

Stochastic Approximation and Newton's Estimate of a Mixing Distribution

Ryan Martin and Jayanta K. Ghosh

Abstract. Many statistical problems involve mixture models and the need for computationally efficient methods to estimate the mixing distribution has increased dramatically in recent years. Newton [*Sankhyā Ser. A* **64** (2002) 306–322] proposed a fast recursive algorithm for estimating the mixing distribution, which we study as a special case of stochastic approximation (SA). We begin with a review of SA, some recent statistical applications, and the theory necessary for analysis of a SA algorithm, which includes Lyapunov functions and ODE stability theory. Then standard SA results are used to prove consistency of Newton's estimate in the case of a finite mixture. We also propose a modification of Newton's algorithm that allows for estimation of an additional unknown parameter in the model, and prove its consistency.

Key words and phrases: Stochastic approximation, empirical Bayes, mixture models, Lyapunov functions.

1. INTRODUCTION

The aim of the present paper is to review the subject of stochastic approximation (SA), along the way highlighting some recent statistical applications, and to explore its relationship with a recent algorithm [26–28] for estimating a mixing distribution.

SA was introduced in [31] as an algorithmic method for finding the root of a function h when only noisy observations on h are available. It has since developed into an important area of systems control and optimization, with numerous applications in statistics. In Section 2 we give a brief introduction to the SA algorithm and review three recent and innovative statistical applications. The first two [6, 16] strengthen the EM and Metropolis algorithms, respectively, and the third is a versatile Monte Carlo integration method, called Stochastic Approximation Monte Carlo (SAMC) [21],

which can be applied in a variety of statistical problems. We demonstrate that combining SAMC with the *energy–temperature duality* [18] provides a method for estimating the normalizing constant of a density. We then state a theorem providing sufficient conditions for almost sure convergence of a SA algorithm, which is used in Section 3 to study the convergence properties of a mixing distribution estimate. For this purpose, the necessary stability theory for ordinary differential equations (ODEs) is developed.

Many statistical problems involve modeling with latent, or unobserved, random variables, for example, cluster analysis [24] and multiple testing or estimation with high-dimensional data [1, 8, 9, 34, 36]. The distribution of the manifest, or observed, random variables then becomes a mixture of the form

$$(1.1) \quad \Pi_f(x) = \int_{\Theta} p(x|\theta) f(\theta) d\mu(\theta),$$

where $\theta \in \Theta$ is the latent variable or parameter, and f is an unknown mixing density with respect to the measure μ on Θ . Estimation of f plays a fundamental role in many inference problems, such as an empirical Bayes approach to multiple testing.

For the deconvolution problem, when $p(x|\theta)$ in (1.1) is of the form $p(x - \theta)$, asymptotic results for estimates of f , including optimal rates of convergence, are known [10]. A nonparametric Bayes approach to

Ryan Martin is Graduate Student, Department of Statistics, Purdue University, 250 North University Street, West Lafayette, Indiana 47907, USA (e-mail: martinrg@stat.purdue.edu). Jayanta K. Ghosh is Professor, Department of Statistics, Purdue University, 150 N. University Street, West Lafayette, 47907, USA and Professor Emeritus, Division of Theoretical Statistics and Mathematics, Indian Statistical Institute, 203 B T Road Kolkata, India 700108 (e-mail: ghosh@stat.purdue.edu).

Gaussian deconvolution is discussed in [13]. For estimating Π_f , a Bayesian might assume an a priori distribution on f , inducing a prior on Π_f via the map $f \mapsto \Pi_f$. Consistency of the resulting estimate of Π_f is considered in [3, 11, 12].

In Section 3, we describe a recursive algorithm of Newton et al. [26–28] for estimating the mixing density f . This estimate is significantly faster to compute than the popular nonparametric Bayes estimate based on a Dirichlet process prior. In fact, the original motivation [27] for the algorithm was to approximate the computationally expensive Bayes estimate. The relative efficiency of the recursive algorithm compared to MCMC methods used to compute the Bayes estimate, coupled with the similarity of the resulting estimates, led Quintana and Newton [29] to suggest the former be used for *Bayesian exploratory data analysis*.

While Newton's algorithm performs well in examples and simulations (see [14, 26–29, 38] and Section 3.3), very little is known about its large-sample properties. A rather difficult proof of consistency, based on an approximate martingale representation of the Kullback–Leibler divergence, is given by Ghosh and Tokdar [14] when Θ is finite; see Section 3.1. In Section 3.2, we show that Newton's algorithm can be expressed as a stochastic approximation and results presented in Section 2.4 are used to prove a stronger consistency theorem than in [14] for the case of finite Θ , where the Kullback–Leibler divergence serves as the Lyapunov function.

The numerical investigations in Section 3.3 consider two important cases when Θ is finite, namely, when f is strictly positive on Θ and when $f(\theta) = 0$ for some $\theta \in \Theta$. In the former case, our calculations show that Newton's estimate is superior, in terms of accuracy and computational efficiency, to both the nonparametric MLE and the Bayes estimate. For the latter case, when only a superset of the support of f is known, the story is completely different. While Newton's estimate remains considerably faster than the others, it is not nearly as accurate.

We also consider the problem where the sampling density $p(x|\theta)$ of (1.1) is of the form $p(x|\theta, \xi)$, where f is a mixing density or prior for θ , and ξ is an additional unknown parameter. Newton's algorithm is unable to handle unknown ξ , and we propose a modified algorithm, called N + P, capable of recursively estimating both f and ξ . We express this algorithm as a general SA and prove consistency under suitable conditions.

In Section 5 we briefly discuss some additional theoretical and practical issues concerning Newton's recursive algorithm and the N + P.

2. STOCHASTIC APPROXIMATION

2.1 Algorithm and Examples

Consider the problem of finding the unique root ξ of a function $h(x)$. If $h(x)$ can be evaluated exactly for each x and if h is sufficiently smooth, then various numerical methods can be employed to locate ξ . A majority of these numerical procedures, including the popular Newton–Raphson method, are iterative by nature, starting with an initial guess x_0 of ξ and iteratively defining a sequence $\{x_n\}$ that converges to ξ as $n \rightarrow \infty$. Now consider the situation where only noisy observations on $h(x)$ are available; that is, for any input x one observes $y = h(x) + \varepsilon$, where ε is a zero-mean random error. This problem arises in situations where $h(x)$ denotes the expected value of the response when the experiment is run at setting x . Unfortunately, standard deterministic methods cannot be used in this problem.

In their seminal paper, Robbins and Monro [31] proposed a *stochastic approximation* algorithm for defining a sequence of design points $\{x_n\}$ targeting the root ξ of h in this noisy case. Start with an initial guess x_0 . At stage $n \geq 1$, use the state x_{n-1} as the input, observe $y_n = h(x_{n-1}) + \varepsilon_n$, and update the guess $(x_{n-1}, y_n) \mapsto x_n$. More precisely, the Robbins–Monro algorithm defines the sequence $\{x_n\}$ as follows: start with x_0 and, for $n \geq 1$, set

$$(2.1) \quad \begin{aligned} x_n &= x_{n-1} + w_n y_n \\ &= x_{n-1} + w_n \{h(x_{n-1}) + \varepsilon_n\}, \end{aligned}$$

where $\{\varepsilon_n\}$ is a sequence of i.i.d. random variables with mean zero, and the weight sequence $\{w_n\}$ satisfies

$$(2.2) \quad w_n > 0, \quad \sum_n w_n = \infty, \quad \sum_n w_n^2 < \infty.$$

While the SA algorithm above works in more general situations, we can develop our intuition by looking at the special case considered in [31], namely, when h is bounded, continuous and monotone decreasing. If $x_n < \xi$, then $h(x_n) > 0$ and we have

$$\begin{aligned} \mathbb{E}(x_{n+1}|x_n) &= x_n + w_{n+1} \{h(x_n) + \mathbb{E}(\varepsilon_{n+1})\} \\ &= x_n + w_{n+1} h(x_n) \\ &> x_n. \end{aligned}$$

Likewise, if $x_n > \xi$, then $\mathbb{E}(x_{n+1}|x_n) < x_n$. This shows that the move $x_n \mapsto x_{n+1}$ will be in the correct direction on average.

Some remarks on the conditions in (2.2) are in order. While $\sum_n w_n^2 < \infty$ is necessary to prove convergence, an immediate consequence of this condition is that $w_n \rightarrow 0$. Clearly $w_n \rightarrow 0$ implies that the effect of the noise vanishes as $n \rightarrow \infty$. This, in turn, has an averaging effect on the iterates y_n . On the other hand, the condition $\sum_n w_n = \infty$ washes out the effect of the initial guess x_0 . For further details, see [25].

We conclude this section with three simple examples of SA to shed light on when and how the algorithm works. Example 2.1, taken from [19], page 4, is an important special case of the Robbins–Monro algorithm (2.1) which further motivates the algorithm as well as the conditions (2.2) on the sequence $\{w_n\}$. Example 2.2 uses SA to find quantiles of a t -distribution, and Example 2.3 illustrates a connection between SA and empirical Bayes, two of Robbins's greatest contributions.

EXAMPLE 2.1. Let F_ξ be the cdf of a distribution with mean ξ . Then estimation of ξ is equivalent to solving $h(x) = 0$ where $h(x) = \xi - x$. If Z_1, \dots, Z_n are i.i.d. observations from F_ξ , then the average \bar{Z}_n is the least squares estimate of ξ . To see that $\{\bar{Z}_n\}$ is actually a SA sequence, recall the computationally efficient recursive expression for \bar{Z}_n :

$$(2.3) \quad \bar{Z}_n = \bar{Z}_{n-1} + n^{-1}(Z_n - \bar{Z}_{n-1}).$$

If we let $x_n = \bar{Z}_n$, $w_n = n^{-1}$ and $y_n = Z_n - \bar{Z}_{n-1}$, then (2.3) is exactly of the form of (2.1), with $\{w_n\}$ satisfying (2.2). Moreover, if $\varepsilon_n = Z_n - \xi$, then we can write $y_n = h(x_{n-1}) + \varepsilon_n$. With this setup, we could study the asymptotic behavior of x_n using the SA analysis below (see Sections 2.3 and 2.4), although the SLLN already guarantees $x_n \rightarrow \xi$ a.s.

EXAMPLE 2.2. Suppose we wish to find the α th quantile of the t_ν distribution; that is, we want to find the solution to the equation $F_\nu(x) = \alpha$, where F_ν is the cdf of the t_ν distribution. While there are numerous numerical methods available (e.g., Newton–Raphson or bisection), we demonstrate below how SA can be used to solve this problem. Making use of the well-known fact that the t_ν distribution is a scale-mixture of normals, we can write

$$F_\nu(x) = \mathbb{E}[\Phi(x|\nu^{-1}Z)], \quad Z \sim \chi_\nu^2,$$

where $\Phi(x|\sigma^2)$ is the cdf of the $N(0, \sigma^2)$ distribution. Now, for $Z_1, Z_2, \dots \stackrel{\text{i.i.d.}}{\sim} \chi_\nu^2$, the sequence $\{y_n\}$ defined

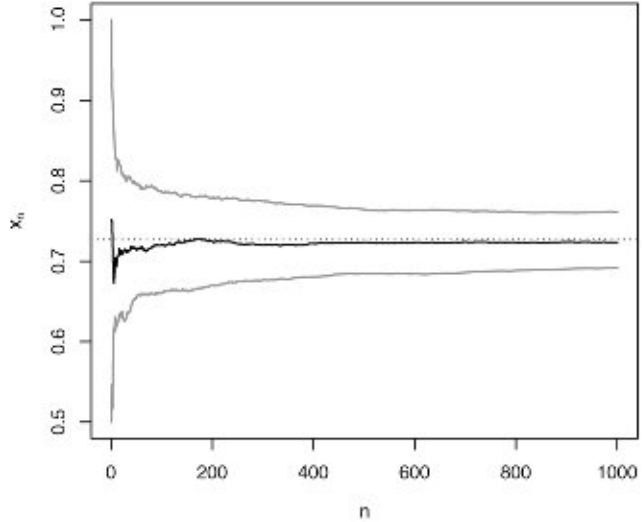


FIG. 1. Sample paths of the three SA sequences $\{x_n\}$ in Example 2.2. The dotted line is the exact 75th percentile of the t_5 distribution.

as $y_n = \alpha - \Phi(x_{n-1}|\nu^{-1}Z_n)$ are noisy observations of $h(x_{n-1}) = \alpha - F_\nu(x_{n-1})$. This h is bounded, continuous and monotone decreasing so the Robbins–Monro theory says that the sequence $\{x_n\}$ defined as (2.1) converges to the true quantile, for any initial condition x_0 . For illustration, Figure 1 shows the first 1000 iterations of the sequence $\{x_n\}$ for $\alpha = 0.75$, $\nu = 5$ and for three starting values $x_0 \in \{0.5, 0.75, 1.0\}$.

