



A Survey of Nonparametric Mixing Density Estimation via the Predictive Recursion Algorithm

Ryan Martin

North Carolina State University, Raleigh, USA

Abstract

Nonparametric estimation of a mixing density based on observations from the corresponding mixture is a challenging statistical problem. This paper surveys the literature on a fast, recursive estimator based on the *predictive recursion* algorithm. After introducing the algorithm and giving a few examples, I summarize the available asymptotic convergence theory, describe an important semiparametric extension, and highlight two interesting applications. I conclude with a discussion of several recent developments in this area and some open problems.

AMS (2000) subject classification. Primary 62G07; Secondary 62C12.
Keywords and phrases. Empirical Bayes, high-dimensional inference, Jayanta K. Ghosh, mixture model, recursive estimation.

1 Introduction

Estimating a mixing distribution based on samples from a mixture is arguably one of the most difficult statistical problems. It boils down to estimating the distribution of a variable based on only indirect or noise-corrupted observations. Nonparametric density estimation is already sufficiently challenging when one has direct observations let alone with only indirect observations. But understanding this latent variable distribution has many important practical consequences so, despite the problem's difficulty, there are now a number of different methods available for estimating that distribution. Here I will focus on a particular method, known as *predictive recursion* (PR), that provides a fast and easy-to-compute nonparametric estimate of a mixing density.

The work on computation for Bayesian nonparametrics—in particular, for the Dirichlet process mixture model—in the late 1990s and early 2000s

provided the original impetus for the development of PR. At that time, Markov chain Monte Carlo (MCMC) for fitting Dirichlet process mixture models was an active area of research, e.g., Escobar (1994), Escobar and West (1995), MacEachern (1994) & MacEachern (1998), MacEachern and Müller (1998), and Neal (2000), but computational power then was nowhere close to what it is now, so there was also an interest in developing alternatives to MCMC which were faster and easier in some sense. At that time, Michael Newton and collaborators, in a series of papers (Newton et al., 1998; Newton and Zhang, 1999; Newton, 2002), developed the predictive recursion algorithm which aimed at providing a fast, MCMC-free approximation of the posterior mean of the mixing distribution under a Dirichlet process mixture model. There was no doubt that the algorithm was fast and produced high-quality estimates in real- and simulated-data examples, but by the mid-2000s it was still unclear what specifically the PR algorithm was doing and what kind of properties the resulting PR estimator had. Jayanta K. Ghosh, or JKG for short, learned of the challenging open questions surrounding the PR algorithm and, naturally, was intrigued. In 2005, he and his then student, Surya Tokdar, published the first fully rigorous investigation of the convergence properties of the PR estimator (Ghosh and Tokdar, 2006). Around that time, I was a PhD student at Purdue University looking for an advisor and a research project. JKG generously shared with me a number of very promising ideas, but the one that stuck—and eventually became the topic of my thesis (Martin, 2009)—was a deeper theoretical and practical investigation into the rather elusive PR algorithm.

Between 2007 and 2012, JKG, Surya, and I were actively working on theory for and methodology based on PR. The three of us eventually shifted our respective research foci to other things, but the developments continued. In particular, James Scott and his collaborators found that PR is a powerful tool for handling the massive data and associated large-scale multiple testing problems arising in real-world applications. I have also recently started working on some new PR-adjacent projects and those results shed light on the PR algorithm itself. More on these efforts below.

Sadly, on September 30th, 2017, JKG passed away, leaving a gaping hole in the scientific community that had once been overflowing with kindness and ingenuity. Aside from his tremendous scholarly impact, JKG also touched the lives of many in a personal way. I had the privilege of participating in several special JKG memorial conference sessions and I was moved by the many fond memories of JKG shared by the participants.¹ To me, JKG was

¹Anirban DasGupta’s “Remembering Professor Jayanta K. Ghosh” is an absolute must-read; see <http://www.stat.purdue.edu/news/2017/jayanta-ghosh.html>.

the epitome of a scientist: his research efforts were fueled by nothing other than an intense curiosity about the world, and his generosity as a teacher and mentor stemmed from an equally intense desire to share all that he knew.

At face value, the goal of this paper is to review the PR algorithm, its theoretical properties, applications, and various extensions. In particular, after a review of mixture models in Section 2, I proceed in Section 3 to define the PR estimator, give some illustrative examples, and summarize its theoretical convergence properties. An important extension of PR is presented in Section 4, one that sets the scene for the applications described in Section 5. At a higher level, however, the goal of this paper is to highlight an interesting albeit lesser-known area of statistics in which JKG had a major influence. With this in mind, I present some recent developments and open problems in Sections 6 and 7, respectively, in hopes of stimulating new research activity in this area and furthering JKG's legacy. Section 8 gives some concluding remarks.

2 Background on Mixture Models

Consider independent and identically distributed (iid) data Y_1, \dots, Y_n with common density function given by the mixture model

$$f(y) = f_p(y) = \int_{\mathbb{U}} k(y | u) p(u) \nu(du), \quad y \in \mathbb{Y} \subseteq \mathbb{R}^d. \quad (2.1)$$

Here $k(\cdot | u)$ is (for now) a fully known kernel, i.e., a density function with respect to, say, Lebesgue measure on \mathbb{Y} for each $u \in \mathbb{U}$, and p is an unknown density with respect to the given measure ν on \mathbb{U} . The goal is estimation of the *mixing* density p based on iid data Y_1, \dots, Y_n from the *mixture* density f . I will assume throughout that p is identifiable, but this is non-trivial; see Teicher (1961, 1963) and San Martin and Quintana (2002). Deconvolution is a special case, where $k(y | u) = k(y - u)$, and special techniques are available for this problem (Stefanski and Carroll, 1990; Zhang, 1990, 1995; Fan, 1991). Here I will focus on methods for general mixture models.

There are a number of approaches to this problem. One is to give p some additional structure, for example, to express p as a discrete distribution. This makes f in Eq. 2.1 a *finite mixture model* and producing maximum likelihood estimates (MLEs) of the parameters that characterize p , namely, the mixture weights and locations, can be readily found via, say, the EM algorithm (Dempster et al., 1977). One can alternatively give a prior distribution for the mixture weights and locations and then use, say, an EM-like

data-augmentation strategy (e.g., van Dyk and Meng, 2001) to sample from the posterior distribution and perform Bayesian inference.

This approach, unfortunately, has some drawbacks. In particular, the methods above can only be easily employed when the number of mixture components is known, which is an unrealistic assumption. One can use model selection techniques to choose the number of components as part of a likelihood-based analysis (e.g., Leroux, 1992). Similarly, the Bayesian can put a prior distribution on the number of mixture components (e.g., Richardson and Green, 1997). Ideally, one could let the data automatically choose the number of components, and there are nonparametric methods that can handle this. Neither the nonparametric MLE (e.g., Lindsay, 1995; Laird, 1978) nor the Dirichlet process mixture model (e.g., Müller and Quintana, 2004; Ghosal, 2010) require the user to choose the number of mixture components. In fact, JKG frequently worked with Dirichlet process mixture models; see Ghosal et al. (1999) and Ghosh and Ramamoorthi (2003).

What makes estimation of p difficult is that there are many different p for which the corresponding mixture closely approximates the empirical distribution of Y_1, \dots, Y_n . That is, even if p is identifiable, it is “just barely so.” Since the above methods are primarily focused on finding a p such that the mixture Eq. 2.1 fits the data well, there is no guarantee that the resulting \hat{p} is a good estimate of p . In fact, the nonparametric MLE is *discrete* almost surely (Lindsay, 1995, Theorem 21), and the posterior mean of p under a Dirichlet process mixture model also has some discrete-like features (e.g., Tokdar et al., 2009, Figs. 1–2). Therefore, if p is assumed to be a smooth density, then a discrete estimator would clearly be unsatisfactory. Smoothing of, say, the nonparametric MLE has been considered, but I will not discuss this here; see Eggermont and LaRiccia (1995). One could also consider maximizing a penalized likelihood, one that encourages smoothness (Liu et al., 2009; Madrid-Padilla et al., 2018), but the computations are non-trivial.

