

Improved EM for Mixture Proportions with Applications to Nonparametric ML Estimation for Censored Data

Yaming Yu

Department of Statistics, University of California, Irvine

February 18, 2010

Abstract

Improved EM strategies, based on the idea of efficient data augmentation (Meng and van Dyk 1997, 1998), are presented for ML estimation of mixture proportions. The resulting algorithms inherit the simplicity, ease of implementation, and monotonic convergence properties of EM, but have considerably improved speed. Because conventional EM tends to be slow when there exists a large overlap between the mixture components, we can improve the speed without sacrificing the simplicity or stability, if we can reformulate the problem so as to reduce the amount of overlap. We propose simple “squeezing” strategies for that purpose. Moreover, for high-dimensional problems, such as computing the nonparametric MLE of the distribution function with censored data, a natural and effective remedy for conventional EM is to add exchange steps (based on improved EM) between adjacent mixture components, where the overlap is most severe. Theoretical considerations show that the resulting EM-type algorithms, when carefully implemented, are globally convergent. Simulated and real data examples show dramatic improvement in speed in realistic situations.

Keywords: AECM; cocktail algorithm; data augmentation; doubly censored data; EM; global convergence; NPMLE; nonparametric mixtures; squeezing; vertex exchange method.

1 Introduction

Several statistical problems give rise to a likelihood function formally equivalent to that of a finite mixture model with known component densities. One example is maximum likelihood (ML) estimation in a saturated multinomial model with ignorable missing data where some units are partially classified (Dempster et al. 1977). Another example is nonparametric ML estimation (NPMLE) of a mixing distribution (Lindsay 1983) when this distribution is assumed to be supported on a finite grid. Closely related is the NPMLE problem for the distribution function for

censored data (Groeneboom and Wellner 1992; Böhning et al. 1996). The EM algorithm (Dempster et al. 1977; Meng and van Dyk 1997) is among the simplest and best known methods for ML computation in such mixture-like problems; Turnbull (1976) used it on censored data before Dempster et al. (1977) laid down the general framework. The potential slow convergence of EM in general, and for NPMLE computation in particular, is also well documented. Other methods of computing the NPMLE include the iterative convex minorant (ICM) algorithm (Aragon and Eberly 1992; Jongbloed 1998), the vertex exchange method (VEM; Böhning et al. 1996), and constrained Newton methods (Wang 2008). EM or EM-like algorithms are also widely used for related problems such as optimal experimental design (Silvey et al. 1978; Yu 2010a), Poisson image reconstruction using positron emission tomography (Vardi et al. 1985), and channel capacity calculations in Shannon theory (Arimoto 1972; Blahut 1972; Csiszár and Tusnády 1984).

This paper is concerned with improved EM strategies for maximizing a finite mixture log-likelihood with known component densities. Possible extensions to more general problems are mentioned in Section 5. Our main motivation is fast computation of the NPMLE for censored data. The NPMLE problem is challenging partly because of the high dimension (there are many mixture components), and the heavy overlap between components, which slows down conventional EM. Our goal is to design algorithms that improve the speed of EM, but preserve its simplicity, ease of implementation, and monotonic convergence properties. First, we introduce “squeezing” strategies that reformulate the problem so as to reduce the overlap between component densities. Such squeezing strategies capitalize on the idea of efficient data augmentation and are inspired by Fessler and Hero (1994). The resulting EM algorithms converge faster because they correspond to augmented data that are less informative. Secondly, we observe that although “squeezing” may not always be effective for the entire collection of mixture components, we can apply it to sub-collections that overlap most severely. Adding such EM-based conditional maximization steps (nearest neighbor exchanges) can improve the speed dramatically. Overall, our algorithms fit in the broad spectrum of alternating-expectation-conditional-maximization (AECM) schemes (Meng and van Dyk 1997, 1998). The simplicity and effectiveness of these algorithms testify to the advantage of working within the general EM framework (Dempster et al. 1977; Wu 1983; Meng and Rubin 1993; Liu and Rubin 1994; Meng and van Dyk 1997; Liu et al. 1998). Also relevant is the work of Pilla and Lindsay (2001), who focus on the nonparametric mixture problem and

propose pairing nearby components and rotating the pairs for fast ML computation.

In Section 2, we introduce the squeezing strategies to improve EM for maximizing a mixture log-likelihood. Intuitively, squeezing yields an equivalent problem where the mixture components have less overlap, and its effect on the speed of EM is explained in terms of efficient data augmentation. Section 3 argues that squeezing strategies can be effectively implemented to sub-collections of the mixture components. This leads to a “cocktail algorithm” with several different moves that complement each other. A global convergence theorem for the cocktail algorithm is proved. A real-data example is included as an illustration. Section 4 focuses on the NPMLE problem for censored data, and demonstrates the effectiveness of our new algorithms through simulation. Section 5 concludes with a discussion on possible extensions (to the bivariate interval censoring problem, for example). Efficient implementations of our EM-type algorithms, which take advantage of the sparsity features of the NPMLE problem for censored data, are collected in the appendix.

2 EM Algorithms for Mixture Proportions

Suppose n observations $y = (y_1, \dots, y_n)$ are taken from a mixture of m known densities with unknown proportions p_1, \dots, p_m . Writing f_{ij} as the j th component density evaluated at y_i , we can express the log-likelihood for $\mathbf{p} = (p_1, \dots, p_m)$ as

$$l(\mathbf{p}) = \sum_{i=1}^n \log \left(\sum_{j=1}^m f_{ij} p_j \right). \quad (2.1)$$

We seek to maximize (2.1) over $\mathbf{p} \in \Theta$ where

$$\Theta = \left\{ \mathbf{p} : \sum_{j=1}^m p_j = 1, p_j \geq 0; l(\mathbf{p}) > -\infty \right\}.$$

2.1 Conventional EM

Conventional EM introduces latent indicators I_{ij} such that $I_{ij} = 1$ if the i th observation is from component j , and $I_{ij} = 0$ otherwise. At iteration t , when the current estimate of \mathbf{p} is $\mathbf{p}^{(t)} = (p_1^{(t)}, \dots, p_m^{(t)})$, the E-step simply computes the conditional expectation of I_{ij} given observed data and $\mathbf{p}^{(t)}$:

$$E(I_{ij} | y, \mathbf{p}^{(t)}) = \frac{f_{ij} p_j^{(t)}}{\sum_{k=1}^m f_{ik} p_k^{(t)}}.$$

The M-step then sets $p_j^{(t+1)} = \sum_{i=1}^n E(I_{ij}|y, \mathbf{p}^{(t)})/n$. Overall each iteration can be written as

$$p_j^{(t+1)} = \frac{1}{n} \sum_{i=1}^n \left(\frac{f_{ij}}{\sum_{k=1}^m f_{ik} p_k^{(t)}} \right) p_j^{(t)}, \quad j = 1, \dots, m. \quad (2.2)$$

The EM algorithm maintains monotone increase in the log-likelihood, i.e., $l(\mathbf{p}^{(t+1)}) \geq l(\mathbf{p}^{(t)})$. Moreover, when started from the interior of the parameter space, i.e., when $p_j^{(0)} > 0$ for all $1 \leq j \leq m$, EM is guaranteed to converge to $\hat{\mathbf{p}}$, the MLE (Csiszár and Tusnády 1984). Convergence is potentially very slow, however, when there exists heavy overlap among the mixture components.