EXAMPLE 2.3. In Section 3 we consider a particular recursive estimate and show that it is of the form of a general SA. It turns out that the problem there can also be expressed as an empirical Bayes (EB) problem [30]. In this simple example, we demonstrate the connection between SA and EB, both of which are theories pioneered by Robbins. Consider the simple hierarchical model

$$\lambda_1, \dots, \lambda_n \stackrel{\text{i.i.d.}}{\sim} \text{Exp}(\xi) \quad \text{and} \quad Z_i|\lambda_i \stackrel{\text{ind}}{\sim} \text{Poi}(\lambda_i)$$

for $i = 1, \dots, n$, where the exponential rate $\xi > 0$ is unknown. EB tries to estimate ξ based on the observed data Z_1, \dots, Z_n . Here we consider a recursive estimate of ξ . Fix an initial guess x_0 of ξ . Assuming ξ is equal to x_0 , the posterior mean of λ_1 is $(Z_1 + 1)/(x_0 + 1)$, which is a good estimate of ξ^{-1} if x_0 is close to ξ . Iterating this procedure, we can generate a sequence

$$(2.4) \quad x_i = x_{i-1} + w_i \left[\frac{1}{x_{i-1}} - \frac{Z_i + 1}{x_{i-1} + 1} \right],$$

where $\{w_i\}$ is assumed to satisfy (2.2). Let y_i denote the quantity in brackets in (2.4) and take its expectation

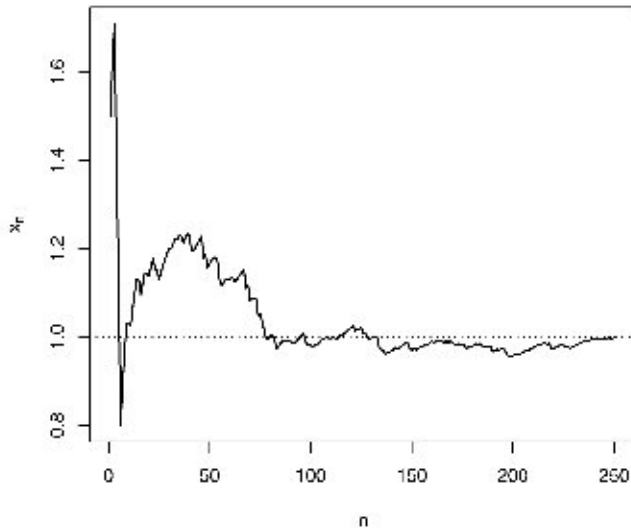


FIG. 2. Sample path of the sequence $\{x_n\}$ in Example 2.3. The dotted line is the value of ξ used for data generation.

with respect to the distribution of Z_i :

$$(2.5) \quad h(x) = \mathbb{E}(y_i | x_{i-1} = x) = \frac{\xi - x}{\xi x(x+1)}.$$

Then the sequence $\{x_n\}$ in (2.4) is a SA targeting a solution of $h(x) = 0$. Since h is continuous, decreasing and $h(x) = 0$ iff $x = \xi$, it follows from the general theory that $x_n \rightarrow \xi$. Figure 2 shows the first 250 steps of such a sequence with $x_0 = 1.5$.

The examples above emphasize one important property that $h(x)$ must satisfy, namely, that it must be easy to “sample” in the sense that there is a function $H(x, z)$ and a random variable Z such that $h(x) = \mathbb{E}[H(x, Z)]$. Another thing, which is not obvious from the examples, is that $h(x)$ must have certain *stability* properties. In general, a SA sequence need not have a unique limit point. However, conditions can be imposed which guarantee convergence to a particular solution ξ of $h(x) = 0$, provided that ξ is a *stable* solution to the ODE $\dot{x} = h(x)$. This is discussed further in Section 2.3.

2.2 Applications

2.2.1 Stochastic approximation EM. The EM algorithm [7] has quickly become one of the most popular computational techniques for likelihood estimation in a host of standard and nonstandard statistical problems. Common to all problems in which the EM can be applied is a notion of “missing data.”

Consider a problem where data Y is observed and the goal is to estimate the parameter θ based on its likelihood function $L(\theta)$. Suppose that the observed

data Y is *incomplete* in the sense that there is a component Z which is *missing*—this could be actual values which are not observed, as in the case of censored data, or it could be latent variables, as in a random effects model. Let $X = (Y, Z)$ denote the *complete* data. Then the likelihood function $f(z, \theta)$ based on the complete data x is related to $L(\theta)$ according to the formula $L(\theta) = \int f(z, \theta) dz$. The EM algorithm produces a convergent sequence of estimates by iteratively filling in the missing data Z in the E-step and then maximizing the simpler complete-data likelihood function $f(z, \theta)$ in the M-step. The E-step is performed by sampling the z -values from the density

$$p(z|\theta) = \begin{cases} f(z, \theta)/L(\theta), & \text{if } L(\theta) \neq 0, \\ 0, & \text{if } L(\theta) = 0, \end{cases}$$

which is the predictive density of Z , given Y and θ .

It is often the case that at least one of the E-step and M-step is computationally difficult, and many variations of the EM have been introduced to improve the rate of convergence and/or simplify the computations. In the case where the E-step cannot be done analytically, Wei and Tanner [39] suggest replacing the expectation in the E-step with a Monte Carlo integration. The resulting MCEM algorithm comes with its own challenges, however; for example, simulating the missing data Z_{nj} , for $j = 1, \dots, m_n$, from $p(z|\theta_n)$ could be quite expensive.

Delyon, Lavielle and Moulines [6] propose, in the case where integration in the E-step is difficult or intractable, an alternative to the MCEM using SA.

SAEM ALGORITHM. At step n , simulate the missing data Z_{nj} from the posterior distribution $p(z|\theta_n)$, $j = 1, \dots, m_n$. Update $\hat{Q}_n(\theta)$ using

$$\hat{Q}_n(\theta) = (1 - w_n)\hat{Q}_{n-1}(\theta) + \frac{w_n}{m_n} \sum_{j=1}^{m_n} \log f(Z_{nj}, \theta),$$

where $\{w_n\}$ is a sequence as in (2.2). Then choose θ_{n+1} such that $\hat{Q}_n(\theta_{n+1}) \geq \hat{Q}_n(\theta)$ for all $\theta \in \Theta$.

Compared to the MCEM, the SAEM algorithm’s use of the simulated data Z_{nj} is much more efficient. At each iteration, the MCEM simulates a new set of missing data from the posterior distribution and forgets the simulated data from the previous iteration. On the other hand, note that the inclusion of $\hat{Q}_{n-1}(\theta)$ in the SAEM update $\theta_n \mapsto \theta_{n+1}$ implies *all* the simulated data points contribute. It is pointed out in [6] that the SAEM performs strikingly better than the MCEM in problems where maximization is much cheaper than simulation.

Delyon, Lavielle and Moulines [6] show, using general SA results, that for a broad class of complete-data likelihoods $f(z, \theta)$ and under standard regularity conditions, the SAEM sequence $\{\theta_n\}$ converges a.s. to the set of stationary points $\{\theta : \nabla L(\theta) = 0\}$ of the incomplete-data likelihood. Moreover, they prove that the only attractive stationary points are local maxima; that is, saddle points of $L(\theta)$ are avoided a.s.

2.2.2 Adaptive Markov Chain Monte Carlo. A random walk Metropolis (RWM) algorithm is a specific MCMC method that can be designed to sample from almost any distribution π . In this particular case, the proposal is $q(x, y) = q(x - y)$, where q is a symmetric density. A popular choice of q is a $N_p(0, \Sigma)$ density. It is well known that the convergence properties of Monte Carlo averages depend on the choice of the proposal covariance matrix Σ , in the sense that it affects the rate at which the generated stochastic process explores the support of π . Trial and error methods for choosing Σ can be difficult and time consuming. One possible solution would be to use the history of the process to suitably tune the proposal. These so-called adaptive algorithms come with their own difficulties, however. In particular, making use of the history destroys the Markov property of the process so nonstandard results are needed in a convergence analysis. For instance, when the state space contains an atom, Gilks, Roberts and Sahu [15] propose an adaptive algorithm that suitably updates the proposal density only when the process returns to the atom. The resulting process is not Markov, but ergodicity is proved using a regeneration argument [15].

An adaptive Metropolis (AM) algorithm is presented by Haario, Saksman and Tamminen [16], which uses previously visited states to update the proposal covariance matrix Σ . Introduce a mean μ and set $\theta = (\mu, \Sigma)$. Let $\{w_n\}$ be a deterministic sequence as in (2.2).

AM ALGORITHM. Fix a starting point z_0 and initial estimates μ_0 and Σ_0 . At iteration $n \geq 1$ draw z_n from $N_p(z_{n-1}, c\Sigma_{n-1})$ and set

$$\begin{aligned}\Sigma_n &= (1 - w_n)\Sigma_{n-1} + w_n(z_n - \mu_{n-1})(z_n - \mu_{n-1})', \\ \mu_n &= (1 - w_n)\mu_{n-1} + w_n z_n.\end{aligned}$$

Note that if $w_n = n^{-1}$, then μ_n and Σ_n are the sample mean and covariance matrix, respectively, of the observations z_1, \dots, z_n . The constant c in the AM is fixed and depends only on the dimension d of the support of π . A choice of c which is, in some sense, optimal is $c = 2.4^2/d$ ([32], page 316).

It is pointed out in [16] that the AM has the advantage of starting the adaptation from the very beginning. This property allows the AM algorithm to search the support of π more effectively *earlier* than other adaptive algorithms. Note that for the algorithm of [15] mentioned above, the adaptation does not begin until the atom is first reached; although the renewal times are a.s. finite, they typically have no finite upper bound.

It is shown in [16] that, under certain conditions, the stationary distribution of the stochastic process $\{z_n\}$ is the target π , the chain is ergodic (even though it is no longer Markovian), and there is almost sure convergence to $\theta_\pi = (\mu_\pi, \Sigma_\pi)$, the mean and covariance of the target π . This implies that, as $n \rightarrow \infty$, the proposal distributions in the AM algorithm will be close to the "optimal" choice. If $H(z, \theta) = (z - \mu, (z - \mu)(z - \mu)' - \Sigma)$, then the AM is a general SA algorithm with $\theta_n = \theta_{n-1} + w_n H(z_n, \theta_{n-1})$, and Andrieu, Moulines and Priouret [2] extend the work in [16] via new SA stability results.

2.2.3 Stochastic approximation Monte Carlo. Let \mathcal{X} be a finite or compact space with a dominating measure ν . Let $p(x) = \kappa p_0(x)$ be a probability density on \mathcal{X} with respect to ν with possibly unknown normalizing constant $\kappa > 0$. We wish to estimate $\int f d\nu$, where f is some function depending on p or p_0 . For example, suppose $p(x)$ is a prior and $g(y|x)$ is the conditional density of y given x . Then $f(x) = g(y|x)p(x)$ is the unnormalized posterior density of x and its integral, the marginal density of y , is needed to compute a Bayes factor.

The following stochastic approximation Monte Carlo (SAMC) method is introduced in [21]. Let A_1, \dots, A_m be a partition of \mathcal{X} and let $\eta_i = \int_{A_i} f d\nu$ for $1 \leq i \leq m$. Take $\hat{\eta}_i(0)$ as an initial guess, and let $\hat{\eta}_i(n)$ be the estimate of η_i at iteration $n \geq 1$. For notational convenience, write

$$\theta_{ni} = \log \hat{\eta}_i(n) \quad \text{and} \quad \theta_n = (\theta_{n1}, \dots, \theta_{nm})'.$$

The probability vector $\pi = (\pi_1, \dots, \pi_m)'$ will denote the *desired* sampling frequency of the A_i 's; that is, π_i is the proportion of time we would like the chain to spend in A_i . The choice of π is flexible and does not depend on the particular partition $\{A_1, \dots, A_m\}$.