The mixture Eq. 2.1 and the desire to estimate the mixing density manifests naturally when the model is expressed hierarchically. That is, if unobservable latent variables U_1, \dots, U_n are iid p and the conditional distribution of Y_i , given $U_i = u$, is $k(y | u)$, then the marginal distribution of Y_i has a density of the form Eq. 2.1. Often, the latent variables are the relevant quantities, e.g., measures of students’ “ability,” so estimating p would be of immediate practical interest. This is a hopeless endeavor with only a few indirect observations from p , but, in the early 2000s, DNA microarray technologies changed this. As Efron (2003) explains, this technology created a plethora of real-life problems where the individual Y_i carries minimal information about

its corresponding U_i , but the collection (Y_1, \dots, Y_n) carries a lot of information about p . One way to take advantage of this information is to model U_1, \dots, U_n as *exchangeable* rather than iid, which amounts to assuming some “similarity” across cases. This similarity suggests that it may be beneficial to share information and, mathematically, the exchangeability assumption results in inference about U_i that depend on all the data, not just on Y_i . This type of “borrowing strength” (e.g., Ghosh et al., 2006, p. 257) was a central theme that emerged in much of JKG’s later work, including Bogdan et al. (2011) & Bogdan et al. (2008), Dutta et al. (2012), and Datta and Ghosh (2013). An attractive alternative to a full hierarchical model, one that retains its “borrowing strength” feature, is an *empirical Bayes* solution, à la Robbins (1956, 1964, 1983), where the data is used to estimate p .

3 Predictive Recursion

3.1. Algorithm The methods described above are all likelihood-based, i.e., either the likelihood is optimized to produce an estimator or the likelihood is used to update a prior via Bayes’s theorem, leading to a posterior distribution. The predictive recursion (PR) algorithm, on the other hand, is not likelihood-based, at least not in its formulation. Instead, PR processes the data points one at a time, using the following fast recursive update.

PREDICTIVE RECURSION ALGORITHM. Initialize the algorithm with a guess p_0 of the mixing density and a sequence $\{w_i : i \geq 1\} \subset (0, 1)$ of weights. Given the data sequence Y_1, \dots, Y_n from the mixture model Eq. 2.1, evaluate

$$p_i(u) = (1 - w_i)p_{i-1}(u) + w_i \frac{k(Y_i | u)p_{i-1}(u)}{f_{i-1}(Y_i)}, \quad i = 1, \dots, n, \quad (3.1)$$

where $f_{i-1}(y) = \int k(y | u)p_{i-1}(u)\nu(du)$ is the mixture corresponding to p_{i-1} . Return p_n and $f_n = f_{p_n}$ as the final estimates.

Motivation for the PR algorithm, as described in Newton et al. (1998), came from the simple and well-known formula for the posterior mean of p , under a Dirichlet process mixture model, based on a single observation. That is, if the mixing distribution is assigned a Dirichlet process prior, with precision parameter $\alpha > 0$ and base measure with density p_0 , then the posterior mean has density

$$\frac{\alpha}{\alpha + 1} p_0(u) + \frac{1}{\alpha + 1} \frac{k(Y_1 | u)p_0(u)}{f_0(Y_1)},$$

which corresponds to the PR update with $w_i = (\alpha + i)^{-1}$. Therefore, PR is exactly the Dirichlet process mixture model posterior mean when $n = 1$; I

refer to this as the *one-step correspondence*. For $n \geq 1$, Newton’s proposal is simply to apply the one-step correspondence in each iteration. That is, the PR algorithm treats the output from the previous iteration like a prior in the next, and the update is just a weighted average of the “prior” and its corresponding posterior based on a single data point. This is an intuitively very reasonable idea, easy to implement, and fast to compute.

Next are several quick observations about the PR algorithm.

- PR can estimate a density with respect to any user-specified dominating measure. That is, if p_0 is a density with respect to ν , then so is p_n for all n . Contrast this to the discrete nonparametric MLE and the “rough” (e.g., Figure 1, Tokdar et al., 2009) Dirichlet process mixture posterior mean. Having control of the dominating measure gives PR some advantages in certain applications; see Section 5.
- The weight sequence (w_i) affects PR’s practical performance. Theory in Section 3.3 gives some guidance about the choice of weights, and examples usually take $w_i = (c + i)^{-\gamma}$ for some constants $c > 0$ and $\gamma \in (\frac{1}{2}, 1]$.
- The PR algorithm takes the form of *stochastic approximation* (Robbins and Monro, 1951), which is designed for root-finding when function evaluations are subject to error. This connection between the two recursive algorithms, fleshed out in Martin and Ghosh (2008), throws light on how the PR algorithm works. Convergence properties for PR can be derived from general results for stochastic approximation (e.g., Martin 2012), but this is so far limited to finite mixture cases.
- One potentially concerning feature of PR is that the final estimate depends on the order in which the data sequence is processed. In other words, at least in the iid case, p_n is not a function of the sufficient statistic and, therefore, is not a Bayes estimate. This order dependence is relatively weak even for moderate n , and can be effectively eliminated by averaging over permutations of the data sequence, resulting in a Rao–Blackwellized version of the original PR estimator (Tokdar et al., 2009). In my experience, averaging over only a relatively small number of permutations, say, 25, is needed for the order-dependence to be negligible; more permutations will not adversely affect the estimate, but it will not lead to any substantial improvements either. An approach that leverages PR’s order-dependence for uncertainty quantification is discussed in Dixit and Martin (2019).

- The only non-trivial computation involved in the PR algorithm is the evaluation of the normalizing constant, $f_{i-1}(Y_i)$, for each $i = 1, \dots, n$ in Eq. 3.1. For one- and two-dimensional integrals, this can be done fast and accurately with quadrature; the higher-dimensional case is discussed in Section 7. Therefore, for such mixtures, PR's computational complexity is $O(n)$ compared to the fully Bayesian methods that require processing the entire data sequence in every MCMC iteration. PR's low computational cost is what allows for averaging over multiple data permutations in just a matter of seconds. And it is also possible to further reduce PR's computation time by processing the different permutations in parallel.

3.2. Illustrations

3.2.1. *Poisson Mixture.* Example 1.2 in Böhning (2000) presents data Y_1, \dots, Y_n on the number of illness spells for $n = 602$ pre-school children in Thailand over a two-week period. The relatively large number of children—120 in total—with no illness spells makes these data *zero-inflated* and, therefore, a Poisson model is not appropriate. This suggests a Poisson mixture model and here I will fit such a model, nonparametrically, using PR.

In the mixture model formulation, $k(y | u)$ denotes a Poisson mass function with rate u , and U_i represents, say, a latent “healthiness” index for the i^{th} child. Panel a in Fig. 1 shows the PR estimate of this density based on a $\text{Unif}(0, 25)$ initial guess, weights as described above with $\gamma = 0.67$, and 25 random permutations of the data sequence. This calculation took just over 1 second on my laptop computer running R without parallelization. The relatively high concentration near 0 is consistent with the zero-inflation seen in the data. Also shown in this panel is the nonparametric MLE, a discrete distribution, as presented in Wang (2007, Table 1). Note that the bump in the PR estimate around $y = 3$ is consistent with the large mass assigned near $u = 3$ by the nonparametric MLE. But while the estimated mixing distributions are dramatically different, the two corresponding mixture distributions in Panel (b) look very similar and both provide a good fit to the data, even the zero-inflation. Interestingly, the likelihood ratio of PR versus the nonparametric MLE is 0.98, very close to 1. Therefore, within the class of mixing *densities*, there is little room to improve upon the PR estimator in terms of its quality of fit to the data; see, also, Chae et al. (2018) and Section 6.2 below.

3.2.2. *Gaussian Mixture.* Gaussian mixture models, where $k(y | u)$ is a normal density with mean u and variance either fixed or estimated from data,

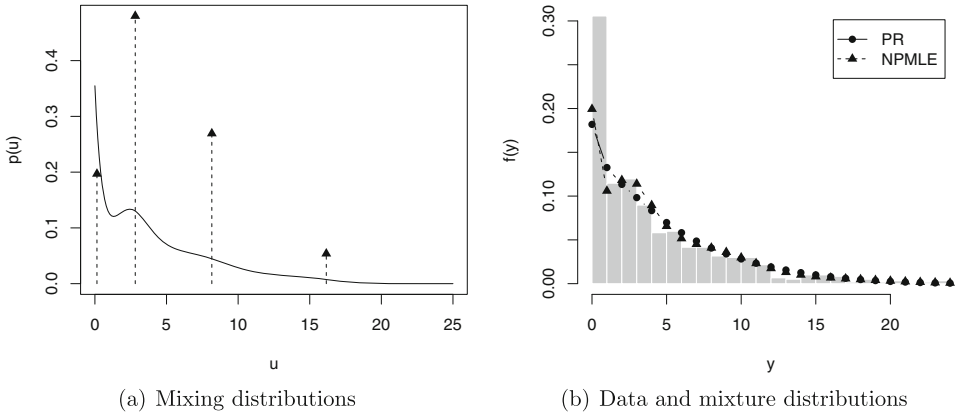


Figure 1: Plots of the estimated mixing and mixture distributions—based on PR and nonparametric maximum likelihood—for the Poisson mixture example in Section 3.2.1

are widely used models for density estimation, clustering, etc. Following Roeder (1990) and many others, I will consider data on the velocities (in thousands of km/sec) of $n = 82$ galaxies moving away from Earth.

