

A deconvolution path for mixtures

Oscar-Hernan Madrid-Padilla

*Department of Statistics
University of California, Berkeley
e-mail: omadrid@berkeley.edu*

Nicholas G. Polson

*Booth School of Business
University of Chicago
e-mail: ngp@chicagobooth.edu*

and

James Scott

*Department of Information, Risk, and Operations Management
and Department of Statistics and Data Sciences
University of Texas at Austin
e-mail: james.scott@mcombs.utexas.edu*

Abstract: We propose a class of estimators for deconvolution in mixture models based on a simple two-step “bin-and-smooth” procedure applied to histogram counts. The method is both statistically and computationally efficient: by exploiting recent advances in convex optimization, we are able to provide a full deconvolution path that shows the estimate for the mixing distribution across a range of plausible degrees of smoothness, at far less cost than a full Bayesian analysis. This enables practitioners to conduct a sensitivity analysis with minimal effort. This is especially important for applied data analysis, given the ill-posed nature of the deconvolution problem. Our results establish the favorable theoretical properties of our estimator and show that it offers state-of-the-art performance when compared to benchmark methods across a range of scenarios.

MSC 2010 subject classifications: Primary 62G05; secondary 62G07.

Keywords and phrases: Deconvolution, mixture models, penalized likelihood, empirical Bayes, sensitivity analysis.

Received May 2017.

1. Deconvolution in mixture models

Suppose that we observe $\mathbf{y} = (y_1, \dots, y_n)$ from the model

$$y_i \mid \mu_i \sim \phi(y_i \mid \mu_i), \quad \mu_i \stackrel{i.i.d.}{\sim} f_0, \quad (1)$$

where $\phi(\cdot \mid \mu)$ is a known distribution with location parameter μ , and f_0 is an unknown mixing distribution. Marginally, we have specified a mixture model for y_i :

$$m(y_i) = \int_{\mathbb{R}} \phi(y_i - \mu_i) f_0(\mu_i) d\mu_i = (\phi * f_0)(y_i). \quad (2)$$

The problem of estimating the mixing distribution f_0 is commonly referred to as *deconvolution*: we observe draws from the convolution $m = \phi * f_0$, rather than from f_0 directly, and we wish to invert this blur operation to recover the distribution of the latent location parameters. Models of this form have been used in a wide variety of applications and have attracted significant attention in the literature (e.g. Kiefer and Wolfowitz, 1956; Ferguson, 1973; Fan, 1991; Newton, 2002; Ghosal and Van Der Vaart, 2001). Yet the estimation of f_0 continues to pose both theoretical and practical challenges, making it an active area of statistical research (e.g. Delaigle and Hall, 2014; Efron, 2016).

In this paper, we propose a nonparametric method for deconvolution that is both statistically and computationally efficient. Our method can be motivated in terms of an underlying Bayesian model incorporating a prior into model (1), but it does not involve full Bayes analysis. Rather, we use a two-step “bin and smooth” procedure. In the “bin” step, we form a histogram of the sample, yielding the number of observations x_j that fall into the j th histogram bin. In the “smooth” step, we use the counts x_j to compute a maximum *a posteriori* (MAP) estimate of f_0 under a prior that encourages smoothness.

We show that this nonparametric empirical-Bayes procedure yields excellent performance for deconvolution, at reduced computational cost compared to full nonparametric Bayesian methods. Our main theorem establishes conditions under which the method yields a consistent estimate of the mixing distribution f_0 . We also provide simulation evidence that the method offers practical improvements over existing state-of-the-art methods.

The structure of the paper is as follows. Section 2 discusses connections of our work with previous efforts in the literature studying the problem of deconvolution. In Section 3.1 we introduce our proposed approach to deconvolution based on imposing smooth constraints on the log-space of the mixing density. We see how this novel perspective can be thought as a MAP estimator for which a Poisson surrogate model is presented in Section 3.2. We then make important distinctions between our proposal and other penalized methods in Section 3.3. The paper next describes solution algorithms for our proposed methods, model selection aspects, and a toy example in Sections 3.4, 3.5 and 3.6 respectively. Section 4.1 then demonstrates a sensitivity analysis example on a real data set illustrating the advantage of having a full path of solutions. We also include, in Section 4.2, a sensitivity analysis to the choice of the number of bins. Section 5 shows almost sure recovery of the mixing density when maximization is carried on a certain class of functions. The practical validity of our methodology is then assessed in Section 6 on both mixing density recovery and normal-means problem. For the first of these tasks, comparisons with other methods on simulated data are evaluated via Earth Mover’s Distance and mean squared error. For the normal-means problem, a similar setting is considered, but measuring performance in terms of mean squared error. Finally, Section 7 summarizes our contributions.

2. Connections with previous work

The goal of this section is to give the reader a sense of some of the historically influential work on deconvolution in the statistics literature, postponing a detailed review until Section 3.3, when we describe in more detail several approaches that have closer connections with our method. We organize our discussion into three categories: likelihood-based, Bayesian, and kernel-based methods.

Likelihood-based deconvolution methods date back at least to the seminal work by Kiefer and Wolfowitz (1956). The Kiefer–Wolfowitz estimator (KW) has some appealing features: it is completely nonparametric and invariant to translations of the data, it requires no tuning parameters, and it is consistent under fairly general conditions. Balanced against these desirable features there is, perhaps, a disadvantage when estimating a smooth true mixing density: the KW estimator is a discrete distribution involving as many as $n + 1$ point masses. This motivates the use of some prior or penalty as a regularizer to the the log-likelihood.

Our approach to deconvolution falls within a penalized likelihood framework. It has important connections with classical ideas of penalized likelihood density estimation (e.g. Good and Gaskins, 1971; Silverman, 1982), and at least in spirit, our work resembles recent deconvolution proposed methods by Lee et al. (2013), Wager (2013), and Efron (2016). A more thorough discussion between these methods and ours will be made in Section 3.3. For now we highlight the feature which distinguishes our approach from these previous methods: it uses a Poisson surrogate likelihood with total variation regularization on the logarithm of the mixing density.

Deconvolution has also been studied using Bayesian methods. In the context of repeated measurements, or multivariate deconvolution, recent work includes Sarkar et al. (2014a) and Sarkar et al. (2014b); Staudenmayer et al. (2008). Moreover, for the one dimensional density estimation problem, a flexible choice is the Dirichlet Process (DP) studied in Ferguson (1973) and Escobar and West (1995). For a Dirichlet prior in deconvolution problems, concentration rates were recently studied in Donnet et al. (2014). Related models were considered by Do et al. (2005) and Muralidharan (2010) for finite mixture of normals.

The DP provides a very general framework for estimating the mixing density f_0 . However, as Martin and Tokdar (2012) argue, fitting a Dirichlet process mixture does not scale well with the number of observations n . For microarray studies, n ranges from thousands to tens of thousands, whereas for more recent studies of fMRI data or single-nucleotide polymorphisms, n can reach several hundreds of thousands (e.g. Tansey et al., 2014). For such large data sets, fitting a DP mixture model can be very time-consuming.

To overcome this difficulty, Newton (2002), Tokdar et al. (2009), Martin and Tokdar (2011), and Martin and Tokdar (2012) studied a predictive recursive (PR) algorithm. The resulting estimator scales well with large data sets while remaining reasonably accurate, thereby solving one the main challenges faced by the fully Bayesian approach. In Section 6, we will provide detailed benchmarks between PR, DP and our approach.

Finally, we note some of the work on kernel based methods for deconvolution. This includes, among others, work by Carroll and Hall (1988); Stefanski and Carroll (1990); Zhang (1990); Fan (1991); Fan and Koo (2002); Hall et al. (2007); Carroll et al. (2012); Delaigle and Hall (2014). Roughly speaking, their idea is motivated by (1) after taking the Fourier transform of the corresponding convolution of densities, then solving for the unknown mixing density using kernel approximations for the Fourier transform of the true marginal density. The resulting kernel estimator enjoys attractive theoretical properties: for each $\mu_0 \in \mathbb{R}$, the estimator has optimal rates of convergence towards $f_0(\mu_0)$ for squared-error loss when the function f_0 belongs to a smooth class of functions (Fan, 1991). We do not intent to provide theoretical comparisons with these kernel methods, but instead we will provide comprehensive numerical comparisons in Section 6.

3. A deconvolution path

3.1. Overview of approach

We now described our proposed approach in detail. We study deconvolution estimators related to the variational problem

$$\underset{f}{\text{minimize}} \quad - \sum_{i=1}^n \log(\phi * f)(y_i) \quad \text{subject to} \quad \int_{\mathbb{R}} f(\mu) d\mu = 1, \quad J(f) \leq t, \quad (3)$$

where $J(f)$ is a known penalty functional. The choices of J we consider include ℓ_1 or ℓ_2 penalties on the derivatives in the log-space to encourage smoothness:

$$\|\log f^{(k)}\|_s^q = \int_{\mathbb{R}} |\log f^{(k)}(\mu)|^s d\mu, \quad (4)$$

with $s = q = 1$ or $s = q = 2$ and where $\log f^{(k)}$ is the derivative of order k of the log prior. The penalty involving the first derivative is an especially interpretable one, as $d \log f(\mu)/d\mu = f'(\mu)/f(\mu)$ is the score function of the mixing density.

Note that an alternative interpretation of our approach is as a MAP estimator. To see this we consider the (possibly improper) prior on the mixing density

$$p(f) \propto \exp(-J(f)) \mathbb{I}(f \in \mathcal{A}),$$

where \mathcal{A} is an appropriate class of density functions. The posterior distribution is $p(f | y) \propto p(y | f)p(f)$, and our MAP estimator therefore solves, for an appropriate $\tau > 0$,

$$\operatorname{argmin}_{f \in \mathcal{A}} - \log p(y | f) + \frac{\tau}{2} J(f). \quad (5)$$

3.2. Binned-counts problem

Throughout this section we assume that ϕ corresponds to the pdf of the standard normal distribution, although the arguments can easily be generalized to other distributions.

To make estimation efficient in scenarios with thousands or even millions of observations, we actually fit a MAP estimator based on binning the data. First, we use the samples to form a histogram $\{I_j, x_j\}_{j=1}^D$ with D bins, where I_j is the j -th interval in the histogram and $x_j = \#\{y_i \in I_j\}$ is the associated count. For ease of exposition, we assume that the intervals take the form $I_j = \xi_j \pm \Delta/2$, i.e. have midpoints ξ_j and width Δ , although this is not essential to our analysis. To arrive at a discrete version of our estimator, instead of Problem (3), we consider an approximation and a reparametrization $g = \log f$. We also put the penalty in the objective function with a regularization parameter $\tau > 0$. This leads to the following optimization problem:

$$\text{minimize}_{g} \quad -\frac{1}{n} \sum_{j=1}^D x_j \log(\phi * e^g(\xi_j)) + \frac{\tau}{2} J(e^g) \quad \text{subject to} \quad \int e^{g(\mu)} d\mu = 1. \tag{6}$$

We then approximate (6) by solving

$$\begin{aligned} \text{minimize}_{\theta \in \mathbb{R}^D} \quad & -\frac{1}{n} \sum_{j=1}^D x_j \log\left(\sum_{i=1}^D \Delta \phi(\xi_j - \xi_i) e^{\theta_i}\right) + \frac{\tau}{2} \|\Delta^{(k+1)} \theta\|_q^s \\ \text{subject to} \quad & \sum_{i=1}^D \Delta e^{\theta_i} = 1, \end{aligned} \tag{7}$$

where $s = q = 1$ or $s = q = 2$, and $\Delta^{(k+1)}$ is the k -th order discrete difference operator. Concretely, when $k = 0$, $\Delta^{(1)}$ is the $(D - 1) \times D$ matrix encoding the first differences of adjacent values:

$$\Delta^{(1)} = \begin{pmatrix} 1 & -1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & -1 & 0 & \cdots & 0 \\ \vdots & & & & \ddots & \vdots \\ 0 & \cdots & & 0 & 1 & -1 \end{pmatrix}. \tag{8}$$

For $k \geq 1$, $\Delta^{(k+1)}$ is defined recursively as $\Delta^{(k+1)} = \Delta^{(1)} \Delta^{(k)}$, where $\Delta^{(1)}$ from (8) is of the appropriate dimension. Thus when $k = 0$, we penalize the total variation of the vector θ (c.f. Rudin et al., 1992; Tibshirani et al., 2005) and should expect estimates that are shrunk towards piecewise-constant functions. When $k \geq 1$, the estimator penalizes higher-order versions of total variation, similar to the polynomial trend-filtering estimators studied by Tibshirani (2014).

