

BAYESIAN NONPARAMETRIC MODELS FOR PEAK IDENTIFICATION IN MALDI-TOF MASS SPECTROSCOPY¹

BY LEANNA L. HOUSE²,
MERLISE A. CLYDE^{2,3} AND ROBERT L. WOLPERT⁴

Virginia Tech, Duke University and Duke University

We present a novel nonparametric Bayesian approach based on Lévy Adaptive Regression Kernels (LARK) to model spectral data arising from MALDI-TOF (Matrix Assisted Laser Desorption Ionization Time-of-Flight) mass spectrometry. This model-based approach provides identification and quantification of proteins through model parameters that are directly interpretable as the number of proteins, mass and abundance of proteins and peak resolution, while having the ability to adapt to unknown smoothness as in wavelet based methods. Informative prior distributions on resolution are key to distinguishing true peaks from background noise and resolving broad peaks into individual peaks for multiple protein species. Posterior distributions are obtained using a reversible jump Markov chain Monte Carlo algorithm and provide inference about the number of peaks (proteins), their masses and abundance. We show through simulation studies that the procedure has desirable true-positive and false-discovery rates. Finally, we illustrate the method on five example spectra: a blank spectrum, a spectrum with only the matrix of a low-molecular-weight substance used to embed target proteins, a spectrum with known proteins, and a single spectrum and average of ten spectra from an individual lung cancer patient.

Received July 2010; revised November 2010.

¹Supported by NSF Grant DMS-04-22400. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

²Supported in part by NSF Grant DMS-03-42172.

³Supported in part by NSF Grant DMS-04-06115 and NIH Grant R01-HL090559-01.

⁴Supported in part by NSF Grants DMS-07-57549 and PHY-09-41373, and NASA Grant NNX09AK60G.

Key words and phrases. Gamma random field, kernel regression, Lévy random fields, reversible jump Markov chain Monte Carlo, wavelets.

<p>This is an electronic reprint of the original article published by the Institute of Mathematical Statistics in <i>The Annals of Applied Statistics</i>, 2011, Vol. 5, No. 2B, 1488–1511. This reprint differs from the original in pagination and typographic detail.</p>
--

1. Introduction. Recent innovations in protein separation methods, ionization procedures and detection algorithms have led mass spectrometry (MS) to play a vital role in the explosive growth of proteomics [Dass (2001), page xxi]. Despite technological advances in data collection, it remains challenging to extract biologically relevant information (such as biomarkers) from MS spectral data [Dass (2001), Chapters 3 and 5; Coombes et al. (2005a); Baggerly, Morris and Coombes (2004, 2006); Clyde, House and Wolpert 2006; Morris et al. (2006)].

Identifying peak locations (which represent proteins) and quantifying protein abundance in spectra is often preceded by a multi-stage analysis involving calibration, normalization, baseline subtraction and filtering of noise [Morris et al. (2005); Tibshirani et al. (2004); Yasui et al. (2003)]. A problem with such an approach is that each individual step may introduce errors, artifacts or biases that may interfere with later stages of the analyses such as classification of subjects or identification of biomarkers. Methods that model background, noise and features simultaneously may lead to improved classification or inferences [Coombes et al. (2005b)]. Nonparametric methods such as wavelet regression have proved successful in simultaneously modeling background and denoising, allowing one to extract features or regions of spectra that differentiate groups [Yasui et al. (2003); Coombes et al. (2005b); Wang, Ray and Mallick (2007); Morris et al. (2008)]. While wavelets are well suited for modeling local features like spectral peaks, neither the scales and locations that index basis functions nor coefficients used in the wavelet representation of expected intensity have any inherent biological interpretation for the typical wavelets used in practice, such as Daubechies' "least asymmetric" $1a8$ wavelet family. As an alternative to nonparametric regression for modeling intensities, Guindani et al. (2006) and Müller et al. (2010) developed a Bayesian mixture model based on beta distributions to estimate a density function for time-of-flight. The parameters of this model are more interpretable than the wavelet regression methods, but the approach does not incorporate information about peak resolution, which we will show allows one to resolve broad peaks into multiple peaks.