2.2 Squeezing Strategy I

To improve conventional EM, let us introduce an auxiliary vector $g = (g_1, \dots, g_n)$ and write the objective function (2.1) as

$$l(\mathbf{p}) = n \log 2 + \sum_{i=1}^n \log \left(\sum_{j=1}^m g_i p_j / 2 + \sum_{j=1}^m (f_{ij} - g_i) p_j / 2 \right) \quad (2.3)$$

$$= n \log 2 + \sum_{i=1}^n \log \left(g_i / 2 + \sum_{j=1}^m (f_{ij} - g_i) p_j / 2 \right). \quad (2.4)$$

We require $g_i \geq 0$ and $f_{ij} - g_i \geq 0$ for all i , i.e.,

$$0 \leq g_i \leq \min_j f_{ij}. \quad (2.5)$$

Conventional EM (2.2) can also be derived from (2.3), viewing it as a mixture log-likelihood with $2m$ components with proportions $p_j/2$, $j = 1, \dots, m$, each appearing twice. Specifically, under this new formulation, we let the density of component j ($j = 1, \dots, 2m$) evaluated at observation i be

$$\tilde{f}_{ij} = \begin{cases} f_{ij} - g_i, & 1 \leq j \leq m, \\ g_i, & m+1 \leq j \leq 2m. \end{cases}$$

Let \tilde{I}_{ij} be the latent indicator of whether the i th observation is from component j , $j = 1, \dots, 2m$. The j th component has proportion $p_j/2$ if $1 \leq j \leq m$, and $p_{j-m}/2$ if $m+1 \leq j \leq 2m$. Then the E-step becomes

$$E\left(\tilde{I}_{ij} \mid y, \mathbf{p}^{(t)}\right) = \frac{\tilde{f}_{ij}}{g_i + \sum_{k=1}^m \tilde{f}_{ik} p_k^{(t)}} \times \begin{cases} p_j^{(t)}, & 1 \leq j \leq m, \\ p_{j-m}^{(t)}, & m+1 \leq j \leq 2m, \end{cases} \quad (2.6)$$

and the M-step becomes

$$p_j^{(t+1)} \propto \sum_{i=1}^n E \left(\tilde{I}_{ij} + \tilde{I}_{i,j+m} \mid y, \mathbf{p}^{(t)} \right), \quad j = 1, \dots, m. \quad (2.7)$$

Routine algebra reveals that the resulting EM iteration is the same as (2.2).

We can also apply EM to maximize (2.4), viewing it as a problem with $m + 1$ mixture components, one of which has proportion $1/2$. Equivalently, in the above derivation of (2.6) and (2.7), instead of \tilde{I}_{ij} , $1 \leq i \leq n$, $1 \leq j \leq 2m$, let us treat

$$\tilde{I}_{ij}, \quad 1 \leq i \leq n, \quad 1 \leq j \leq m, \quad \text{and} \quad \tilde{I}_{i0} \equiv \sum_{j=m+1}^{2m} \tilde{I}_{ij}, \quad 1 \leq i \leq n,$$

as the set of latent indicators. The latter, being a collapsed version of the former, contains less information about \mathbf{p} . Note that \tilde{I}_{i0} is the indicator of a mixture component with proportion $1/2$, whose density evaluated at observation i is g_i . The E-step proceeds to calculate the conditional expectation of \tilde{I}_{ij} , $j = 0, 1, \dots, m$, resulting in the same formula as (2.6) for $j = 1, \dots, m$. The M-step becomes

$$p_j^{(t+1)} \propto \sum_{i=1}^n E \left(\tilde{I}_{ij} \mid y, \mathbf{p}^{(t)} \right), \quad j = 1, \dots, m. \quad (2.8)$$

The conditional expectation of \tilde{I}_{i0} does not appear in (2.8). Combining (2.6) (for $j = 1, \dots, m$) with (2.8), we obtain an EM iteration as

$$p_j^{(t+1)} = \left(n - \sum_{i=1}^n \frac{g_i}{\sum_{k=1}^m f_{ik} p_k^{(t)}} \right)^{-1} \sum_{i=1}^n \left(\frac{f_{ij} - g_i}{\sum_{k=1}^m f_{ik} p_k^{(t)}} \right) p_j^{(t)}, \quad j = 1, \dots, m. \quad (2.9)$$

The iteration (2.2) corresponds to (2.9) with $g \equiv 0$.

Because (2.9) is derived in the EM framework, it inherits nearly all the desirable properties of (2.2). For example, each iteration of (2.9) increases the log-likelihood (2.4). Furthermore, because the convergence rate of EM is determined by the fraction of missing information, we know that (2.9) converges faster than (2.2) because it is based on a reduced set of latent variables. This is an example of efficient data augmentation (Meng and van Dyk 1997): we speed up EM by augmenting less. This strategy of improving (2.2) is called a squeezing strategy because the key equivalent formula (2.4) is obtained by subtracting (“squeezing out”) a nonnegative vector g from each component density.

A slight extension of the above discussion shows that the convergence rate of (2.9) is monotonic in the squeezing parameter g . Because of the restriction (2.5), the optimal g is therefore the upper bound $g_i = \min_j f_{ij}$, that is, we perform as much squeezing as allowed. This optimal g may be viewed as the overlap among all component densities. If this overlap is small, i.e., g is close to a vector of zeros, then (2.9) is not very different from (2.2). Hence we may expect significant speedup only if the overlap is large enough.

This squeezing strategy (as well as the strategy of Section 2.3) is inspired by the work of Fessler and Hero (1994) on efficient EM for Poisson image problems. Similar strategies also work for the Arimoto-Blahut algorithm for calculating the Shannon capacity of a discrete memoryless channel (see Yu 2010b).

2.3 Squeezing Strategy II

The squeezing strategy of Section 2.2 can be improved by further manipulation of the log-likelihood. Let us introduce another auxiliary vector $\beta = (\beta_1, \dots, \beta_m)$, and rewrite (2.4) as

$$l(\mathbf{p}) = n \log(2 + \beta_+) + \sum_{i=1}^n \log \frac{g_i - \sum_{j=1}^m \tilde{f}_{ij} \beta_j + \sum_{j=1}^m \tilde{f}_{ij} p_j + \sum_{j=1}^m \tilde{f}_{ij} \beta_j}{2 + \beta_+} \quad (2.10)$$

$$= n \log(2 + \beta_+) + \sum_{i=1}^n \log \frac{g_i - \sum_{j=1}^m \tilde{f}_{ij} \beta_j + \sum_{j=1}^m \tilde{f}_{ij} (p_j + \beta_j)}{2 + \beta_+}, \quad (2.11)$$

where $\tilde{f}_{ij} = f_{ij} - g_i$ as before, and $\beta_+ = \sum_{j=1}^m \beta_j$. It is required that

$$\beta_j \geq 0, \quad 1 \leq j \leq m; \quad g_i - \sum_{j=1}^m \tilde{f}_{ij} \beta_j \geq 0, \quad 1 \leq i \leq n. \quad (2.12)$$

When $m = 2$ and $g_i = \min_j f_{ij}$, (2.12) is equivalent to

$$0 \leq \beta_1 \leq \min_{i: f_{i1} > f_{i2}} \frac{f_{i2}}{f_{i1} - f_{i2}}, \quad 0 \leq \beta_2 \leq \min_{i: f_{i2} > f_{i1}} \frac{f_{i1}}{f_{i2} - f_{i1}}. \quad (2.13)$$

In general, however, it is not clear how to reduce (2.12) to an explicit range for β . Although we do address the choice of β in this section, further study is desired.