SAMC ALGORITHM. Starting with initial estimate θ_0 , for $n \geq 0$ simulate a sample z_{n+1} using a RWM algorithm with target distribution

$$(2.6) \quad p(z|\theta_n) \propto \sum_{i=1}^m f(z) e^{-\theta_{ni}} I_{A_i}(z), \quad z \in \mathcal{X}.$$

Then set $\theta_{n+1} = \theta_n + w_{n+1}(\zeta_{n+1} - \pi)$, where the deterministic sequence $\{w_n\}$ is as in (2.2), and $\zeta_{n+1} = (I_{A_1}(z_{n+1}), \dots, I_{A_m}(z_{n+1}))'$.

The normalizing constant in (2.6) is generally unknown and difficult to compute. However, $p(z|\theta_n)$ is only used at the RWM step where it is only required that the target density be known up to a proportionality constant.

It turns out that, in the case where no A_i are empty, the observed sampling frequency $\hat{\pi}_i$ of A_i converges to π_i . This shows that $\hat{\pi}_i$ is independent of its probability $\int_{A_i} p \, d\nu$. Consequently, the resulting chain will not get stuck in regions of high probability, as a standard Metropolis chain might.

The sequence $\{\theta_n\}$ is a general stochastic approximation and, using the convergence results of [2], Liang, Liu and Carroll [21] show that if no A_i is empty and suitable conditions are met, then

$$(2.7) \quad \theta_{ni} \rightarrow C + \log \int_{A_i} f \, d\nu - \log \pi_i \quad \text{a.s.,}$$

for $1 \leq i \leq m$ as $n \rightarrow \infty$, for some arbitrary constant C . Liang, Liu and Carroll [21] point out a *lack of identifiability* in the limit (2.7); that is, C cannot be determined from $\{\theta_n\}$ alone. Additional information is required, such as $\sum_{i=1}^m \hat{\eta}_i(n) = c$ for each n and for some known constant c .

In Example 2.4, we apply SAMC to estimate the partition function in the one-dimensional Ising model. In this simple situation, a closed-form expression is available, which we can use as a baseline for assessing the performance of the SAMC estimate.

EXAMPLE 2.4. Consider a one-dimensional Ising model, which assumes that each of the d particles in a system has positive or negative spin. The Gibbs distribution on $\mathcal{X} = \{-1, 1\}^d$ has density (with respect to counting measure ν)

$$p_T(x) = \frac{1}{Z(T)} e^{-E(x)/T}, \quad Z(T) = \sum_{x \in \mathcal{X}} e^{-E(x)/T},$$

where T is the temperature, and E is the energy function defined, in this case, as $E(x) = -\sum_{i=1}^{d-1} x_i x_{i+1}$. The partition function $Z(T)$ is of particular interest to physicists: the thermodynamic limit $F(T) = \lim_{d \rightarrow \infty} d^{-1} \log Z(T)$ is used to study phase transitions [4]. In this simple case, a closed-form expression for $Z(T)$ is available. There are other more complex systems, however, where no analytical solution is available and $\nu(\mathcal{X}) = 2^d$ is too large to allow for naïve calculation of $Z(T)$.

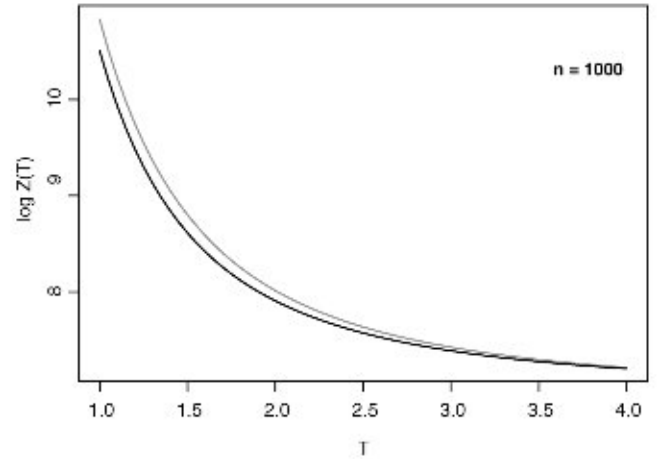


FIG. 3. $\log Z(T)$ (gray) and SAMC estimate $\log \hat{Z}(T)$ (black) in Example 2.4.

Our jumping-off point is the *energy–temperature duality* [18] $Z(T) = \sum_u \Omega(u) e^{-u/T}$, where $\Omega(u) = \nu\{x : E(x) = u\}$ is the density of states. We will apply SAMC to first estimate $\Omega(u)$ and then estimate $Z(T)$ with a plug-in:

$$\hat{Z}(T) = \sum_u \hat{\Omega}(u) e^{-u/T}.$$

Note here that a single estimate of Ω can be used to estimate the partition function for any T , eliminating the need for simulations at multiple temperatures. Furthermore, $\sum_u \Omega(u) = \nu(\mathcal{X}) = 2^d$ is known so, by imposing this condition on the estimate $\hat{\Omega}$ we do not fall victim to the lack of identifiability mentioned above. Figure 3 shows the true partition function $Z(T) = 2^d \cosh^{d-1}(1/T)$ for $d = 10$ as well as the SAMC estimate $\hat{Z}(T)$ as a function of $T \in [1, 4]$, on the log-scale, based on $n = 1000$ iterations. Clearly, \hat{Z} performs quite well in this example, particularly for large T .

2.3 ODE Stability Theory

The asymptotic theory of ODEs plays an important role in the convergence analysis of a SA algorithm. After showing the connection between SA and ODEs, we briefly review some of the ODE theory that is necessary in the sequel.

Recall the general SA algorithm in (2.1) given by $x_n = x_{n-1} + w_n y_n$. Assume there is a measurable function h such that $h(x_{n-1}) = \mathbb{E}[y_n | x_{n-1}]$ and rewrite this algorithm as

$$x_n = x_{n-1} + w_n h(x_{n-1}) + w_n \{y_n - h(x_{n-1})\}.$$

Define $M_n = y_n - h(x_{n-1})$. Then $\{M_n\}$ is a zero-mean *martingale* sequence and, under suitable conditions, the martingale convergence theorem guarantees that M_n becomes negligible as $n \rightarrow \infty$, leaving us with

$$\begin{aligned} x_n &= x_{n-1} + w_n h(x_{n-1}) + w_n M_n \\ &\approx x_{n-1} + w_n h(x_{n-1}). \end{aligned}$$

But this latter “mean trajectory” is deterministic and essentially a finite difference equation with small step sizes. Rearranging the terms gives us

$$\frac{x_n - x_{n-1}}{w_n} = h(x_{n-1}),$$

which, for large n , can be approximated by the ODE $\dot{x} = h(x)$. It is for this reason that the study of SA algorithms is related to the asymptotic properties of solutions to ODEs.

Consider a general autonomous ODE $\dot{x} = h(x)$, where $h: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a bounded and continuous, possibly nonlinear, function. A solution $x(t)$ of the ODE is a trajectory in \mathbb{R}^d with a given initial condition $x(0)$. Unfortunately, in many cases, a closed-form expression for a solution $x(t)$ is not available. For that reason, other methods are necessary for studying these solutions and, in particular, their properties as $t \rightarrow \infty$.

Imagine a physical system, such as an orbiting celestial body, whose state is being governed by the ODE $\dot{x} = h(x)$ with initial condition $x(0) = x_0$. Then, loosely speaking, the system is stable if choosing an alternative initial condition $x(0) = x'_0$ in a neighborhood of x_0 has little effect on the asymptotic properties of the resulting solution $x(t)$. The following definition makes this more precise.

DEFINITION 2.5. A point $\xi \in \mathbb{R}^d$ is said to be *locally stable* for $\dot{x} = h(x)$ if for each $\varepsilon > 0$ there is a $\delta > 0$ such that if $\|x(0) - \xi\| < \delta$, then $\|x(t) - \xi\| < \varepsilon$ for all $t \geq 0$. If ξ is locally stable and $x(t) \rightarrow \xi$ as $t \rightarrow \infty$, then ξ is *locally asymptotically stable*. If this convergence holds for all initial conditions $x(0)$, then the asymptotic stability is said to be *global*.

Points ξ for which stability is of interest are *equilibrium points* of $\dot{x} = h(x)$. Any point ξ such that $h(\xi) = 0$ is called an equilibrium point, since the constant solution $x(t) \equiv \xi$ satisfies $\dot{x} = h(x)$.

EXAMPLE 2.6. Let $\dot{x} = Ax$, where A is a fixed $d \times d$ matrix. For an initial condition $x(0) = x_0$, we can write an explicit formula for the particular solution: $x(t) = e^{At}x_0$ for $t \geq 0$. Suppose, for simplicity, that A has a spectral decomposition $A = U\Lambda U'$, where U is

orthogonal and Λ is a diagonal matrix of the eigenvalues $\lambda_1, \dots, \lambda_d$ of A . Then the matrix exponential can be written as $e^{At} = Ue^{\Lambda t}U'$, where $e^{\Lambda t}$ is diagonal with i th element $e^{\lambda_i t}$. Clearly, if $\lambda_i < 0$, then $e^{\lambda_i t} \rightarrow 0$ as $t \rightarrow \infty$. Therefore, if A is negative definite, then the origin $x = 0$ is globally asymptotically stable.

When explicit solutions are not available, proving asymptotic stability for a given equilibrium point will require a so-called *Lyapunov function* [20].

DEFINITION 2.7. Let $\xi \in \mathbb{R}^d$ be an equilibrium point of the ODE $\dot{x} = h(x)$ with initial condition $x(0) = x_0$. A function $\ell: \mathbb{R}^d \rightarrow \mathbb{R}$ is called a *Lyapunov function* (at ξ) if:

- ℓ has continuous first partial derivatives in a neighborhood of ξ ;
- $\ell(x) \geq 0$ with equality if and only if $x = \xi$;
- the time derivative of ℓ along the path $x(t)$, defined as $\dot{\ell}(x) = \nabla \ell(x)'h(x)$, is ≤ 0 .

A Lyapunov function is said to be *strong* if $\dot{\ell}(x) = 0$ implies $x = \xi$.

Lyapunov functions are a generalization of the potential energy of a system, such as a swinging pendulum, and Lyapunov's theory gives a formal extension of the stability principles of such a system. Theorem 2.8 is very powerful because it does not require an explicit formula for the solution. See [20] for a proof and various extensions of the Lyapunov theory.

THEOREM 2.8. *If there exists a (strong) Lyapunov function in a neighborhood of an equilibrium point ξ of $\dot{x} = h(x)$, then ξ is (asymptotically) stable.*

There is no general recipe for constructing a Lyapunov function. In one important special case, however, a candidate Lyapunov function is easy to find. Suppose $h(x) = -\nabla g(x)$, for some positive definite, sufficiently smooth function g . Then $\ell(x) = g(x)$ is a Lyapunov function since $\dot{\ell}(x) = -\|\nabla g(x)\|^2 \leq 0$.

EXAMPLE 2.9. Consider again the linear system $\dot{x} = Ax$ from Example 2.6, where A is a $d \times d$ negative definite matrix. Here we will derive asymptotic stability by finding a Lyapunov function and applying Theorem 2.8. In light of the previous remark, we choose $\ell(x) = -\frac{1}{2}x'Ax$. Then $\dot{\ell}(x) = -\|Ax\|^2 \leq 0$ so ℓ is a strong Lyapunov function for $\dot{x} = Ax$ and the origin is asymptotically stable by Theorem 2.8.

Of interest is the stronger conclusion of *global* asymptotic stability. Note, however, that Theorem 2.8

does not tell us how far x_0 can be from the equilibrium in question and still get asymptotic stability. For the results that follow, we will prove the *global* part directly.

2.4 SA Convergence Theorem

Consider, for fixed x_0 and $\{w_n\}$ satisfying (2.2), the general SA algorithm

$$(2.8) \quad x_n = \text{Proj}_X\{x_{n-1} + w_n y_n\}, \quad n \geq 1,$$

where $X \subset \mathbb{R}^d$ is compact and $\text{Proj}_X(x)$ is a projection of x onto X . The projection is necessary when boundedness of the iterates cannot be established by other means. The *truncated* or *projected* algorithm (2.8) is often written in the alternative form [19]

$$(2.9) \quad x_n = x_{n-1} + w_n y_n + w_n z_n,$$

where z_n is the “minimum” z such that $x_{n-1} + w_n y_n + w_n z$ belongs to X .

Next we state the main stochastic approximation result used in the sequel, a special case of Theorem 5.2.3 in [19]. Define the filtration sequence $\mathcal{F}_n = \sigma(y_1, \dots, y_n)$.