Interestingly, following the proof of Theorem 1 in Padilla and Scott (2015) we find that (7) is equivalent to

$$\text{minimize}_{\theta \in \mathbb{R}^D} \quad l(\theta) + \frac{\tau}{2} \|\Delta^{(k+1)} \theta\|_q^s, \tag{9}$$

where

$$l(\theta) = \sum_{j=1}^D \{\lambda_j(\theta) - x_j \log \lambda_j(\theta)\},$$

with $\lambda_j(\theta) = \sum_{i=1}^D G_{ij} e^{\theta_i}$, $G_{ij} = \Delta \phi(\xi_j - \xi_i)$, and $\hat{\theta}$ solves (9) if only if $\hat{\theta} - \log(n \Delta) \mathbf{1}$ solves (7). Hence, in practice we solve the unconstrained optimization Problem (9). This is due to the fact that a constrained problem, as in (7), is typically more difficult than an unconstrained problem. In particular, in the next subsections we will see how (9) is amenable to efficient optimization algorithms. Moreover, as stated above, the equivalence between (7) and (9) can be seen as in Padilla and Scott (2015). However, it is worth mentioning that this is similar in spirit to Theorem 3.1 in Silverman (1982) which is in the context of a variational density estimation problem.

3.3. Other related approaches

The estimator defined in (5) belongs to a general class of MAP estimators that have been studied in Good and Gaskins (1971) and Silverman (1982) for the classical problem of density estimation. For deconvolution problems we note the recent work by Wager (2013), which penalizes the marginal density rather than the derivatives of the mixing density as we propose. Moreover, the penalty in Wager (2013) is not designed to encourage smoothness. Rather, it is an ℓ_2 projection on to the space of acceptable marginal densities (i.e. those that are valid Gaussian convolutions)—a criterion that our estimates satisfy by construction. An alternative penalized likelihood method was studied in Lee et al. (2013) in the different context where the marginal density has atoms. There the authors use the roughness penalty $J(f) = \int |f'(\mu)|^2 d\mu$. This differs from our approach to penalizing the log-mixing density, thereby ensuring that the solution will be positive. Also, we allow different degrees of smoothness depending on the choice of k . Moreover, while Lee et al. (2013) only considered sample sizes in the order of hundreds, we show in the next sections that our estimator can scale to much larger data sets while still enjoying attractive computational and statistical properties.

More recently, Efron (2016) proposed a penalized approach to deconvolution that is also based on regularization but differs from ours. This approach proceeds by assuming that the mixing density is discrete with support $\{\theta_1, \dots, \theta_N\}$. Then Efron (2016) specified a parametric model on the mixing density of the form

$$f(\theta_j) = \exp(Q_{j,\cdot}^T \alpha - c(\alpha)), \quad j = 1, \dots, N,$$

where $Q_{j,\cdot}$ is the j -th row of the matrix $Q \in \mathbb{R}^{N \times p}$, for some $p > 0$, which is used to encourage structure on the mixing density. Moreover, $c(\alpha)$ is a normalizing constant satisfying

$$c(\alpha) = \log \left(\sum_{j=1}^N \exp(Q_{j,\cdot}^T \alpha) \right).$$

Then summing over all the $\{\theta_j\}_{j=1}^N$ and considering the contribution of the

different samples, Efron (2016) arrives to the minus log-likelihood

$$l(\alpha) = - \sum_{i=1}^n \log \left(\sum_{j=1}^N N(y_i | \theta_j, 1) \exp(Q_j^T \alpha - c(\alpha)) \right).$$

The estimator from Efron (2016) results from solving

$$\underset{\alpha}{\text{minimize}} \quad l(\alpha) + c_0 \left(\sum_{h=1}^p \alpha_h^2 \right)^{1/2}, \tag{10}$$

where c_0 is either 1 or 2.

We note that the estimator (10) is possibly limited by the following features. First, the choice of the matrix Q , which Efron (2016) recommends to be a spline basis representation, might produce estimates that suffer from local-adaptivity problems. Moreover, there is no theoretical support of choosing $c_0 \in \{1, 2\}$ —although to be fair, this can be addressed by performing a somewhat more exhaustive model-selection process. In our experiments section we will present experimental comparisons between the estimator in (10) and our approach.

Finally, we emphasize that our approach is not, in any sense, related to the ridge parameter deconvolution estimator from Hall et al. (2007). This estimator does not penalize the log-likelihood. Rather is designed to avoid the need to choose a kernel function in the original kernel estimator from Fan (1991), which it accomplishes by “ridging” the integral in its definition with a positive function.

3.4. Solution algorithms

We now discuss implementation details for solving (9) in the case $s = q = 1$. To solve this problem, motivated by the work on trend filtering for regression by Ramdas and Tibshirani (2014), we rewrite the problem as

$$\underset{\theta \in \mathbb{R}^D}{\text{minimize}} \quad l(\theta) + \frac{\tau}{2} \|\Delta^{(1)}\alpha\|_1 \quad \text{subject to} \quad \alpha = \Delta^{(k)}\theta. \tag{11}$$

Next we proceed to solve (11) via the alternating-direction method of multipliers (ADMM), as in Ramdas and Tibshirani (2014). (See Boyd et al., 2011, for an overview of ADMM.) By exploiting standard results we arrive at the scaled augmented Lagrangian corresponding to the constrained problem (11):

$$L_\rho(\theta, \alpha, u) = l(\theta) + \frac{\tau}{2} \|\Delta^{(1)}\alpha\|_1 + \rho u^T (\alpha - \Delta^{(k)}\theta) + \frac{\rho}{2} \|\alpha + u - \Delta^{(k)}\theta\|_2^2.$$

This leads to the following ADMM updates at each iteration j :

$$\begin{aligned} \theta^{j+1} &\leftarrow \underset{\theta}{\text{argmin}} \left(l(\theta) + \frac{\rho}{2} \|\alpha^j + u^j - \Delta^{(k)}\theta\|_2^2 \right), \\ \alpha^{j+1} &\leftarrow \underset{\alpha}{\text{argmin}} \left(\frac{1}{2} \|\alpha - \Delta^{(k)}\theta^{j+1} + u^j\|_2^2 + \frac{\tau}{2\rho} \|\Delta^{(1)}\alpha\|_1 \right), \\ u^{j+1} &\leftarrow u^j + \alpha^{j+1} - \Delta^{(k)}\theta^{j+1}. \end{aligned} \tag{12}$$

Note that in (12) the update for θ involves solving a sub-problem whose solution is not analytically available. To deal with this, we use the well known Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm, which is very efficient because the gradient of the θ sub-problem objective is available in closed form. The update for α can be computed in linear time by appealing to the dynamic programming algorithm from Johnson (2013).

In the case $p = q = 2$, both components of the objective function in (9) have closed-form gradients; see the appendix. Thus we can solve the problem using any algorithm that can use function and gradient calls. In our experiments, we use BFGS.

3.5. Solution path and model selection

One of the major advantages of our approach is that it easily yields an entire deconvolution path, comprising a family of estimates $\hat{f}(\tau)$ over a grid of smoothness parameters. Although in principle any deconvolution algorithm can yield such a path by solving the problem for many smoothing parameters, our path is generated in a highly efficient manner, using warm starts. We initially solve (11) for a large value of τ , for which the resulting estimate is nearly constant. We then use this solution to initialize the ADMM at a slightly smaller value of τ , which dramatically reduces the computation time compared to an arbitrarily chosen starting point. We proceed iteratively until solutions have been found across a decreasing grid of τ values (which are typically spaced uniformly in $\log \tau$).

The resulting deconvolution path can be used to inspect a range of plausible estimates for f_0 , with varying degrees of smoothness. This allows the data analyst to bring to bear any further prior information (such as the expected number of modes in f_0) that was not formally incorporated into the loss function. It also enables sensitivity analysis with respect to different plausible assumptions about the smoothness of the mixing distribution. We illustrate this approach with a real-data example in Section 4.1.

However, in certain cases—for example, in our simulation studies—it is necessary to select a particular value of τ using a default rule. We now briefly describe heuristics for doing so based on ℓ_1 and ℓ_2 penalties with $k = 1$. These heuristics are used in our simulation studies. For the case of ℓ_1 regularization, motivated by Tibshirani and Taylor (2012), we consider a surrogate AIC approach by computing

$$\text{AIC}_\tau = l(\hat{\theta}_\tau) + k + 1 + \left| \left\{ i : (\Delta^{(k+1)}\hat{\theta}_\tau)_i \neq 0 \right\} \right|,$$

and choosing the value of τ that minimizes this expression. Here, $\hat{\theta}_\tau$ denotes the solution with ℓ_1 penalty and regularization parameter τ .

In the case of ℓ_2 regularization the situation is more difficult, since there is not an intuitive notion of the number of parameters of the model. Instead, we consider an ad-hoc procedure based on cross validation. This solves the problem

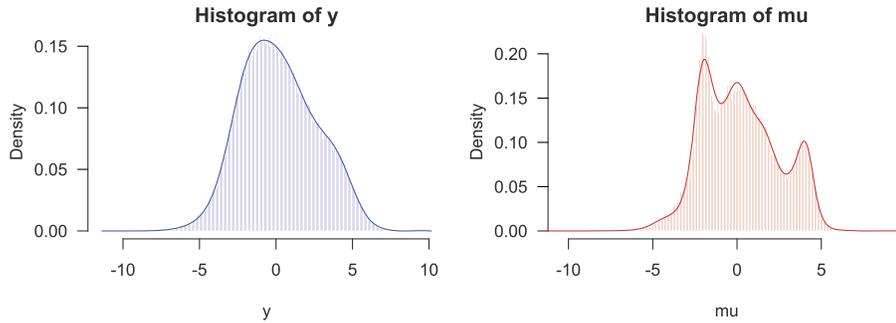


Fig 1: Example of deconvolution with an ℓ^2 penalty on the discrete first derivative ($k = 1$). The left panel shows the data histogram together with the fitted marginal density as a solid curve. The right panel shows the histogram of the μ_i 's together with the estimated mixing measure as a solid curve.

for a grid of regularization parameters and chooses the parameter that minimizes $l(\hat{\theta}_\tau^{\text{held out}}) + \|\Delta^{(k+1)}\hat{\theta}_\tau^{\text{held out}}\|_1$, where $\hat{\theta}_\tau^{\text{held out}}$ is defined as

$$\hat{\theta}_\tau^{\text{held out}} = \hat{\theta}_\tau - \log(n\Delta) + \log(n_{\text{held out}}\Delta)$$

with $\hat{\theta}_\tau$ the solution obtained by fitting the model on the training set which consists of 75% of the data. Here, $l(\hat{\theta}_\tau^{\text{held out}})$ is evaluated using the counts from the held out set which has 25% of the data, and $n_{\text{held out}}$ is the number of observations in such set. Our motivation for using the additional term $\|\Delta^{(k+1)}\hat{\theta}_\tau^{\text{held out}}\|_1$ is that ℓ_0 works well when the problem is formulated with ℓ_1 regularization. However, when (9) is formulated using ℓ_2 , the penalty ℓ_0 is not suitable so instead we use ℓ_1 . Our simulations in the experiments section will show that this rule works well in practice.

3.6. A toy example

We conclude this section by illustrating the accuracy of our regularized deconvolution approach on a toy example. In this example we draw 10^5 samples $\{y_i\}$ with the corresponding $\{\mu_i\}$ drawn from a mixture of three normal distributions. Figure 1 shows the samples of both the observations y_i (left panel) and the means μ_i (right panel), together with the reconstructions provided by our method. Here, we solve the ℓ^2 version of problem (9) by using the BFGS algorithm and choosing τ with the heuristic just described.

It is clear that regularizing with an ℓ_2 penalty provides an excellent fit of the marginal density. Surprisingly, it can also capture all three modes of the true mixing density, a feature which is completely obscured in the marginal). Our experiments in Section 6 will show in a more comprehensive way that our method far outperforms other approaches in its ability to provide accurate estimates for multi-modal mixing distributions.

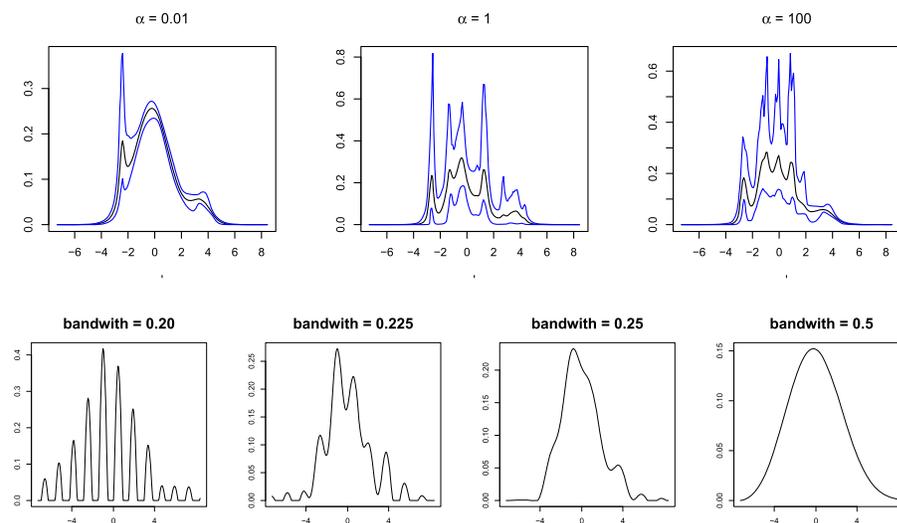


Fig 2: The first three panels show 95% confidence bands and posterior mean from 15000 posterior samples from a mixture of 10 normals prior on the latent variables μ . Panels 4-7 then shows the estimated mixing density using the kernel estimator with different bandwidth choices.