Iteration (2.9) can also be derived from (2.10), viewing it as a mixture log-likelihood with $2m + 1$ components. The density of the j th component ($0 \leq j \leq 2m$) evaluated at observation i

is

$$f_{ij}^\# = \begin{cases} g_i - \sum_{k=1}^m \tilde{f}_{ik} \beta_k, & j = 0, \\ \tilde{f}_{ij}, & 1 \leq j \leq m, \\ \tilde{f}_{i,j-m}, & m+1 \leq j \leq 2m. \end{cases}$$

The proportion of the j th component is $(2 + \beta_+)^{-1}$ if $j = 0$, $p_j/(2 + \beta_+)$ if $1 \leq j \leq m$, and $\beta_{j-m}/(2 + \beta_+)$ if $m+1 \leq j \leq 2m$. The factor $2 + \beta_+$ makes these proportions sum to one. Let $I_{ij}^\#$ denote the indicator of whether observation i is from component j . Similar to Section 2.2, if we treat $I_{ij}^\#$, $1 \leq i \leq n$, $0 \leq j \leq 2m$, as latent variables, then the resulting EM iteration is precisely (2.9).

On the other hand, we can derive an EM iteration based on (2.11), viewing it as a mixture log-likelihood with $m+1$ components, one of which has proportion $(2 + \beta_+)^{-1}$, and the others have proportions $(p_j + \beta_j)/(2 + \beta_+)$, $j = 1, \dots, m$. Equivalently, we treat

$$I_{i0}^\#, \quad 1 \leq i \leq n, \quad \text{and} \quad I_{ij}^\# + I_{i,j+m}^\#, \quad 1 \leq i \leq n, \quad 1 \leq j \leq m, \quad (2.14)$$

as latent indicators instead of the entire collection $I_{ij}^\#$, $1 \leq i \leq n$, $0 \leq j \leq 2m$. The E-step, as before, is to calculate the conditional expectations. We have

$$E \left(I_{ij}^\# + I_{i,j+m}^\# \mid y, \mathbf{p}^{(t)} \right) = \frac{\tilde{f}_{ij} (p_j^{(t)} + \beta_j)}{g_i + \sum_{k=1}^m \tilde{f}_{ik} p_k^{(t)}} \equiv K_{ij}.$$

The M-step seeks to maximize the function

$$\sum_{j=1}^m \sum_{i=1}^n K_{ij} \log(p_j + \beta_j). \quad (2.15)$$

By checking the Karush-Kuhn-Tucker conditions, it can be shown that (2.15) is maximized by choosing p_j as

$$p_j^{(t+1)} = \max \left\{ 0, \delta \sum_{i=1}^n K_{ij} - \beta_j \right\}, \quad 1 \leq j \leq m, \quad (2.16)$$

where δ is determined by the constraint $\sum_{j=1}^m p_j^{(t+1)} = 1$. Iteration (2.9) corresponds to (2.16) with $\beta_j \equiv 0$.

The new EM iteration (2.16) is more complicated than (2.9) or (2.2), but only slightly so. First, we observe that the right-hand side of (2.16) is a continuous and increasing function of δ ,

and hence a δ exists to ensure $\sum_{j=1}^m p_j^{(t+1)} = 1$, as long as $\sum_i K_{ij} > 0$ for some j . Moreover, such a δ can be found efficiently (in $O(m \log m)$ time) using a “waterfilling” algorithm (see Appendix A in Yu 2010b).

The convergence rate of (2.16) is no worse than that of (2.9), because the latent variables (2.14) are less informative than the entire collection $I_{ij}^\#$, $1 \leq i \leq n$, $0 \leq j \leq 2m$, which lead to (2.9). A slight extension of this argument reveals that, for fixed g , the convergence rate of (2.16) is monotonic in $\beta = (\beta_1, \dots, \beta_m)$. Such results resemble those of Yu (2010b) on improved Arimoto-Blahut algorithms for channel capacity calculations. In Yu (2010b), the convergence rate comparison results are derived by calculating the matrix rate and analyzing its eigenvalues. Because we work within the EM framework, however, the convergence rate comparison results here are obtained automatically once we specify the appropriate latent variables.

Combined with the convergence rate comparisons of Section 2.2, the above considerations suggest the following guideline for choosing the squeezing parameters g and β .

- Choose $g_i = \min_j f_{ij}$, which satisfies the upper bound in (2.5).
- Choose β_j to be as large as possible, subject to (2.12).

When $m = 2$, the condition (2.12) reduces to (2.13). Hence we recommend setting β at the upper bounds in (2.13).

It is helpful to write down an explicit formula for (2.16) with these optimal choices of g and β in the $m = 2$ case. Actually, for later convenience, we present a slightly more general iterative formula for maximizing ($m = 2$)

$$\tilde{l}(\mathbf{p}) = \sum_{i=1}^n \log(r_i + f_{i1}p_1 + f_{i2}p_2), \quad (2.17)$$

subject to $p_j \geq 0$, $j = 1, 2$, and $p_1 + p_2 = \beta_0$. Here r_i , $i = 1, \dots, n$, are nonnegative constants, and $\beta_0 > 0$ is fixed. Define $g_i = \min\{f_{i1}, f_{i2}\}$, and

$$\beta_1 = \min_{i: f_{i1} > f_{i2}} \frac{r_i + \beta_0 f_{i2}}{f_{i1} - f_{i2}}, \quad \beta_2 = \min_{i: f_{i2} > f_{i1}} \frac{r_i + \beta_0 f_{i1}}{f_{i2} - f_{i1}}.$$

Suppose the current parameter estimate is $\mathbf{p}^{(t)}$. We compute

$$S_j = \left(p_j^{(t)} + \beta_j\right) \sum_{i=1}^n \frac{f_{ij} - g_i}{r_i + f_{i1}p_1^{(t)} + f_{i2}p_2^{(t)}}, \quad j = 1, 2.$$

Then we update $\mathbf{p}^{(t)}$ as

$$p_j^{(t+1)} = \max\{0, \min\{\beta_0, (\beta_0 + \beta_1 + \beta_2)S_j/(S_1 + S_2) - \beta_j\}\}, \quad j = 1, 2, \quad (2.18)$$

which is a slight generalization of (2.16) for $m = 2$. The iteration (2.18) is uniquely defined if $S_1 + S_2 \neq 0$, for which it suffices to have $f_{i1} - f_{i2} \neq 0$. (Inspection shows that $p_j^{(t)} + \beta_j > 0$, $j = 1, 2$, as long as $\tilde{l}(\mathbf{p}^{(t)}) > -\infty$.) If $f_{i1} - f_{i2} \equiv 0$, then the M-step is not unique, but it is convenient to set $\mathbf{p}^{(t+1)} = \mathbf{p}^{(t)}$. Because of the max and min operations, (2.18) can potentially transfer all the mass from one component to the other in a single step. This will be especially useful in later sections after we introduce nearest neighbor exchanges.

