, fastest asymptotic local convergence rates are obtained for ϵ near 1, and, for the optimal ϵ, $E(\nabla\Phi_\epsilon(\Theta^o)) = I - \epsilon QR \approx 0$. Thus for mixtures whose component populations are "widely separated," the optimal ϵ is only slightly greater than 1, and rapid first-order local convergence of the iterative procedure (4) to Θ can be expected asymptotically for this ϵ.

Now suppose that the component populations in the mixture are such that at least two pairs (μ_i^o, Σ_i^o) and (μ_j^o, Σ_j^o), $i \neq j$, are nearly identical. Then $\beta_i(x) \approx \beta_j(x)$, $\beta_i(x)\gamma_i(x) \approx \beta_j(x)\gamma_j(x)$ and $\beta_i(x)\delta_i(x) \approx \beta_i(x)\delta_j(x)$, and it follows that R is nearly singular and, hence, that τ is near zero. One concludes that the optimal ϵ cannot be much smaller than 2. In fact, if ρ is near 1, as is the case when all pairs (μ_i^o, Σ_i^o) are nearly identical, then the optimal ϵ must lie near 2. Furthermore, the spectral radius of $E(\nabla\Phi_\epsilon(\Theta^o))$ is near 1, even for the optimal ϵ; therefore, slow first-order convergence can be expected asymptotically in this case.

5. Maximum-likelihood estimates of the a priori probabilities and the means.

It happens that, if the covariance matrices Σ_i, $i = 1,\ldots,m$, are held fixed, then, under certain conditions, an appropriately restricted version of the iterative procedure (4) converges locally with probability 1 to a maximum-likelihood estimate of the parameters α_i^o and μ_i^o, $i = 1,\ldots,m$, whenever the number of observations in the sample reaches a certain finite size. To be more specific, we introduce the following notation: For $\tilde{\Theta} = \begin{pmatrix}\alpha\\ \mu\end{pmatrix} \in \mathcal{A} \oplus \mathcal{M}$ and $\bar{\Sigma} \in \mathcal{S}$, denote $\tilde{\theta} = \begin{pmatrix}\alpha\\ \mu\\ \Sigma\end{pmatrix} \in \mathcal{A} \oplus \mathcal{M} \oplus \mathcal{S}$ by $(\tilde{\Theta},\bar{\Sigma})$. Then, for given $\bar{\Sigma}$, the likelihood equations for the parameters α_i and μ_i, $i = 1,\ldots,m$ can be written as

$$\tilde{\Theta} = \begin{pmatrix} A(\tilde{\Theta},\bar{\Sigma}) \\ M(\tilde{\Theta},\bar{\Sigma}) \end{pmatrix}$$

or, equivalently, as

$$\tilde{\Theta} = \tilde{\Phi}_\epsilon(\tilde{\Theta},\bar{\Sigma}) \equiv (1-\epsilon)\tilde{\Theta} + \epsilon \begin{pmatrix} A(\tilde{\Theta},\bar{\Sigma}) \\ M(\tilde{\Theta},\bar{\Sigma}) \end{pmatrix} \tag{7}$$

for any ϵ. The appropriate iterative procedure to consider is the following: Beginning with some starting value $\tilde{\Theta}^{(1)}$, define successive iterates inductively by

$$\tilde{\Theta}^{(k+1)} = \tilde{\Phi}_\epsilon(\tilde{\Theta}^{(k)},\bar{\Sigma}) \tag{8}$$

for $k = 1,2,3,\ldots$. Our result concerning this procedure is given by the theorem and its corollary below.

THEOREM: If $N \geq m(n+1)$ and if $(\tilde{\Theta},\bar{\Sigma})$ is a solution of (7) which lies sufficiently near a solution of (3), then, with probability 1, $\tilde{\Phi}_\epsilon$ is a locally contractive operator (in some norm on $\mathcal{A}\oplus\mathcal{M}$) near $\tilde{\Theta}$ whenever $0 < \epsilon < 2$.

COROLLARY: If $N \geq m(n+1)$ and if $(\tilde{\Theta},\bar{\Sigma})$ is a solution of (7) which lies sufficiently near a solution of (3), then, with probability 1, the iterative procedure (8) converges locally to $\tilde{\Theta}$ whenever $0 < \epsilon < 2$.