4. Sensitivity

4.1. Sensitivity analysis across the path

In this section, we provide an example of a sensitivity analysis using our deconvolution path estimator. We examine data originally collected and analyzed by Singh et al. (2002) on gene expression for 12,600 genes across two samples: 9 healthy patients and 25 patients with prostate tumors. The data come as a set of 12,600 t -statistics computed from gene-by-gene tests for whether the mean gene-expression score differs between the two groups. After turning these 12,600 t -statistics into z -scores via a CDF transform, we estimate a deconvolution path assuming a Gaussian convolution kernel. We use an ℓ^2 penalty and a grid of τ values evenly spaced on the logarithmic scale between 10^7 and 10^{-3} .

Each row of Figure 3 shows five points along the deconvolution path; the regularization parameter is largest in Row A and gets smaller in each successive row. Within each row, the left column shows the estimated mixing distribution \hat{f} for the given value of τ . The middle column shows the histogram of the data together with the fitted marginal density $\hat{n} = \phi * \hat{f}$. The right column shows the fitted marginal density on the log scale, with a regular grid to facilitate comparison of the results across different values of τ .

The figure shows that, while the estimate of the mixing distribution changes dramatically across the deconvolution path, the estimate of the marginal density is much more stable. Even on the log scale (right column), the differences among

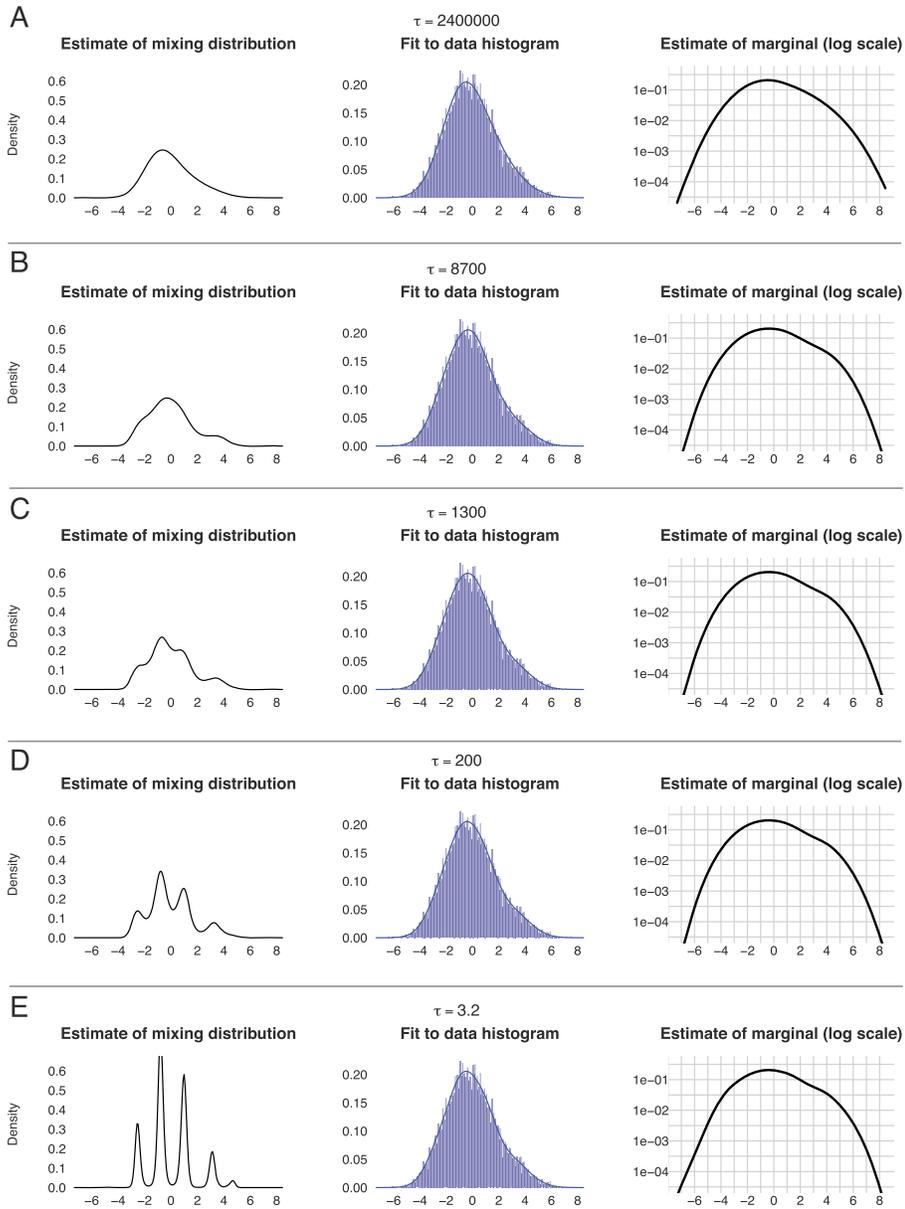


Fig 3: Rows A–E show five points along the deconvolution path for the prostate cancer gene-expression data. The regularization parameter is largest in Row A and gets smaller in each succeeding row.

the fitted marginal densities are not visually apparent in Panels B through E, even as the regularization parameter varies across three orders of magnitude.

This vividly demonstrates the well-known fact that deconvolution, especially of a Gaussian kernel, is a very ill-posed inverse problem. There is little information in the data to distinguish a smooth mixing distribution from a highly multimodal one, and the model-selection heuristics described earlier are imperfect. A decision to prefer Panel B to Panel E, for instance, is almost entirely due to the effect of the prior. Yet for most common deconvolution methods, the mapping between prior assumptions and the smoothness of the estimate is far from intuitive. By providing a full deconvolution path, our method makes this mapping visually explicit.

For reference, it is interesting to compare our deconvolution path to the results of other methods. Figure 2 shows the result of using MCMC to fit a 10-component mixture of normals to the mixing distribution. The weights in the Gaussian mixture were assigned three different symmetric Dirichlet priors, with concentration parameter $\alpha \in \{0.01, 1, 100\}$. Panels 1-3 in Figure 2 show the posterior mean and posterior 95% credible envelopes for f_0 ; these settings span a wide range of expected degrees of smoothness for f_0 , and they yield a correspondingly wide range of posterior estimates. Comparing Figures 3 and 2, we see that the deconvolution path spans essentially the entire range of plausible posterior estimates for f_0 arising under any of the concentration parameters. In contrast, Panels 4-7 in Figure 2 show that the kernel estimates are either overly smooth or wiggly.

4.2. Sensitivity to the number of bins

We now proceed to assess the sensitivity of our approach to the choice of the number of bins. We focus on our ℓ_2 regularization deconvolution as the conclusions for the ℓ_1 based method are similar.

To construct our example, we choose a mixing density (f_0) from our experiment section, and generate $n = 20000$ samples $\{\mu_i\}_{i=1}^n$ and then $y_i \sim N(\mu_i, 1)$, for $i = 1, \dots, n$. Using the data $\{y_i\}_{i=1}^n$, we apply the ℓ_2 deconvolution method with different choices of bin size and regularization parameter. The results in Figure 4 show that, for a fixed choice of τ , the estimated mixing distribution changes as we change d . This is not surprising, as the parameter d controls how many terms are involved in the likelihood, and increasing d implies that a fixed choice of τ becomes a weaker constraint. In practice, we have noticed excellent empirical results by fixing the parameter d and then performing model selection to choose τ , for instance as described in Section 3.5.

5. Theoretical properties

In this section we establish some important theoretical properties of our estimators by thinking of them as approximations to problems involving sieves. We

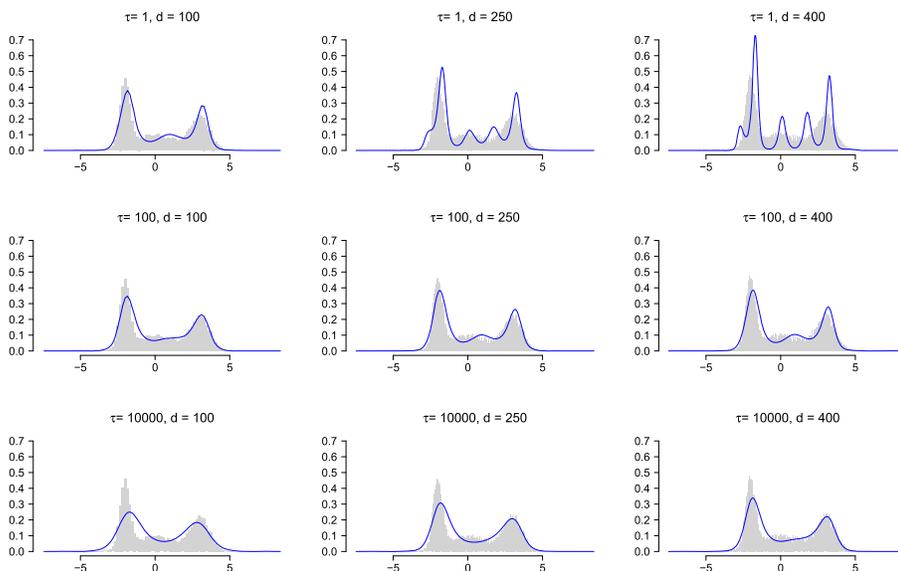


Fig 4: Each plot depicts the estimated mixing density using ℓ_2 regularization deconvolution with the tuning parameter τ and the bin size d as specified.

start by showing consistency of the mixing density in L1 norm. We do not provide convergence rates since, unlike the kernel estimator from Fan (1991) and the predictive recursion from Newton (2002), our method cannot be expressed in analytical form. Moreover, while the method from Fan (1991) remarkably attains minimax rates under squared error loss for point estimates of the true mixing, in an earlier work Carroll and Hall (1988) suggested that convergence rates for Gaussian deconvolution might be too pessimistic given the unbounded support nature of the classes of functions considered. This is out of the scope of our paper, but we do provide evidence in the later sections that our estimator can outperform existing non-parametric methods.

Throughout we consider $k \in \mathbb{N} - \{0\}$ and $q > 0$ to be fixed. We also denote by \mathcal{P} the set of densities in \mathbb{R} , thus $\mathcal{P} := \{f : \int_{\mathbb{R}} f(\mu) d\mu = 1; f \geq 0\}$, where $d\mu$ denotes Lebesgue measure. Moreover, given any non-negative function f we say that $b \in T_f$ if

$$\max \left(\|f\|_{\infty}, \|(\log f)^{(k+1)}\|_{\infty}, |(\log f)^{(k)}(0)|, \dots, |(\log f)(0)| \right) \leq b,$$

and $(\log f)^{(k+1)}$ is b -Lipschitz. Here, given an arbitrary function g , we use the notation $\|\cdot\|_{\infty}$ to indicate the usual supremum norm on the support of g . Moreover g is called T_m -Lipschitz if it satisfies $|g(x) - g(y)| \leq T_m |x - y|$, for all x and y .

In this section two metrics of interest will be repeatedly used. The first one is

the usual ℓ_1 distance $d(f, g) = \int_{\mathbb{R}} |f - g|$. The other metric of interest will be the Hellinger distance whose square is given as $H^2(f, g) := \int_{\mathbb{R}} |\sqrt{f}(\mu) - \sqrt{g}(\mu)|^2 d\mu$. We also use the notation $D_{\text{KL}}(f|g) = \int f(\mu) \log(f(\mu)/g(\mu)) d\mu$. Finally, for $q \in \mathbb{N}$, we define the functional $J_{k,q}$, a generalization of the usual total variation, as $J_{k,q}(f) := \int_{\mathbb{R}} |f^{(k+1)}(\mu)|^q d\mu$.

Next we state some assumptions for our first consistency result. Our approach is to consider the objective function in (3) restricted to a smaller domain than that of its original formulation. This will then allow to prove that the new problem is not ill defined and also its solutions enjoy asymptotic properties of convergence towards the true mixing density. We refer to Geman and Hwang (1982) for a general perspective on sieves.

Assumptions and definitions Let A be a set of functions that satisfies the following.

Assumption 1. If $f \in A$ then $f \in \mathcal{P}$, $f > 0$, $J_{k,q}(\log f) < \infty$, and there exists a constant t_f such that $t_f \in T_f$.

Assumption 2. For all $m \in \mathbb{N}$, there exists a set $S_m \subset A$ and constants $T_m, K_m > 0$ such that for all $f \in S_m$ it holds that $t_f = T_m$ and $J_{k,q}(\log f) \leq K_m$. Moreover, for all m , the set S_m induces a tight set of probability measures in $(\mathbb{R}, \mathbb{B}(\mathbb{R}))$ satisfying $S_m \subset S_{m+1}$. In addition, $\cup_m S_m$ is dense in A with respect to the metric d .