3 Nearest Neighbor Exchanges and the Cocktail Algorithm

3.1 Nearest Neighbor Exchanges

The intuition that conventional EM tends to be slow when there exists heavy overlap between mixture components is used to our advantage in Section 2 for designing faster EM schemes via squeezing. However, the strategies so far begin by squeezing out a common vector g from *each* of the components. When there exist many components, it is conceivable that squeezing applied to all components may not be effective, even though a sub-collection may have severe overlap. Then it is worthwhile to apply some form of “local squeezing” to a sub-collection of components.

For example, consider a nonparametric mixture problem where the observations y_i are assumed to be drawn independently from a mixture of normals

$$y_i \sim \sum_{j=1}^m p_j \mathcal{N}(\mu_j, 1).$$

The variance is fixed for simplicity. We assume the mixing distribution puts mass p_j on $\mathcal{N}(\mu_j, 1)$, where μ_j is obtained by discretizing an interval, say $\mu_j = Uj/m$, and $U > 0$ denotes the largest of the normal means. As m increases, the overlap between adjacent densities $\mathcal{N}(\mu_j, 1)$ and $\mathcal{N}(\mu_{j+1}, 1)$ increases, and conventional EM slows down. The global squeezing strategies may not be effective, however, because the overlap between the left-most density $\mathcal{N}(U/m, 1)$ and the right-most density $\mathcal{N}(U, 1)$ can still be small.

A natural remedy, therefore, is to exchange the mass between each pair of nearby components in turn, holding the other components fixed. This is similar to (but somewhat simpler than) the paired and rotated EM of Pilla and Lindsay (2001). In general, given the current parameter estimate $\mathbf{p}^{(t)}$, let $j_1 < \dots < j_{q+1}$ be the elements of $\{j : p_j^{(t)} > 0\}$ where $q + 1$ is the number of support points of $\mathbf{p}^{(t)}$. We perform mass exchanges between j_k and j_{k+1} for $k = 1, \dots, q$ in turn, i.e.,

$$\mathbf{p}^{(t+k/q)} = VE(j_k, j_{k+1}, \mathbf{p}^{(t+(k-1)/q)}), \quad k = 1, \dots, q. \quad (3.1)$$

We use $\tilde{\mathbf{p}} = VE(u, v, \mathbf{p})$, $u \neq v$, to denote an update of the form

$$\tilde{p}_j = \begin{cases} p_j, & j \notin \{u, v\}, \\ p_j + \delta, & j = u, \\ p_j - \delta, & j = v, \end{cases}$$

where $\delta \in [-p_u, p_v]$ is chosen so that $l(\tilde{\mathbf{p}}) \geq l(\mathbf{p})$. To choose the step-length δ , we naturally use (2.18). The iteration (2.18) is applicable because, when other components are held fixed, the log-likelihood for p_u, p_v is exactly in the form of (2.17). Because (2.18) is an EM iteration, the log-likelihood is automatically monotonic. Moreover, (2.18) is easy to implement, and very amenable to theoretical analysis. These all add to the appeal of (2.18) when compared with standard tools such as Newton's method. We refer to the composite mapping $\mathbf{p}^{(t)} \rightarrow \mathbf{p}^{(t+1)}$ given by (3.1) as the set of nearest neighbor exchanges (NNEs).

We have found that by adding nearest neighbor exchanges based on (2.18) to conventional EM can lead to considerably improved speed. There is one caveat, however. Conventional EM is usually started at the interior of the parameter space, i.e., $p_j^{(0)} > 0$ for all j , because once a component receives zero mass, it does so in all subsequent iterations. By adding nearest neighbor exchanges, certain p_j may be set to zero. While this has the desirable effect of eliminating bad support points, it may accidentally eliminate a good one, and yield a suboptimal solution. The problem is easily remedied, however, by adding in the following step, known as the vertex direction method (VDM; Fedorov 1972). Given the current parameter estimate \mathbf{p} , we first calculate the derivatives

$$d_j = \frac{\partial l(\mathbf{p})}{\partial p_j} = \sum_{i=1}^n \frac{f_{ij}}{\eta_i},$$

where $\eta_i = \sum_{j=1}^m f_{ij} p_j$. Let $j^\#$ denote any index such that d_j , $j = 1, \dots, m$, is maximized. Then we update \mathbf{p} to $\tilde{\mathbf{p}}$ with

$$\tilde{p}_j = \begin{cases} (1 - \delta)p_j, & j \neq j^\#, \\ (1 - \delta)p_j + \delta, & j = j^\#, \end{cases} \quad (3.2)$$

where $\delta \in [0, 1]$ is chosen so that $l(\tilde{\mathbf{p}}) \geq l(\mathbf{p})$. We use iteration (2.18) for choosing δ because $l(\tilde{\mathbf{p}})$ as a function of $(\delta, 1 - \delta)$ is again in the form of (2.17).

Let us denote the mapping (3.2) with δ chosen by (2.18) as $\tilde{\mathbf{p}} = VDM(j^\#, \mathbf{p})$. In Section 3.2 we show that VDM based on (2.18), when added to conventional EM and nearest neighbor exchange iterations, results in a globally convergent algorithm. That is, starting from any $\mathbf{p}^{(0)} \in \Theta$, all limit points of the resulting algorithm are global maxima of the log-likelihood function on Θ . The proof actually shows that by adding VDM to any monotonic algorithm we obtain a globally convergent algorithm.

3.2 The Cocktail Algorithm

We summarize a “cocktail algorithm” based on VDM, nearest neighbor exchanges, and EM steps. A convergence proof is then provided. Empirical evaluation of such a strategy is presented in Section 3.3. Yu (2009) applies this strategy to the D-optimal design problem, and reports dramatic improvement in speed. We show similar performance for the mixture problem.

Cocktail Algorithm

- 1 At iteration t , first perform a VDM step (3.2) where δ is chosen using (2.18).
- 2 Then use the output of VDM, say $\tilde{\mathbf{p}}$, as input for the nearest neighbor exchanges, i.e., (3.1), again based on (2.18).
- 3 Finally, update the output of (3.1) using (2.2), i.e., conventional EM, to obtain the next iterate $\mathbf{p}^{(t+1)}$.