THEOREM 2.10. *For $\{x_n\}$ in (2.8) with $\{w_n\}$ satisfying (2.2), assume*

- (SA1) $\sup_n \mathbb{E}\|y_n\|^2 < \infty$.
- (SA2) *There exists a continuous function $h(\cdot)$ and a random vector β_n such that $\mathbb{E}(y_n | \mathcal{F}_{n-1}) = h(x_{n-1}) + \beta_n$ a.s. for each n .*
- (SA3) $\sum_n w_n \|\beta_n\|$ converges a.s.

If ξ is globally asymptotically stable for $\dot{x} = h(x)$, then $x_n \rightarrow \xi$ a.s.

3. NEWTON’S RECURSIVE ESTIMATE

Let Θ and \mathcal{X} be the parameter space and sample space, equipped with σ -finite measures μ and ν , respectively. Typically, Θ and \mathcal{X} are subsets of Euclidean space and ν is Lebesgue or counting measure. The measure μ varies depending on the inference problem: for estimation, μ is usually Lebesgue or counting measure, but for testing, μ is often something different (see Example 3.1).

Consider the following model for pairs of random variables $(X_i, \theta^i) \in \mathcal{X} \times \Theta$:

$$(3.1) \quad \theta^i \stackrel{\text{i.i.d.}}{\sim} f, \quad X_i | \theta^i \stackrel{\text{ind}}{\sim} p(\cdot | \theta^i), \quad i = 1, \dots, n,$$

where $\{p(\cdot | \theta) : \theta \in \Theta\}$ is a parametric family of probability densities with respect to ν on \mathcal{X} and f is a probability density with respect to μ on Θ . In the present

case, the variables (parameters) $\theta^1, \dots, \theta^n$ are not observed. Therefore, under model (3.1), X_1, \dots, X_n are i.i.d. observations from the marginal density Π_f in (1.1). We call f the mixing density (or prior, in the Bayesian context) and the inference problem is to estimate f based on the data observed from Π_f . The following example gives a very important special case of this problem—the analysis of DNA microarray data.

EXAMPLE 3.1. A microarray is a tool that gives researchers the ability to simultaneously investigate the effects of numerous genes on the occurrence of various diseases. Not all of the genes will be *expressed*—related to the disease in question—so the problem is to identify those which are. Let θ^i represent the expression level of the i th gene, with $\theta^i = 0$ indicating the gene is not expressed. After some reduction, the data X_i is a measure of θ^i , and the model is of the form (3.1) with f being a prior density with respect to $\mu = \lambda_{\text{Leb}} + \delta_{\{0\}}$. Consider the multiple testing problem

$$H_{0i} : \theta^i = 0, \quad i = 1, \dots, n.$$

The number n of genes under investigation is often in the thousands so, with little information about θ^i in X_i , choosing a fixed prior f would be problematic. On the other hand, the data contain considerable information about the prior f so the *empirical Bayes* approach—using the data to *estimate* the prior—has been quite successful [9].

In what follows, we focus our attention on a particular estimate of the mixing density f . Let $x_1, \dots, x_n \in \mathcal{X}$ be i.i.d. observations from the mixture density Π_f in (1.1). Newton [26] suggests the following algorithm for estimating f .

NEWTON’S ALGORITHM. Choose a positive density f_0 on Θ and weights $w_1, \dots, w_n \in (0, 1)$. Then for $i = 1, \dots, n$, compute

$$(3.2) \quad f_i(\theta) = (1 - w_i)f_{i-1}(\theta) + w_i \frac{p(X_i | \theta) f_{i-1}(\theta)}{\Pi_{i-1}(X_i)},$$

where $\Pi_j(x) = \int p(x | \theta) f_j(\theta) d\mu(\theta)$, and report $f_n(\theta)$ as the final estimate.

In the following subsections we establish some asymptotic properties of f_n as $n \rightarrow \infty$ and we show the results of several numerical experiments that demonstrate the finite-sample accuracy of Newton’s estimate (3.2) in both the discrete and continuous cases. First, a few important remarks.

- The update $f_{i-1} \mapsto f_i$ in (3.2) is similar to a Bayes estimate based on a Dirichlet process prior (DPP), given the information up to, and including, time $i - 1$. That is, after observing X_1, X_2, \dots, X_{i-1} , a Bayesian might model f with a DPP $\mathcal{D}(\frac{1-w_i}{w_i}, f_{i-1})$. In this case, the posterior expectation is exactly the f_i in (3.2).
- Because f_n depends on the ordering of the data and not simply on the sufficient statistic $(X_{(1)}, \dots, X_{(n)})$, it is *not* a posterior quantity.
- The algorithm is very fast: if one evaluates (3.2) on a grid of m points $\theta_1, \dots, \theta_m$ and calculates the integral in Π_{i-1} using, say, a trapezoid rule, then the computational complexity is mn .

3.1 Review of Convergence Results

In this section, we give a brief review of the known convergence results for Newton's estimate f_n in the case of a finite parameter space. The case of a compact Θ is quite different and, until very recently [38], nothing was known about the convergence of f_n in such problems; see Section 5.

Newton [26], building on the work in [27, 28], states the following convergence theorem.

THEOREM 3.2. *Assume the following:*

- (N1) Θ is finite and μ is counting measure.
- (N2) $\sum_n w_n = \infty$.
- (N3) $p(x|\theta)$ is bounded away from 0 and ∞ .

Then surely there exists a density f_∞ on Θ such that $f_n \rightarrow f_\infty$ as $n \rightarrow \infty$.

Newton [26] presents a proof of Theorem 3.2 based on the theory of nonhomogeneous Markov chains. He proves that f_n represents the n -step marginal distribution of the Markov chain $\{Z_n\}$ given by

$$Z_0 \sim f_0, \quad Z_n = \begin{cases} Z_{n-1}, & \text{with prob } 1 - w_n, \\ Y_n, & \text{with prob } w_n, \end{cases}$$

where Y_n has density $\propto p(X_n|\theta) f_{n-1}(\theta)$. However, the claim that this Markov chain admits a stationary distribution is incomplete—N2 implies the chain $\{Z_n\}$ is weakly ergodic but the necessary strong ergodicity property does not follow, even when Θ is finite. Counterexamples are given in [14, 17]. Ghosh and Tokdar [14] prove consistency of f_n along quite different lines. For probability densities ψ and φ with respect to μ , define the Kullback–Liebler (KL) divergence,

$$(3.3) \quad K(\psi, \varphi) = \int_{\Theta} \psi \log(\psi/\varphi) d\mu.$$

The following theorem is proved in [14] using an approximate martingale representation of $K(f, f_n)$.

THEOREM 3.3. *In addition to N1–N3, assume*

- (GT1) $\sum_n w_n^2 < \infty$.
- (GT2) f is identifiable; that is, $f \mapsto \Pi_f$ is injective.

Then $K(f, f_n) \rightarrow 0$ a.s. as $n \rightarrow \infty$.

Part of the motivation for the use of the KL divergence lies in the fact that the ratio f_n/f_{n-1} has a relatively simple form. More important, however, is the Lyapunov property shown in the proof Theorem 3.4. Sufficient conditions for GT2 in the case of finite Θ are given in, for example, [22, 37]. San Martin and Quintana [33] also discuss the issue of identifiability in connection with the consistency of f_n .

3.2 Newton's Estimate as a SA

Here we show that Newton's algorithm (3.2) is a special case of SA. First, note that if f is viewed as a prior density, then estimating f is an *empirical Bayes* (EB) problem. The ratio in (3.2) is nothing but the posterior distribution of θ , given x_i , and assuming that the prior f is equal to f_{i-1} . This, in fact, is exactly the approach taken in Example 2.3 to apply SA in an EB problem.

Let μ be counting measure and $d = \mu(\Theta)$. We can think of $f_n(\theta)$ as a vector $f_n = (f_n^1, \dots, f_n^d)'$ in the probability simplex Δ^d , defined as

$$\Delta^d = \left\{ (\varphi^1, \dots, \varphi^d)' \in [0, 1]^d : \sum_{i=1}^d \varphi^i = 1 \right\}.$$

Define $H : \mathcal{X} \times \Delta^d \rightarrow \mathbb{R}^d$ with k th component

$$(3.4) \quad H_k(x, \varphi) = \frac{p(x|\theta_k)\varphi^k}{\Pi_\varphi(x)} - \varphi^k, \quad k = 1, \dots, d,$$

where $\Pi_\varphi(x) = \sum_k p(x|\theta_k)\varphi^k$ is the marginal density on \mathcal{X} induced by $\varphi \in \Delta^d$. Then (3.2) becomes

$$(3.5) \quad f_n = f_{n-1} + w_n H(X_n, f_{n-1}).$$

Let $P_x = \text{diag}\{p(x|\theta_k) : k = 1, \dots, d\}$ be the diagonal matrix of the sampling density values and define the mapping $h : \Delta^d \rightarrow \mathbb{R}^d$ to be the conditional expectation of $H(x, f_n)$, given $f_n = \varphi$:

$$(3.6) \quad \begin{aligned} h(\varphi) &= \int_{\mathcal{X}} H(x, \varphi) \Pi_f(x) d\nu(x) \\ &= \int_{\mathcal{X}} \frac{\Pi_f(x)}{\Pi_\varphi(x)} P_x \varphi d\nu(x) - \varphi, \end{aligned}$$

where $f = (f^1, \dots, f^d)'$ is the true mixing/prior distribution. From (3.6), it is clear that f solves the equation $h(\varphi) = 0$ which implies (i) f is an equilibrium

point of the ODE $\dot{\varphi} = h(\varphi)$, and (ii) that f is a fixed point of the map

$$T(\varphi) = h(\varphi) + \varphi = \int \frac{\Pi_f(x)}{\Pi_\varphi(x)} P_x \varphi d\nu(x).$$

Newton [26], page 313, recognized the importance of this map in relation to the limit of f_n . Also, the use of T in [5, 35] for the I -projection problem is closely related to the SA approach taken here.

We have shown that (3.5) can be considered as a general SA algorithm, targeting the solution $\varphi = f$ of the equation $h(\varphi) = 0$ in Δ^d . Therefore, the SA results of Section 2.4 can be used in the convergence analysis. The following theorem is proved in Appendix A.1.

THEOREM 3.4. *Assume N1, N2, GT1 and GT2. If $p(\cdot|\theta) > 0$ v.a.e. for each θ , then $f_n \rightarrow f$ a.s.*

REMARK 3.5. Removal of the boundedness condition N3 on $p(x|\theta)$ in Theorem 3.4 extends the consistency result of [14] to many important cases, such as mixtures of normal or gamma densities.

REMARK 3.6. Theorem 3.4 covers the *interior* case (when f is strictly positive) as well as the *boundary* case (when $f^i = 0$ for some i). The fact that $f_0^i > 0$ implies $f_n^i > 0$ for all n suggests that convergence may be slow in the boundary case.

3.3 Simulations

Here we provide numerical illustrations comparing the performance of Newton's estimate with that of its competitors. We consider a location-mixture of normals; that is, $p(\cdot|\theta)$ is a $N(\theta, \sigma^2)$ density. The weights are set to be $w_i = (i+1)^{-1}$ and the initial estimate f_0 is taken to be a $\text{Unif}(\Theta)$ density. For the Bayes estimate, we assume a Dirichlet process prior $f \sim \mathcal{D}(1, f_0)$ in each example.

EXAMPLE 3.7 (Finite Θ). In this example, we compare Newton's recursive (NR) estimate with the nonparametric maximum likelihood (NPML) estimate and the nonparametric Bayes (NPB) estimate. Computation of NR and NPML (using the EM algorithm) is straightforward. Here, in the case of finite Θ , we use sequential imputation [23] to calculate NPB. Take $\Theta = \mathbb{Z} \cap [-4, 4]$, and set $\sigma = 1$ in $p(x|\theta)$. We consider two different mixing distributions on Θ :

- I. $f = \text{Bin}(8, 0.6)$,
- II. $f = 0.5\delta_{\{-2\}} + 0.5\delta_{\{2\}}$.

We simulate 50 data sets of size $n = 100$ from the models corresponding to mixing densities I, II and computing the three estimates for each. Figure 4 shows the resulting estimates for a randomly chosen data set from each model. Notice that NR does better for model I than both NPML and NPB. The story is different for model II—both NPML and NPB are considerably better than NR. This is further illustrated in Figure 5 where the KL divergence $K(\Pi_f, \hat{\Pi}_n)$ on $\mathcal{X} = \mathbb{R}$ is summarized over the 50 samples. We see that NR has a slightly smaller KL number than NPML and NPB for model I, but they clearly dominate NR for model II. This discrepancy is at least partially explained by Remark 3.6; see Section 5 for further discussion. We should point out, however, that both NPML and NPB take significantly longer to compute than NR, about 100 times longer on average.