Assumption 3. Data model: we assume that y_1, \dots, y_n are independent draws from the density $\phi * f_0$, $f_0 \in A$, with ϕ being an arbitrary density function satisfying $\max(\|\phi\|_{\infty}, \|\phi'\|_{\infty}) < \infty$ and $\int_{\mathbb{R}} \log(\phi * f_0(\mu)) \phi * f_0(\mu) d\mu < \infty$.

Assumption 4. The set

$$A_m = \left\{ \gamma \in S_m : D_{\text{KL}}(\phi * f_0 | \phi * \gamma) = \inf_{\beta \in S_m} D_{\text{KL}}(\phi * f_0 | \phi * \beta) \right\},$$

satisfies $d(f_0, \gamma) \rightarrow 0$ as $m \rightarrow \infty$ for all $\gamma \in A_m$, where the convergence is uniform in A_m .

Assumption 5. We assume that the y_1, \dots, y_n are binned into D_n different intervals with frequency counts $\{x_j\}_{j=1, \dots, D_n}$ such that $n^{-1} \|x\|_{\infty} \rightarrow 0$ a.s., and we denote by ξ_j an arbitrary point in interval j . Note that this trivially holds for the case where $D_n = n$ and $x_j = 1$ for all j .

Assumption 6. There exists $f_m \in A_m$ such that

$$\sum_{j=1}^{D_n} \frac{x_j}{n} \log(\phi * f_m(\xi_j)) \rightarrow \int_{\mathbb{R}} \log(\phi * f_m(\xi)) \phi * f_0(\xi) d\xi \text{ a.s. as } n \rightarrow \infty.$$

If the $x_j = 1$ and $\xi_j = y_j$ for all $j = 1, \dots$, this condition can be disregarded.

Assumptions 1–3 are natural for the original variational problem proposed earlier. The Lipschitz condition, the bounds on the behavior of the functions at zero, and the tightness of distributions are merely used to ensure that the sieves will indeed be compact sets with respect to the metric d . An example of a set A satisfying these assumptions will be provided below.

With regards to the remaining conditions, Assumption 4 tell us that the sieves S_m are rich enough to approximate the true mixing density sufficiently well. In particular, this condition holds if the true mixing density f_0 belongs to S_m for some S_m . On the other hand, the last two assumptions can be disregarded when the counts in the bins are all one.

Assumption 5 imposes the constraint that $n^{-1} \|x\|_\infty \rightarrow 0$ a.s., where x is the vector of counts across bins. Note that $\mathbb{E}(x_j) = n p_j$, where p_j is the probability under the marginal density, of y_i falling in bin j . Hence by the Dominated Convergence Theorem, Assumption 5 implies that p_j goes, uniformly in j , to zero as n goes to infinity.

Before stating our first result, we consider a generic example for which the assumptions above hold. The example below basically states that the set A can be taken as the set of probability distributions consisting of finite mixtures of normals, with a fixed number of components. We could also consider the set of mixtures of normals with any number of components but with locations and scales on a compact set.

Example 1. Define the set A as

$$A := \left\{ f : f(x) = \sum_{i=1}^K w_i^f \frac{1}{\sqrt{2\pi (\sigma_i^f)^2}} \exp\left(-\frac{(x - \mu_i^f)^2}{2 (\sigma_i^f)^2}\right), \sigma_i^f > 0, w_i^f \in \Pi_K \right\},$$

where Π_K is the K -th standard simplex for some $K \in \mathbb{N}$. Moreover, we construct the sieves S_m as

$$S_m =: \left\{ f \in A : t_f = T_m, J_{k,q}(\log f) \leq K_m, \sigma_i^f \in I_m^\sigma, w_i^f \geq 0, \mu_i^f \in I_m^\mu \right\}$$

with $T_m, K_m \uparrow \infty$, and I_m^σ an increasing sequence of compact sets in $(0, \infty)$ for which

$$(0, \infty) := \cup_m I_m^\sigma,$$

and the sets I_m^μ are also increasing with

$$(-\infty, \infty) := \cup_m I_m^\mu,$$

Then, for $k \geq 2$, the sets A and S_m satisfy the Assumptions (1) and (2) from above. If in $\int_{\mathbb{R}} |(\log f)^{(k+1)}|^q$ we replace \mathbb{R} by a bounded set, then the same holds for $k = 1$.

We are now ready to state our first consistency result. Its proof generalizes ideas from Theorem 1 in Geman and Hwang (1982).

Theorem 1. *If Assumptions (1-6) hold, then, the problem*

$$\underset{f \in S_m}{\text{minimize}} \quad - \sum_{j=1}^{D_n} x_j \log(\phi * f)(\xi_j),$$

has solution set $M_m^n \neq \emptyset$. Moreover, for any sequence m_n increasing slowly enough it holds that

$$\sup_{\beta \in M_{m_n}^n} d(\beta, f_0) \rightarrow 0 \quad \text{a.s.} \quad (13)$$

In Theorem 1, the sequence T_{m_n} is arbitrary and can grow as fast as desired. Moreover, the a.s statement is on the probability space $(\mathbb{R}^\infty, \mathcal{F}, F_0 \times F_0 \times F_0 \dots)$ with F_0 the measure on $(\mathbb{R}, \mathbb{B}(\mathbb{R}))$ induced by $\phi * f_0$, and with \mathcal{F} the completion of $\mathbb{B}(\mathbb{R})^\infty$. Here $\mathbb{B}(\mathbb{R})$ refers to the class of Borel sets in \mathbb{R} .

We also observe that Theorem 1 can be restated in terms of the Hellinger distance. To see this, we recall the well known inequality between densities p and q : $H^2(p, q) \leq d(p, q)$. Therefore, we can replace d in Equation (13) by the Hellinger distance $(\sqrt{H^2})$.

Finally, we emphasize that Theorem 1 only ensures convergence but provides no information regarding speed of convergence. We leave for future work the characterization of the sequence m_n .

6. Experiments

6.1. Mixing density estimation

In this section we show the potential gain given by our penalized approaches. We start by considering the task of recovering the true mixing distribution. We evaluate the performance of our methods described in Section 3 which we call L1-deconvolution (L1-D) and L2-deconvolution (L2-D) depending on the regularization penalty used in the estimation. As competitors we consider a mixture of normals model (MN), the predictive recursion algorithm (PR) from Newton (2002), and the Fourier transform kernel deconvolution method (FTKD) from Fan (1991). Our comparisons are based on four examples which are shown in Figure 5. These examples are intended to illustrate the performance under different scenarios involving smooth and sharp densities. Next we describe the simulation setting and as well as the implementation details of the competing methods.

As a flexible Bayesian model we decided to use a prior for the mixing density based on a mixture of 10 normals (MN). Here, the weights of the mixture components are drawn from a Dirichlet prior with concentration parameter 1. This is done in order to have a uniform prior on the simplex. For the locations of the mixture we consider non-informative priors given as $N(0, 10^2)$ while for the variances of the mixture components we place a inverse gamma prior with shape parameter 0.01 and rate 0.01. The complete model can be then thought as a weak limit approximation to the Dirichlet process, (Ishwaran and Zarepour,

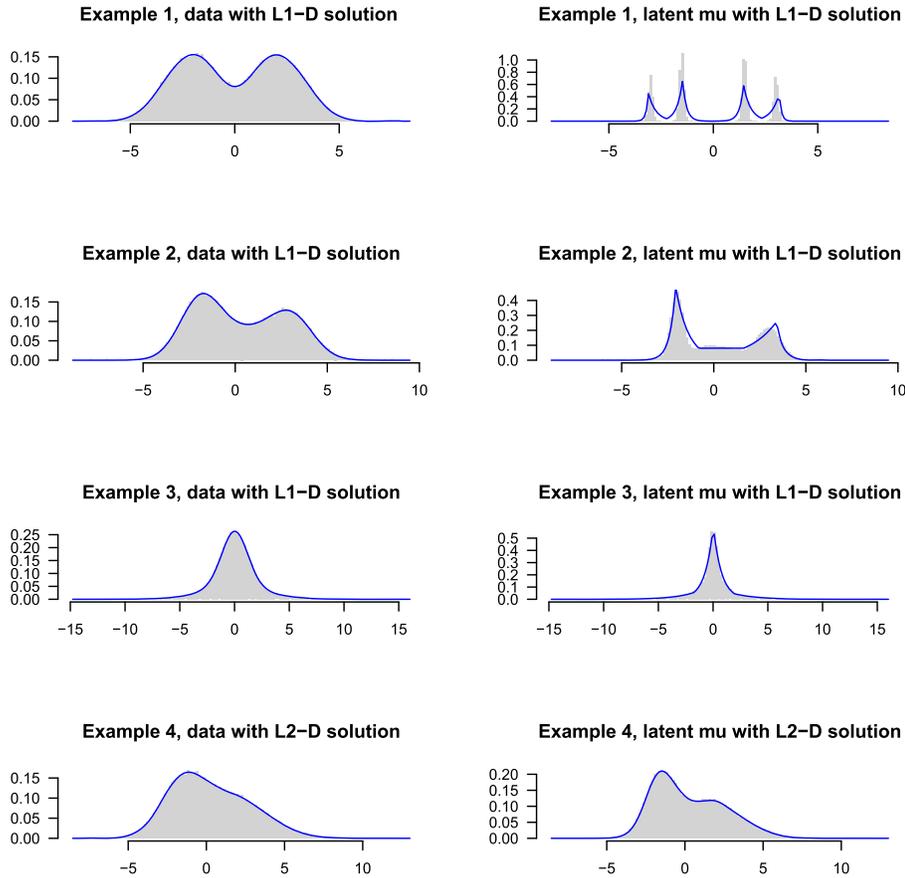


Fig 5: The first panel shows a histogram of observed data $\{y_i\}_{i=1}^n$ for our first example, and the L1-D marginal density estimate plotted on top of the histogram. Here the data has been generated as $y_i \sim N(\mu_i, 1)$ where μ_i is a draw from the mixing density. The second panel shows, for this same example, the histogram of $\{\mu_i\}_{i=1}^n$ (unobserved draws from the mixing density) and the L1-D estimate of the mixing density plotted on top of it. Panels 3-6 show the respective cases of Examples 2 and 3. The last two panels show the corresponding plots for the L2-D solution and Example 4.

2002). Also, Gibbs sampling is accomplished straightforwardly by introducing a data augmentation with a variable z_i indicating the component to which μ_i belongs.

The next competing model is the predictive recursion algorithm from Newton (2002) for which we choose the weights w_i as in Martin and Tokdar (2012), close to the limit of the upper bound of the convergence rate for PR given in

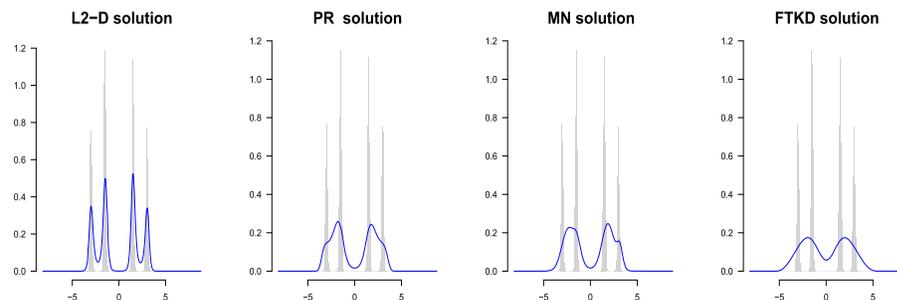


Fig 6: For the mixing density illustrated in Example 1 of Figure 5 we show the estimated mixing densities with different methods. From left to right the first two panels correspond to the estimated mixing densities using L2-D and PR algorithms along with the latent μ_i 's. The third and four panels show the estimated density using MN and FTKD both with the latent μ_i 's. In all cases the estimation was done with the same data where $n = 10^5$

Tokdar et al. (2009). Moreover we average the PR estimator over 10 different permutations of the input data in order to obtain a smaller expected error Tokdar et al. (2009).

On the other hand, for the Fourier transform kernel deconvolution method, we consider different choices of bandwidth: the rule of thumb from Fan (1991), the plug in bandwidth from Stefanski and Carroll (1990), and the 2-stage plug-in bandwidth from Delaigle and Gijbels (2002). Our estimates are obtained using the R package fDKDE available at <http://www.ms.unimelb.edu.au/~aurored/links.html>, which addresses the main concerns associated with the R package decon; see Delaigle (2014).

For the final competitor, the “g-modeling” approach from Efron (2016) (g-M), we use the newly released R package deconvolveR.