Note that this is only one of the potential algorithms based on reducing the overlap between component densities. There is much room for further exploration. One could consider, for example, an algorithm that uses only Steps 1 and 2 above at each iteration. That is, a VDM step is combined

with nearest neighbor exchanges. We call this algorithm NNE+. We design the cocktail algorithm in the hope that the nearest neighbor steps and conventional EM can complement each other, since NNE+ focuses on purely local modifications, whereas EM focuses on purely global ones. As we shall illustrate in empirical examples, the performance of NNE+ or conventional EM (each by itself) can be poor, but the cocktail algorithm is very fast.

Because the cocktail algorithm consists of many EM sub-steps, the log-likelihood is guaranteed to increase at each iteration. Further analysis yields the following convergence theorem.

Theorem 1. *The cocktail algorithm is globally convergent. That is, if $\mathbf{p}^{(t)}$ is a sequence generated by the cocktail algorithm starting from any $\mathbf{p}^{(0)} \in \Theta$, then all limit points of $\mathbf{p}^{(t)}$ are global maxima of $l(\mathbf{p})$ on $\mathbf{p} \in \Theta$.*

Proof. The proof is similar to that of Theorem 1 in Yu (2009). Let $\tilde{\mathbf{p}}^{(t)}$ denote the output of the VDM step at iteration t . By monotonicity,

$$l(\mathbf{p}^{(t)}) \leq l(\tilde{\mathbf{p}}^{(t)}) \leq l(\mathbf{p}^{(t+1)}).$$

Hence the two sequences $l(\mathbf{p}^{(t)})$ and $l(\tilde{\mathbf{p}}^{(t)})$ tend to the same (finite) limit. Let \mathbf{p}^* be a limit point of $\mathbf{p}^{(t)}$, and let $\mathbf{p}^{(t_j)}$ be a subsequence converging to \mathbf{p}^* . Without loss of generality, we may assume that the VDM steps $\mathbf{p}^{(t_j)} \rightarrow \tilde{\mathbf{p}}^{(t_j)}$ are all performed on the same index $k = j^\#$ as in (3.2), since at least one of the m indices will appear infinitely often. If $f_{ik} \equiv \sum_j f_{ij} p_j$ for some $\mathbf{p} = \mathbf{p}^{(t_j)}$ or $\mathbf{p} = \mathbf{p}^*$, then we can show directly that $\partial l(\mathbf{p}) / \partial p_k = n$. By the choice of k , we have $\partial l(\mathbf{p}) / \partial p_j \leq n$ for all $1 \leq j \leq m$, and hence \mathbf{p} is already a global maximum by the general equivalence theorem (Lindsay 1983). Assume $f_{ik} \not\equiv \sum_j f_{ij} p_j$ for all $\mathbf{p} = \mathbf{p}^{(t_j)}$ and $\mathbf{p} = \mathbf{p}^*$. Then inspection of (3.2) and (2.18) reveals that all VDM steps $\mathbf{p}^{(t_j)} \rightarrow \tilde{\mathbf{p}}^{(t_j)}$ are uniquely defined. Moreover, when k is considered fixed, the VDM mapping is continuous at \mathbf{p}^* . Hence $\tilde{\mathbf{p}}^{(t_j)}$ converges to $VDM(k, \mathbf{p}^*) = \tilde{\mathbf{p}}$, say, and $l(\tilde{\mathbf{p}}) = l(\mathbf{p}^*)$ as a result. Since this VDM step $\mathbf{p}^* \rightarrow \tilde{\mathbf{p}}$ is derived in the EM framework, $\tilde{\mathbf{p}}$ being uniquely defined means that it is the unique maximizer at the M-step. If $\tilde{\mathbf{p}} \neq \mathbf{p}^*$, then the expected complete-data log-likelihood increases strictly, and so does the observed log-likelihood, which contradicts $l(\tilde{\mathbf{p}}) = l(\mathbf{p}^*)$. It follows that $\tilde{\mathbf{p}} = \mathbf{p}^*$, i.e., \mathbf{p}^* is a fixed point of the VDM mapping. Inspection of (3.2) and (2.18), however, shows that this fixed point must satisfy $\partial l(\mathbf{p}^*) / \partial p_k \leq n$, which implies that \mathbf{p}^* is a global maximum. \square

3.3 Numerical Illustration

This section gives a numerical illustration of the effectiveness of the cocktail algorithm. The cocktail algorithm is compared with conventional EM, i.e., iteration (2.2), the algorithm NNE+ mentioned in Section 3.2, and the vertex exchange method (VEM) of Böhning (1985). To describe VEM, suppose the parameter estimate at iteration t is $\mathbf{p}^{(t)}$. Define $d_j = \partial l(\mathbf{p}^{(t)}) / \partial p_j$, and let $j^\#$ and $j_\#$ be indices between 1 and m such that

$$d_{j_\#} = \min_{j: p_j^{(t)} > 0} d_j, \quad d_{j^\#} = \max_{1 \leq j \leq m} d_j.$$

VEM sets $\mathbf{p}^{(t+1)}$ as

$$\mathbf{p}^{(t+1)} = VE(j^\#, j_\#, \mathbf{p}^{(t)}).$$

That is, we exchange the mass between the indices $j^\#$ and $j_\#$ so as to increase the log-likelihood. We again employ (2.18) to choose the step-length.

We run each of (conventional) EM, NNE+, VEM, and the cocktail algorithm until convergence and record the number of iterations and computing time. An iteration of NNE+ consists of one iteration of VDM and the set of nearest neighbor exchanges. An iteration of the cocktail algorithm consists of one iteration each of NNE+ and conventional EM. We shall concentrate on the computing time as a more objective measure of performance. All calculations are performed on the same Sun Solaris 10 machine, and the computing time is recorded using the R function `system.time()`. The program is written in C and is available, together with the R interface, upon request from the author.

Each algorithm is started from the same uniform probability vector, i.e., $p_j^{(0)} = 1/m$, $j = 1, \dots, m$. We use the common convergence criterion

$$\max_{1 \leq j \leq m} d_j - n \leq \epsilon, \tag{3.3}$$

where $d_j = \partial l(\mathbf{p}) / \partial p_j$. The theoretical basis for (3.3) is that for any \mathbf{p} that satisfies this criterion we have

$$l(\hat{\mathbf{p}}) - l(\mathbf{p}) \leq \epsilon,$$

where $\hat{\mathbf{p}}$ is the MLE (Lindsay 1983, Böhning et al. 1996). We choose $\epsilon = 10^{-6}$ in our experiments.

EM, NNE+, VEM, and the cocktail algorithm are tested on data taken from Roeder (1990) concerning the velocities of 82 galaxies. Following Pilla and Lindsay (2001), we fit a normal finite

Table 1: Iteration count and computing time (in seconds) until convergence for four algorithms on the galaxy data.

| Iteration count | | | | Computing time | | | |
|-----------------|------|-----|----------|----------------|------|------|----------|
| EM | NNE+ | VEM | Cocktail | EM | NNE+ | VEM | Cocktail |
| 21777 | 74 | 974 | 36 | 87 | 0.02 | 0.13 | 0.02 |

mixture model to these data. The means of the normal components lie on a grid of 64 equally-spaced points from 10.0 to 33.94, and the common standard deviation is $\sigma = 0.95$. The algorithms deliver the same MLE as reported by Pilla and Lindsay (2001), and their performance is recorded in Table 1.