Figure 2 shows the data histogram along with the PR estimates of the mixing and mixture distributions, in Panels a and b, respectively. Here the PR algorithm uses a kernel with standard deviation set at $\sigma = 1$; the initial guess is $\text{Unif}(5, 40)$ and the weights and permutation averaging is as in the previous example. The mixing density identifies four well-separated modes, but these are arguably not separated enough since the mixture appears to be a bit too smooth. This is likely due to fixing the kernel scale parameter at $\sigma = 1$. The PR formulation can be extended naturally to semiparametric mixtures—see Section 4—and, here, I use this generalization to simultaneously estimate p and the scale parameter σ . The estimate in this case is $\hat{\sigma} = 0.82$ and, as expected, the estimated mixing density has sharper peaks, leading to a less smooth and arguably better estimate of the mixture density.

3.2.3. Binomial Mixture and Empirical Bayes. In basketball, shots made from long distance count for 50% more points than those from shorter distance. These *three-point shots* can have a substantial effect on the outcome of a game, so three-point shooting performance strongly influences teams' offensive and defensive strategies. I downloaded data from www.nba.com that lists the three-point shots made, Y_i , and attempted, N_i , for all $n = 427$ NBA players in the last 10 games of the 2017–2018 season. To study three-point shooting performance, I take $Y_i \sim \text{Bin}(N_i, U_i)$, independent, where

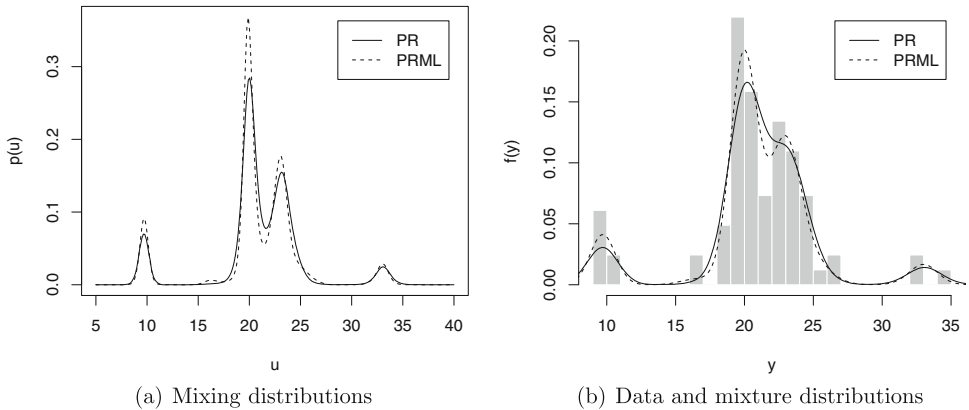


Figure 2: Plots of the estimated mixing and mixture distributions—based on PR and its semiparametric extension, PRML, described in Section 4—for the Gaussian mixture example in Section 3.2.2

N_i is treated like a fixed covariate and U_i represents the latent three-point shooting ability of player $i = 1, \dots, n$ during that crucial series of games at season’s end. Here I want to estimate the latent ability density, p , as part of an “empirical Bayesketball” analysis like in Brown (2008) and elsewhere for hitting in baseball.

The solid black line in Fig. 3 shows the PR estimate of the prior density p based on a $\text{Unif}(0, 1)$ initial guess and weights and permutation averaging as in the previous examples. This is unimodal, with mode 0.36, and concentrates about all its mass in the interval $(0.2, 0.6)$. The other lines in the plot show the corresponding empirical Bayes posterior densities for three selected players, namely, LeBron James, Jarret Allen, and Nikola Vucevic, whose proportion of three-point shots made for this series of games was 19/52, 2/3, and 1/18, respectively. James’s proportion is very close to the estimated prior mode and his number of attempts is high, so his estimated posterior is a more-concentrated version of the prior. Allen’s proportion of makes is high compared to the prior mode, but the number of attempts is low, hence strong shrinkage towards the prior mode. Finally, Vucevic’s proportion is very low but based on a moderate number of attempts, so only a moderate amount of shrinkage towards the prior mode.

3.3. Theoretical Properties Since the PR output p_n is neither a maximum likelihood nor a Bayesian estimator, its convergence properties do not follow immediately from the standard asymptotic theory, so something different is needed. Ghosh and Tokdar (2006) gave the first rigorous results on

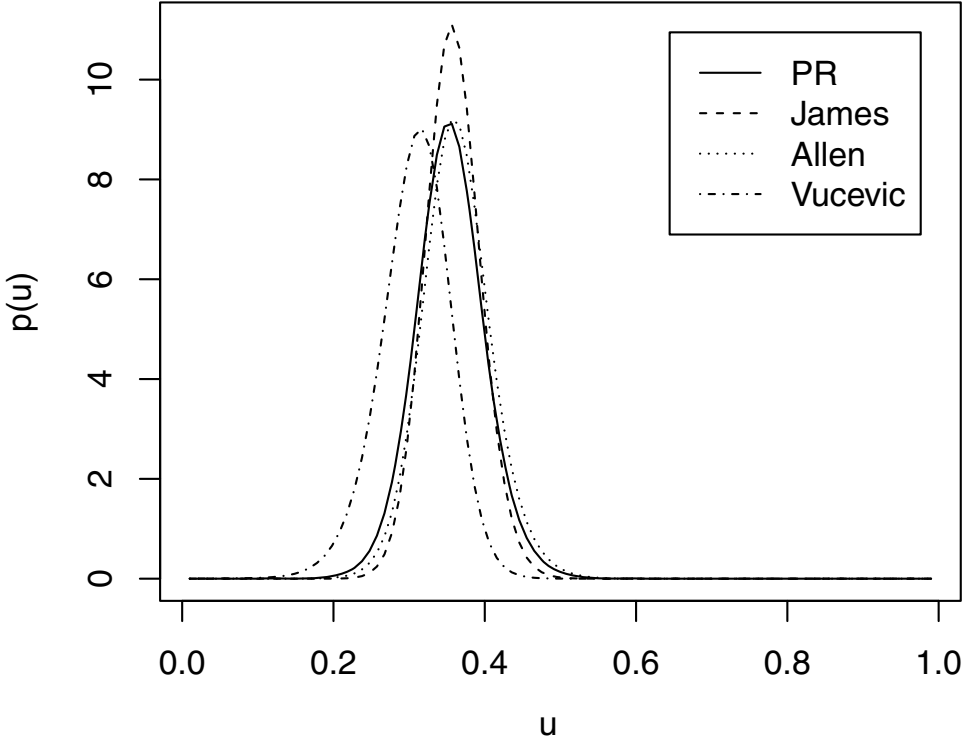


Figure 3: Plot of the PR estimate of the prior and the corresponding posterior for three selected NBA players in the three-point shooting example of Section 3.2.3

PR convergence, using martingale techniques, which were later extended in Tokdar et al. (2009) and again in Martin and Tokdar (2009).

As before, let Y_1, \dots, Y_n be iid samples from a density f^* , but allow the possibility that the posited mixture model is *misspecified*, that is, the common marginal density f^* may not have a mixture representation as in Eq. 2.1. In this misspecified case, since there may not be a “true” mixing density, it is not entirely clear what it means for the PR estimator to converge. The best one could hope for is that the PR estimate, f_n , of the marginal density would converge to the “best possible” mixture of the specified form Eq. 2.1. More specifically, if K denotes the Kullback–Leibler divergence, then, ideally, $K(f^*, f_n)$ would converge to $\inf_f K(f^*, f)$, where the infimum is over the set of mixtures in Eq. 2.1 for the given kernel, etc. Conditions under which the infimum is attained for a mixture $f^\dagger = f_{p^\dagger}$, with corresponding mixing

density p^\dagger , are given in, e.g., Martin and Tokdar (2009, Lemma 3.1) and Kleijn and van der Vaart (2006, Lemma 3.1); recall that I assume the mixture model is identifiable, so this p^\dagger is unique. Of course, if the mixture model is well-specified, then $f^\dagger = f^*$ and p^\dagger equals the true mixing density, p^* .

Naturally, the PR convergence theorem requires some assumptions. There are two sets of conditions, one on the posited mixture model and the other on the PR algorithm's inputs. I briefly summarize each in turn.