For the cases of the true mixing density we consider four densities of the form

$$f_0(\mu) = \sum_{i=1}^K w_i N(\mu | \theta_i, \sigma_i^2).$$

In all cases considered here, the observations arise as in (1) with a standard normal sampling model. In our first example we evaluate performance for samples of a density that has four peaks or explicitly $K = 4$, $w = (0.2, 0.3, 0.3, 0.2)$, $\theta = (-3, -1.5, 1.5, 3)$ and with variances $\sigma^2 = (0.01, 0.01, 0.01, 0.01)$. For the second example we consider a mixture of three normals two of which are smooth while the other has a peak. The true parameters in this case are $K = 3$, $w = (1/3, 1/3, 1/3)$, $\theta = (0, -2, 3)$ and $\sigma^2 = (2, .1, .4)$. The next example is a mixture of $K = 3$ normals, one of which has very high variance. The true parameters chosen are $w = (0.3, 0.4, 0.3)$, $\theta = (0, 0, 0)$ and $\sigma^2 = (0.1, 1, 9)$. Our final example

TABLE 1

Comparisons between the true and estimated mixing densities for different methods given samples from density Example 1. The acronyms here are given the text. The Mean squared error (MSE) is multiplied by 10^2 and reported over an interval containing 95% the mass of the mixing density. Moreover, the Earth Mover's Distance is multiplied by 10. For each performance measure the results reported are obtained averaging over 100 Monte Carlo simulations.

n	MSE						Earth Mover's Distance					
	MN	PR	L2-D	L1-D	FTKD	g-M	MN	PR	L2-D	L1-D	FTKD	g-M
	95%	95%	95%	95%	95%	95%						
2000	9.47	9.12	9.39	9.69	8.89	9.27	3.71	3.52	3.78	3.79	4.35	4.76
10000	8.43	8.72	8.64	7.44	8.87	9.22	3.29	3.31	3.37	2.77	4.26	4.20
25000	8.34	8.46	7.27	5.54	8.88	9.32	3.12	3.18	2.84	2.11	4.24	4.05
50000	8.21	8.23	5.80	4.15	8.85	9.40	3.13	3.07	2.29	1.69	4.22	3.99
100000	8.34	8.05	4.79	3.38	8.69	9.50	3.12	2.98	1.85	1.36	3.91	3.95

is a mixture, with $K = 3$, giving raise to a very smooth density, the parameters are $w = (0.5, 0.4, .1)$, $\theta = (-1.5, 1.5, 4)$ and $\sigma^2 = (1, 2, 2)$. A visualization of these examples is shown in Figure 5.

Given these densities, we consider scenarios with varying number of samples n and for each fixed n we run 100 Monte Carlo simulations. Moreover, for our methods we set D , the number of evenly space points in the grid, to 250. In all cases we set the parameter k to be 1 for our deconvolution methods.

The results on Table 1 illustrate a clear advantage of our penalized likelihood approaches over MN, PR and FTKD which seems even more significant for larger samples size. The estimated mixing density by L1-D is shown in Figure 5 where we can clearly see that L1-D can capture the peaks of the unknown mixing density. Moreover, Figure 6 shows that L2-D can also capture the structure of the true density. In contrast, MN, PR and FTKD all fail to provide reliable estimators.

For our example density 2, we observe from Table 2 that in general L2-D and L1-D offer the best performance. In the case of example 3, we observe that the L1-D again provides better results than the competitors in all the scenarios of sample sizes considered. Even with only 10000, samples L1-D is closer to the true density than all the other methods with more samples. Moreover, L2-D performs much better than PR and FTKD. Also, L2-D seems to be a clear competitor to MN. In the final example density 4, we observe that L2-D is the best method in all the scenarios considered.

Overall, we have shown that for estimating the mixing density, L1-D and L2-D can perform well under different settings, even when other methods exhibit notable deficiencies. The advantage is amplified by the fact that both of our methods are less computationally intensive than MN, with L2-D requiring around 40 seconds to handle problems with $D = 250$, and L1-D under the same problem conditions typically requires around 5 minutes for a full solution path across 50 values of the tuning parameter.

TABLE 2

Simulation results for Examples 2, 3 and 4. Comparisons between the true and estimated mixing densities for different methods given samples from each example density. The acronyms here are given the text. The Mean squared error (MSE) is multiplied by a constant and reported over an interval containing 95% the mass of the mixing density. Moreover, the Earth Mover's Distance is multiplied by a constant. For each performance measure the results reported are obtained averaging over 100 Monte Carlo simulations.

(a) Example 2 results, here each MSE entry is multiplied by 10^3 . Moreover, each Earth Mover's Distance is multiplied by 10^2 .

n	MSE						Earth Mover's Distance					
	MN 95%	PR 95%	L2-D 95%	L1-D 95%	FTKD 95%	g-M 95%	MN	PR	L2-D	L1-D	FTKD	g-M
2000	6.20	2.54	2.74	2.38	6.07	8.23	13.3	12.0	12.9	11.3	24.6	29.8
10000	3.45	1.75	1.60	1.68	5.98	5.82	8.98	8.14	7.92	7.22	22.8	19.4
25000	2.31	1.46	1.19	1.35	5.99	5.01	5.94	6.69	5.87	5.96	22.0	15.7
50000	1.24	1.28	0.89	1.18	5.89	4.68	5.05	6.19	5.01	5.04	21.6	13.3
100000	0.78	1.07	0.74	0.87	4.97	4.03	4.10	5.45	4.35	4.09	18.3	12.9

(b) Example 3. MSE is multiplied by 10^3 and each Earth Mover's Distance is multiplied by 10^2 .

n	MSE						Earth Mover's Distance					
	MN 95%	PR 95%	L2-D 95%	L1-D 95%	FTKD 95%	g-M 95%	MN	PR	L2-D	L1-D	FTKD	g-M
2000	5.28	1.83	2.09	0.96	4.95	7.16	11.2	11.7	12.0	9.39	24.2	35.1
10000	3.06	1.38	1.46	0.61	4.86	1.45	6.98	7.69	7.21	5.38	22.1	44.3
25000	1.51	1.16	1.18	0.47	4.61	1.18	4.89	6.12	5.51	3.81	21.0	41.0
50000	0.95	1.06	1.00	0.42	3.77	1.13	4.04	5.69	4.77	3.23	17.3	43.2
100000	0.72	0.95	0.86	0.38	3.48	2.42	3.71	5.24	4.26	2.92	20.0	42.4

(c) Example 4, each MSE entry is multiplied by 10^4 . Moreover, each Earth Mover's Distance is multiplied by 10^2 .

n	MSE						Earth Mover's Distance					
	MN 95%	PR 95%	L2-D 95%	L1-D 95%	FTKD 95%	g-M 95%	MN	PR	L2-D	L1-D	FTKD	g-M
2000	20.6	4.75	1.82	3.48	3.25	7.06	11.1	10.0	8.07	8.32	15.2	17.4
10000	7.64	1.89	0.65	1.93	2.87	4.20	6.35	6.22	4.34	5.27	14.2	10.2
25000	2.04	1.10	0.48	2.19	2.57	2.34	3.82	4.33	3.14	3.76	13.0	7.75
50000	1.03	0.69	0.36	1.20	2.02	1.95	2.61	3.42	2.12	2.87	10.9	6.37
100000	0.50	0.55	0.39	0.90	1.36	1.49	1.99	2.82	1.71	2.53	8.90	5.29

6.2. Normal means estimation

After evaluating our proposed methodology for the task of estimating the mixing density, we now, for the case of standard normal kernel, focus on the esti-

TABLE 3

Mean squared error, of the normal means estimates, times 100, averaging over 100 Monte Carlo simulations, for different methods given samples from example 1.

n	L2-D	L1-D	PR	MN	Efron	GMLEBIP	SC	BG	NLP
2000	64.31	64.29	64.16	67.50	70.27	64.48	68.24	65.57	64.11
10000	63.89	63.68	63.86	63.18	70.00	63.80	65.56	64.06	63.27
25000	63.52	63.37	63.69	63.84	69.96	63.39	64.66	63.65	63.60
50000	63.27	63.21	63.55	65.20	69.85	63.23	64.15	63.44	63.26
100000	63.27	63.23	63.59	63.79	69.89	63.21	63.86	63.39	63.18

mation of the normals means $\{\mu_i\}$. For this, we consider comparisons using the best four among the methods used before in addition to other procedures that we briefly discuss next.

As it is well known (e.g Efron (2011) for description and references), assuming that the marginal density is known, one can use Tweedie's formula to estimate $\{\mu_i\}$. For all the methods here this is the approach that we take, except for MN in which case we use the posterior means resulting from Gibbs sampling inference. For the methods depending on grid estimator, the number of bins is set to 250.

For the method of Efron (2011), we set to 5 the degree of the polynomial approximation to the logarithm of the marginal true density (we found larger values to be less numerically stable). The Poisson surrogate model is then fitted in R using the command `glm`. We also compare against the general maximum likelihood empirical-Bayes estimator (GMLEB) from Jiang and Zhang (2009), which is a discretized version of the original Kiefer-Wolfowitz estimator. For our comparisons we use the algorithm proposed in Koenker and Mizera (2014) based on an interior point method algorithm (GMLEBIP). We use the R package `REBayes` in order to obtain this estimator (Koenker, 2013). On the other hand, for the shape constrained (SC) estimator from Koenker and Mizera (2014), we rely on a binned count approach based on a weighted likelihood using R code provided by the authors. Moreover, we consider the estimator from Brown and Greenshtein (2009) using the default choice of bandwidth $h_n = (\log n)^{-1/2}$, which we refer to as BG. The final competitor is the non-linear projection (NLP) estimator from Wager (2013).

From Table 3 it is clear that the best methods for example 1 are L1-D, L2-D, GMLEBIP, and NLP. Moreover, it is not surprising that GMLEBIP provides good estimates given that the true mixing density has mixture components that have small variance.

For example 2, we can see from Table 4 that again L2-D and L1-D provide competitive estimates. The other suitable methods for this example seem to be PR and GMLEBIP. With slightly worse estimates MN, BG and SC provide results that are still competitive, with SC being particularly attractive given its computational speed to provide solutions.

Finally, for examples 3 and 4 we can see in Tables 5 and 6 respectively that L1-D and L2-D are the best or among the best methods in terms of mean

TABLE 4

Mean squared error, of the normal means estimates, times 100, averaging over 100 Monte Carlo simulations, for different methods given samples from example 2.

n	L2-D	L1-D	PR	MN	Efron	GMLEBIP	SC	BG	NLP
2000	65.42	65.46	65.36	64.33	69.97	66.20	69.60	66.99	65.75
10000	64.98	65.06	65.08	65.66	69.75	65.29	67.10	65.54	65.95
25000	65.19	65.08	65.32	65.21	69.94	65.12	66.42	65.49	65.09
50000	64.99	65.08	65.13	65.44	69.93	65.03	65.97	65.19	65.24
100000	65.02	64.95	65.14	65.03	69.84	65.02	65.69	65.14	64.96

TABLE 5

Mean squared error, of the normal means estimates, times 100, averaging over 100 Monte Carlo simulations, for different methods given samples from example 3.

n	L2-D	L1-D	PR	MN	Efron	GMLEBIP	SC	BG	NLP
2000	64.99	64.96	65.41	69.05	70.54	65.74	69.70	66.99	65.77
10000	64.73	64.76	64.96	64.21	71.34	64.85	66.92	65.36	64.81
25000	64.52	64.57	64.75	64.97	71.42	64.65	66.62	64.82	64.62
50000	64.51	64.61	64.73	65.38	71.52	64.64	66.60	64.67	64.57
100000	64.54	64.41	64.76	64.54	71.96	64.56	65.17	64.62	64.46

TABLE 6

Mean squared error, of the normal means estimates, times 100, averaging over 100 Monte Carlo simulations, for different methods given samples from example 4.

n	L2-D	L1-D	PR	MN	Efron	GMLEBIP	SC	BG	NLP
2000	79.63	80.20	79.89	78.68	80.00	80.97	85.47	81.58	80.01
10000	79.32	79.35	79.42	79.34	79.99	79.74	82.18	79.89	79.64
25000	79.39	79.31	79.48	78.79	79.96	79.30	80.98	79.65	79.39
50000	79.21	79.25	79.29	79.85	79.82	79.40	80.58	79.36	79.39
100000	79.29	79.22	79.37	79.51	79.91	79.30	80.15	79.37	79.36

squared distance when recovering the unknown means μ_i . Table 6 also suggests that Efron's estimator is more suitable when the true mixing density is very smooth with no sharp peaks.

Code An R package has been placed in Github with the proposed methods. To install the package simply call in R:

```
install_github('RRDecon', 'omadrid1')
```

The source code is available at <https://github.com/omadrid1/RRDecon>.

7. Discussion

In many problems in statistics and machine learning, we observe a blurred version of an unknown mixture distribution which we would like to recover via deconvolution. The main challenge is to find an approach that is computationally

fast but still possesses nice statistical guarantees in the form of rates of convergence. We propose a two-step “bin-and-smooth” procedure that achieves both of these goals. This reduces the deconvolution problem to a Poisson-regularized model which can be solved either via standard methods for smooth optimization, or with a fast version of the alternating-direction method of multipliers (ADMM). Our approach reduces the computational cost compared to a fully Bayesian method and yields a full deconvolution path to illustrate the sensitivity of our solution to the specification of the amount of regularization. We provide theoretical guarantees for our procedure. In particular, under suitable regularity conditions, we establish the almost-sure convergence of our estimator towards the mixing density.