In this paper we propose a novel nonparametric method employing Lévy Adaptive Kernel Regression (LARK) models [Wolpert, Clyde and Tu (2011); Clyde and Wolpert (2007)], which retains the adaptivity and flexibility that make wavelet and other nonparametric methods appealing but, in contrast to these other methods, uses model parameters with direct biological interpretations. This offers the opportunity to elicit meaningful expert opinion to guide the selection of prior distributions, and a posteriori to provide posterior distributions on model parameters that are meaningful to the expert. The model presented in this article is intended for use with a "single" spectrum, although an entire set of exchangeable spectra may be analyzed by applying

the method to the mean spectrum, similar to the mean-spectrum undecimated wavelet thresholding (MUDWT) method of Morris et al. (2005).

The paper is arranged as follows. We begin in Section 2 with a brief overview of MALDI-TOF mass spectrometry. In Section 3 we develop a statistical model for protein abundance as a function of time-of-flight based on the recently developed nonparametric LARK models. Prior distributions for the model parameters are developed in Section 4 from elicited expert knowledge about the MALDI-TOF procedure and from exploratory analysis of MALDI-TOF data from related experiments. Inference about parameters of clinical interest are obtained from posterior distributions and are described in Section 5. We then illustrate the methodology in Section 6 using a series of experiments with real data. We use laboratory data from three experiments with known sources: blank spectra, which help to characterize noise in MALDI-TOF experiments; matrix spectra, which help to characterize background; and a known protein mixture, which illustrates the model’s ability to resolve masses. The LARK model is also applied to single and mean spectra from a recent lung cancer study conducted at the Duke University Medical Center. In Section 7 we compare the LARK model to the wavelet method of Morris et al. (2005) using mean spectra from the TOF simulator of Coombes et al. (2005b). We conclude with a discussion and suggestions for future work in Section 8.

2. MALDI-TOF data. In Matrix Assisted Laser Desorption Time-of-Flight Mass Spectrometry, or MALDI-TOF MS, inference about the molecular composition of a sample is based on indirect measurement of molecular masses. Molecules are initially embedded in a *matrix* of low molecular weight substance, such as sinapinic acid, and placed on a metal plate. The molecules are then simultaneously dislodged (by vaporizing the substrate) and ionized (by removing one or more electrons from molecules) by a series of laser pulses. The now-charged molecules are accelerated by a strong electric field toward a detector. In the Applied Biosystems Voyager DE Biospectrometry Workstation [Applied Biosystems (2001)], a linear detector measures ion abundance over time, then sends a signal at regular time intervals (*clock ticks*) to a digitizer for conversion to measured intensities; for the examples considered below, samples are taken at regular 4 ns or 16 ns intervals. The reported intensity in each time interval is typically the aggregate sum over several repeated laser “shots,” leading to what we will refer to as a single spectrum, with response ion intensity at corresponding time-of-flights (TOFs).

Figure 1 illustrates serum protein spectra from a single individual with lung cancer from a study conducted at the Duke Medical Center Radiology Department [Wang et al. (2003)]. Each serum sample was separated into 20 fractions along a pH gradient prior to the MALDI-TOF analysis to reduce