Proof of the Theorem: Suppose that $N \geq m(n+1)$ and $0 < \epsilon < 2$. As in the proof of the preceding theorem, it suffices to show that, with probability 1, $\nabla_{\tilde{\Theta}}\tilde{\Phi}_\epsilon(\tilde{\Theta},\bar{\Sigma})$ has operator norm less than 1 with respect to some vector norm on $\mathcal{A}\oplus\mathcal{M}$. Since $\nabla_{\tilde{\Theta}}\tilde{\Phi}_\epsilon$ depends continuously on $\tilde{\Theta}$ and $\bar{\Sigma}$, this need only be shown when $(\tilde{\Theta},\bar{\Sigma})$ is a solution of (3).

From the proof of the preceding theorem, one sees that is $(\tilde{\Theta},\bar{\Sigma})$ is a solution of (3), then

$$\nabla_{\tilde{\Theta}}\tilde{\Phi}_\epsilon(\tilde{\Theta},\bar{\Sigma}) = \begin{pmatrix} I & 0 \\ 0 & I \end{pmatrix} - \begin{pmatrix} (\operatorname{diag}\alpha_i) & 0 \\ 0 & I \end{pmatrix} \{\frac{1}{N}\sum_{k=1}^{N}\tilde{V}(x_k)\langle\tilde{V}(x_k),\cdot\rangle\},$$

where

$$\tilde{V}(x) = \begin{pmatrix} \beta_1(x) \\ \vdots \\ \beta_m(x) \\ \\ \beta_1(x)\gamma_1(x) \\ \vdots \\ \beta_m(x)\gamma_m(x) \end{pmatrix}$$

and the inner product $\langle \cdot, \cdot \rangle$ is now the inner product induced on $\mathcal{A} \oplus \mathcal{M}$ by scalar multiplication on $\mathbb{R}^1$ and the inner products $\langle \cdot, \cdot \rangle_i'$ on $\mathbb{R}^n$. As before, $\nabla_{\tilde{\Theta}} \Phi_\epsilon(\tilde{\Theta}, \tilde{\Sigma})$ is of the form $I - \epsilon \tilde{Q}\tilde{R}$, where

$$\tilde{Q} = \begin{pmatrix} (\text{diag } \alpha_i) & 0 \\ \\ 0 & I \end{pmatrix}$$

and

$$\tilde{R} = \{\frac{1}{N} \sum_{k=1}^{N} \tilde{V}(x_k) \langle \tilde{V}(x_k), \cdot \rangle\}.$$

We observe that $\tilde{Q}\tilde{R}$ is symmetric and positive semi-definite with respect to the inner product $\langle \cdot, \tilde{Q}^{-1} \cdot \rangle$. In fact, it is shown in Appendix 2 that, with probability 1, $\tilde{Q}\tilde{R}$ is positive-definite with respect to this inner product. Consequently, the theorem will be proved if it can be shown that

$$\langle \tilde{W}, \tilde{Q}^{-1}[\tilde{Q}\tilde{R}]\tilde{W} \rangle = \langle \tilde{W}, \tilde{R}\tilde{W} \rangle \leq \langle \tilde{W}, \tilde{Q}^{-1}\tilde{W} \rangle$$

for all $\tilde{W} \in \mathcal{A} \oplus \mathcal{M}$.

For

$$\tilde{W} = \begin{pmatrix} y_1 \\ \vdots \\ y_m \\ v_1 \\ \vdots \\ v_m \end{pmatrix} \in \mathcal{A} \oplus \mathcal{M},$$

one has

$$\begin{aligned} \langle \tilde{W}, \tilde{R}\tilde{W} \rangle &= \frac{1}{N} \sum_{k=1}^{N} [\sum_{i=1}^{m} y_i \beta_i(x_k) + \sum_{i=1}^{m} v_i^T(\alpha_i \Sigma_i^{-1}) \beta_i(x_k) \gamma_i(x_k)]^2 \\ &= \frac{1}{N} \sum_{k=1}^{N} [\sum_{i=1}^{m} (\alpha_i^{-1} y_i + v_i^T \Sigma_i^{-1} \gamma_i(x_k)) \alpha_i \beta_i(x_k)]^2 \\ &\leq \frac{1}{N} \sum_{k=1}^{N} \sum_{i=1}^{m} [\alpha_i^{-1} y_i + v_i^T \Sigma_i^{-1} \gamma_i(x_k)]^2 \alpha_i \beta_i(x_k) \end{aligned}$$

by Schwarz's Inequality. Since $(\tilde{\Theta}, \bar{\Sigma})$ is a solution of (3), this easily yields

$$\langle \tilde{W}, \tilde{R}\tilde{W} \rangle \leq \sum_{i=1}^{m} \alpha_i^{-1} y_i^2 + \sum_{i=1}^{m} v_i^T(\alpha_i \bar{\Sigma}_i) v_i = \langle \tilde{W}, \tilde{Q}^{-1}\tilde{W} \rangle,$$

and the proof is complete.