- For the mixture model, more general results are available, but here I will assume that the mixing densities are all fully supported on a compact set \mathbb{U} . I will also assume that the kernel is such that $u \mapsto k(y | u)$ is bounded and continuous for almost all y . Finally, certain integrability of density ratios is needed in the proof, so it will be assumed that

$$\sup_{u_1, u_2 \in \mathbb{U}} \int \left\{ \frac{k(y | u_1)}{k(y | u_2)} \right\}^2 f^*(y) dy < \infty. \quad (3.2)$$

This is a strong condition, but, since \mathbb{U} is assumed to be compact, it holds if $k(y | u)$ is an exponential family and f^* has Gaussian-like tails.

- For the PR algorithm's inputs, namely, the initial guess p_0 and the weight sequence (w_i) , the assumptions are quite mild. First, it is necessary that the support of p_0 contain that of p^\dagger . If the compact support \mathbb{U} is known, then this is trivially satisfied. Second, the weights must satisfy

$$\sum_{i=1}^{\infty} w_i = \infty \quad \text{and} \quad \sum_{i=1}^{\infty} w_i^2 < \infty. \quad (3.3)$$

The suggested class of weights, $w_i = (c + i)^{-\gamma}$, for $\gamma \in (\frac{1}{2}, 1]$ satisfy this.

The following theorem summarizes what is currently known about the convergence properties of the PR estimators p_n and f_n . A version of the consistency result below, in the well-specified case, is also presented in Section 5.4 of Ghosal and van der Vaart (2017).

THEOREM 1. *Assume that Y_1, Y_2, \dots are iid samples from density f^* and that the aforementioned conditions are met. Set $K_n = K(f^*, f_n) - \inf_p K(f^*, f_p)$.*

1. *Then $K_n \rightarrow 0$ almost surely.*
2. *If $\sum_n a_n w_n^2 < \infty$, where $a_n = \sum_{i=1}^n w_i$, then $a_n K_n \rightarrow 0$ almost surely.*

3. If the kernel is tight in the sense of Martin and Tokdar (2009, Condition A6) then p_n converges weakly to p^\dagger almost surely.

An interesting by-product of the proof of Theorem 1 creates an asymptotic link between PR and the nonparametric MLE. That is, the PR estimator, p_n , is converging to a solution P^\dagger , which may or may not have a density, such that

$$\int \frac{k(y | u)}{f_{P^\dagger}(y)} f^*(y) dy = 1 \quad \text{for } P^\dagger\text{-almost all } u.$$

But according to Lindsay (1995, p. 115), the nonparametric maximum likelihood estimator \hat{P} is characterized as a solution to

$$\frac{1}{n} \sum_{i=1}^n \frac{k(Y_i | u)}{f_{\hat{P}}(Y_i)} = 1 \quad \text{for } \hat{P}\text{-almost all } u.$$

When n is large, the above average is approximately equal to the expectation with respect to f^* , hence a link between the PR algorithm's target and the nonparametric MLE.

The first and third claims in Theorem 1 establish consistency of the PR estimates. The compactness condition eluded to in the third claim holds for all the standard kernels so it imposes no practical constraints.

The second claim in Theorem 1 gives a bound on the PR rate of convergence. That condition is satisfied for $w_i = (c+i)^{-\gamma}$ for $\gamma \in (\frac{2}{3}, 1]$, and gives a corresponding Kullback–Leibler convergence rate for f_n of about $n^{-1/3}$. Unfortunately, this leaves something to be desired. For example, Ghosal and van der Vaart (2001) showed that, with a Gaussian kernel and a Dirichlet process prior on the mixing distribution, the Bayes posterior concentrates around a true Gaussian mixture at nearly a n^{-1} rate in Kullback–Leibler divergence. But the PR rate above makes no assumptions about the true density f^* so it is interesting to understand the nature of that rate. Martin and Tokdar (2009) showed that PR's $n^{-1/3}$ rate is “minimax” in the sense that it is the rate PR attains when $f^*(y) = k(y | u^*)$ for some fixed u^* value, the “most extreme” kind of mixture where “ p^* ” is a point mass at u^* .

4 Semiparametric Mixture Extension

So far, I have assumed that the kernel k in the mixture model is fixed. However, there are cases in which it would make sense to allow the kernel to depend on some other parameters, say, θ , that do not get mixed over. The standard example would be to allow a Gaussian kernel to depend on

some scale parameter while being mixed over the mean; see below. That is, here I am concerned with a semiparametric mixture model where the goal is to simultaneously estimate both the mixing density p and the non-mixing structural parameter θ . For this, it turns out that the asymptotic theory for PR under model misspecification plays an important role.

Write the θ -dependent kernel as $k_\theta(y | u)$, and let $p_{i,\theta}$ denote the PR estimate of the mixing density based on Y_1, \dots, Y_i , with kernel k_θ fixed throughout. Also write $f_{i,\theta}(y) = \int k_\theta(y | u) p_{i,\theta}(u) \nu(du)$ for the corresponding mixture. Next, define a sort of “likelihood function” based on the PR output, that is,

$$L_n(\theta) = \prod_{i=1}^n f_{i-1,\theta}(Y_i). \quad (4.1)$$

Martin and Tokdar (2011) motivated this choice of likelihood by showing that $L_n(\theta)$ had features resembling that of the marginal likelihood for θ under a fully Bayesian Dirichlet process mixture model. I will refer to Eq. 4.1 as the *PR marginal likelihood*, and I proceed to estimate the structural parameter by maximizing this function.

For a quick example, consider a kernel $k_\theta(y | u) = \mathbf{N}(y | u, \theta^2)$. The likelihood function in Eq. 4.1 can be readily evaluated and maximized numerically to simultaneously estimate p and θ . This approach was carried out in the galaxy data example of Section 3.2.2 and the additional flexibility of being able to estimate the kernel scale parameter via PR marginal likelihood optimization resulted in an estimated mixture density that fit the data histogram better compared to that from the original PR.

Maximizing $L_n(\theta)$ is equivalent to minimizing $n^{-1} \sum_{i=1}^n \log\{f^*(Y_i)/f_{i-1,\theta}(Y_i)\}$, and it follows from Theorem 1 that this latter function converges pointwise, as $n \rightarrow \infty$, to $\inf_p K(f^*, f_{p,\theta})$, where the infimum is over all mixing densities. Therefore, at least intuitively, one would expect that

$$\hat{\theta} \rightarrow \arg \min_{\theta} \left\{ \inf_p K(f^*, f_{p,\theta}) \right\}, \quad n \rightarrow \infty. \quad (4.2)$$

It turns out, however, that this consistency property is quite difficult to demonstrate in general; see Section 7. But numerical results in Martin and Tokdar (2011) and elsewhere suggest that Eq. 4.2 does hold and, moreover, so does asymptotic normality.

5 Applications

There are a number of applications of the PR algorithm and its semiparametric extension in the literature. See Tao et al. (1999), Newton and Zhang

(1999), the example in Newton (2002) based on the genetics application in Newton et al. (2001), Todem and Williams (2009), and the very recent work by Woody and Scott (2018) on valid Bayesian post-selection inference. Here I only highlight two specific applications, one in large-scale significance testing, an area in which JKG worked, and one in robust regression.

5.1. Large-Scale Significance Testing In the hierarchical model formulation at the end of Section 2, consider a large collection U_1, \dots, U_n of latent variables where case i is said to be “null” if $U_i = 0$ and “non-null” otherwise. An example is DNA microarray experiments where the cases correspond to genes and “null” means that the gene is not differentially expressed. Of course, only noisy measurements Y_1, \dots, Y_n of U_1, \dots, U_n are available, so the goal is to test the sequence of hypotheses, $H_{0i} : U_i = 0$ versus $H_{1i} : U_i \neq 0$, $i = 1, \dots, n$. What makes this an interesting statistical problem is that n is large and most of the cases are null, e.g., most genes are not associated with a particular phenotype, so it is beneficial to share information across cases. Brad Efron wrote extensively on empirical Bayes solutions this problem in the early 2000s (e.g., Efron, 2010), and here I will summarize a PR-based implementation of Efron’s approach presented in Martin and Tokdar (2012). Recent extensions of this proposal to handle covariates and certain spatial dependence are presented in Scott et al. (2015) and Tansey et al. (2018), respectively.