EXAMPLE 3.8 (Compact Θ). We consider a one- and a two-component mixture of beta densities on $\Theta = [0, 1]$ as the true f :

- I. $f = \text{Beta}(2, 7)$,
- II. $f = 0.33 \text{Beta}(3, 30) + 0.67 \text{Beta}(4, 4)$.

Let $\sigma = 0.1$ be the normal sampling variance. Again, computation of NR is straightforward. To compute NPB, the importance sampling algorithm in [38] that makes use of a collapsing of the Polyá Urn scheme is used. Figure 6 shows a typical realization of NR and NPB, based on a sample of size $n = 100$ from each of the corresponding marginals. Note that the Bayes estimate does a rather poor job here, being much too spiky in both cases. This is mainly because the posterior for f sits on discrete distributions. On the other hand, Newton's estimate has learned the general shape of f after only 100 iterations and results in a much better estimate than NPB. Furthermore, on average, the computation time for NR is again about 100 times less than that of NPB.

4. N+P ALGORITHM

Suppose that the sampling distribution on \mathcal{X} is parametrized not only by θ but by an additional parameter ξ . An example of this is the normal distribution with mean θ and variance $\xi = \sigma^2$. More specifically, we replace the sampling densities $p(x|\theta)$ of Section 3 with $p(x|\theta, \xi)$ where θ is the latent variable, and ξ is also unknown. Newton's algorithm cannot be used in this situation since θ does not fully specify the sampling density.

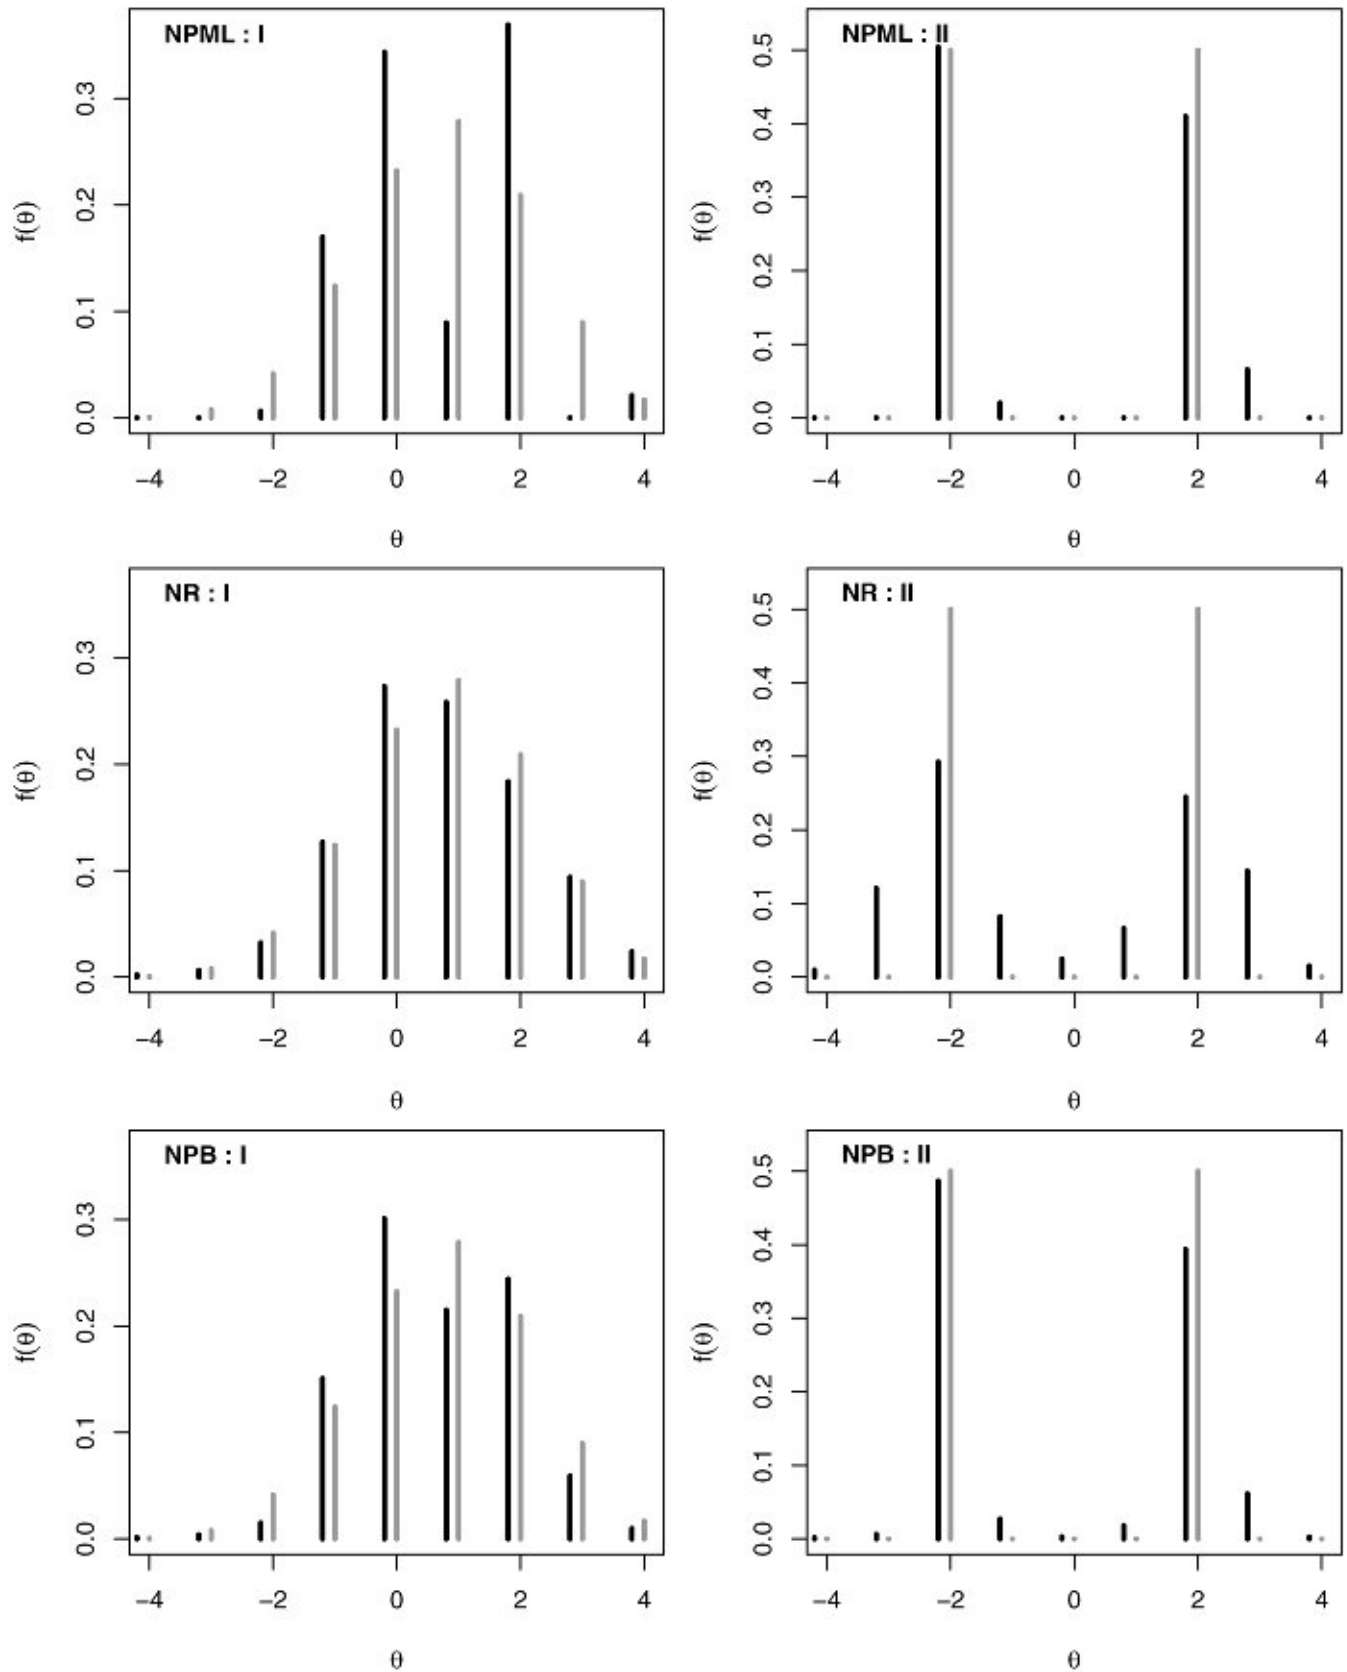


FIG. 4. Estimates of mixing densities I and II in Example 3.7. Left column: True f (gray) for model I and the three estimates (black). Right column: True f (gray) for model II and the three estimates (black).

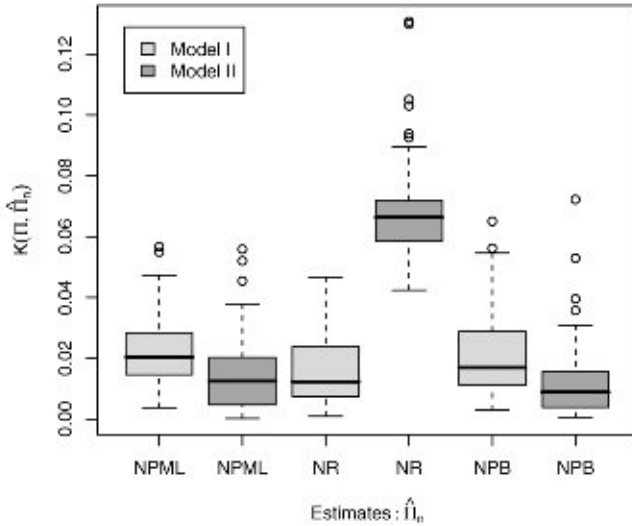


FIG. 5. Summary of the KL divergence $K(\Pi_f, \hat{\Pi}_n)$ for the three estimates $\hat{\Pi}_n$ in models I and II in Example 3.7.

In this section we introduce a modification of Newton’s algorithm to simultaneously and recursively estimate both a mixing distribution and an additional unknown parameter. This modification, called the

Newton+Plug-in (N+P), is actually quite simple—at each step we use a plug-in estimate of ξ in the update (3.2). We show that the N+P algorithm can be written as a general SA algorithm and, under certain conditions, prove its consistency.

Let $p(x|\theta, \xi)$ be a two-parameter family of densities on \mathcal{X} , and consider the model

$$(4.1) \quad \begin{aligned} \theta^1, \dots, \theta^n &\stackrel{i.i.d.}{\sim} f, \\ X_{i1}, \dots, X_{ir} &\stackrel{i.i.d.}{\sim} p(\cdot|\theta^i, \xi), \quad i = 1, \dots, n, \end{aligned}$$

where f is an unknown density on Θ and the parameter $\xi \in \Xi$ is also unknown. The number of replicates $r \geq 2$ is assumed fixed. Note that (4.1) is simply a *non-parametric* random effects model.

Assume, for simplicity, that $\Xi \subseteq \mathbb{R}$; the more general case $\Xi \subseteq \mathbb{R}^p$ is a natural extension of what follows. Let $\Theta = \{\theta_1, \dots, \theta_d\}$ be a *finite* set and take μ to be counting measure on Θ . Recall that Δ^d is the probability simplex. Assume:

- (NP1) $\xi \in \text{int}(\Xi_0)$, where Ξ_0 is a compact and convex subset of Ξ .