If the conclusion of the theorem holds for some solution $(\tilde{\Theta}, \bar{\Sigma})$ of (7), then, as in the preceding section, a particular value of ϵ can be determined which yields the fastest uniform rate of local convergence of (8) near $\tilde{0}$. With respect to the inner product $\langle \cdot, \tilde{Q}^{-1} \cdot \rangle$, $\tilde{Q}\tilde{R}$ is positive-definite and symmetric on $\mathcal{E} \oplus \mathcal{M}$. Denoting the largest and smallest eigenvalues of the

restriction of $\mathfrak{QR}$ to $\mathfrak{E} \oplus \mathfrak{M}$ by ρ and τ, respectively, one sees that the optimal ϵ is again given by $\epsilon = \frac{2}{\rho+\tau}$. Since the restriction of $\mathfrak{QR}$ has operator norm no greater than 1 with respect to the inner product $\langle \cdot, Q^{-1} \cdot \rangle$, ρ must be no greater than 1. Hence, $\epsilon \geq \frac{1}{1+\tau}$, where τ lies between 0 and 1. Reasoning as before, one sees that the optimal ϵ lies near 1 if the component populations are "widely separated," and cannot be much less than 2 if two or more of the populations have nearly identical means and covariance matrices. In the former case, rapid first-order local convergence of (8) can be expected for the optimal ϵ. In the latter case, if ρ is near 1, then the optimal ϵ must be near 2, and slow first-order convergence of (8) can be expected, even for the optimal ϵ.

6. Concluding remarks.

A number of numerical techniques for obtaining maximum-likelihood estimates of the parameters for a mixture of normal distributions have been discussed in the literature. In addition to the usual steepest-ascent method for obtaining a local maximum of the log-likelihood function, we mention in particular Newton's method, the method of scoring, and the modifications of these procedures investigated by Kale [8] for obtaining solutions of the likelihood equations. It is our feeling that the iterative procedure (4) offers considerable computational advantages over these procedures in many cases of practical interest.

Even though the partial derivatives of the log-likelihood function are not appreciably more difficult to evaluate than the expressions used in defining the function Φ_ϵ, the procedure (4), which is a generalized steepest-ascent

(deflected gradient) method appears to have two particular advantages over the usual steepest-ascent method. First, the major practical implication of this note is that the iterative procedure (4) converges whenever the step-size ϵ lies in an interval which is completely independent of the particular mixture problem at hand. It is readily ascertained that this cannot be said for the regular steepest-ascent procedure

$$\alpha_i^{(q+1)} = \alpha_i^{(q)} + \epsilon\left[\frac{1}{N}\sum_{k=1}^{N}\frac{p_i(x_k)}{p(x_k)} - \frac{1}{mN}\sum_{j=1}^{m}\sum_{k=1}^{N}\frac{p_j(x_k)}{p(x_k)}\right]$$

$$\mu_i^{(q+1)} = \mu_i^{(q)} + \epsilon\left[\frac{1}{N}\sum_{k=1}^{N}\frac{\alpha_i^{(q)}p_i(x_k)}{p(x_k)}\Sigma_i^{(q)-1}(x_k-\mu_i^{(q)})\right]$$

$$\Sigma_i^{(q+1)} = \Sigma_i^{(q)} + \epsilon\left[\frac{1}{2N}\sum_{k=1}^{N}\frac{\alpha_i^{(q)}p_i(x_k)}{p(x_k)}\left[-\Sigma_i^{(q)-1} + \Sigma_i^{(q)-1}(x_k-\mu_i^{(q)})(x_k-\mu_i^{(q)})^T\Sigma_i^{(q)-1}\right.\right.$$

Second, if ϵ is no greater than 1 the successive iterates defined by (4) automatically satisfy the requisite constraints on the parameters, i.e., the successive Σ_i's are, with probability 1 for large N, positive-definite and the successive α_i's are positive and sum to 1.

Although Newton's method and the method of scoring offer quadratic and near-quadratic convergence, respectively, for large sample sizes, they require at each iteration the inversion of a square matrix whose dimension is equal to the number of independent variables among the parameters, namely $\frac{m(n+1)(n+2)}{2} - 1$. Thus these methods may be less efficient computationally than the iterative procedure (4) if m and n are large, even though they may yield a satisfactory approximate solution after fewer iterations. The modified versions of Newton's method and the method of scoring do not require the re-calculation of the inverse of a

large matrix at each step. However, quadratic convergence is not achieved with these modified methods, and multiplication by a large matrix must still be carried out at each iteration.