There are a number of directions for future inquiry, including multivariate extensions and extensions to multiple hypothesis testing. These are active areas of current research.

Appendix A: Technical supplement

A.1. Example 1

Consider the density

$$f(x) = \sum_{i=1}^K w_i N(x | \mu_i, \sigma_i^2)$$

where $\sum_{i=1}^K w_i = 1$, $w_i > 0$. Then

$$f'(x) = \sum_{i=1}^K \frac{-w_i(x - \mu_i)}{\sigma_i^2} N(x | \mu_i, \sigma_i^2)$$

$$f''(x) = \sum_{i=1}^K \frac{w_i(x - \mu_i)^2}{\sigma_i^4} N(x | \mu_i, \sigma_i^2) - \sum_{i=1}^K \frac{w_i}{\sigma_i^2} N(x | \mu_i, \sigma_i^2).$$

Moreover,

$$(\log f)'' = \frac{f f'' - (f')^2}{f^2}.$$

Therefore,

$$\begin{aligned} (\log f)''(x) &= \frac{\left[\sum_{i=1}^K w_i N(x | \mu_i, \sigma_i^2) \right] \left[\sum_{i=1}^K \frac{w_i(x - \mu_i)^2}{\sigma_i^4} N(x | \mu_i, \sigma_i^2) - \sum_{i=1}^K \frac{w_i}{\sigma_i^2} N(x | \mu_i, \sigma_i^2) \right]}{\left[\sum_{i=1}^K \frac{-w_i(x - \mu_i)}{\sigma_i^2} \phi_{0,1} \left(\frac{x - \mu_i}{\sigma_i} \right) \right]^2} \frac{f(x)^2}{f(x)^2} \\ &= - \frac{\sum_{i=1}^K \frac{w_i}{\sigma_i^2} N(x | \mu_i, \sigma_i^2)}{f(x)} + \frac{\left[\sum_{i=1}^K w_i N(x | \mu_i, \sigma_i^2) \right] \left[\sum_{i=1}^K \frac{w_i(x - \mu_i)^2}{\sigma_i^4} N(x | \mu_i, \sigma_i^2) \right]}{f(x)^2} \\ &\quad - \frac{\left[\sum_{i=1}^K \frac{-w_i(x - \mu_i)}{\sigma_i^2} N(x | \mu_i, \sigma_i^2) \right]^2}{f(x)^2}. \end{aligned}$$

Hence, there exists constants c_{ij} and polynomials satisfying

$$(\log f)''(x) = -\frac{\sum_{i=1}^K \frac{w_i}{\sigma_i^2} N(x | \mu_i, \sigma_i^2)}{\sum_{i=1}^K w_i N(x | \mu_i, \sigma_i^2)} + \frac{\sum_{i \neq j} c_{ij} T_{ij}(x) \exp\left(-\frac{(x-\mu_i)^2}{2\sigma_i^2} - \frac{(x-\mu_j)^2}{2\sigma_j^2}\right)}{\left[\sum_{i=1}^K w_i N(x | \mu_i, \sigma_i^2)\right]^2}.$$

Hence,

$$\begin{aligned} & (\log f)'''(x) \\ &= -\frac{\left[-\sum_{i=1}^K \frac{w_i(x-\mu_i)}{\sigma_i^4} N(x|\mu_i, \sigma_i^2)\right] \left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]}{\left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]^2} \\ & \quad - \frac{\left[-\sum_{i=1}^K \frac{w_i(x-\mu_i)}{\sigma_i^2} N(x|\mu_i, \sigma_i^2)\right] \left[\sum_{i=1}^K \frac{w_i}{\sigma_i^2} N(x|\mu_i, \sigma_i^2)\right]}{\left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]^2} \\ & \quad + \frac{\left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]^2 \sum_{i \neq j} \left(c_{ij} T'_{ij}(x) + c_{ij} T_{ij}(x) \left(-\frac{x-\mu_i}{\sigma_i^2} - \frac{x-\mu_j}{\sigma_j^2}\right)\right) \exp\left(-\frac{(x-\mu_i)^2}{2\sigma_i^2} - \frac{(x-\mu_j)^2}{2\sigma_j^2}\right)}{\left[\sum_{i=1}^K \frac{w_i}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(x-\mu_i)^2}{2\sigma_i^2}\right)\right]^4} \\ & \quad + \frac{2 \left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right] \left[\sum_{i=1}^K -\frac{w_i(x-\mu_i)}{\sigma_i^2} N(x|\mu_i, \sigma_i^2)\right] \sum_{i \neq j} c_{ij} T_{ij}(x) \exp\left(-\frac{(x-\mu_i)^2}{2\sigma_i^2} - \frac{(x-\mu_j)^2}{2\sigma_j^2}\right)}{\left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]^4} \\ &= \frac{\sum_{i \neq j} r_{ij}(x) \exp\left(-\frac{(x-\mu_i)^2}{\sigma_i^2} - \frac{(x-\mu_j)^2}{2\sigma_j^2}\right)}{\left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]^2} \\ & \quad + \frac{\sum_{|\{i_1, i_2, i_3, i_4\}| \geq 2} P_{i_1, i_2, i_3, i_4}(x) \exp\left(\sum_{j=1}^4 -\frac{(x-\mu_{i_j})^2}{2\sigma_{i_j}^2}\right)}{\left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]^4} \\ &= \frac{\sum_{|\{i_1, i_2, i_3, i_4\}| \geq 2} \bar{P}_{i_1, i_2, i_3, i_4}(x) \exp\left(\sum_{j=1}^4 -\frac{(x-\mu_{i_j})^2}{2\sigma_{i_j}^2}\right)}{\left[\sum_{i=1}^K w_i N(x|\mu_i, \sigma_i^2)\right]^4} \end{aligned}$$

with r_{ij} , P_{i_1, i_2, i_3, i_4} and $\bar{P}_{i_1, i_2, i_3, i_4}$ polynomials. Consequently,

$$(\log f)''''(x) = \frac{\sum_{|\{i_1, \dots, i_8\}| \geq 2} S_{i_1, \dots, i_8}(x) \exp\left(-\sum_{j=1}^8 \frac{(x-\mu_{i_j})^2}{2\sigma_{i_j}^2}\right)}{\left[\sum_{i=1}^K w_i N(x | \mu_i, \sigma_i^2)\right]^8}.$$

Now we show some properties of $(\log f)^{(3)}$ and $(\log f)^{(4)}$.

First, assume that $|\{\sigma_1, \dots, \sigma_K\}| > 1$ and choose $i^* = \arg \min_i \frac{1}{\sigma_i^2}$. Then

$$(\log f)'''(x) = \frac{\sum_{|\{i_1, i_2, i_3, i_4\}| \geq 2} \bar{P}_{i_1, i_2, i_3, i_4}(x) \exp\left(\sum_{j=1}^4 -\frac{(x-\mu_{i_j})^2}{2\sigma_{i_j}^2} + 4\frac{(x-\mu_{i^*})^2}{2\sigma_{i^*}^2}\right)}{\left(\sum_{i=1}^K \frac{w_i}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(x-\mu_i)^2}{2\sigma_i^2} + \frac{(x-\mu_{i^*})^2}{2\sigma_{i^*}^2}\right)\right)^4}.$$

Hence $(\log f)'''(x) \rightarrow 0$ as $x \rightarrow \pm\infty$. Also,

$$\begin{aligned} \int_{\mathbb{R}} |(\log f)'''(x)| dx &\leq \sum_{|\{i_1, i_2, i_3, i_4\}| \geq 2} \int_{\mathbb{R}} \frac{|\bar{P}_{i_1, i_2, i_3, i_4}(x)|}{\left[\frac{\sqrt{2\pi}\sigma_{i^*}}{w_{i^*}}\right]^4} \\ &\quad \times \exp\left(\sum_{j=1}^4 -\frac{(x-\mu_{i_j})^2}{2\sigma_{i_j}^2} + 4\frac{(x-\mu_{i^*})^2}{2\sigma_{i^*}^2}\right) \\ &\leq \infty, \end{aligned}$$

and we can also easily show that $\int_{\mathbb{R}} |(\log f)'''(x)|^q dx < \infty$ for $q > 2$. The same argument also holds if we replace $(\log f)'''$ by $(\log f)''''$.

If $|\{\sigma_1, \dots, \sigma_K\}| = 1$ then we can assume without loss of generality that $|\{\mu_1, \dots, \mu_K\}| > 1$ (otherwise the proof is trivial). Then taking $i^* = \arg \max_i \mu_i$ we have

$$\begin{aligned} &(\log f)'''(x) \\ &= \frac{\sum_{|\{i_1, i_2, i_3, i_4\}| \geq 2} \bar{P}_{i_1, i_2, i_3, i_4}(x) \exp\left(\sum_{j=1}^4 \left[-\frac{x(\mu_{i^*} - \mu_{i_j})}{\sigma^2} - \frac{\mu_{i_j}^2}{2} + \frac{\mu_{i^*}^2}{2}\right]\right)}{\left(\sum_{i=1}^K \frac{w_i}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{x(\mu_{i^*} - \mu_i)}{\sigma^2} - \frac{\mu_i^2}{2} + \frac{\mu_{i^*}^2}{2}\right)\right)^4}. \end{aligned}$$

From this, we can see that $\int_0^\infty |(\log f)'''(x)| dx < \infty$ and $(\log f)'''(x) \rightarrow 0$ as $x \rightarrow \infty$. Moreover, if we take $i^* = \arg \min_i \mu_i$ then we see that $\int_{-\infty}^0 |(\log f)'''(x)| dx < \infty$ and $(\log f)'''(x) \rightarrow 0$ as $x \rightarrow -\infty$. The same result also holds for $(\log f)''''$.

On the other hand, we can proceed by induction and show the same properties for $(\log f)^{(l)}$ for $l > 4$.

It remains to check that S_m is tight, and dense in A . The former follows immediately by the Prokhorov's theorem, and the latter is immediate by the fact that $A = \cup_m S_m$.

A.2. Gradient expression for ℓ_2 regularization

Here we write the mathematical expressions for the gradient of the objective function when performing L2 de-convolution, As in Section 3.3 of the main

document. Using the notation there, we have that

$$[\nabla l(\theta)]_j = \sum_{i=1}^D G_{ij} e^{\theta_j} \left(\frac{x_i}{\lambda_i(\theta)} - 1 \right),$$

and

$$\nabla \|\Delta^{(k+1)}\theta\|_2^2 = 2 \left(\Delta^{(k+1)} \right)^T \Delta^{(k+1)} \theta.$$

A.3. Proof of Theorem 1

Proof. Motivated by Geman and Hwang (1982), given $\alpha \in A$ we define the function $F(\xi, \alpha) = (\phi * \alpha)(\xi)$ for $\xi \in \mathbb{R}$. Clearly, $F(\xi, \alpha)$ is a density that induces a measure in \mathbb{R} that is absolutely continuous with respect to the Lebesgue measure in \mathbb{R} . Also, we observe that if $\alpha, \beta \in A$, then, for any Borel measurable set E , we have by Tonelli's theorem that

$$\begin{aligned} \left| \int_E \phi * \alpha(\mu) d\mu - \int_E \phi * \beta(\mu) d\mu \right| &= \left| \int_{\mathbb{R}} \left(\int_E \phi(\mu - y) d\mu \right) (\alpha(y) - \beta(y)) dy \right| \\ &\leq d(\alpha, \beta). \end{aligned}$$

Hence $d(\alpha, \beta) = 0$ implies that $\phi * \alpha$ and $\phi * \beta$ induce the same probability measures in $(\mathbb{R}, \mathbb{B}(\mathbb{R}))$.

Next we verify the assumptions in Theorem 1 from Geman and Hwang (1982). This is done into different steps below. Steps 1-4 verify the assumptions B1-B4 in Theorem 1 from Geman and Hwang (1982). Steps 5-6 are needed in the general case in which the data is binned. These are also related to ideas from Wald (1949).