Efron (2008) describes the *two-groups model* where Y_1, \dots, Y_n are assumed to have a common density function of the form

$$f(y) = \pi f^{(0)}(y) + (1 - \pi) f^{(1)}(y), \quad (5.1)$$

where $f^{(0)}$ and $f^{(1)}$ correspond to the densities under null and non-null settings, respectively, and π represents the proportion of null cases. He argues that, basically without generality, one can take $f^{(0)}(y) = \mathbf{N}(y \mid \mu, \sigma^2)$, but perhaps with parameters (μ, σ^2) that need to be estimated, i.e., an *empirical null* (Efron, 2004). Assuming, for the moment, that all the pieces in Eq. 5.1 are known, one can show that the Bayes test of H_{0i} would reject if $\text{fdr}(Y_i) \leq c$, for $c = 0.1$, say, where fdr —the *local false discovery rate*—is given by

$$\text{fdr}(y) = \pi f^{(0)}(y) / f(y). \quad (5.2)$$

Efron’s insight was that, since n is large, nonparametric estimation of the marginal density is straightforward and, likewise, since most of the cases are null, (π, μ, σ^2) could also be estimated. Plugging these estimates into the expression Eq. 5.2 and carrying out the sequence of tests with the corresponding estimate of fdr is Efron’s empirical Bayes solution. Details can be

found in, Efron (2004), and alternative strategies are given in Jin and Cai (2007), Muralidharan (2010), Jin et al. (2010), and Jeng et al. (2018).

An advantage of Efron’s approach is that it is apparently not necessary to directly model the possibly complicated non-null density $f^{(1)}$. However, it is possible that the independent estimates of f and $f^{(0)}$ are incompatible in the sense that, e.g., $\hat{\pi}\hat{f}^{(0)}(y) > \hat{f}(y)$ for some y . To avoid such issues, a model for all the ingredients in Eq. 5.1—one that is sufficiently flexible in $f^{(1)}$ —is needed. Toward this, Martin and Tokdar (2012) embed Eq. 5.1 into the general mixture formulation Eq. 2.1 by taking the dominating measure

$$\nu(du) = \delta_0(du) + \lambda_{[-1,1]}(du),$$

a point-mass at 0 plus Lebesgue measure on $[-1, 1]$, and kernel

$$k_\theta(y | u) = \mathbf{N}(y | \mu + \tau\sigma u, \sigma^2), \quad \theta = (\mu, \tau, \sigma).$$

With these choices, the mixture in Eq. 2.1 takes the form

$$f(y) = \pi\mathbf{N}(y | \mu, \sigma^2) + (1 - \pi) \int_{-1}^1 \mathbf{N}(y | \mu + \tau\sigma u, \sigma^2) p(u) du, \quad (5.3)$$

which can immediately be identified as a model-based version of Eq. 5.1. Intuitively, the non-null cases, which correspond to “signals,” should tend to be larger magnitude, so it makes sense that $f^{(1)}$ have heavier tails than $f^{(0)}$. The normal location mixture in Eq. 5.3 can achieve this, and the parameter τ controls roughly how much heavier the normal the tails need to be. Since the PR algorithm respects the specified dominating measure, the combined discrete-continuous form of the mixing distribution can be handled easily, and (a minor modification of) the semiparametric extension of PR in Section 4 can be applied to fit the model in Eq. 5.3 and define the corresponding empirical Bayes testing procedure based on the plug-in estimate of fdr.

For illustration, I consider data from the study in van’t Wout et al. (2003) that compares the genetic profiles of four healthy and four HIV-positive patients. The goal is to determine which, if any, of the $n = 7680$ genes are differentially expressed between the two groups. This example is described in Efron (2010, Section 6.1D). Figure 4 shows the results of the PR model fit; in particular, $\hat{\mu} = 0.07$, $\hat{\sigma} = 0.74$, and $\hat{\pi} = 0.88$. The estimated f clearly fits the data histogram, which is wide enough to leave room for the normal $f^{(0)}$ and the heavier-tailed bimodal estimate of $f^{(1)}$. The inverted scale shows the estimated fdr and the “ $\widehat{\text{fdr}} \leq 0.1$ ” cutoffs are also show. Finally, the plot indicates that 121 genes are identified by the test as differentially expressed,

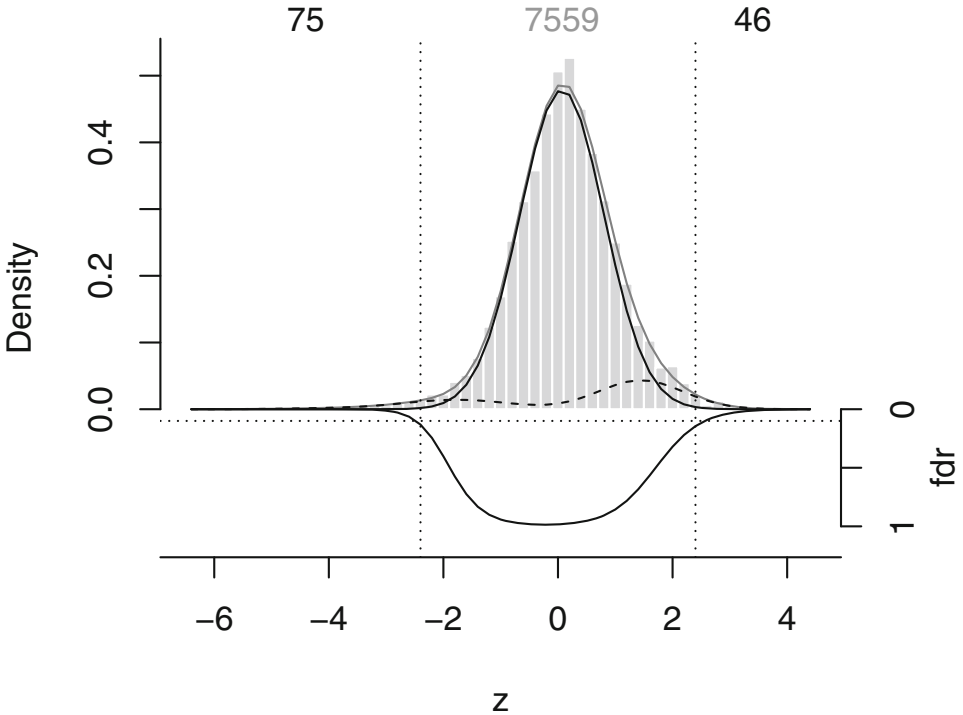


Figure 4: Histogram of the z-scores for the HIV data in van’t Wout et al. (2003), along with estimates of f , $\pi f^{(0)}$, and $(1 - \pi)f^{(1)}$ overlaid. The inverted scale shows the estimated fdr and the $\widehat{\text{fdr}} \leq 0.1$ cutoffs for the test

46 are up- and 75 are down-regulated. The conclusions here are similar to those obtained by Efron, but this is not always the case; cf. Martin and Tokdar (2012).

5.2. *Robust Regression* Consider a linear regression model where

$$y_i = x_i^\top \beta + \varepsilon_i, \quad i = 1, \dots, n,$$

where y_i is a real-valued response, x_i is a d -vector of predictor variables, β is a d -vector of regression coefficients, and ε_i are measurement errors, assumed to be iid. Quantification of uncertainty about estimates or predictions in this setting requires specification of a distribution for the errors. A standard choice is to assume the errors are normal, leading to simple closed-form expressions for the MLEs with straightforward sampling distribution properties. However, if the normal error assumption is questionable, e.g., if there are “outliers,” then the MLEs will suffer. Therefore, it is of interest

to develop procedures that can handle different error distribution assumptions, especially those with heavier-than-normal tails. It is indeed possible to introduce a heavy-tailed error distribution, such as Student-t with small degrees of freedom, and work out the corresponding MLEs and their properties, but this is still an assumption that may not be appropriate for the given problem. A more flexible, nonparametric choice of error distribution would be desirable. Motivated by the fact that scale mixtures of normals produce heavy-tailed distributions, Martin and Han (2016) introduce a mixture model formulation and propose to estimate both the mixing density and β using the semiparametric extension of PR described in Section 4. Here I briefly summarize their approach.

With a slight abuse of my previous notation, let me write f for the density function of the measurement errors, $\varepsilon_1, \dots, \varepsilon_n$. Expressing f as the mixture

$$f(\varepsilon) = \int_0^\infty \mathbf{N}(\varepsilon \mid 0, u^2) p(u) du,$$

for some unknown mixing density p , is one way to induce a flexible, heavy-tailed distribution for the errors. For any fixed β , by writing $\varepsilon_i = y_i - x_i^\top \beta$, the PR algorithm can be used to estimate the mixing and mixture densities, p and f , respectively. Of course, those estimates would depend on β so, like in Section 4, I could define a marginal likelihood in β to be maximized, leading to a simultaneous estimate of β and p . Optimization of this marginal likelihood is non-trivial, but Martin and Han (2016) propose a hybrid PR–EM algorithm wherein they introduce latent variables U_i from the mixing distribution to make the “complete-data” likelihood of a simple Gaussian form. Details are in their paper and an R code implementation is available at my website. Pastpipatkul et al. (2017) used a similar PR–EM strategy in a time series application.