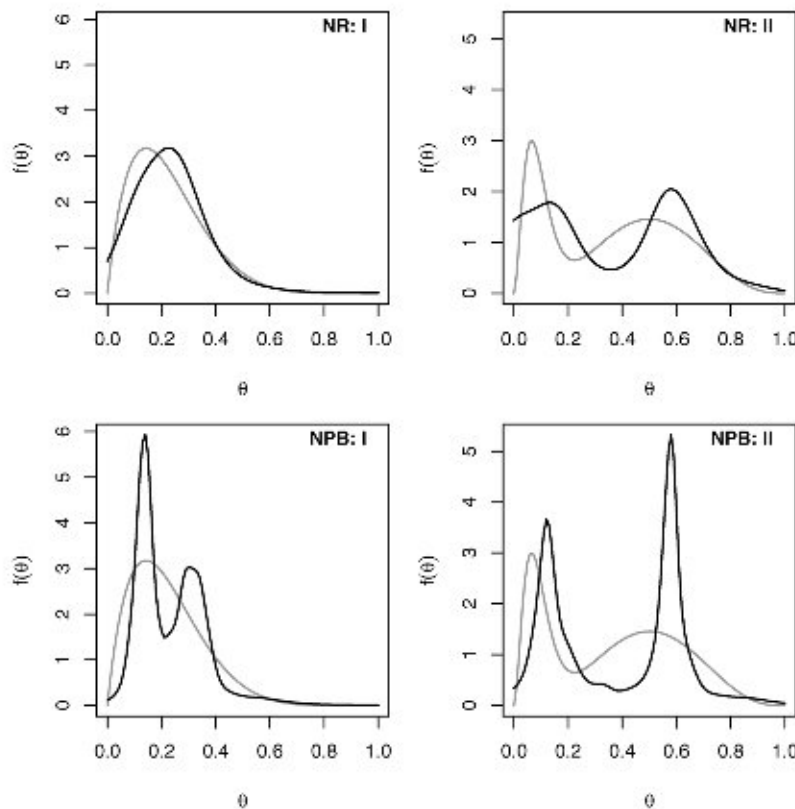


FIG. 6. Estimates of the mixing densities I and II in Example 3.8. Top row: true f (gray) and NR (black). Bottom row: true f (gray) and NPB (black).

(NP2) $f \in \text{int}(\Delta_0)$ where $\Delta_0 \subset \Delta^d$ is compact and, for each $\varphi \in \Delta_0$, the coordinates $\varphi^1, \dots, \varphi^d$ are bounded away from zero.

The subset Ξ_0 can be arbitrarily large so assumption NP1 causes no difficulty in practice. Assumption NP2 is somewhat restrictive in that f must be strictly positive. While NP2 seems necessary to prove consistency (see Appendix A.3), simulations suggest that this assumption can be weakened.

The N+P algorithm uses an estimate of ξ at each step in Newton's algorithm. We assume here that an unbiased estimate is available:

(NP3) There exists an unbiased estimate $T_{\text{UBE}}(x)$, $x \in \mathcal{X}^r$, of ξ with variance $v^2 < \infty$.

Later we will replace the unbiased estimate with a Bayes estimate. This will require replacing NP3 with another assumption.

At time $i = 1, \dots, n$, we observe an r -vector $X_i = (X_{i1}, \dots, X_{ir})'$ and we compute $\hat{\xi}^{(i)} = T_{\text{UBE}}(X_i)$. An unbiased estimate of ξ based on the entire data X_1, \dots, X_n would be the average $\xi_n = n^{-1} \sum_{i=1}^n \hat{\xi}^{(i)}$, which has a convenient recursive expression

$$(4.2) \quad \xi_i = i^{-1}[(i-1)\xi_{i-1} + \hat{\xi}^{(i)}], \quad i = 1, \dots, n.$$

More importantly, by construction, $\hat{\xi}^{(1)}, \dots, \hat{\xi}^{(n)}$ are i.i.d. random variables with mean ξ and finite variance. It is, therefore, a consequence of the SLLN that ξ_n , as defined in (4.2), converges a.s. to ξ . While this result holds for any unbiased estimate T , an unbiased estimate T' with smaller variance is preferred, since it will have better finite-sample performance.

Define the mapping $H: \mathcal{X}^r \times \Delta_0 \times \Xi_0 \rightarrow \mathbb{R}^d$ with k th component

$$(4.3) \quad H_k(x, \varphi, \psi) = \frac{\bar{p}(x|\theta_k, \psi)\varphi^k}{\sum_j \bar{p}(x|\theta_j, \psi)\varphi^j} - \varphi^k,$$

for $k = 1, \dots, d$, where φ and ψ denote generic elements in Δ_0 and Ξ_0 , respectively, and $\bar{p}(\cdot|\theta, \psi)$ is the joint density of an i.i.d. sample of size r from $p(\cdot|\theta, \psi)$.

N+P ALGORITHM. Choose an initial estimate $f_0 \in \Delta_0$, weights $w_1, \dots, w_n \in (0, 1)$, and an arbitrary $\xi_0 \in \Xi_0$. Then for $i = 1, \dots, n$ compute

$$\begin{aligned} \xi_i &= \text{Proj}_{\Xi_0} \{i^{-1}[(i-1)\xi_{i-1} + \hat{\xi}^{(i)}]\}, \\ f_i &= \text{Proj}_{\Delta_0} \{f_{i-1} + w_i H(X_i; f_{i-1}, \xi_i)\}, \end{aligned}$$

and produce (f_n, ξ_n) as the final estimate.

We claim that the N+P algorithm for estimating f can be written as a general SA involving the true but unknown ξ plus an additional perturbation. Define the quantities

$$(4.4) \quad h(f_{n-1}) = \mathbb{E}[H(X_n, f_{n-1}, \xi)|\mathcal{F}_{n-1}],$$

$$(4.5) \quad \begin{aligned} \beta_n &= \mathbb{E}[H(X_n, f_{n-1}, \xi_n)|\mathcal{F}_{n-1}] \\ &\quad - \mathbb{E}[H(X_n, f_{n-1}, \xi)|\mathcal{F}_{n-1}], \end{aligned}$$

where $\mathcal{F}_{n-1} = \sigma(X_1, \dots, X_{n-1})$, so that

$$\mathbb{E}[H(X_n, f_{n-1}, \xi_n)|\mathcal{F}_{n-1}] = h(f_{n-1}) + \beta_n.$$

Now the update $f_{n-1} \mapsto f_n$ can be written as

$$(4.6) \quad f_n = f_{n-1} + w_n \{h(f_{n-1}) + M_n + \beta_n + z_n\},$$

where z_n is the "minimum" z keeping f_n in Δ_0 , and

$$M_n = H(X_n, f_{n-1}, \xi_n) - h(f_{n-1}) - \beta_n$$

is a martingale adapted to \mathcal{F}_{n-1} . Notice that (4.6) is now in a form in which Theorem 2.10 can be applied. We will make use of the Law of Iterated Logarithm so define $u(t) = (2t \log \log t)^{1/2}$. The consistency properties of the N+P algorithm are summarized in the following theorem.

THEOREM 4.1. Assume N1, GT1, GT2, NP1–NP3. In addition, assume

(NP4) $\frac{\partial}{\partial \psi} H(x; \varphi, \psi)$ is bounded on $\mathcal{X}^r \times \Delta_0 \times \Xi_0$.

(NP5) $\sum_n w_n n^{-1} u(n)$ converges.

Then $(f_n, \xi_n) \rightarrow (f, \xi)$ a.s. as $n \rightarrow \infty$.

We now remove the restriction to unbiased estimates of ξ , focusing primarily on the use of a Bayes estimate in place of the unbiased estimate. But first, let $\tilde{\xi}_i = T(X_1, \dots, X_i)$ be any suitable estimate of ξ based on only X_1, \dots, X_i . Then replace the N+P update $f_{i-1} \mapsto f_i$ with

$$\tilde{f}_i = \text{Proj}_{\Delta_0} \{\tilde{f}_{i-1} + w_i H(X_i, \tilde{f}_{i-1}, \tilde{\xi}_i)\}.$$

While this adaptation is more flexible with regard to the choice of estimate, this additional flexibility does not come for free. Notice that the algorithm is no longer recursive. That is, given a new data point x_{n+1} , we need more information than just the pair $(\tilde{f}_n, \tilde{\xi}_n)$ to obtain $(\tilde{f}_{n+1}, \tilde{\xi}_{n+1})$.

COROLLARY 4.2. If assumptions NP3 and NP5 in Theorem 4.1 are replaced by

(NP3') $|\tilde{\xi}_n - \xi| = O(\rho_n)$ a.s. as $n \rightarrow \infty$,

(NP5') $\sum_n w_n \rho_n < \infty$,

then $(\tilde{f}_n, \tilde{\xi}_n) \rightarrow (f, \xi)$ a.s. as $n \rightarrow \infty$.

Typically, for Bayes and ML estimates, the rate is $\rho_n = n^{-1/2}$. Then NP5' holds if, e.g., $w_n \sim n^{-1}$.

To illustrate the N+P and its modified version, consider the special case where $p(\cdot|\theta, \xi)$ in (4.1) is a normal density with mean θ and $\xi = \sigma^2$ is the unknown variance. That is,

$$X_{i1}, \dots, X_{ir} \stackrel{\text{i.i.d.}}{\sim} N(\theta^i, \sigma^2), \quad i = 1, \dots, n.$$

Moreover, the statistic $S_i = \bar{X}_i$ is sufficient for the mean and the density $g(\cdot|\theta, \sigma^2)$ of S_i is known. Therefore, H in (4.3) can be written as

$$(4.7) \quad H_k(s, \varphi, \psi) = \frac{g(s|\theta_k, \psi)\varphi^k}{\sum_j g(s|\theta_j, \psi)\varphi^j} - \varphi^k$$

for $k = 1, \dots, d$, where $g(s|\theta, \psi)$ is the $N(\theta, \psi/r)$ density. Even in this simple example, it is not obvious that the function H in (4.7) satisfies NP4. A proof of the following proposition is in Appendix A.3.

PROPOSITION 4.3. NP4 holds for H in (4.7).

Let Σ_0 be the Ξ_0 defined in the general setup. For the N+P, we choose $T_{\text{UBE}}(x)$ to be the sample variance of x , resulting in the recursive estimate

$$(4.8) \quad \sigma_i^2 = \frac{1}{i(r-1)} \sum_{k=1}^i \sum_{j=1}^r (X_{kj} - \bar{X}_k)^2.$$

For σ^2 , take the standard noninformative prior $p(\sigma^2) = (\sigma^2)^{-1}$. Under squared-error loss, the Bayes estimate of σ^2 based on X_1, \dots, X_i is

$$(4.9) \quad \begin{aligned} \tilde{\sigma}_i^2 &= \mathbb{E}(\sigma^2 | \mathcal{F}_i) \\ &= \frac{1}{i(r-1) - 2} \sum_{k=1}^i \sum_{j=1}^r (X_{kj} - \bar{X}_k)^2. \end{aligned}$$

Note that $|\tilde{\sigma}_n^2 - \sigma^2| = O(n^{-1/2})$ a.s. so the conclusion of Corollary 4.2 holds if $w_n \sim n^{-1}$.

The following example compares three resulting estimates for this location mixture of normals problem: when σ^2 is known, when (4.8) is used with the N+P and when (4.9) is used in the modified N+P. Convergence of the iterates holds in each case by Theorems 3.4 and 4.1 and Corollary 4.2.

EXAMPLE 4.4. Let $\Theta = \mathbb{Z} \cap [-4, 4]$ and take f to be a Bin(8, 0.5) density on Θ . Suppose $r = 10$, $n = 100$, $w_i = (i + 1)^{-1}$ and set $\sigma^2 = 1.5$. For each of 100 simulated data sets, the three estimates of f are computed using Newton's algorithm, the N+P and

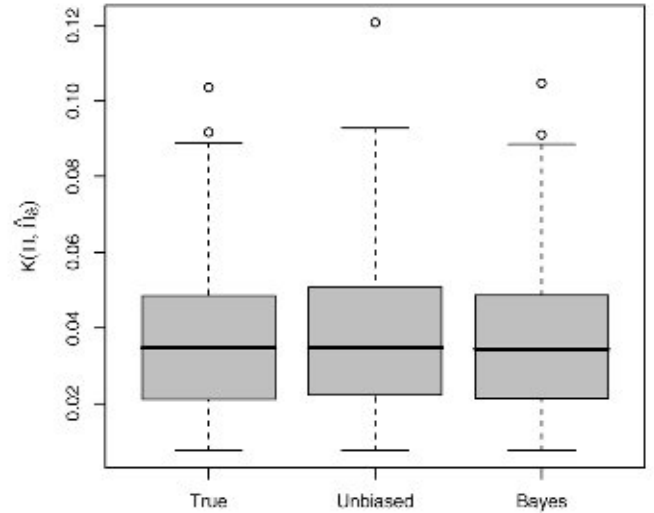


FIG. 7. Summary of KL divergences $K(\Pi, \hat{\Pi}_{\hat{\sigma}})$ for the three algorithms in Example 4.4.

the Bayes modification. Each algorithm produces estimates \hat{f} and $\hat{\sigma}^2$ with which we compute $\hat{\Pi}_{\hat{\sigma}}(s) = \sum_{j=1}^d g(s|\theta_j, \hat{\sigma}^2) \hat{f}^j$. Figure 7 summarizes the 100 KL divergences $K(\Pi, \hat{\Pi}_{\hat{\sigma}})$ for each of the three estimates. Surprisingly, little efficiency is lost when an estimate of σ^2 is used rather than the true value. Also, the N+P and the Bayes modification perform comparably, with the Bayes version performing perhaps slightly better on average. Note that no projections onto $\Sigma_0 = [10^{-4}, 10^4]$ or

$$\Delta_0 = \{\varphi \in \Delta : \varphi^k \geq 10^{-4}, k = 1, \dots, d\}$$

were necessary in this example.