Step 1

Given $\alpha \in A$ and $\epsilon > 0$, the function

$$\xi \rightarrow \sup_{\beta \in S_m: d(\alpha, \beta) < \epsilon} (\phi * \beta)(\xi)$$

is continuous and therefore measurable on ξ . To see this, simply note that for any $\beta \in A$ we have that

$$\|(\phi * \beta)'\|_{\infty} = \|(\phi' * \beta)\|_{\infty} \leq \|\phi'\|_{\infty} \int_{\mathbb{R}} \beta(\mu) d\mu = \|\phi'\|_{\infty}.$$

Hence all the functions $\beta \in S_m$ are $(\|\phi'\|_{\infty} + 1)$ -Lipschitz and the claim follows. Also, we note that

$$\lim_{\epsilon \rightarrow 0} \sup_{\beta \in S_m: d(\alpha, \beta) < \epsilon} (\phi * \beta)(\xi) = \phi * \alpha(\xi).$$

This follows by noticing that

$$\begin{aligned} \left| \sup_{\beta \in S_m: d(\alpha, \beta) < \epsilon} (\phi * \beta)(\xi) - \phi * \alpha(\xi) \right| &\leq \sup_{\beta \in S_m: d(\alpha, \beta) < \epsilon} |(\phi * (\beta - \alpha))(\xi)| \\ &\leq \|\phi\|_{\infty} \sup_{\beta: d(\alpha, \beta) < \epsilon} d(\alpha, \beta) \\ &\leq \epsilon \|\phi\|_{\infty}. \end{aligned}$$

Step 2

Define $E_\alpha(g) := \int_{\mathbb{R}} g(\xi) (\phi * \alpha)(\xi) d\xi$ for any function g . Then for any $\alpha \in A$ and $\epsilon > 0$ we have

$$\begin{aligned} E_{f_0} \left(\log \left(\sup_{\beta \in S_m: d(\alpha, \beta) < \epsilon} (\phi * \beta)(\xi) \right) \right) &\leq E_{f_0} \left(\log \left(\sup_{\beta: d(\alpha, \beta) < \epsilon} (\phi * \beta)(\xi) \right) \right) \\ &\leq \int_{\mathbb{R}} \log(\|\phi\|_\infty) \phi * f_0(\xi) d\xi \\ &< \infty. \end{aligned}$$

Step 3

Next we show that S_m is compact on (A, d) . Throughout, we use the notation \rightarrow_u to indicate uniform convergence. To show the claim, choose $\{\alpha_l\}$ a sequence in S_m . Then since $\{(\log(\alpha_l))^{(k+1)}\}$ are T_m -Lipschitz and uniformly bounded it follows by Arzela-Ascoli Theorem that there exists a sub-sequence $\{\alpha_{1,l}\} \subset \{\alpha_l\}$ such that $(\log(\alpha_{1,l}))^{(k+1)} \rightarrow_u g_{k+1}$ in $[-1, 1]$ for some function $g_{k+1} : [-1, 1] \rightarrow \mathbb{R}$ which is also T_m -Lipschitz. Note that we can again use Arzela-Ascoli Theorem applied to the sequence $\{\alpha_{1,l}\}$ to ensure that there exists a sub-sequence $\{\alpha_{2,l}\} \subset \{\alpha_{1,l}\}$ such that $(\log(\alpha_{2,l}))^{(k+1)} \rightarrow_u g_{k+1}$, in $[-2, 2]$. Thus we extend the domain of g_{k+1} if necessary.

Proceeding by induction we conclude that for every $N \in \mathbb{N}$ there exists a sequence $\{\alpha_{N,l}\}_{l \in \mathbb{N}} \subset \{\alpha_{N-1,l}\}_{l \in \mathbb{N}}$ such that

$$(\log(\alpha_{N,l}))^{(k+1)} \rightarrow_u g_{k+1}, \text{ as } l \rightarrow \infty,$$

in $[-N, N]$ as $l \rightarrow \infty$. Hence with Cantor's diagonal argument we conclude that there exists a sub-sequence $\{\alpha_{l_j}\} \subset \{\alpha_l\}$ such that

$$(\log(\alpha_{l_j}))^{(k+1)} \rightarrow_u g_{k+1} \text{ as } j \rightarrow \infty,$$

in $[-N, N]$ for all $N \in \mathbb{N}$. Since $|(\log(\alpha_{l_j}))^{(k)}(0)| \leq T_m$ for all j . Then without loss of generality, we can assume that

$$(\log(\alpha_{l_j}))^{(k)} \rightarrow_u g_k \text{ as } j \rightarrow \infty,$$

in $[-N, N]$ for all $N \in \mathbb{N}$ and where the function g_k satisfies $g'_k = g_{k+1}$. Continuing with this process we can assume, without loss of generality, that

$$\log(\alpha_{l_j}) \rightarrow_u g \text{ as } j \rightarrow \infty,$$

in $[-N, N]$ for all $N \in \mathbb{N}$ for some function g satisfying $g^{(j)} = g_j$ for all $0 \leq j \leq k + 1$. Therefore,

$$\alpha_{l_j} \rightarrow_u \exp(g) \text{ as } j \rightarrow \infty, \tag{14}$$

in $[-N, N]$ for all $N \in \mathbb{N}$.

Let us now prove that $\exp(g) \in S_m$. First, we observe by the Fatou's lemma e^g is integrable in \mathbb{R} with respect to the Lebesgue measure. Since S_m is tight, we obtain

$$d(\exp(g), \alpha_{I_j}) \rightarrow 0.$$

This clearly also implies that $\exp(g)$ integrates to 1 or $\exp(g) \in \mathcal{P}$. Note that also by Fatou's lemma we have that $J_{k,q}(g) \leq K_m$ and by construction,

$$\max\left(\|\exp(g)\|_\infty, \|g^{(k+1)}\|_\infty, |g^{(k)}(0)|, \dots, |g(0)|\right) \leq T_m.$$

Finally, combining all of this with g^{k+1} being T_m -Lipschitz, we arrive to $\exp(g) \in S_m$.

Step 4 By assumption (4), we have that

$$\sup_{\alpha \in A_m} d(f_0, \alpha) \rightarrow 0, \quad \text{as } m \rightarrow \infty.$$

Step 5

Let us show that

$$\lim_{\epsilon \rightarrow 0} E_{f_0} \left(\sup_{d(\alpha, \beta) < \epsilon, \beta \in S_m} \log(\phi * \beta) \right) = E_{f_0}(\log(\phi * \alpha))$$

for all $\alpha \in S_m$. First, note that for all ξ

$$0 \leq \max \left\{ 0, \sup_{d(\alpha, \beta) < \epsilon, \beta \in S_m} \log(\phi * \beta)(\xi) \right\} \leq \max \{0, \log(\|\phi\|_\infty)\}.$$

Hence, by Step 1 we obtain

$$0 \leq \lim_{\epsilon \rightarrow 0} E_{f_0} \left(\max \left\{ 0, \sup_{d(\alpha, \beta) < \epsilon, \beta \in S_m} \log(\phi * \beta) \right\} \right) = E_{f_0}(\max \{\log(\phi * \alpha), 0\}) < \infty.$$

Now we observe that

$$0 \leq -\min \left\{ 0, \sup_{d(\alpha, \beta) < \epsilon, \beta \in S_m} \log(\phi * \beta)(\xi) \right\} \leq -\min \{0, \log(\phi * \alpha(\xi))\},$$

and the claim follows from the monotone convergence theorem.

If $x_j = 1$ and $\xi_j = y_j$ for all $j = 1, \dots, D_n$, the claim of Theorem 1 follows from Theorem 1 from Geman and Hwang (1982). Otherwise, we continue the proof below. In either case we can see that the solution set M_m^n is not empty given that the map $\alpha \rightarrow \phi * \alpha(\xi)$ is continuous with respect to the metric d for any ξ .

Step 6

Note that, by Glivenko-Cantelli Theorem and our assumption on the maximum number of bins, we have that, almost surely, the random distribution

$$G_n(\xi) = \sum_{j=1}^{D_n} \frac{x_j}{n} I_{(-\infty, \xi]}(\xi_j)$$

converges weakly to the distribution function associated with $\phi * f_0$. Hence, almost surely, from the Portmanteau theorem we have for any $\alpha \in S_{m_n}$ and any $\delta > 0$ it holds that

$$\begin{aligned} \limsup_{l \rightarrow \infty} \sum_{j=1}^{D_l} \frac{x_j}{l} \sup_{d(\alpha, \beta) < \delta, \beta \in S_m} \log(\phi * \beta)(\xi_j) \\ \leq E_{f_0} \left(\sup_{d(\alpha, \beta) < \delta, \beta \in S_m} \log(\phi * \beta)(\xi) \right), \end{aligned} \tag{15}$$

since the function

$$\xi \rightarrow \sup_{d(\alpha, \beta) < \delta, \beta \in S_m} \log(\phi * \beta)(\xi),$$

is continuous and bounded by above.

Next we define

$$m_1 = \min \left\{ m : \sup_{\alpha \in A_m} d(\alpha, f_0) < \frac{1}{2} \right\}.$$

Clearly, $\beta_1, \beta_2 \in A_{m_1}$ implies $d(\beta_1, \beta_2) < 1$. Also, we see that the set $\Pi_1 := \{\alpha \in S_{m_1} : d(\alpha, f_0) \geq 1\} \subset S_{m_1} - A_{m_1}$ is d-compact. Hence, there exists $\alpha_1^1, \dots, \alpha_{h_1}^1$ in Π_1 such that $\Pi_1 \subset \cup_{l=1}^{h_1} \{\alpha \in \Pi_1 : d(\alpha, \alpha_l^1) < \delta_{1,l}\}$ for positive constants $\{\delta_{1,l}\}$ satisfying that

$$E_{f_0} \left(\sup_{d(\alpha, \alpha_l^1) < \delta_{1,l}, \alpha \in \Pi_1} \log(\phi * \alpha) \right) < E_{f_0}(\log(\phi * f_{m_1}))$$

for $l = 1, \dots, h_1$. Therefore from our assumptions on the sets A_m and also from (15), we arrive at

$$\begin{aligned} \limsup_{r \rightarrow \infty} \sum_{j=1}^{D_r} \frac{x_j}{r} \left(\sup_{d(\alpha, \alpha_l^1) < \delta_{1,l}, \alpha \in \Pi_1} \log(\phi * \alpha)(\xi_j) - \log(\phi * f_{m_1})(\xi_j) \right) \\ \leq E_{f_0} \left(\sup_{d(\alpha, \alpha_l) < \delta_l, \alpha \in \Pi_1} \log(\phi * \alpha) \right) - E_{f_0}(\log(\phi * f_{m_1})) \\ < 0, \quad \text{a.s.} \end{aligned}$$

Hence

$$\limsup_{r \rightarrow \infty} \sum_{j=1}^{D_r} x_j \left(\sup_{d(\alpha, \alpha_l^1) < \delta_{1,l}, \alpha \in \Pi_1} \log(\phi * \alpha)(\xi_j) - \log(\phi * f_{m_1}(\xi_j)) \right) = -\infty \text{ a.s.}$$

This implies

$$\lim_{r \rightarrow \infty} \frac{\sup_{\alpha \in \Pi_1} \prod_{j=1}^{D_r} \phi * \alpha(\xi_j)^{x_j}}{\prod_{j=1}^{D_r} \phi * f_{m_1}(\xi_j)^{x_j}} = 0 \text{ a.s.}$$

Next we define

$$k_1 = \min \left\{ k_0 : r \geq k_0 \text{ implies } \frac{\sup_{\alpha \in \Pi_1} \prod_{j=1}^{D_r} \phi * \alpha(\xi_j)^{x_j}}{\prod_{j=1}^{D_r} \phi * f_{m_1}(\xi_j)^{x_j}} < 1 \right\}$$

and we set $m_{k_1} = m_1$. Therefore,

$$\sup_{\alpha \in M_{m_{k_1}}^{k_1}} d(\alpha, f_0) < 1.$$

Next we define m_2 as

$$m_2 = \min \left\{ m \geq m_1 : \sup_{\alpha \in A_m} d(\alpha, f_0) < \frac{1}{4} \right\} + 1$$

Then, $\beta_1, \beta_2 \in A_{m_2}$ implies $d(\beta_1, \beta_2) < 1/2$. We also see that $\Pi_2 := \{\alpha \in S_{m_2} : d(\alpha, f_0) \geq 1/2\} \subset S_{m_2} - A_{m_2}$ is d-compact. Hence there exists $\alpha_1^2, \dots, \alpha_{h_2}^2$ in Π_2 such that $\Pi_2 \subset \cup_{l=1}^{h_2} \{\alpha : d(\alpha, \alpha_l^2) < \delta_{2,l}\}$ for positive constants $\{\delta_{2,l}\}$ satisfying

$$E_{f_0} \left(\sup_{d(\alpha, \alpha_l^2) < \delta_{2,l}, \alpha \in \Pi_2} \log(\phi * \alpha) \right) < E_{f_0}(\log(\phi * f_{m_2}))$$

for $l = 1, \dots, h_2$. Therefore,

$$\begin{aligned} & \limsup_{r \rightarrow \infty} \sum_{j=1}^{D_r} \frac{x_j}{r} \left(\sup_{d(\alpha, \alpha_l^2) < \delta_{2,l}, \alpha \in \Pi_2} \log(\phi * \alpha)(\xi_j) - \log(\phi * f_{m_2}(\xi_j)) \right) \\ & \leq E_{f_0} \left(\sup_{d(\alpha, \alpha_l^2) < \delta_{2,l}, \alpha \in \Pi_2} \log(\phi * \alpha) \right) - E_{f_0}(\log(\phi * f_{m_2})) \\ & < 0 \text{ a.s.} \end{aligned}$$