As an example, I consider data on mathematics proficiency presented in Table 11.4 of Kutner et al. (2005). The response variable, y , is the students’ average mathematics proficiency exam score for 37 U.S. states, the District of Columbia, Guam, and the Virgin Islands; hence, $n = 40$. The predictor variable, x , is the percentage of students in each state with at least three types of reading materials at home. This is an interesting example because D.C. and Virgin Islands are outliers in y and Guam is an outlier in both x and y . The general trend suggests a quadratic model,

$$y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \varepsilon_i, \quad i = 1, \dots, n,$$

and the plot in Fig. 5 shows the data and the results of three fits of the above model, namely, ordinary least squares, Huber’s robust least squares,

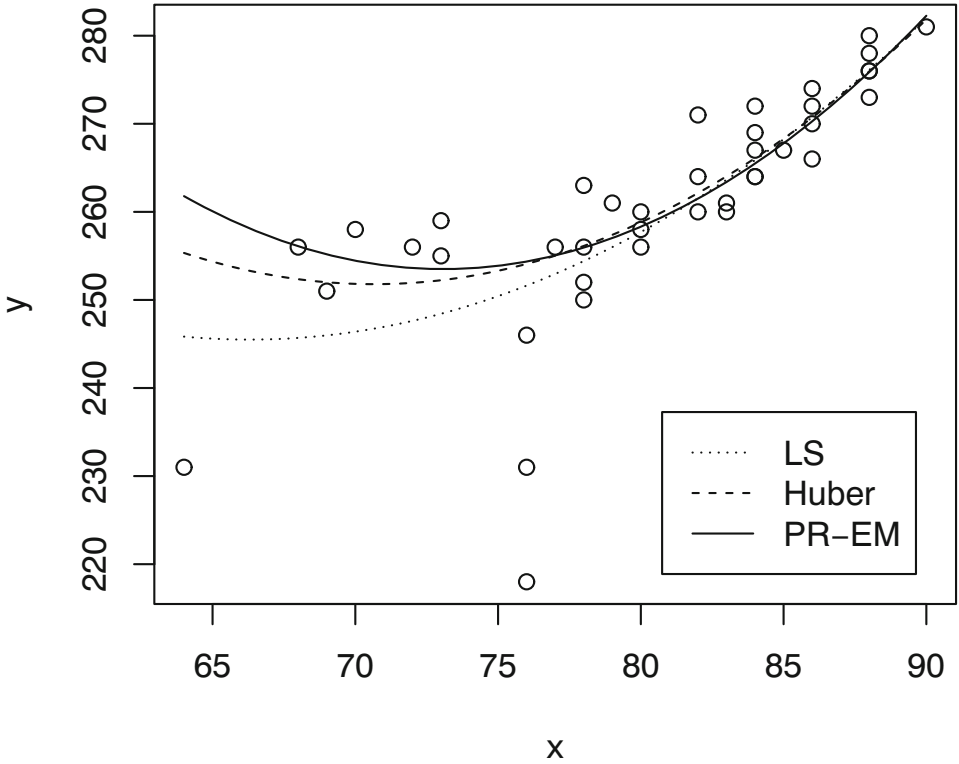


Figure 5: Scatter plot of the mathematics proficiency score data and the three different quadratic model fits

and PR-EM. Here the former two methods are both more influenced by the outliers than the PR-EM method, suggesting that the latter puts lesser weight on those extremes in the model fit.

6 Recent Developments

6.1. PR for the Mixture The PR algorithm is designed for estimating the mixing distribution but, of course, it is at least conceptually straightforward to produce a corresponding estimate for the mixture distribution. However, the PR algorithm requires numerical integration at each iteration, which itself requires that the mixing density support be known, compact, and no more than two dimensions. If the sole purpose of the mixture model was to facilitate density estimation, as is often the case, then the above requirements are a hindrance. It is, therefore, natural to ask if it is possible to formulate a numerical integration-free version of the PR algorithm directly

on the mixture density. Hahn et al. (2018) happened upon an affirmative answer to this question while investigating a seemingly unrelated updating property of Bayesian predictive distributions.

Given a prior density $p_0(u)$ and a kernel $k(y | u)$, let f_0 be the prior predictive density, with corresponding distribution function F_0 . For a sequence of data Y_1, Y_2, \dots , let f_i denote the Bayesian posterior predictive distribution for Y_{i+1} , given Y_1, \dots, Y_i . Then Hahn et al. (2018) showed that there exists a sequence of bivariate copula densities c_i such that $f_i(y) = c_i(F_{i-1}(y), F_{i-1}(Y_i))f_{i-1}(y)$. That is, if this sequence of copula densities were known, then one could recursively update the Bayesian predictive distribution without any posterior sampling, MCMC, etc. For simple Bayesian models, the closed-form expressions for the copula densities can be derived, but not in general.

Indeed, for a Dirichlet process mixture model, only the first in the sequence of copula densities can be derived in closed-form. This suggests following Newton's strategy, capitalizing on the one-step exactness of the recursive update, to derive a new algorithm. The specific proposal in Hahn et al. (2018) is the update

$$f_n(y) = (1 - w_n)f_{n-1}(y) + w_n g_\rho(F_{n-1}(y), F_{n-1}(Y_n)) f_{n-1}(y),$$

where g_ρ is the Gaussian copula density with correlation parameter ρ . Those authors show that this algorithm is fast to compute and provides accurate estimate in finite-sample simulation experiments. They also prove consistency under tail conditions on the true density. It would be interesting to investigate convergence rates and to extend this method to handle multivariate and dependent data sequences.

6.2. A Variation on PR A potentially troubling feature of the PR algorithm is its dependence on the order of the data sequence. Averaging over permutations reduces this dependence, but is not a fully satisfactory fix. Therefore, other similar algorithms might be of interest.

One that has made an appearance in numerous places across the literature, but has yet to be systematically studied, is as follows. Start by assuming that the mixture density f in Eq. 2.1 is *known*. Then the algorithm

$$p_t(u) = \int \frac{k(y | u)f(y)}{\int k(y | v)p_{t-1}(v) dv} dy, \quad t \geq 1, \quad (6.1)$$

will converge to a solution of the inverse problem defined in Eq. 2.1; see Chae et al. (2019). In a statistical context, where f is unknown but data Y_1, \dots, Y_n is available, there are a number of ways one can modify the above algorithm.

One is to replace f in Eq. 6.1 with the empirical distribution of Y_1, \dots, Y_n . This produces a smooth mixing density estimate at every finite t , but Chae et al. (2018) show that it converges, as $t \rightarrow \infty$, to the discrete nonparametric MLE. It would be interesting to determine a stopping criterion such that the corresponding estimator could be called a smooth nonparametric *near-MLE*. Alternatively, one can pick any suitable density estimate \hat{f} and plug in to Eq. 6.1. Numerical results indicate that this procedure will produce high-quality estimates of the mixing density, but its theoretical properties are still under investigation.

7 Open Problems

PROBLEM 1. This one goes all the way back to Newton’s original development. What is PR doing? Is there any precise sense in which PR, or its permutation-averaged version, gives an approximation to the Dirichlet process mixture Bayes estimator? Newton et al. (1998) showed that the connection is exact for $n = 1$ and they also investigated the case of $n = 2$. In particular, for two observations, Y_1 and Y_2 , they show that both PR and the Dirichlet process mixture posterior mean take the form

$$a_0 p_0(u) + a_1 p_0(u | Y_1) + a_2 p_0(u | Y_2) + a_{12} p_0(u | Y_1, Y_2),$$

where $p_0(u | Y_i) \propto k(Y_i | u) p_0(u)$ and

$$p_0(u | Y_1, Y_2) \propto k(Y_1 | u) k(Y_2 | u) p_0(u),$$

the only difference being in the coefficients a_1, a_2, a_{12} . It can also be shown that the permutation-averaged PR estimator is of the same form, again with different coefficients, but I will not list these here. For the general n case, if one imagines averaging the PR expression in Proposition 12 of Ghosh and Tokdar (2006) over different permutations of the data sequence, one can vaguely see something reminiscent of the Dirichlet process mixture posterior mean expression given in Lo (1984). So it seems like something interesting could be there, but the details have eluded me so far.