Appendix 1.

We now give a brief proof of the existence and uniqueness of the strongly consistent maximum-likelihood estimate. For the sake of generality, this is done in a somewhat broader context than is necessary for this paper.

Let $p(x,\Theta)$ be a probability density function of a vector variable $x \in \mathbb{R}^n$ and a vector parameter $\Theta \in \mathbb{R}^\nu$. If $\{x_k\}_{k=1,\ldots,N}$ is an independent sample of observations on a random variable $x \in \mathbb{R}^n$ whose probability density function is $p(x,\Theta^o)$ for some $\Theta^o \in \mathbb{R}^\nu$, then a maximum-likelihood estimate of Θ^o is a choice of Θ which locally maximizes the log-likelihood function

$$L = \sum_{k=1}^{N} \log p(x_k,\Theta).$$

If p is a differentiable function of Θ, then a necessary condition for a maximum-likelihood estimate is that the likelihood equations

$$\frac{\partial L}{\partial \theta_i} = 0, \quad i = 1,\ldots,\nu,$$

be satisfied, where θ_i is the $i^{\underline{th}}$ component of Θ. In the following, our objective is to show that if p satisfies certain conditions, then, given any sufficiently small neighborhood of Θ^o, there is, with probability 1 as N approaches infinity, a unique solution of the likelihood equations in that neighborhood, and this solution is a maximum-likelihood estimate of Θ^o.

We assume that $p(x,\Theta)$ satisfies the following conditions of Chanda [1]:
(a) There is a neighborhood Ω of Θ^o such that for all $\Theta \in \Omega$, for almost

all $x \in \mathbb{R}^n$, and for $i,j,k=1,\ldots,\nu$, $\frac{\partial p}{\partial \theta_i}$, $\frac{\partial^2 p}{\partial \theta_i \partial \theta_j}$, and $\frac{\partial^3 p}{\partial \theta_i \partial \theta_j \partial \theta_k}$ exist and satisfy

$$\left|\frac{\partial p}{\partial \theta_i}\right| \le f_i(x), \quad \left|\frac{\partial \partial^2 p}{\partial \theta_i \partial \theta_j}\right| \le f_{ij}(x), \quad \left|\frac{\partial^3 \log p}{\partial \theta_i \partial \theta_j \partial \theta_k}\right| \le f_{ijk}(x),$$

where f_i and f_{ij} are integrable and f_{ijk} satisfies

$$\int_{\mathbb{R}^n} f_{ijk}(x) p(x,\Theta^o)\,dx < \infty.$$

(b) The matrix $J(\Theta) = \left(\int_{\mathbb{R}^n} \frac{\partial \log p}{\partial \theta_i} \frac{\partial \log p}{\partial \theta_j} p\,dx\right)$ is positive-definite at Θ^o.

$$\text{Let } \mathcal{L}(\Theta) = \begin{pmatrix} \frac{1}{N}\frac{\partial L}{\partial \theta_1} \\ \vdots \\ \frac{1}{N}\frac{\partial L}{\partial \theta_\nu} \end{pmatrix}.$$

It is immediately seen that $\mathcal{L}(\Theta) = 0$ if and only if the likelihood equations are satisfied, and that, by the Law of Large Numbers, $\mathcal{L}(\Theta^o)$ converges with probability 1 to zero. Furthermore, it follows from assumptions (a) and (b) above that there exists a neighborhood Ω^o of Θ^o (contained in Ω and, for convenience, convex) and a positive ϵ such that, with probability 1 as N approaches infinity, $\nabla\mathcal{L}(\Theta) \le -\epsilon I$ for all $\Theta \in \Omega^o$. (The inequality is with respect to the usual ordering on symmetric matrices.) Denoting the spherical neighborhood of radius δ about Θ^o by Ω_δ, we establish the following

Lemma: With probability 1 as N approaches infinity,

(1) is one-to-one on Ω^o,

(ii) (Ω_δ) contains the ball of radius $\epsilon\delta$ about (θ^o) whenever $\Omega_\delta \subseteq \Omega^o$.