Clearly the nearest neighbor exchanges are very effective, since both NNE+ and the cocktail algorithm improve conventional EM dramatically, reducing its computing time by orders of magnitude. Adding conventional EM to NNE+ appears to have increased the computing time per iteration, but decreased the number of iterations, so that the cocktail algorithm and NNE+ have similar overall computing time. We remark that the cocktail algorithm is also appealing because it is easy to implement and requires virtually no tuning.

4 Efficient EM for Computing the NPMLE for Censored Data

We show that the EM strategies designed in Sections 2 and 3, in particular the cocktail algorithm of Section 3, are also effective for the NPMLE problem for censored data. Section 4.1 briefly reviews how this problem can be viewed as a problem of mixture proportions. Section 4.2 highlights fast implementations of our algorithms. Section 4.3 contains numerical illustrations using simulated data.

4.1 NPMLE for Censored Data

Assume that failure time data collected from n units are independent and identically distributed according to a distribution function F , except that they are subject to censoring. Following Gentleman and Geyer (1994), assume there is an inspection time process Q which is independent of the failure times, and suppose each unit is subject to inspections governed by Q independently. The observed data then consist of n observation intervals $(l_i, r_i]$, where l_i is the last inspection time prior to failure and r_i is the first inspection time after failure for subject i . Right censoring may be represented by $r_i = \infty$, and left censoring by $l_i = 0$; the exact failure time is observed when l_i coincides with r_i , i.e., when the individual is subject to continuous inspections.

Doubly censored data arise when units are subject to both right and left censoring. The inspection time process is given by a pair of random variables (L, U) , $L < U$. If the failure time T satisfies $L < T \leq U$, then T is observed; if $T \leq L$, then the observation is left censored at L ; if $T > U$, then the observation is right censored at U . For (case 2) interval censored data, the failure time T is not observed, but only known to fall within a random interval.

In either case, let us order the distinct observation times (i.e., end points of the intervals $(l_i, r_i]$) as $0 = z_0 < z_1 < \dots < z_{m-1} < z_m = \infty$. Denote

$$\mathbf{p} = (p_1, \dots, p_m), \quad p_j = F(z_j) - F(z_{j-1}),$$

and define the $n \times m$ matrix (f_{ij}) by

$$f_{ij} = \begin{cases} 1, & z_j \in (l_i, r_i], \\ 0, & \text{otherwise.} \end{cases}$$

For notational convenience we assume that observation intervals are open at the left, and closed at the right, end points, though extension to the general case is trivial; if the i th observation is exact, for example, we simply set $f_{ij} = 1$ if $z_j = l_i (= r_i)$ and $f_{ij} = 0$ otherwise. Seeking the maximizer \hat{F} that jumps only at observed time points, we note that the log-likelihood function is exactly (2.1). Thus all the EM-type algorithms in Sections 2 and 3 can be applied to compute the NPMLE, defined as the $\hat{\mathbf{p}}$ that maximizes (2.1).

4.2 Efficient Implementations

Straightforward implementation of conventional EM, NNE+, VEM, or the cocktail algorithm, as described in Sections 2 and 3, each requires $O(nm)$ time per iteration. For univariate censored data, however, one can take advantage of the special structure of the matrix (f_{ij}) to derive $O(n)$ implementations. This is a substantial reduction since m is often of the same order of magnitude as n . It is especially relevant in situations such as bootstrap resampling, when repeated use of the algorithms is needed.

Noted only briefly by Jongbloed (1998), who does not give the technical details concerning computational complexity, this possibility of fast EM implementation has remained largely unnoticed. Zhang and Jamshidian (2004) present fast implementations for doubly censored data, but not (case 2) interval censored data. In Appendix A, we show that all four algorithms admit efficient implementations for univariate interval censoring in general.

4.3 Evaluation of Algorithms on Doubly Censored Data

In this section we evaluate the effectiveness of EM, NNE+, VEM, and the cocktail algorithm for computing the NPMLE for doubly censored data. Simulations are performed under conditions similar to those of Wellner and Zhan (1997), and Zhang and Jamshidian (2004). Wellner and Zhan (1997) propose an effective algorithm which combines EM and iterative convex minorant (ICM; Jongbloed 1998) iterations. We have decided to focus on evaluating the cocktail algorithm relative to EM, NNE+ and VEM because they are easy to describe and easy to implement. A full evaluation, including the case of bivariate censoring, is work in progress.

Our simulation setting is as follows. The failure time T_i for unit i is generated as an independent exponential random variable with mean 1, and the censoring variables L_i, U_i are generated as the q_1 th and q_2 th order statistics of 20 independent uniform(0, 1) random variables. The resulting observation is

$$\begin{aligned} T_i, & \text{ if } L_i < T_i \leq U_i; \\ (0, L_i], & \text{ if } T_i \leq L_i; \\ (U_i, \infty], & \text{ if } T_i > U_i. \end{aligned}$$

By adjusting q_1 and q_2 we obtain varying degrees of censoring.

As in Section 3.3, all algorithms are started at the uniform probability vector, and the com-

Table 2: Means and standard deviations of the number of iterations and computing time (in seconds) until convergence for four algorithms for doubly censored data. Input data are generated with $q_1 = 3$ and $q_2 = 18$.

| | Iteration count | | | | Computing time | | | |
|------------|-----------------|---------|-------|----------|----------------|------|------|----------|
| | EM | NNE+ | VEM | Cocktail | EM | NNE+ | VEM | Cocktail |
| $n = 1000$ | | | | | | | | |
| mean | 5076 | 15456 | 5714 | 46.2 | 3.22 | 20.5 | 0.80 | 0.07 |
| s.d. | 194 | 488 | 206 | 7.7 | 0.20 | 1.2 | 0.04 | 0.01 |
| $n = 2000$ | | | | | | | | |
| mean | 9920 | 44044 | 12008 | 67.3 | 17.0 | 157 | 3.77 | 0.28 |
| s.d. | 892 | 1645 | 269 | 7.3 | 2.2 | 9 | 0.15 | 0.08 |
| $n = 4000$ | | | | | | | | |
| mean | 20347 | 100000+ | 24924 | 93.3 | 85.6 | 740+ | 19.1 | 0.77 |
| s.d. | 1056 | | 369 | 7.9 | 5.2 | | 0.7 | 0.06 |

mon convergence criterion is (3.3) with $\epsilon = 10^{-6}$. Our limited experience suggests that VEM may benefit from a starting value with fewer support points, but the cocktail algorithm is relatively insensitive to the initial number of support points. Again, an iteration of the cocktail algorithm consists of an iteration of VDM, the set of nearest neighbor exchanges, and an iteration of conventional EM.

Based on 10 replications, Tables 2 and 3 display the means and standard deviations of the number of iterations and computing time (in seconds) until convergence for EM, NNE+, VEM and the cocktail algorithm. The input data are generated with $q_1 = 3$ and $q_2 = 18$ (moderate censoring) for Table 2, and with $q_1 = 8$ and $q_2 = 12$ (heavy censoring) for Table 3.