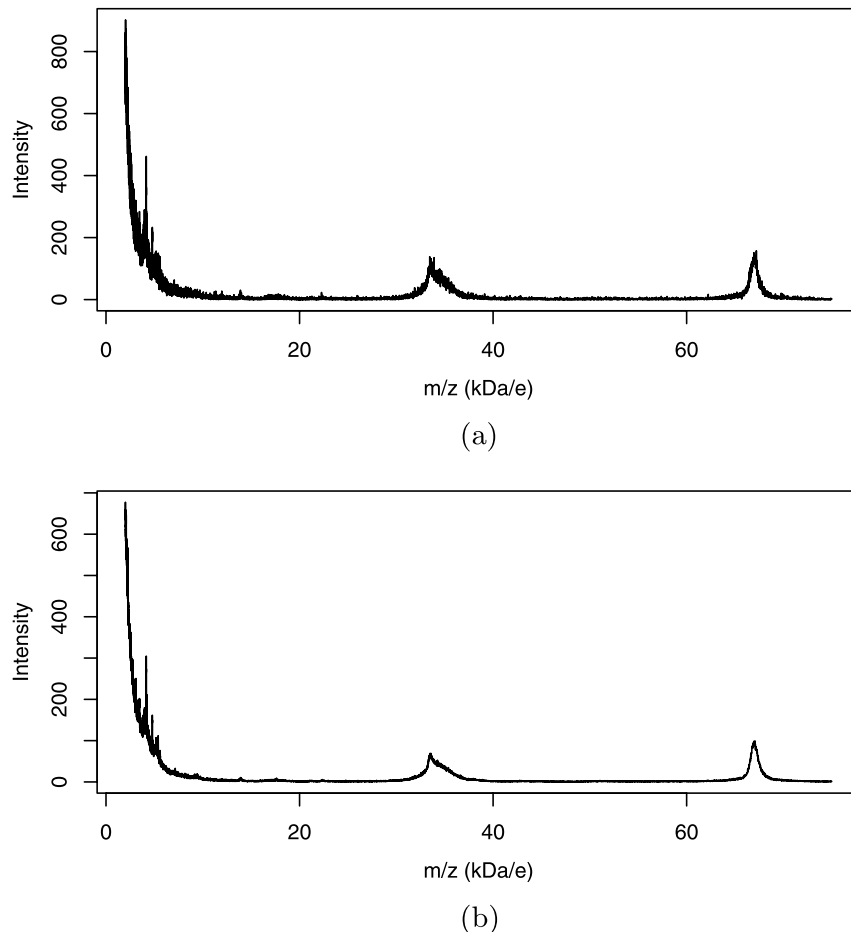


FIG. 1. *Single spectrum (a) and mean of ten spectra (b) from a lung cancer patient.*

saturation of the signal. Ten replicated spectra were obtained for each fraction, each with ten laser shots, using Voyager with a sinapinic acid matrix. For our analyses we randomly selected one fraction from one subject; Figure 1 shows the single spectrum from the chosen subject-fraction (a), and the mean spectrum obtained by averaging the intensities of all ten replicates (b) from the same fraction. Morris et al. (2005) suggest using the mean spectrum as a way to reduce noise from various sources as discussed below.

Distance traveled under constant acceleration is a quadratic function of time, leading to a simple but nonlinear relationship between TOF and the molecules' masses and ionic charge (the latter two enter only through their quotient, the *mass to charge ratio* m/z [Coombes et al. (2005a)]). Under ideal conditions the TOF spectrum would show a spike at the TOF corre-

sponding to each molecular species present. In actual MALDI-TOF spectra (Figure 1) we observe irregular peaks rather than one-dimensional spikes because molecules of equal size and charge do not all reach the detector at the same time. The most important of the many causes of TOF dispersion is variability in the amount of ionizing laser energy received by molecules of varying location within the matrix; those further from the matrix surface or from the center of the laser pulse may receive less kinetic energy and thus have lower initial velocities than similarly-sized molecules located closer to the center, delaying their arrival at the detector. Molecules may exchange energy in collisions, and may lose or gain mass through fragmentation and agglomeration, respectively. All these lead to TOF variation for each molecular species [Coombes et al. (2005a); Zhigilei and Garrison (1998); Franzen (1997)]. Furthermore, while the abundance of a protein in the sample ideally corresponds to the area under the TOF distribution curve, factors such as ion suppression, multiple charged ions, adducts (the addition of other molecules to the protein), protein-protein interactions and isotope distributions may result in a protein being represented by more than one peak in the spectrum.