So proceeding as before we obtain

$$\lim_{r \rightarrow \infty} \frac{\sup_{\alpha \in \Pi_2} \prod_{j=1}^{D_r} \phi * \alpha(\xi_j)^{x_j}}{\prod_{j=1}^{D_r} \phi * f_{m_2}(\xi_j)^{x_j}} = 0 \text{ a.s.}$$

Finally we define

$$k_2 = \min \left\{ k_0 : k_0 \geq k_1 \text{ and } r \geq k_0 \text{ implies } \frac{\sup_{\alpha \in \Pi_2} \prod_{j=1}^{D_r} \phi * \alpha(\xi_j)^{x_j}}{\prod_{j=1}^{D_r} \phi * f_{m_2}(\xi_j)^{x_j}} < 1 \right\},$$

and we set $m_k = m_1$ for all $k_1 \leq k < k_2$ and $m_{k_2} = m_2$. By construction, we have that

$$\sup_{\alpha \in M_{m_{k_2}}^{k_2}} d(\alpha, f_0) < 1/2.$$

Thus an induction argument allow us to conclude the proof. □

Appendix B: Simulation details

For the cases of the true mixing density we consider four densities of the form

$$f_0(\mu) = \sum_{i=1}^K w_i N(\mu \mid \theta_i, \sigma_i^2).$$

In all cases considered here, the observations arise as in (1) with a standard normal sampling model. In our first example we evaluate performance for samples of a density that has four peaks or explicitly $K = 4$, $w = (0.2, 0.3, 0.3, 0.2)$, $\theta = (-3, -1.5, 1.5, 3)$ and with small variance $\sigma^2 = (0.01, 0.01, 0.01, 0.01)$. For the second example we consider a mixture of three normals two of which are smooth while the other has a peak. The true parameters in this case are $K = 3$, $w = (1/3, 1/3, 1/3)$, $\theta = (0, -2, 3)$ and $\sigma^2 = (2, .1, .4)$. The next example is a mixture of $K = 3$ normals, one of which has very high variance. The true parameters chosen are $w = (0.3, 0.4, 0.3)$, $\theta = (0, 0, 0)$ and $\sigma^2 = (0.1, 1, 9)$. Our final example is a mixture, with $K = 3$, giving raise to a very smooth density, the parameters are $w = (0.5, 0.4, .1)$, $\theta = (-1.5, 1.5, 4)$ and $\sigma^2 = (1, 2, 2)$. A visualization of these examples is shown in Figure 5.

B.1. Sensitivity to the number of bins

Figure 7 shows the performance of both L1-D and L2-D generally improves as we increase D . However, based on our experience, $D = 250$ is a reasonable choice. Specially for L1-D whose computational burden increases more rapidly. For $D = 250$ it typically takes around 5 min to compute the solution path for L1-D with 50 values of the regularization parameter. In contrast, L2-D only requires around 40 seconds under the same setting.

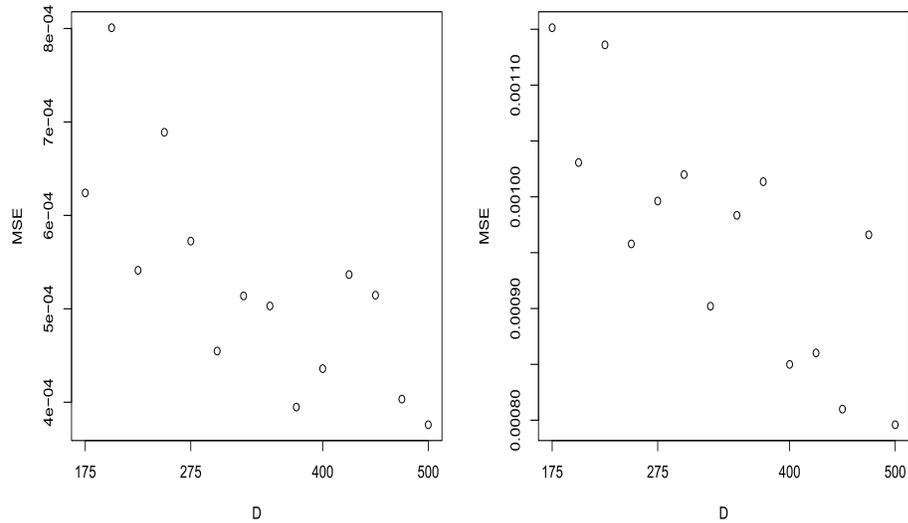


Fig 7: Sensitivity of our methods to the parameter D . We consider $n = 10^6$ samples from Example 1 in the paper. The first panel corresponds to L2-D while the second on to L1-D. The results show the average MSE on a 95% mass interval of the true mixing density. The average is taken over 20 MC simulations.

References

- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends[®] in Machine Learning*, 3(1):1–122, 2011.
- L. D. Brown and E. Greenshtein. Nonparametric empirical bayes and compound decision approaches to estimation of a high-dimensional vector of normal means. *The Annals of Statistics*, pages 1685–1704, 2009. [MR2533468](#)
- R. Carroll, A. Delaigle, and P. Hall. Deconvolution when classifying noisy data involving transformations. *Journal of the American Statistical Association*, 107(499):1166–1177, 2012. [MR3010903](#)
- R. J. Carroll and P. Hall. Optimal rates of convergence for deconvolving a density. *Journal of the American Statistical Association*, 83(404):1184–1186, 1988. [MR0997599](#)
- A. Delaigle. Nonparametric kernel methods with errors-in-variables: Constructing estimators, computing them, and avoiding common mistakes. *Australian & New Zealand Journal of Statistics*, 56(2):105–124, 2014. [MR3226432](#)
- A. Delaigle and I. Gijbels. Estimation of integrated squared density derivatives from a contaminated sample. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):869–886, 2002. [MR1979392](#)
- A. Delaigle and P. Hall. Parametrically assisted nonparametric estimation of a density in the deconvolution problem. *Journal of the American Statistical*

- Association*, 109(506):717–729, 2014. [MR3223745](#)
- K.-A. Do, P. Muller, and F. Tang. A Bayesian mixture model for differential gene expression. *Journal of the Royal Statistical Society, Series C*, 54(3):627–44, 2005. [MR2137258](#)
- S. Donnet, V. Rivoirard, J. Rousseau, and C. Scricciolo. Posterior concentration rates for empirical bayes procedures, with applications to dirichlet process mixtures. *arXiv preprint arXiv:1406.4406*, 2014. [MR3706755](#)
- B. Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–14, 2011. [MR2896860](#)
- B. Efron. Empirical bayes deconvolution estimates. *Biometrika*, 103(1):1–20, 2016. [MR3465818](#)
- M. D. Escobar and M. West. Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90:577–88, 1995. [MR1340510](#)
- J. Fan. On the optimal rates of convergence for nonparametric deconvolution problems. *The Annals of Statistics*, pages 1257–1272, 1991. [MR1126324](#)
- J. Fan and J.-Y. Koo. Wavelet deconvolution. *Information Theory, IEEE Transactions on*, 48(3):734–747, 2002. [MR1889978](#)
- T. S. Ferguson. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1:209–30, 1973. [MR0350949](#)
- S. Geman and C.-R. Hwang. Nonparametric maximum likelihood estimation by the method of sieves. *The Annals of Statistics*, 10(2):401–14, 1982. [MR0653512](#)
- S. Ghosal and A. W. Van Der Vaart. Entropies and rates of convergence for maximum likelihood and bayes estimation for mixtures of normal densities. *The Annals of Statistics*, pages 1233–1263, 2001. [MR1873329](#)
- I. J. Good and R. A. Gaskins. Nonparametric roughness penalties for probability densities. *Biometrika*, 58(2):255–77, 1971. [MR0319314](#)
- P. Hall, A. Meister, et al. A ridge-parameter approach to deconvolution. *The Annals of Statistics*, 35(4):1535–1558, 2007. [MR2351096](#)
- H. Ishwaran and M. Zarepour. Exact and approximate sum representations for the dirichlet process. *The Canadian Journal of Statistics/La Revue Canadienne de Statistique*, pages 269–283, 2002. [MR1926065](#)
- W. Jiang and C.-H. Zhang. General maximum likelihood empirical bayes estimation of normal means. *The Annals of Statistics*, 37(4):1647–1684, 2009. [MR2533467](#)
- N. A. Johnson. A dynamic programming algorithm for the fused lasso and l0-segmentation. *Journal of Computational and Graphical Statistics*, 22(2):246–260, 2013. [MR3173713](#)
- J. Kiefer and J. Wolfowitz. Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *The Annals of Mathematical Statistics*, 27:887–906, 1956. [MR0086464](#)
- R. Koenker. Rebayes: empirical bayes estimation and inference in r. *R package version 0.41*, 2013.
- R. Koenker and I. Mizera. Convex optimization, shape constraints, compound decisions, and empirical bayes rules. *Journal of the American Statistical As-*

- sociation, 109(506):674–685, 2014. [MR3223742](#)
- M. Lee, P. Hall, H. Shen, J. S. Marron, J. Tolle, and C. Burch. Deconvolution estimation of mixture distributions with boundaries. *Electronic journal of statistics*, 7:323, 2013. [MR3020423](#)
- R. Martin and S. T. Tokdar. Semiparametric inference in mixture models with predictive recursion marginal likelihood. *Biometrika*, 98(3):567–582, 2011. [MR2836407](#)
- R. Martin and S. T. Tokdar. A nonparametric empirical Bayes framework for large-scale multiple testing. *Biostatistics*, 13(3):427–39, 2012.
- O. Muralidharan. An empirical bayes mixture method for effect size and false discovery rate estimation. *The Annals of Applied Statistics*, pages 422–438, 2010. [MR2758178](#)
- M. A. Newton. On a nonparametric recursive estimator of the mixing distribution. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 306–322, 2002. [MR1981761](#)
- O. H. M. Padilla and J. G. Scott. Nonparametric density estimation by histogram trend filtering. *arXiv preprint arXiv:1509.04348*, 2015.
- A. Ramdas and R. J. Tibshirani. Fast and flexible ADMM algorithms for trend filtering. Technical report, Carnegie Mellon University, <http://www.stat.cmu.edu/~ryantibs/papers/fasttf.pdf>, 2014.
- L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(259–68), 1992. [MR3363401](#)
- A. Sarkar, B. K. Mallick, J. Staudenmayer, D. Pati, and R. J. Carroll. Bayesian semiparametric density deconvolution in the presence of conditionally heteroscedastic measurement errors. *Journal of Computational and Graphical Statistics*, 23(4):1101–1125, 2014a. [MR3270713](#)
- A. Sarkar, D. Pati, B. K. Mallick, and R. J. Carroll. Bayesian semiparametric multivariate density deconvolution. *arXiv preprint arXiv:1404.6462*, 2014b.
- B. W. Silverman. On the estimation of a probability density function by the maximum penalized likelihood method. *The Annals of Statistics*, pages 795–810, 1982. [MR0663433](#)
- D. Singh, P. G. Febbo, K. Ross, D. G. Jackson, J. Manola, C. Ladd, P. Tamayo, A. A. Renshaw, A. V. D’Amico, J. P. Richie, E. S. Lander, M. Loda, P. W. Kantoff, T. R. Golub, and W. R. Sellers. Gene expression correlates of clinical prostate cancer behavior. *Cancer Cell*, 1(2):203–9, 2002.
- J. Staudenmayer, D. Ruppert, and J. P. Buonaccorsi. Density estimation in the presence of heteroscedastic measurement error. *Journal of the American Statistical Association*, 103(482):726–736, 2008. [MR2524005](#)
- L. A. Stefanski and R. J. Carroll. Deconvolving kernel density estimators. *Statistics*, 21(2):169–184, 1990. [MR1054861](#)
- W. Tansey, O. Koyejo, R. A. Poldrack, and J. G. Scott. False discovery rate smoothing. Technical report, University of Texas at Austin, 2014. <http://arxiv.org/abs/1411.6144>.
- R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society (Series B)*, 67:91–108, 2005. [MR2136641](#)

- R. J. Tibshirani. Adaptive piecewise polynomial estimation via trend filtering. *The Annals of Statistics*, 42(1):285–323, 2014. [MR3189487](#)
- R. J. Tibshirani and J. Taylor. Degrees of freedom in lasso problems. *The Annals of Statistics*, 40(2):1198–1232, 2012. [MR2985948](#)
- S. T. Tokdar, R. Martin, and J. K. Ghosh. Consistency of a recursive estimate of mixing distributions. *The Annals of Statistics*, pages 2502–2522, 2009. [MR2543700](#)
- S. Wager. A geometric approach to density estimation with additive noise. *Statistica Sinica*, 2013. [MR3235388](#)
- A. Wald. Note on the consistency of the maximum likelihood estimate. *The Annals of Mathematical Statistics*, pages 595–601, 1949. [MR0032169](#)
- C.-H. Zhang. Fourier methods for estimating mixing densities and distributions. *The Annals of Statistics*, pages 806–831, 1990. [MR1056338](#)