Proof: We may assume that $\nabla\mathcal{L}(\theta) \leq -\epsilon I$ for all $\theta \in \Omega^o$, since the probability that this is the case is 1 as N approaches infinity. To prove (i), suppose that $\mathcal{L}(\theta^1) = \mathcal{L}(\theta^2)$ for θ^1 and θ^2 in Ω^o. Then

$$0 = (\theta^1 - \theta^2)^T[\mathcal{L}(\theta^1) - \mathcal{L}(\theta^2)]$$
$$= (\theta^1 - \theta^2)^T\{\int_o^1 \nabla\mathcal{L}(\theta^2 + t[\theta^1 - \theta^2])dt\}(\theta^1 - \theta^2).$$

The negative-definiteness of $\nabla\mathcal{L}$ implies that $\theta^1 = \theta^2$, and (i) is proved.

To prove (ii), suppose that $\Omega_\delta \subseteq \Omega^o$, and let θ^1 be a boundary point of Ω_δ. Then

$$\mathcal{L}(\theta^1) - \mathcal{L}(\theta^o) = \{\int_o^1 \nabla\mathcal{L}(\theta^o + t[\theta^1 - \theta^o])dt\}(\theta^1 - \theta^o).$$

After left-multiplying this equation by $(\theta^1 - \theta^o)^T$, one verifies using Schwarz's inequality and the negative-definiteness of $\nabla\mathcal{L}$ that

$$\left\| \mathcal{L}(\theta^1) - \mathcal{L}(\theta^o) \right\| \geq \epsilon \left\| \theta^1 - \theta^o \right\| = \epsilon\,\delta,$$

where $\|\ \|$ denotes the usual Euclidean norm on $\mathbb{R}^\nu$. Since all boundary points of $\mathcal{L}(\Omega_\delta)$ are images under $\mathcal{L}$ of boundary points of Ω_δ, the proof of (ii) is complete.

The desired result of this appendix follows immediately from this lemma and

the remarks preceding it. Indeed, if Ω^1 is any neighborhood of Θ^o which is contained in Ω^o, then one can find a δ for which $\Omega_\delta \subseteq \Omega^1 \subseteq \Omega^o$. By the lemma, the probability is 1 as N tends to infinity that $\mathcal{L}$ is one-to-one on Ω^1 and that $\mathcal{L}(\Omega_\delta)$ and, hence, $\mathcal{L}(\Omega^1)$ contain the ball of radius $\epsilon\delta$ about $\mathcal{L}(\Theta^o)$. Since $\mathcal{L}(\Theta^o)$ converges with probability 1 to zero, one concludes that, with probability 1 as N approaches infinity, there exists a unique $\Theta \in \Omega^1$ for which $\mathcal{L}(\Theta) = 0$. Since the probability also is 1 as N approaches infinity that $\nabla\mathcal{L}$ is negative-definite on Ω^1, this Θ is, with probability 1, a maximum-likelihood estimate.

Appendix 2.

We now prove that the operator $\tilde{\tilde{QR}}$ is positive-definite on $\mathcal{A} \oplus \mathcal{M}$ with probability 1 whenever $N \geq m(n+1)$. Since

$$\tilde{\tilde{QR}} = \begin{pmatrix} (\text{diag } \alpha_i) & 0 \\ 0 & I \end{pmatrix} \{\frac{1}{N} \sum_{k=1}^{N} \tilde{V}(x_k) \langle \tilde{V}(x_k), \cdot \rangle\},$$

it suffices to show that the vectors

$$\tilde{V}(x_k) = \begin{pmatrix} \frac{p_1(x_k)}{p(x_k)} \\ \vdots \\ \frac{p_m(x_k)}{p(x_k)} \\ \\ \frac{p_1(x_k)}{p(x_k)}(x_k - \mu_1) \\ \vdots \\ \frac{p_m(x_k)}{p(x_k)}(x_k - \mu_m) \end{pmatrix}, \quad k = 1, \ldots, N,$$

span $\mathcal{A} \oplus \mathcal{M}$ with probability 1 whenever $N \geq m(n+1)$. This follows from the more general result below.

Lemma. Let $\{x_k\}_{k=1,\ldots,N}$ be an independent sample of observations on a random variable x in $\mathbb{R}^s$ which is distributed with a probability density function p. If V is a real-analytic function from $\mathbb{R}^s$ to $\mathbb{R}^t$ whose component functions are linearly independent, then the vectors $V(x_k)$, $k=1,\ldots,N$, span $\mathbb{R}^t$ with

probability 1 whenever $N \geq t$.