The interpretation and analysis of MALDI-TOF data are complicated by several other sources of variation described by Morris et al. (2005) and Coombes et al. (2005a). In addition to *measurement error* which may mask or distort protein peaks, at least three other sources complicate the comparison or synthesis of multiple spectra: *calibration* (uncertainty in the conversion of TOF to m/z , including variable latency that affects time registration); *background* (a constant or even time-varying trend in the overall level); and *scale* (caused by many things including variability of laser intensity). One way to accommodate this is to construct models for peak identification and quantification that incorporate these recognized sources of variability, as in the wavelet approach of Coombes et al. (2005b), Morris et al. (2005) developed for calibrated spectra. Our approach, based on kernel models, has the added advantage that model parameters have direct physical interpretations.

3. A model for MALDI-TOF. To reduce the variability attributable to differing numbers of laser shots and differing baselines, we first standardize the spectrum and model

$$(3.1) \quad Y_t \equiv \frac{Y_t^{\text{ob}} - \min(\mathbf{Y}^{\text{ob}})}{l}$$

for a raw spectrum $\mathbf{Y}^{\text{ob}} \equiv \{Y_t^{\text{ob}}\}$ for $T_0 \leq t \leq T_1$, where T_0 and T_1 correspond to the range of TOFs of scientific interest and l is the number of laser shots summarized by \mathbf{Y}^{ob} . The initial molecular velocities are expected to be approximately Gaussian in distribution [Dass (2001), page 75]. This and the dynamics of the MALDI-TOF process [Coombes et al. (2005a)] suggest that

TOFs for a single isotopic peak will also have symmetric bell-shaped distributions in the time domain, leading us [and others—see Morris et al. (2005); Malyarenko et al. (2005)] to prefer time of flight (TOF, in μs) rather than mass-to-charge ratios (m/z , in Da/e) for spectral modeling (although we follow convention in reporting and plotting results in m/z). Because ions from the matrix may saturate the detector at initial TOFs [Malyarenko et al. (2005)] and masses less than 2 kDa were not of scientific interest to our collaborators, T_0 will correspond to the TOF of a mass of 2 kDa throughout, unless otherwise noted. While the nonparametric model that we propose can accommodate an arbitrary lower bound (even $T_0 = 0$), modeling these extra initial peaks will increase the running time of the algorithm, with little or no improvement in peak identification or model fit for the rest of the spectrum.

3.1. *Peak shape.* The shape of a symmetric isotopic peak may be represented by a probability density function for TOF t with parameters governing the protein peak’s location τ and width ω . Examples include the Gaussian

$$(3.2) \quad k(t; \tau, \omega) = \frac{1}{\sqrt{2\pi}\omega} \exp(-|t - \tau|^2/2\omega^2)$$

and Cauchy (sometimes called Lorentzian in the MS literature)

$$(3.3) \quad k(t; \tau, \omega) = \frac{\omega}{\pi(\omega^2 + |t - \tau|^2)},$$

as in Dass (2001), page 75, Kempka et al. (2004), and Applied Biosystems (2001), pages 6–30.

A protein signature associated with J peaks may now be represented as a sum

$$(3.4) \quad f(t) = \sum_{j=1}^J k(t; \tau_j, \omega_j) \eta_j,$$

where $\{\tau_j\}$, $\{\omega_j\}$ and $\{\eta_j\}$ represent the unknown location (TOF), peak width and abundance of the j th protein (or molecule), respectively. This is a special case of the Lévy Adaptive Regression Kernel (LARK) models of Wolpert, Clyde and Tu (2011), Clyde and Wolpert (2007), which generalize classical kernel regression [Wand and Jones (1995)] by allowing the number of kernels and the “smoothing parameters” (ω) of the kernel k to adapt to the unknown degree of smoothness in the data, as in wavelet models [Morris et al. (2005)]. While both methods lead to excellent function reconstructions, the parameters in the kernels (τ_j, ω_j) and kernel coefficients η_j of the LARK model have direct biological interpretations which aide in prior specification (detailed below) and posterior interpretation.