PROBLEM 2. Existing implementations of PR have used numerical integration to evaluate the normalizing constant at each iteration. So even though the theory puts no restriction on the dimension of the latent variable space, the reliance on quadrature methods makes it difficult to handle mixture over more than one or two dimensions. Is it possible to use Monte Carlo methods to compute this integral? A strategy that works with a fixed set of particles with weights that are updated at each iteration seems particularly promising, but these weights would need to be monitored carefully.

PROBLEM 3. A bound on the convergence rate of the PR estimator was stated in Theorem 1, but I noted that this bound is conservative in the sense that it seems to be attained when the mixing distribution is a point mass, not a smooth density. So a relevant question is how one could incorporate smoothness assumptions about the true density to improve upon this rate?

PROBLEM 4. The PR algorithm is naturally sequential and would be ideal in cases where the data ordering matters, e.g., dependent data problems. However, currently nothing is known about PR in such cases; in fact, even how to define the PR algorithm in such cases is not clear. A suggestion is made in Ghosal and Roy (2009) but, to my knowledge, there are no results in this direction currently available.

PROBLEM 5. For the PR-based estimate of the structural parameter θ described in Section 4, currently very little is known about its theoretical properties. I indicated there that simulation experiments suggest an asymptotic normality result holds, but this has yet to be rigorously demonstrated. In classical iid problems, the log-likelihood is additive and the central limit theorem can be used after linearization. For the PR likelihood, however, the i^{th} term depends—in a complicated way—on all of Y_1, \dots, Y_i . Martingale laws of large numbers and central limit theorems seem promising but, unfortunately, no progress has been made along these lines yet.

PROBLEM 6. I mentioned a few high-dimensional empirical Bayes applications here in this review and, for these problems, I always felt that there should, at least in some cases, be a theoretical benefit to plugging in a smooth estimate of the prior density compared to, say, a discrete estimate like in Jiang and Zhang (2009). Unfortunately, I have not been able to identify a theoretical benefit, but I still believe that one exists.

8 Conclusion

In this paper, I have reviewed the work on theory and applications of the PR algorithm for estimating mixing distributions, along the way highlighting some new developments and some open problems. This is only one of the many areas that JKG had an impact so, naturally, my review here made connections to a number of adjacent topics on which JKG worked, including Bayesian nonparametrics, density estimation, and high-dimensional inference. That these are still “hot topics” in statistics research is surely no coincidence, it is a testament to JKG’s incredible foresight and influence. I was so tremendously lucky to have had the opportunity to know and to work with JKG, and it is an honor to dedicate this work to him.

Acknowledgments. The author is grateful to the guest editors of this special issue of *Sankhya, Series B* honoring JKG, and to the anonymous reviewers for their helpful feedback. This work is partially supported by the National Science Foundation, DMS-1737929.

References

- BOGDAN, M., GHOSH, J.K. and TOKDAR, S.T. (2008). *A comparison of the Benjamini-Hochberg procedure with some Bayesian rules for multiple testing*. IMS, Beachwood, BALAKRISHNAN, N., PEÑA, E. and SILVAPULLE, M. (eds.), p. 211–230.
- BOGDAN, M., CHAKRABARTI, A., FROMMLET, F. and GHOSH, J.K. (2011). Asymptotic Bayes-optimality under sparsity of some multiple testing procedures. *Ann. Statist.* **39**, 1551–1579.
- BÖHNING, D. (2000). *Computer-assisted Analysis of Mixtures and Applications: Meta-analysis, Disease Mapping and Others*. Chapman and Hall–CRC, Boca Raton.
- BROWN, L. (2008). In-season prediction of batting averages: a field test of empirical Bayes and Bayes methodologies. *Ann. Appl. Stat.* **2**, 113–152.
- CHAE, M., MARTIN, R. and WALKER, S.G. (2018). Convergence of an iterative algorithm to the nonparametric MLE of a mixing distribution. *Statist. Probab. Lett.* **140**, 142–146.
- CHAE, M., MARTIN, R. and WALKER, S.G. (2019). On an algorithm for solving Fredholm integrals of the first kind. *Stat. Comput.* **29**, 645–654.
- DATTA, J. and GHOSH, J.K. (2013). Asymptotic properties of Bayes risk for the horseshoe prior. *Bayesian Anal.* **8**, 111–131.
- DEMPSTER, A., LAIRD, N. and RUBIN, D. (1977). Maximum-likelihood from incomplete data via the EM algorithm (with discussion). *J. Roy. Statist. Soc. Ser. B* **39**, 1–38.
- DIXIT, V. and MARTIN, R. (2019). Permutation-based uncertainty quantification about a mixing distribution. Unpublished manuscript, arXiv:1906.05349.
- DUTTA, R., BOGDAN, M. and GHOSH, J.K. (2012). Model selection and multiple testing—a Bayes and empirical Bayes overview and some new results. *J. Indian Statist. Assoc.* **50**, 105–142.
- EFRON, B. (2003). Robbins, empirical Bayes and microarrays. *Ann. Statist.* **31**, 366–378.
- EFRON, B. (2004). Large-scale simultaneous hypothesis testing: the choice of a null hypothesis. *J. Amer. Statist. Assoc.* **99**, 96–104.
- EFRON, B. (2008). Microarrays, empirical Bayes and the two-groups model. *Statist. Sci.* **23**, 1–22.
- EFRON, B. (2010). *Large-Scale Inference Volume 1 of Institute of Mathematical Statistics Monographs*. Cambridge University Press, Cambridge.
- EGGERMONT, P.P.B. and LARICCIA, V.N. (1995). Maximum smoothed likelihood density estimation for inverse problems. *Ann. Statist.* **23**, 199–220.
- ESCOBAR, M.D. (1994). Estimating normal means with a Dirichlet process prior. *J. Amer. Statist. Assoc.* **89**, 268–277.
- ESCOBAR, M.D. and WEST, M. (1995). Bayesian density estimation and inference using mixtures. *J. Amer. Statist. Assoc.* **90**, 577–588.
- FAN, J. (1991). On the optimal rates of convergence for nonparametric deconvolution problems. *Ann. Statist.* **19**, 1257–1272.
- GHOSAL, S. (2010). *The Dirichlet process, related priors and posterior asymptotics*. Cambridge Univ. Press, Cambridge, p. 35–79.

- GHOSAL, S. and ROY, A. (2009). *Bayesian nonparametric approach to multiple testing*. World Scientific Press, Singapore, SASTRY, N.S.N., RAO, T.S.S.R.K., DELAMPADY, M. and RAJEEV, B. (eds.), p. 139–164.
- GHOSAL, S. and VAN DER VAART, A.W. (2001). Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities. *Ann. Statist.* **29**, 1233–1263.
- GHOSAL, S. and VAN DER VAART, A. (2017). *Fundamentals of Nonparametric Bayesian Inference Volume 44 of Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.
- GHOSH, J.K. and RAMAMOORTHY, R.V. (2003). *Bayesian Nonparametrics*. Springer, New York.
- GHOSH, J.K. and TOKDAR, S.T. (2006). *Convergence and consistency of Newton's algorithm for estimating mixing distribution*. Imp. Coll. Press, London, FAN, J. and KOUL, H. (eds.), p. 429–443.
- GHOSAL, S., GHOSH, J.K. and RAMAMOORTHY, R.V. (1999). Posterior consistency of Dirichlet mixtures in density estimation. *Ann. Statist.* **27**, 143–158.
- GHOSH, J.K., DELAMPADY, M. and SAMANTA, T. (2006). *An Introduction to Bayesian Analysis*. Springer, New York.
- HAHN, P.R., MARTIN, R. and WALKER, S.G. (2018). On recursive Bayesian predictive distributions. *J. Amer. Statist. Assoc.* **113**, 1085–1093.
- JENG, X.J., ZHANG, T. and TZENG, J.-Y. (2018). Efficient signal inclusion with genomic applications. *J. Amer. Statist. Assoc.*, to appear; arXiv:[1805.10570](https://arxiv.org/abs/1805.10570).
- JIANG, W. and ZHANG, C.-H. (2009). General maximum likelihood empirical Bayes estimation of normal means. *Ann. Statist.* **37**, 1647–1684.
- JIN, J. and CAI, T.T. (2007). Estimating the null and the proportional of nonnull effects in large-scale multiple comparisons. *J. Amer. Statist. Assoc.* **102**, 495–506.
- JIN, J., PENG, J. and WANG, P. (2010). A generalized Fourier approach to estimating the null parameters and proportion of nonnull effects in large-scale multiple testing. *J. Statist. Res.* **44**, 103–127.
- KLEIJN, B.J.K. and VAN DER VAART, A.W. (2006). Misspecification in infinite-dimensional Bayesian statistics. *Ann. Statist.* **34**, 837–877.
- KUTNER, M.I., NACHTSHEIM, C.J., NETER, J. and LI, W. (2005). *Applied Linear Statistical Models*, 5th edn. McGraw-Hill/Irwin.
- LAIRD, N. (1978). Nonparametric maximum likelihood estimation of a mixed distribution. *J. Amer. Statist. Assoc.* **73**, 805–811.
- LEROUX, B.G. (1992). Consistent estimation of a mixing distribution. *Ann. Statist.* **20**, 1350–1360.
- LINDSAY, B.G. (1995). *Mixture Models, Theory, Geometry and Applications*. Haywood, IMS.
- LIU, L., LEVINE, M. and ZHU, Y. (2009). A functional EM algorithm for mixing density estimation via nonparametric penalized likelihood maximization. *J. Comput. Graph. Statist.* **18**, 481–504.
- LO, A.Y. (1984). On a class of Bayesian nonparametric estimates. I. Density estimates. *Ann. Statist.* **12**, 351–357.
- MACEACHERN, S.N. (1994). Estimating normal means with a conjugate style Dirichlet process prior. *Comm. Statist. Simulation Comput.* **23**, 727–741.
- MACEACHERN, S.N. (1998). *Computational methods for mixture of Dirichlet process models*. Springer, New York, DEY, D., MÜLLER, P. and SINHA, D. (eds.), p. 23–43.
- MACEACHERN, S. and MÜLLER, P. (1998). Estimating mixture of Dirichlet process models. *J. Comput. Graph. Statist.* **7**, 223–238.