5. DISCUSSION

In this paper, we have used general results in the area of stochastic approximation to prove a consistency theorem for a recursive estimate of a mixing distribution/prior in the case of a finite parameter space Θ . It is natural to wonder if this theorem can be extended to the case where f is an infinite-dimensional parameter on an uncountable space Θ . Very recently, Tokdar, Martin and Ghosh [38] have proved consistency of f_n in the infinite-dimensional case, under mild conditions. Their argument is based on the approximate martingale representation used in [14] but applied to the KL divergence $K(\Pi_f, \Pi_n)$ between the induced marginals. Again, there is a connection between their approach and the SA approach taken here, namely, $K(\Pi_f, \Pi_\varphi)$ is also a Lyapunov function for the associated ODE $\dot{\varphi} = h(\varphi)$.

In addition to convergence, there are some other interesting theoretical and practical questions to consider. First and foremost, there is the question of *rate of convergence* which, from a practical point of view, is much more important than convergence alone. We expect that, in general, the rate of convergence will depend on the support of f_0 , the weights w_n and, in the case of an uncountable Θ , the smoothness of f . Whatever the true rate of convergence might be, Example 3.7 (model II) demonstrated that this rate is unsatisfactory when the support of f is misspecified. For this reason, a modification of the algorithm that better handles such cases would be desirable.

Another question of interest goes back to the original motivation for Newton's recursive algorithm. To an orthodox Bayesian, any method which performs well should be at least *approximately Bayes*. Stemming from the fact that the recursive estimate and the nonparametric Bayes estimate, with the appropriate Dirichlet process prior, agree when $n = 1$, Newton et al. [26–28] claim that the former should serve as a suitable approximation to the latter. Our calculations in Section 3.3 disagree. In particular, we see two examples in the finite case, one where the recursive estimate is significantly better and the other where the Bayes estimate is significantly better. A new question arises: if it is not an approximation of the Dirichlet process prior Bayes estimate, *for what prior does the recursive estimate approximate the corresponding Bayes estimate?*

Finally, it should be pointed out that our approach to the finite mixture problem is somewhat less general than would be desirable. In particular, we are assuming that the support of f is within a *known* finite set of points. In general, however, what is known is that the support of f is contained in, say, a bounded interval. In this case, a set of grid points $\Theta = \{\theta_1, \dots, \theta_m\}$ are chosen to approximate the *unknown* support $\Theta^* = \{\theta_1^*, \dots, \theta_M^*\}$ of f . Newton's algorithm will produce an estimate f_n on Θ in this case, but it is impossible to directly compare f_n to f since their supports Θ and Θ^* may be entirely different. There is no problem comparing the marginals, however. This leads us to the following important conjecture, closely related to the so-called *I*-projections in [5, 35].

CONJECTURE. *Let Π_{f_n} and Π_f be the marginal densities corresponding to f_n on Θ and f on Θ^* , respectively. Then, as $n \rightarrow \infty$,*

$$K(\Pi_f, \Pi_{f_n}) \rightarrow \inf_{\varphi} K(\Pi_f, \Pi_{\varphi}) \quad \text{a.s.,}$$

where φ ranges over all densities on Θ .

Despite these unanswered practical and theoretical questions, the strong performance of Newton's algorithm and the N+P algorithm in certain cases and, more importantly, their computational cost-effectiveness, make them very attractive compared to the more expensive nonparametric Bayes estimate or the nonparametric MLE, and worthy of further investigation.

APPENDIX: PROOFS

A.1 Proof of Theorem 3.4

To prove the theorem, we need only show that the algorithm (3.5) satisfies the conditions of Theorem 2.10. First note that f_n is, for each n , a convex combination of points in the *interior* of Δ^d so no projection as in (2.8) is necessary. Second, the random variables β_n in assumption SA2 are identically zero so SA3 is trivially satisfied.

Let $\{u_n\}$ be a convergent sequence in Δ^d , where $u_n = (u_n^1, \dots, u_n^d)'$. The limit $u = (u^1, \dots, u^d)' = \lim_{n \rightarrow \infty} u_n$ also belongs to Δ so $h(u)$ is well defined. To prove that $h = (h_1, \dots, h_d)'$ is continuous, we show that $h_k(u_n) \rightarrow h_k(u)$ for each $k = 1, \dots, d$ as $n \rightarrow \infty$. Consider

$$h_k(u_n) = \int \frac{p(x|\theta_k)u_n^k}{\Pi_{u_n}(x)} \Pi_f(x) dv(x) - u_n^k.$$

The integrand $p(\cdot|\theta_k)u_n^k/\Pi_{u_n}(\cdot)$ is nonnegative and bounded *v*-a.e. for each k . Then by the bounded convergence theorem we get

$$\lim_{n \rightarrow \infty} h_k(u_n) = h_k(u), \quad k = 1, \dots, d.$$

But $\{u_n\} \subset \Delta^d$ was arbitrary so h is continuous.

Next, note that $H(x, f_n)$ is the difference of two points in Δ^d and is thus bounded independent of x and n . Then SA1 holds trivially.

Finally, we show that f is globally asymptotically stable for the ODE $\dot{\varphi} = h(\varphi)$ in Δ^d . Note that $\sum_{i=1}^d \dot{\varphi}^i = \sum_{i=1}^d h_i(\varphi) = 0$ so the trajectories lie on the connected and compact Δ^d . Let $\ell(\varphi)$ be the KL divergence, $\ell(\varphi) = \sum_{k=1}^d f^k \log(f^k/\varphi^k)$. We claim that ℓ is a strong Lyapunov function for $\dot{\varphi} = h(\varphi)$ at f . Certainly $\ell(\varphi)$ is positive definite. To check the differentiability condition, we must show that $\ell(\varphi)$ has a well-defined gradient around f , even when f is on the boundary of Δ^d . Suppose, without loss of generality, that f^1, \dots, f^s are positive, $1 \leq s \leq d$, and the remaining f^{s+1}, \dots, f^d are zero. By definition, $\ell(\varphi)$ is constant in $\varphi^{s+1}, \dots, \varphi^d$ and, therefore, the partial derivatives with respect to those φ 's are zero. Thus, for

any $1 \leq s \leq d$ and for any φ such that $\ell(\varphi) < \infty$, the gradient can be written as

$$(A.1) \quad \nabla \ell(\varphi) = -(r^1, r^2, \dots, r^d)' + r^s I'_s,$$

where $r^k = f^k/\varphi^k$ and I_s is a vector whose first s coordinates are one and last $d - s$ coordinates are zero. The key point here is that the gradient of $\ell(\varphi)$, for φ restricted to the boundary which contains f , is exactly (A.1). We can, therefore, extend the definition of $\nabla \ell(\varphi)$ continuously to the boundary if need be.

Given that $\nabla \ell(\varphi)$ exists on all of Δ^d , the time derivative of ℓ along φ is

$$(A.2) \quad \begin{aligned} \dot{\ell}(\varphi) &= \nabla \ell(\varphi)' h(\varphi) \\ &= \int \frac{\Pi_f(x) - \Pi_\varphi(x)}{\Pi_\varphi(x)} \nabla \ell(\varphi)' P_x \varphi \, d\nu(x) \\ &= 1 - \int \frac{\Pi_f}{\Pi_\varphi} \Pi_f \, d\nu. \end{aligned}$$

It remains to show that $\dot{\ell}(\varphi) = 0$ iff $\varphi = f$. Applying Jensen's inequality to $y \mapsto y^{-1}$ in (A.2) gives

$$(A.3) \quad \begin{aligned} \dot{\ell}(\varphi) &= 1 - \int_{\mathcal{X}} \left(\frac{\Pi_\varphi}{\Pi_f} \right)^{-1} \Pi_f \, d\nu \\ &\leq 1 - \left(\int_{\mathcal{X}} \frac{\Pi_\varphi}{\Pi_f} \Pi_f \, d\nu \right)^{-1} = 0, \end{aligned}$$

where equality can hold in (A.3) iff $\Pi_\varphi = \Pi_f$ ν -a.e. We assume the mixtures are identifiable, so this implies $\varphi = f$. Therefore, $\dot{\ell}(\varphi) = 0$ iff $\varphi = f$, and we have shown that ℓ is a strong Lyapunov function on Δ^d . To prove that f is a globally asymptotically stable point for $\dot{\varphi} = h(\varphi)$, suppose that $\varphi(t)$ is a solution, with $\varphi(0) = f_0$, that does not converge to f . Since ℓ is a strong Lyapunov function, the sequence $\ell(\varphi(t))$, as $t \rightarrow \infty$, is bounded, strictly decreasing and, thus, has a limit $\lambda > 0$. Then the trajectory $\varphi(t)$ must fall in the set

$$\Delta^* = \{\varphi \in \Delta^d : \lambda \leq \ell(\varphi) \leq \ell(f_0)\}$$

for all $t \geq 0$. In the case $f \in \text{int}(\Delta^d)$, $\ell(\varphi) \rightarrow \infty$ as $\varphi \rightarrow \partial\Delta$, so the set Δ^* is compact (in the relative topology). If $f \in \partial\Delta^d$, then Δ^* is not compact but, as shown above, $\dot{\ell}(\varphi)$ is well defined and continuous there. In either case, $\dot{\ell}$ is continuous and bounded away from zero on Δ^* , so

$$\sup_{\varphi \in \Delta^*} \dot{\ell}(\varphi) = -L < 0.$$

Then, for any $\tau \geq 0$, we have

$$\ell(\varphi(\tau)) = \ell(f_0) + \int_0^\tau \dot{\ell}(\varphi(s)) \, ds \leq \ell(f_0) - L\tau.$$

If $\tau > \ell(f_0)/L$, then $\ell(\varphi(\tau)) < 0$, which is a contradiction. Therefore, $\varphi(t) \rightarrow f$ for all initial conditions $\varphi(0) = f_0$, so f is globally asymptotically stable. Theorem 2.10 then implies $f_n \rightarrow f$ a.s.

A.2 Proof of Theorem 4.1

The proof of the theorem requires the following lemma, establishing a Lipschitz-type bound on the error terms β_n in (4.5). Its proof follows immediately from NP4 and the Mean Value Theorem.

LEMMA A.1. *Under the assumptions of Theorem 4.1, there exists a number $A \in (0, \infty)$ such that*

$$\|\beta_n\| \leq A\mathbb{E}(|\xi_n - \xi| \mid \mathcal{F}_{n-1}).$$

PROOF OF THEOREM 4.1. The map h in (4.4) has k th component

$$\begin{aligned} h_k(\varphi) &= \int H(x; \varphi, \xi) \Pi_{f,\xi}(s) \, d\nu^r(x) \\ &= \int \frac{\Pi_{f,\xi}(x)}{\Pi_{\varphi,\xi}(x)} p(x|\theta_k, \xi) \varphi^k \, d\nu^r(x) - \varphi^k, \end{aligned}$$

where $\Pi_{f,\xi}(x) = \sum_k \bar{p}(x|\theta_k, \xi) f^k$ is the marginal density of x and ν^r is the product measure on \mathcal{X}^r . Notice that this h , which does not depend on the estimate ξ_n , is exactly the same as the h in (3.6). Therefore, the continuity and stability properties derived in the proof of Theorem 3.4 are valid here as well. All that remains is to show that the β_n 's in (4.5) satisfy SA3 of Theorem 2.10.