3.2. *Peak width and resolution.* Protein peaks tend to be broader for late-arriving molecules than for earlier ones, with width nearly proportional to arrival time [Siuzdak (2003), page 44]; for this reason it is conventional in mass spectrometry to quantify the precision (narrowness) of a kernel $k(\cdot; \tau, \omega)$ not by the scale ω , but by the *resolution*

$$(3.5) \quad \rho \equiv \tau / \Delta\tau,$$

where $\Delta\tau$, the so-called *full width at half mass* or FWHM, is the width of the kernel $k(\cdot; \tau, \omega)$ at half its height [Dass (2001), page 120]. For a symmetric kernel, $\Delta\tau$ is the solution of the equation

$$k(\tau \pm \frac{1}{2}\Delta\tau; \tau, \omega) = \frac{1}{2}k(\tau; \tau, \omega).$$

For the Gaussian and Cauchy kernels we have $\Delta\tau = 2\omega\sqrt{\log 4}$ and $\Delta\tau = 2\omega$, respectively, leading to $\omega = \omega(\tau, \rho)$ with

$$(3.6) \quad \omega(\tau, \rho) = \frac{\tau}{2\rho\sqrt{\log 4}} \quad \text{and} \quad \omega(\tau, \rho) = \frac{\tau}{2\rho},$$

for the Gaussian and Cauchy kernels, respectively. Prior knowledge about resolution can be used to resolve the ambiguity illustrated in Figure 2, where the observed spectrum may arise from either a single wide peak or a pair of nearby narrower peaks. As depicted later in Figure 6(d)–(f), we illustrate how the model is able to “deconvolve” a wide peak into several individual protein peaks in real data.

3.3. *Background noise sources.* Even in the absence of any protein molecules [i.e., with $f(t) \equiv 0$], the MALDI-TOF spectrum does not vanish. Figure 4(a) shows the nearly-constant level of thermal noise from a run with an empty plate, while Figure 4(b) illustrates the rapidly-decreasing signal

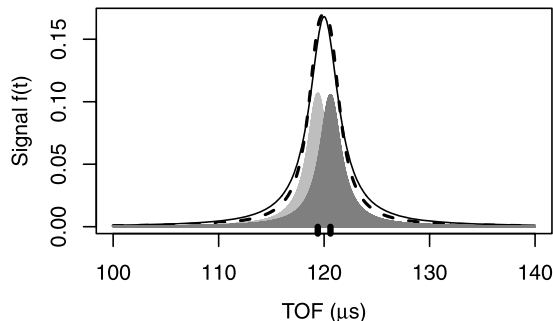


FIG. 2. The (nearly indistinguishable) lines represent simulated protein signals from a sample with either one wide peak (solid) centered at 120 μs or a mixture (dashed) of two narrow peaks centered at bold ticks.

with only the sinapinic acid matrix, showing the early arrival at the detector of ionized matrix molecules (far lighter than typical proteins under study). Since the signal from matrix ions dominates the thermal noise or detector “ringing,” together these sources contribute a background that falls off nearly exponentially to a nonzero asymptote.

Exploratory analysis suggests that the matrix molecular signal $\beta_0(t)$ can be modeled adequately as a constant (see below) plus an exponential function,

$$(3.7) \quad \beta_0(t) = k_0(t; \omega_0)\eta_0 = \frac{\eta_0}{\omega_0} \exp\{-t/\omega_0\}\mathbf{1}_{\{t>0\}},$$

with a characteristic decay time of $\omega_0 > 0$ and intensity $\eta_0 > 0$.

3.4. Expected spectral intensity. To reflect all of these features, we model the expected spectral intensity as

$$(3.8) \quad \mu(t) = \zeta\{(1 - S) + S[f(t) + \beta_0(t)]\}$$

for an overall scale ζ , a dimensionless signal-to-background ratio $S \in [0, 1]$, the protein signal $f(t)$ from equation (3.4), and the matrix molecular signature $\beta_0(t)$ from equation (3.7). The term S represents the proportion of observed intensity produced by molecular signal (both matrix and protein), rather than by thermal noise.