In either situation we see that the cocktail algorithm is a dramatic improvement; it reduces the computing time of (conventional) EM or NNE+ by large factors, the reduction being more significant as n , the number of units, becomes larger. The improvement is especially remarkable

Table 3: Means and standard deviations of the iteration count and computing time (in seconds) until convergence for four algorithms for doubly censored data. Input data are generated with $q_1 = 8$ and $q_2 = 12$.

| | Iteration count | | | | Computing time | | | |
|------------|-----------------|-------|------|----------|----------------|------|------|----------|
| | EM | NNE+ | VEM | Cocktail | EM | NNE+ | VEM | Cocktail |
| $n = 1000$ | | | | | | | | |
| mean | 5793 | 2768 | 2170 | 65.3 | 7.77 | 1.70 | 0.38 | 0.05 |
| s.d. | 877 | 421 | 170 | 8.3 | 1.76 | 0.34 | 0.11 | 0.01 |
| $n = 2000$ | | | | | | | | |
| mean | 11034 | 8669 | 4411 | 103 | 38.0 | 13.0 | 1.56 | 0.20 |
| s.d. | 2022 | 491 | 192 | 8 | 8.3 | 1.0 | 0.08 | 0.01 |
| $n = 4000$ | | | | | | | | |
| mean | 20397 | 27247 | 9176 | 145 | 163 | 89.8 | 8.76 | 0.61 |
| s.d. | 5481 | 1957 | 336 | 17 | 51 | 8.8 | 0.57 | 0.08 |

because the cocktail algorithm is a direct combination of EM and NNE+, each of which is very slow. NNE+ performs much worse for moderately censored data than for heavily censored data. In the case of $q_1 = 3$, $q_2 = 18$ and $n = 4000$, each of the 10 runs of NNE+ takes more than 100000 iterations. As expected, each algorithm takes more iterations as n increases. Somewhat unexpectedly, for the same n , EM has similar iteration counts for moderately versus heavily censored data. Yet the computing time of EM in the heavily censored case is significantly higher. Another peculiarity is that, for the same n , VEM takes fewer iterations and less time for heavily censored data than for moderately censored data. But the main feature in Tables 2 and 3 is the clear superiority of the cocktail algorithm in this example.

5 Discussion

We have shown how to use efficient data augmentation to design fast EM-type algorithms for maximizing a mixture log-likelihood with known component densities. Squeezing strategies are presented that take advantage of the overlap between components. A cocktail algorithm that combines conventional EM with a nearest neighbor exchange strategy is found to perform very well for computing the NPMLE for censored data, which is the intended application area of this work.

The nearest neighbor exchange strategy works well with conventional EM when there is a natural ordering of the mixture components, as in the case of univariate censored data. It would be interesting to extend such algorithms to bivariate censoring (Betensky and Finkelstein 1999), where we observe a pair of possibly censored random variables for each unit. Bivariate censoring presents many inferential and computational challenges. Work on extending the effective nearest neighbor strategy is in progress, with encouraging preliminary results. Extensions that accommodate truncation in addition to censoring, or that facilitate semi-parametric estimation, would also be desirable.

It would be interesting to extend the squeezing strategies of Section 2 to mixture problems with unknown parameters in the component densities. One approach is to again adopt AECM (Meng and van Dyk 1997), and perform two types of EM-based maximization steps, one for the mixture proportions given the other parameters, and one for the other parameters given the mixture proportions. The squeezing strategies can be used at the former maximization step. It

would be worthwhile to investigate the potential gain of using such strategies.

Acknowledgment

The author would like to thank Cliff Anderson-Bergman for helpful discussions on the squeezing strategies of Section 2.

Appendix A: Fast Implementations for the NPMLE Problem with Censored Data

The conventional EM mapping, (2.2), contains a summation of n terms for each $j = 1, \dots, m$. However, one can take advantage of the special structure of the matrix (f_{ij}) for fast computation. This is a zero-one matrix whose non-zero entries are consecutive in each row. Equivalently, if we let

$$a(i) = \min\{j : z_j \in (l_i, r_i]\}, \quad b(i) = \max\{j : z_j \in (l_i, r_i]\}, \quad i = 1, \dots, n, \quad (.1)$$

then $f_{ij} = 1$ if $a(i) \leq j \leq b(i)$ and $f_{ij} = 0$ otherwise. For convenience, we drop the superscripts in (2.2) and focus on how to efficiently compute

$$p_j^{new} = \frac{1}{n} \sum_{i=1}^n \left(\frac{f_{ij}}{\sum_{k=1}^m f_{ik} p_k} \right) p_j, \quad j = 1, \dots, m,$$

for any $\mathbf{p} = (p_1, \dots, p_m) \in \Theta$. Note that computing $\eta_i = \sum_{k=1}^m f_{ik} p_k$, $i = 1, \dots, n$, can be done in $O(m+n)$ time (or equivalently $O(n)$ time because $m \leq 2n+1$) using Algorithm 1.

Algorithm 1

Step 1 Calculate the cumulative sums $s_j = \sum_{k=1}^j p_k$, $j = 1, \dots, m$. Set $s_0 = 0$.

Step 2 Set $\eta_i = s_{b(i)} - s_{a(i)-1}$, $i = 1, \dots, n$.

If we can calculate

$$d_j = \sum_{i=1}^n f_{ij} / \eta_i = \sum_{i: a(i) \leq j \leq b(i)} 1 / \eta_i, \quad (.2)$$

then $p_j^{new} = d_j p_j / n$. To calculate d_j efficiently, we rely on the following algorithm.

Algorithm 2

Step 1 Initialize $d_j = 0$, $j = 1, \dots, m$.

Step 2 For $i = 1, \dots, n$, add $1/\eta_i$ to $d_{a(i)}$, and, if $b(i) + 1 \leq m$, subtract $1/\eta_i$ from $d_{b(i)+1}$.

Step 3 Replace (d_1, \dots, d_m) by its cumulative sum, i.e., for $j = 2, \dots, m$, add d_{j-1} to d_j .

Obviously, Algorithm 2 costs $O(n)$ time. We have

Proposition 1. *Algorithm 2 is valid, i.e., its output agrees with (.2).*

Proof. This can be shown by induction. First, in the output of Algorithm 2, $d_1 = \sum_{i: a(i)=1} 1/\eta_i$, which agrees with (.2). Assume Algorithm 2 gives the correct answer for d_{j-1} , $j > 1$. Then Algorithm 2 computes d_j using

$$d_j = d_{j-1} + \sum_{i: a(i)=j} 1/\eta_i - \sum_{i: b(i)+1=j} 1/\eta_i,$$

which again agrees with (.2) if we consider the difference $d_j - d_{j-1}$. By the induction principle, all d_j , $j = 1, \dots, m$, are correctly computed. \square

Algorithms 1 & 2 clearly give an $O(n)$ implementation of a conventional EM iteration. Because η_i , $i = 1, \dots, n$ and d_j , $j = 1, \dots, m$, are also the key quantities for VDM and VEM, the same efficient implementation applies to VDM and VEM. Specifically, the underlying iteration (2.18) is done in $O(n)$ time by keeping track of η_i .