Proof: Denoting the j^{th} component function of V by v_j, we define a real-analytic function V_j from $\mathbb{R}^s$ to $\mathbb{R}^j$ by

$$V_j(x) = \begin{pmatrix} v_1(x) \\ \vdots \\ v_j(x) \end{pmatrix}$$

for $j = 1,\ldots,t$. Our proof of the lemma consists of showing inductively that, for $j = 1,\ldots,t$, the set $\{V_j(x_k)\}_{k=1,\ldots,j}$ spans $\mathbb{R}^j$ with probability 1. We make the preliminary observation that, since the real-analytic functions v_j are assumed to be linearly independent, any non-zero linear combination of them vanishes only on a set of Lebesque measure zero in $\mathbb{R}^s$.

From the observation above, $V_1(x_1)$ is non-zero with probability 1; hence $V_1(x_1)$ spans $\mathbb{R}^1$ with probability 1. Suppose now that, for some j, $1 \leq j < t$, the set $\{V_j(x_k)\}_{k=1,\ldots,j}$ spans $\mathbb{R}^j$ with probability 1. Then, with probability 1, the set $\{V_{j+1}(x_k)\}_{k=1,\ldots,j+1}$ fails to span $\mathbb{R}^{j+1}$ if and only if

$$V_{j+1}(x_{j+1}) = \sum_{k=1}^{j} c_k V_{j+1}(x_k) \tag{9}$$

for some set of constants $\{c_k\}_{k=1,\ldots,j}$. If (9) holds, the constants c_k are determined by

$$\begin{pmatrix} c_1 \\ \vdots \\ c_j \end{pmatrix} = \tilde{V}_j^{-1} V_j(x_{j+1})$$

with probability 1, where $\tilde{V}_j$ is the $j \times j$ matrix whose k^{th} column is $V_j(x_k)$. Thus, with probability 1, (9) holds if and only if

$$[\tilde{V}_j^{-1} V_j(x_{j+1})]^T \begin{pmatrix} v_{j+1}(x_1) \\ \vdots \\ v_{j+1}(x_j) \end{pmatrix} - v_{j+1}(x_{j+1}) = 0.$$

Now

$$[\tilde{V}_j^{-1} V_j(x)]^T \begin{pmatrix} v_{j+1}(x_1) \\ \vdots \\ v_{j+1}(x_j) \end{pmatrix} - v_{j+1}(x)$$

is a non-zero linear combination of the functions $v_1, \ldots, v_{j+1}$ and, hence, vanishes only a set of Lebesque measure zero in $\mathbb{R}^s$. One concludes that $\{V_{j+1}(x_k)\}_{k=1,\ldots,j+1}$ fails to span $\mathbb{R}^{j+1}$ with probability zero. This completes the induction, and the lemma is proved.

BIBLIOGRAPHY

1. K.C. Chanda, "A note on the consistency and maxima of the roots of the likelihood equations," Biometrika 41 (1954), pp. 56-61.

2. H. Cramér, Mathematical Methods of Statistics, Princeton University Press, Princeton, 1946.

3. N.E. Day, "Estimating the components of a mixture of normal distributions," Biometrika 56 (1969), pp. 463-474.

4. R.O. Duda and P.E. Hart, Pattern Classification and Scene Analysis, John Wiley and Sons, Inc., New York, 1973.

5. V. Hasselblad, "Estimation of parameters for a mixture of normal distributions," Technometrics 8 (1966), pp. 431-446.

6. A.S. Householder, Theory of Matrices and Numerical Analysis, Blaisdell Publishing Co., New York, 1964.

7. V.S. Huzurbazar, "The likelihood equation, consistency and the maxima of the likelihood function," Ann. Eugen., Lond. 14 (1948), pp. 185-200.

8. B.K. Kale, "On the solution of the likelihood equations by iteration processes. The multiparametric case," Biometrika 49 (1962), pp. 479-486.

9. M. Loéve, Probability Theory, D. Van Nostrand Co., New York, 1963.

10. D.G. Luenberger, Optimization by Vector Space Methods, John Wiley and Sons, Inc., New York, 1969.

11. B.C. Peters and W.H. Coberly, "The numerical evaluation of the maximum-likelihood estimate of mixture proportions," to appear.

12. A. Wald, "Note on the consistency of the maximum-likelihood estimate," Ann. Math. Stat. 20 (1949), p. 595.

13. J.H. Wolfe, "Pattern clustering by multivariate mixture analysis," Multivariate Behavorial Research 5 (1970), pp. 329-350.

14. S. Zacks, The Theory of Statistical Inference, John Wiley and Sons, Inc., New York, 1971.