3.5. Likelihood. Both gamma and log-normal distributions are commonly used to model positive data like the standardized responses Y_t . The variance is proportional to the mean for gamma distributions and to the square of the mean for log-normals. Exploratory data analysis (from both a Box–Cox approach and a robust regression illustrated in Figure 3) suggests that the conditional variance of standardized MS data Y_t , given the mean, is nearly proportional to the first power of the mean, supporting the gamma model

$$(3.9) \quad Y_t | \mu(\cdot), \varphi \stackrel{\text{ind}}{\sim} \text{Ga}(\varphi\mu(t), \varphi),$$

with mean $\mu(t)$ and relative precision parameter φ (similar relationships hold for the other data sets, although the slopes vary). This leads to a measurement-error model with likelihood function

$$(3.10) \quad \mathcal{L}(\boldsymbol{\theta}; \mathbf{Y}) = \prod_{i=1}^n \text{Ga}(Y_i; \varphi\mu(t_i), \varphi)$$

for the parameter vector $\boldsymbol{\theta}$ comprising the conditional mean function $\mu(\cdot)$ [or, equivalently from equation (3.8), all of ζ , J , $\{\tau_j, \omega_j, \eta_j\}_{1 \leq j \leq J}$, S , ω_0 , and η_0] and φ . Here $\mathbf{Y} = \{Y(t_i)\}_{1 \leq i \leq n}$ represents the vector of standardized intensities from (3.1), and $\text{Ga}(y; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} y^{\alpha-1} e^{-\beta y} \mathbf{1}_{\{y>0\}}$ is the probability density function at $y \in \mathbb{R}$ for the gamma $\text{Ga}(\alpha, \beta)$ distribution.

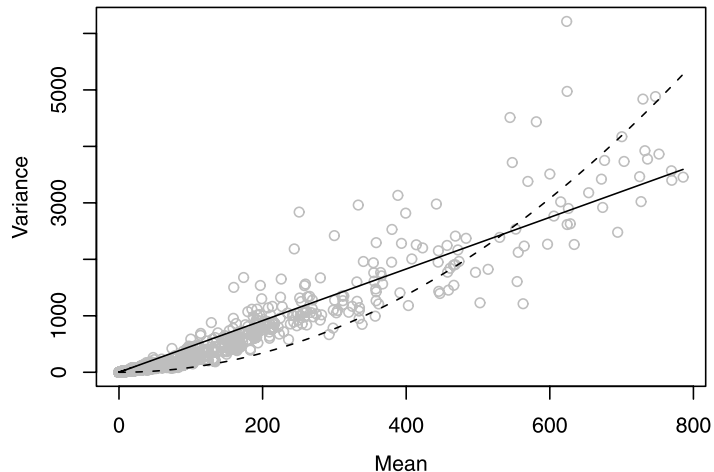


FIG. 3. *Linear (solid) and quadratic (dashed) fits of variance versus mean intensity of intensity for 200 μ s blocks of observations from a single spectrum from a lung cancer patient.*

Typically the likelihood function of equation (3.10) has many modes because it is difficult to distinguish wide peaks from clusters of narrow ones, or small peaks from noise, from the data alone. Estimating θ (and in particular J , the number of protein peaks) by direct maximization of the likelihood leads to over-fitting the data and to over-estimating J . This can be overcome by regularization or by a Bayesian approach, in which prior distributions effectively penalize overly complex models.