- MADRID-PADILLA, O.-H., POLSON, N.G. and SCOTT, J. (2018). A deconvolution path for mixtures. *Electron. J. Stat.* **12**, 1717–1751.
- MARTIN, R. (2009). Fast Nonparametric Estimation of a Mixing Distribution with Application to High-Dimensional Inference. PhD thesis, Purdue University Department of Statistics, West Lafayette, IN.
- MARTIN, R. (2012). Convergence rate for predictive recursion estimation of finite mixtures. *Statist. Probab. Lett.* **82**, 378–384.
- MARTIN, R. and GHOSH, J.K. (2008). Stochastic approximation and Newton’s estimate of a mixing distribution. *Statist. Sci.* **23**, 365–382.
- MARTIN, R. and HAN, Z. (2016). A semiparametric scale-mixture regression model and predictive recursion maximum likelihood. *Comput. Statist. Data Anal.* **94**, 75–85.
- MARTIN, R. and TOKDAR, S.T. (2009). Asymptotic properties of predictive recursion: robustness and rate of convergence. *Electron. J. Stat.* **3**, 1455–1472.
- MARTIN, R. and TOKDAR, S.T. (2011). Semiparametric inference in mixture models with predictive recursion marginal likelihood. *Biometrika* **98**, 567–582.
- MARTIN, R. and TOKDAR, S.T. (2012). A nonparametric empirical Bayes framework for large-scale multiple testing. *Biostatistics* **13**, 427–439.
- MÜLLER, P. and QUINTANA, F.A. (2004). Nonparametric Bayesian data analysis. *Statist. Sci.* **19**, 95–110.
- MURALIDHARAN, O. (2010). An empirical Bayes mixture method for effect size and false discovery rate estimation. *Ann. Appl. Statist.* **4**, 422–438.
- NEAL, R.M. (2000). Markov chain sampling methods for Dirichlet process mixture models. *J. Comput. Graph. Statist.* **9**, 249–265.
- NEWTON, M.A. (2002). On a nonparametric recursive estimator of the mixing distribution. *Sankhyā Ser. A* **64**, 306–322.
- NEWTON, M.A. and ZHANG, Y. (1999). A recursive algorithm for nonparametric analysis with missing data. *Biometrika* **86**, 15–26.
- NEWTON, M.A., QUINTANA, F.A. and ZHANG, Y. (1998). *Nonparametric Bayes methods using predictive updating*. Springer, New York, DEY, D., MÜLLER, P. and SINHA, D. (eds.), p. 45–61.
- NEWTON, M., KENDZIORSKI, C., RICHMOND, C., BLATTNER, F. and TSUI, K. (2001). On differential variability of expression ratios: Improving statistical inference about gene expression changes from microarray data. *J. Comput. Biology* **8**, 37–52.
- PASTPIPATKUL, P., YAMAKA, W. and SRIBOONCHITTA, S. (2017). *Predictive recursion maximum likelihood of threshold autoregressive model*. Springer, KREINOVICH, V., SRIBOONCHITTA, S. and HUYNH, V.-N. (eds.), p. 349–362.
- RICHARDSON, S. and GREEN, P.J. (1997). On Bayesian analysis of mixtures with an unknown number of components. *J. Roy. Statist. Soc. Ser. B* **59**, 731–792.
- ROBBINS, H. (1956). *An empirical Bayes approach to statistics, I*. University of California Press, Berkeley, p. 157–163.
- ROBBINS, H. (1964). The empirical Bayes approach to statistical decision problems. *Ann. Math. Statist.* **35**, 1–20.
- ROBBINS, H. (1983). Some thoughts on empirical Bayes estimation. *Ann. Statist.* **11**, 713–723.
- ROBBINS, H. and MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statistics* **22**, 400–407.
- ROEDER, K. (1990). Density estimation with confidence sets exemplified by superclusters and voids in the galaxies. *J. Amer. Statist. Assoc.* **411**, 617–624.

- SAN MARTIN, E. and QUINTANA, F. (2002). Consistency and identifiability revisited. *Braz. J. Probab. Stat.* **16**, 99–106.
- SCOTT, J.G., KELLY, R.C., SMITH, M.A., ZHOU, P. and KASS, R.E. (2015). False discovery rate regression: An application to neural synchrony detection in primary visual cortex. *J. Amer. Statist. Assoc.* **110**, 459–471.
- STEFANSKI, L. and CARROLL, R.J. (1990). Deconvoluting kernel density estimators. *Statistics* **21**, 169–184.
- TANSEY, W., OLUWASANMI, K., POLDRACK, R.A. and SCOTT, J.G. (2018). False discovery rate smoothing. *J. Amer. Statist. Assoc.* **113**, 1156–1171.
- TAO, H., PALTA, M., YANDELL, B.S. and NEWTON, M.A. (1999). An estimation method for the semiparametric mixed effects model. *Biometrics* **55**, 102–110.
- TEICHER, H. (1961). Identifiability of mixtures. *Ann. Math. Statist.* **32**, 244–248.
- TEICHER, H. (1963). Identifiability of finite mixtures. *Ann. Math. Statist.* **34**, 1265–1269.
- TODEM, D. and WILLIAMS, K.P. (2009). A hierarchical model for binary data with dependence between the design and outcome success probabilities. *Stat. Med.* **28**, 2967–2988.
- TOKDAR, S.T., MARTIN, R. and GHOSH, J.K. (2009). Consistency of a recursive estimate of mixing distributions. *Ann. Statist.* **37**, 2502–2522.
- VAN DYK, D.A. and MENG, X.-L. (2001). The art of data augmentation. *J. Comput. Graph. Statist.* **10**, 1, 1–111. With discussions, and a rejoinder by the authors.
- VAN’T WOUT, A., LEHRMA, G., MIKHEEVA, S., O’KEEFE, G., KATZE, M., BUMGARNER, R., GEISS, G. and MULLINS, J. (2003). Cellular gene expression upon human immunodeficiency virus type 1 injection of cd4+T-Cell lines. *J. Virol.* **77**, 1392–1402.
- WANG, Y. (2007). On fast computation of the non-parametric maximum likelihood estimate of a mixing distribution. *J. R. Stat. Soc. Ser. B* **69**, 185–198.
- WOODY, S. and SCOTT, J.G. (2018). Optimal post-selection inference for sparse signals: a nonparametric empirical-Bayes approach. Unpublished manuscript, arXiv:[1810.11042](https://arxiv.org/abs/1810.11042).
- ZHANG, C.-H. (1990). Fourier methods for estimating mixing densities and distributions. *Ann. Statist.* **18**, 806–831.
- ZHANG, C.-H. (1995). On estimating mixing densities in discrete exponential family models. *Ann. Statist.* **23**, 929–945.

Publisher’s Note. Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

RYAN MARTIN
DEPARTMENT OF STATISTICS,
NORTH CAROLINA STATE UNIVERSITY,
RALEIGH, NC, USA
E-mail: rgmarti3@ncsu.edu