By the SLLN, ξ_n belongs to Ξ_0 for large enough n so we can assume, without loss of generality, that no projection is necessary. Let $S_n = Z_1 + \dots + Z_n$, where the $Z_i = \nu^{-1}(\hat{\xi}^{(i)} - \xi)$ and ν^2 is the variance of $\hat{\xi}^{(i)}$. Then $|\xi_n - \xi| = cn^{-1}|S_n|$, where $c > 0$ is a constant independent of n . Since S_n is a sum of i.i.d. random variables with mean zero and unit variance, the Law of Iterated Logarithm states that

$$(A.4) \quad \limsup_{n \rightarrow \infty} \{|S_n|/u(n)\} = 1 \quad \text{a.s.}$$

Now, by Lemma A.1 and (A.4) we have

$$\|\beta_n\| \leq Acn^{-1}\mathbb{E}(|S_n| \mid \mathcal{F}_{n-1}) = O(n^{-1}u(n))$$

and, therefore, $\sum_n w_n \|\beta_n\|$ converges a.s. by NP5. Condition SA3 is satisfied, completing the proof. \square

A.3 Proof of Proposition 4.3

To prove that the case of a location-mixture of normals with unknown variance is covered by Theorem 4.1, we must show that the function H , defined

in (4.7), satisfies NP4, that is, that the partial derivatives $\frac{\partial}{\partial \psi} H_k(s; \varphi, \psi)$ are bounded.

PROOF OF PROPOSITION 4.3. Clearly each component H_k of H , defined in (4.7), is differentiable with respect to $\psi \in \Sigma_0$ and, after simplification,

$$\frac{\partial}{\partial \psi} H_k(s, \varphi, \psi) = \frac{\varphi^k e^{-r\theta_k^2/2\psi} e^{rs\theta_k/\psi}}{2\psi^2} \cdot \frac{\sum_j u_{kj}(s) \varphi^j e^{-r\theta_j^2/2\psi} e^{rs\theta_j/\psi}}{[\sum_j \varphi^j e^{-r\theta_j^2/2\psi} e^{rs\theta_j/\psi}]^2},$$

where (as $|s| \rightarrow \infty$)

$$(A.5) \quad u_{kj}(s) = \theta_k^2 - \theta_j^2 + 2s(\theta_j - \theta_k) = O(|s|).$$

This derivative is continuous on $s(\mathcal{X}^r) \times \Delta_0 \times \Sigma_0$ and, since Δ_0 and Σ_0 are compact, we know that

$$(A.6) \quad A_k(s) := \sup_{\varphi \in \Delta_0} \sup_{\psi \in \Sigma_0} \left| \frac{\partial}{\partial \psi} H_k(s; \varphi, \psi) \right|$$

is finite for all $s \in s(\mathcal{X}^r)$ and for all k . By the Mean Value Theorem,

$$|H_k(s; \varphi, \psi) - H_k(s; \varphi, \sigma^2)| \leq A_k(s) |\psi - \sigma^2|.$$

It remains to show that $A_k(s)$ is bounded in s . For notational simplicity, assume that φ and ψ are the values for which the suprema in (A.6) are attained. Making a change of variables $y = rs/\psi$ we can, with a slight abuse of notation, write

$$A_k(y) \leq \frac{C_k \varphi^k e^{y\theta_k} \sum_j |u_{kj}(y)| \varphi^j e^{y\theta_j}}{[\sum_j \varphi^j e^{y\theta_j}]^2}.$$

We must show that $A_k(y)$ is bounded as $|y| \rightarrow \infty$. Assume, without loss of generality, that the θ 's are arranged in ascending order: $\theta_1 < \theta_2 < \dots < \theta_d$. Factoring out, respectively, $e^{y\theta_1}$ and $e^{y\theta_d}$, we can write

$$A_k(y) \leq \frac{C_k \varphi^k e^{y(\theta_k - \theta_1)} \sum_j |u_{kj}(y)| \varphi^j e^{y(\theta_j - \theta_1)}}{(\varphi^1)^2 + \sum_{j \neq 1} \sum_{i \neq 1} \varphi^j \varphi^i e^{y(\theta_j - \theta_1) + y(\theta_i - \theta_1)},}$$

$$A_k(y) \leq \frac{C_k \varphi^k e^{y(\theta_k - \theta_d)} \sum_j |u_{kj}(y)| \varphi^j e^{y(\theta_j - \theta_d)}}{(\varphi^d)^2 + \sum_{j \neq d} \sum_{i \neq d} \varphi^j \varphi^i e^{y(\theta_j - \theta_d) + y(\theta_i - \theta_d)}.$$

Note that since $\varphi \in \Delta_0$, each φ^j is bounded away from 0. If $y \rightarrow -\infty$, then the term $e^{y(\theta_k - \theta_1)} \rightarrow 0$ dominates the numerator of the first inequality, while the denominator is bounded. Similarly, if $y \rightarrow +\infty$, then the term $e^{y(\theta_k - \theta_d)} \rightarrow 0$ dominates the numerator in the second inequality, while the denominator is bounded. For the case $k = 1$ or $k = d$, note that $|u_{11}(y)| = |u_{dd}(y)| = 0$, so the two inequalities can

still be applied and a similar argument shows A_1 and A_d are also bounded. Therefore, $A_k(y)$ is bounded for each k and the claim follows by taking A to be $\max\{\sup_y A_k(y) : 1 \leq k \leq d\}$. \square

ACKNOWLEDGMENTS

The authors thank Professors Chuanhai Liu and Surya T. Tokdar for numerous fruitful discussions, as well as Professors Jim Berger and Mike West, the Associate Editor and the two referees for their helpful comments.

REFERENCES

- [1] ALLISON, D., GADBURY, G., HEO, M., FERNÁNDEZ, J., LEE, C., PROLLA, T. and WEINDRUCH, R. (2002). A mixture model approach for the analysis of microarray gene expression data. *Comput. Statist. Data Anal.* **39** 1–20. MR1895555
- [2] ANDRIEU, C., MOULINES, E. and PRIOURET, P. (2005). Stability of stochastic approximation under verifiable conditions. *SIAM J. Control Optim.* **44** 283–312. MR2177157
- [3] BARRON, A., SCHERVISH, M. and WASSERMAN, L. (1999). The consistency of posterior distributions in nonparametric problems. *Ann. Statist.* **27** 536–561. MR1714718
- [4] CIPRA, B. (1987). Introduction to the Ising model. *Amer. Math. Monthly* **94** 937–959. MR0936054
- [5] CSISZÁR, I. (1975). I -divergence geometry of probability distributions and minimization problems. *Ann. Probab.* **3** 146–158. MR0365798
- [6] DELYON, B., LAVIELLE, M. and MOULINES, E. (1999). Convergence of a stochastic approximation version of the EM algorithm. *Ann. Statist.* **27** 98–128. MR1701103
- [7] DEMPSTER, A., LAIRD, N. and RUBIN, D. (1977). Maximum-likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B* **39** 1–38. MR0501537
- [8] EFRON, B. and TIBSHIRANI, R. (2002). Empirical Bayes methods and false discovery rates for microarrays. *Genet. Epidemiol.* **23** 70–86.
- [9] EFRON, B., TIBSHIRANI, R., STOREY, J. and TUSHER, V. (2001). Empirical Bayes analysis of a microarray experiment. *J. Amer. Statist. Assoc.* **96** 1151–1160. MR1946571
- [10] FAN, J. (1991). On the optimal rates of convergence for nonparametric deconvolution problems. *Ann. Statist.* **19** 1257–1272. MR1126324
- [11] GHOSAL, S., GHOSH, J. K. and RAMAMOORTHY, R. V. (1999). Posterior consistency of Dirichlet mixtures in density estimation. *Ann. Statist.* **27** 143–158. MR1701105
- [12] GHOSAL, S. and VAN DER VAART, A. (2001). Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities. *Ann. Statist.* **29** 1233–1263. MR1873329
- [13] GHOSH, J. K. and RAMAMOORTHY, R. V. (2003). *Bayesian Nonparametrics*. Springer, New York. MR1992245
- [14] GHOSH, J. and TOKDAR, S. (2006). Convergence and consistency of Newton's algorithm for estimating a mixing distribution. In *Frontiers in Statistics* 429–443. Imp. Coll. Press, London. MR2326012

- [15] GILKS, W., ROBERTS, G. and SAHU, S. (1998). Adaptive Markov chain Monte Carlo through regeneration. *J. Amer. Statist. Assoc.* **93** 1045–1054. MR1649199
- [16] HAARIO, H., SAKSMAN, E. and TAMMINEN, J. (2001). An adaptive Metropolis algorithm. *Bernoulli* **7** 223–242. MR1828504
- [17] ISAACSON, D. L. and MADSEN, R. W. (1976). *Markov Chains: Theory and Applications*. Wiley, New York. MR0407991
- [18] KOU, S. C., ZHOU, Q. and WONG, W. H. (2006). Equi-energy sampler with applications in statistical inference and statistical mechanics. *Ann. Statist.* **34** 1581–1619. MR2283711
- [19] KUSHNER, H. J. and YIN, G. G. (2003). *Stochastic Approximation and Recursive Algorithms and Applications*, 2nd ed. Springer, New York. MR1993642
- [20] LASALLE, J. and LEFSCHETZ, S. (1961). *Stability by Liapunov's Direct Method with Applications*. Academic Press, New York. MR0132876
- [21] LIANG, F., LIU, C. and CARROLL, R. J. (2007). Stochastic approximation in Monte Carlo computation. *J. Amer. Statist. Assoc.* **102** 305–320. MR2345544
- [22] LINDSAY, B. (1995). *Mixture Models: Theory, Geometry and Applications*. IMS, Haywood, CA.
- [23] LIU, J. S. (1996). Nonparametric hierarchical Bayes via sequential imputations. *Ann. Statist.* **24** 911–930. MR1401830
- [24] MCLACHLAN, G., BEAN, R. and PEEL, D. (2002). A mixture model-based approach to the clustering of microarray expression data. *Bioinformatics* **18** 413–422.
- [25] NEVEL'SON, M. B. and HAS'MINSKII, R. Z. (1973). *Stochastic Approximation and Recursive Estimation*. Amer. Math. Soc., Providence, RI. Translations of Mathematical Monographs, Vol. 47. MR0423714
- [26] NEWTON, M. A. (2002). A nonparametric recursive estimator of the mixing distribution. *Sankhyā Ser. A* **64** 306–322. MR1981761
- [27] NEWTON, M. A., QUINTANA, F. A. and ZHANG, Y. (1998). Nonparametric Bayes methods using predictive updating. In *Practical Nonparametric and Semiparametric Bayesian Statistics* (D. Dey, P. Muller and D. Sinha, eds.) 45–61. Springer, New York. MR1630075
- [28] NEWTON, M. A. and ZHANG, Y. (1999). A recursive algorithm for nonparametric analysis with missing data. *Biometrika* **86** 15–26. MR1688068
- [29] QUINTANA, F. A. and NEWTON, M. A. (2000). Computational aspects of nonparametric Bayesian analysis with applications to the modeling of multiple binary sequences. *J. Comput. Graph. Statist.* **9** 711–737. MR1821814
- [30] ROBBINS, H. (1964). The empirical Bayes approach to statistical decision problems. *Ann. Math. Statist.* **35** 1–20. MR0163407
- [31] ROBBINS, H. and MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statist.* **22** 400–407. MR0042668
- [32] ROBERT, C. P. and CASELLA, G. (2004). *Monte Carlo Statistical Methods*, 2nd ed. Springer, New York. MR2080278
- [33] SAN MARTIN, E. and QUINTANA, F. (2002). Consistency and identifiability revisited. *Braz. J. Probab. Stat.* **16** 99–106. MR1962497
- [34] SCOTT, J. G. and BERGER, J. O. (2006). An exploration of aspects of Bayesian multiple testing. *J. Statist. Plann. Inference* **136** 2144–2162. MR2235051
- [35] SHYAMALKUMAR, N. (1996). Cyclic I_0 projections and its applications in statistics. Technical Report 96-24, Dept. Statistics, Purdue Univ., West Lafayette, IN.
- [36] TANG, Y., GHOSAL, S. and ROY, A. (2007). Nonparametric Bayesian estimation of positive false discovery rates. *Biometrics* **63** 1126–1134.
- [37] TEICHER, H. (1963). Identifiability of finite mixtures. *Ann. Math. Statist.* **34** 1265–1269. MR0155376
- [38] TOKDAR, S. T., MARTIN, R. and GHOSH, J. K. (2008). Consistency of a recursive estimate of mixing distributions. *Ann. Statist.* To appear.
- [39] WEI, G. C. G. and TANNER, M. A. (1990). A Monte Carlo implementation of the EM algorithm and the Poor Man's data augmentation algorithm. *J. Amer. Statist. Assoc.* **85** 699–704.