For NNE+ and the cocktail algorithm, we notice that each sub-step of nearest neighbor exchange in (3.1) affects only a limited number of terms in the log-likelihood. Specifically, to implement sub-step k , only observation intervals $(l_i, r_i]$ such that either $a(i) \leq j_k \leq b(i)$ and $j_{k+1} > b(i)$, or $j_k < a(i)$ and $a(i) \leq j_{k+1} \leq b(i)$, need be considered. Define the set

$$V_k = \{i : a(i) \leq j_k \leq b(i), j_{k+1} > b(i)\} \cup \{i : j_k < a(i), a(i) \leq j_{k+1} \leq b(i)\}.$$

The time cost of sub-step k is proportional to $|V_k|$, the number of entries in V_k . However, because each i belongs to at most two of V_k , $k = 1, \dots, q$, the total number of entries satisfy $\sum_{k=1}^q |V_k| \leq 2n$. Hence, with a bit of bookkeeping, the entire set of nearest neighbor exchanges can be implemented in $O(n)$ time per iteration.

Remark. Algorithms 1 and 2 take $a(i)$, $b(i)$, $i = 1, \dots, n$, given by (.1), as input. Setting these up requires sorting the end points of the observation intervals $(l_i, r_i]$, which costs $O(n \log n)$ time. Although slightly higher than the $O(n)$ per-iteration cost, this is typically a small fraction of the total computing time because of the required number of iterations. Setting up the full matrix (f_{ij}) , on the other hand, costs $O(mn)$ time, which is $O(n^2)$ in the worst case.

References

- [1] Aragon, J. and Eberly, D. (1992). On convergence of convex minorant algorithms for distribution estimation with interval-censored data. *J. Computational and Graphical Statistics* **1**, 129-140.
- [2] Arimoto, S. (1972). An algorithm for computing the capacity of arbitrary discrete memoryless channels, *IEEE Trans. Inform. Theory* **18**, 14–20.
- [3] Betensky, R. A. and Finkelstein, D. M. (1999). A non-parametric maximum likelihood estimator for bivariate interval censored data. *Statistics in Medicine* **18**, 3089-3100.
- [4] Blahut, R.E. (1972). Computation of channel capacity and rate-distortion functions, *IEEE Trans. Inform. Theory* **18**, 460–473.
- [5] Böhning, D. (1985). Numerical estimation of a probability measure. *J. Statist. Plan. Inf.* **11**, 57–69.
- [6] Böhning, D., Schlattmann, P. and Dietz E. (1996). Interval censored data: a note on the nonparametric maximum likelihood estimator of the distribution function. *Biometrika* **83**, 462–466.
- [7] Csiszár, I. and Tusnády, G. (1984). Information geometry and alternating minimization procedures. *Statistics & Decisions* Supplement Issue **1**, 205–237.
- [8] Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977). Maximum likelihood estimation from incomplete data via the EM algorithm (with discussion). *J. Roy. Statist. Soc. B* **39**, 1–38.
- [9] Fedorov, V.V. (1972). *Theory of Optimal Experiments*, Academic Press, New York.

- [10] Fessler, J.A. and Hero, A.O. (1994). Space-alternating generalized expectation-maximisation algorithm, *IEEE Trans. Signal Processing* **42**, 2664–2677.
- [11] Gentleman, R. and Geyer, C. J. (1994). Maximum likelihood for interval censored data: consistency and computation. *Biometrika* **81**, 618–623.
- [12] Groeneboom, P. and Wellner, J. (1992). *Information Bounds and Nonparametric Maximum Likelihood Estimation*. Basel: Birkhäuser.
- [13] Jongbloed, G. (1998). The iterative convex minorant algorithm for nonparametric estimation. *J. Computational and Graphical Statistics* **7**, 310–321.
- [14] Lindsay, B. G. (1983). The geometry of mixture likelihoods: a general theory. *Ann. Statist.* **11**, 86–94.
- [15] Liu, C. H. and Rubin, D. B. (1994). The ECME algorithm: a simple extension of EM and ECM with faster monotone convergence. *Biometrika* **81**, 633–648.
- [16] Liu, C. H., Rubin, D. B. and Wu, Y. N. (1998). Parameter expansion to accelerate EM—the PX-EM algorithm. *Biometrika* **85**, 755–770.
- [17] Meng, X.-L. and Rubin, D. B. (1993). Maximum likelihood estimation via the ECM algorithm: a general framework. *Biometrika* **80**, 267–278.
- [18] Meng, X.-L. and van Dyk, D.A. (1997). The EM algorithm – an old folk-song sung to a fast new tune (with discussion), *J. Roy. Statist. Soc. B* **59**, 511–567.
- [19] Meng, X.-L. and van Dyk, D.A. (1998). Fast EM-type implementations for mixed effects models, *J. Roy. Statist. Soc. B* **60**, 559–578.
- [20] Pilla, R.S. and Lindsay, B.G. (2001). Alternative EM methods for nonparametric finite mixture models, *Biometrika* **88**, 535–550.
- [21] Roeder, K. (1990). Density estimation with confidence sets exemplified by superclusters and voids in the galaxies. *J. Am. Statist. Assoc.* **85**, 617–624.
- [22] Silvey, S.D., Titterton, D.M. and Torsney, B. (1978). An algorithm for optimal designs on a finite design space, *Commun. Stat. Theory Methods* **14**, 1379–1389.

- [23] Turnbull, B. W. (1976). The empirical distribution function with arbitrarily grouped, censored and truncated data. *J. Roy. Statist. Soc. B* **38**, 290–295.
- [24] Vardi, Y., Shepp, L.A. and Kaufman, L. (1985). A statistical model for positron emission tomography (with discussion), *J. Amer. Statist. Assoc.* **80**, 8–37.
- [25] Wang, Y. (2008). Dimension-reduced nonparametric maximum likelihood computation for interval-censored data, *Computational Statistics & Data Analysis* **52**, 2388–2402.
- [26] Wellner, J. A. and Zhan, Y. (1997). A hybrid algorithm for computation of the nonparametric maximum likelihood estimator from censored data. *J. Amer. Statist. Assoc.* **92**, 945-959.
- [27] Wu, C. F. J. (1983). On the convergence properties of the EM algorithm. *Annals of Statistics* **11**, 95–103.
- [28] Yu, Y. (2009). D-optimal designs via a cocktail algorithm. *Preprint*, arXiv:0911.0108
- [29] Yu, Y. (2010a). Monotonic convergence of a general algorithm for computing optimal designs, To appear, *Annals of Statistics*. arXiv:0905.2646
- [30] Yu, Y. (2010b). Squeezing the Arimoto-Blahut algorithm for faster convergence. *Preprint*, arXiv:0906.3849
- [31] Zhang, Y. and Jamshidian, M. (2004). On algorithms for NPMLE of the failure function with censored data. *J. Computational and Graphical Statistics* **13**, 123–140.