4. Prior distributions for MALDI-TOF. We now address the problem of constructing a joint prior distribution for all the unknown parameters of the model of Section 3,

$$\begin{aligned}
 Y_t | \mu(\cdot), \varphi &\stackrel{\text{ind}}{\sim} \text{Ga}(\varphi \mu(t), \varphi), \\
 \mu(t) &= \zeta \{ (1 - S) + S[f(t) + \beta_0(t)] \}, \\
 f(t) &= \sum_{j=1}^J k(t; \tau_j, \omega_j) \eta_j, \\
 \beta_0(t) &= \frac{\eta_0}{\omega_0} \exp\{-(t - T_0)/\omega_0\} \mathbf{1}_{\{t > T_0\}}.
 \end{aligned}
 \tag{4.1}$$

We discuss the approach we used to construct the distributions, employing public knowledge where possible or default procedures otherwise. Parameter values used for the different data sets are summarized in Table 1.

TABLE 1
Prior hyperparameters for each of the real and simulated data sets

Prior hyper-parameters	Data set					
	Blank	Matrix	Known	Single lung	Mean lung	Simulation study
μ_J	20	20	20	100	100	150
λ	0.03	0.03	0.03	0.11	0.11	0.65
ε	0.79	0.79	0.79	0.21	0.21	0.03
T_0	5.58	5.58	5.58	0.03	0.03	13.47
T_1	217.14	217.14	217.14	278.04	278.04	82.78
σ_ρ^2	0.1225	0.1225	0.1225	0.1225	0.1225	0.1225
μ_e	200	200	200	200	200	300
σ_e^2	0.49	0.49	0.49	0.49	0.49	0.49
a_φ	0.25	0.25	0.25	0.25	0.25	–
b_φ	0.02	0.11	25.91	1.14	0.17	–
a_S	1.36	8.33	8.71	9.46	7.29	6.33
b_S	1	1	1	1	1	1
λ_0	0.0009	0.0005	0.0012	0.0006	0.0004	0.0124
$\hat{\omega}_0$	171.68	290.14	127.30	210.96	268.52	11.22
$\sigma_{\omega_0}^2$	0.25	0.25	0.25	0.25	0.25	0.25

4.1. *Measurement error φ and overall level ζ .* The exploratory data analysis of experimental spectra (see Section 3.5) suggests that sample variances of $\{Y_t\}$ are nearly proportional to the mean. We chose a gamma prior distribution $\varphi \sim \text{Ga}(a_\varphi, b_\varphi)$ for the mean-to-variance ratio φ , centered at a data-based value but supporting a wide range of prior uncertainty. We binned the observations $\{Y_t\}$ of each data set in 50 μs -wide blocks and calculated their block-specific means and variances. The prior mean a_φ/b_φ was set to the slope of a regression of the block means on block variances, with $a_\varphi = 0.25$ to attain a coefficient of variation of 2.

The parameter ζ may be interpreted as the mean level or scale for Y_t , since $E[f(t)] \approx 1$ (see Section 4.2). Since experimental levels depend on a wide range of exogenous variables and vary widely among trials, it is difficult to elicit a subjective prior distribution for this quantity. We instead employed a data-dependent rescaling and set $\zeta \equiv \bar{Y}$, the overall mean intensity. This is comparable to rescaling the raw data by the average or total intensity. Sensitivity analysis showed that this gave results very similar to those under a tight data-dependent prior distribution.

4.2. *Prior distribution for protein signature $f(\cdot)$.* We specify a prior distribution for the protein signature

$$f(t) = \sum_{j=1}^J k(t; \tau_j, \omega_j) \eta_j$$

by first specifying a distribution for J and then, conditional on J , taking $\{\tau_j, \omega_j, \eta_j\}$ to be i.i.d. from a specified joint distribution [as in the infinitely divisible construction of the LARK models of Wolpert, Clyde and Tu (2011)]. To reflect uncertainty in the possible number of peaks, Wolpert, Clyde and Tu used a negative binomial distribution $\text{NB}(\alpha_J, \mu_J)$ for J with mean and shape parameters μ_J and α_J . To simplify elicitation while providing robust inference over a range of spectra, we set $\alpha_J = 1$ throughout, leading to a geometric distribution:

$$P[J = j | \mu_J] = \left(\frac{1}{1 + \mu_J} \right) \left(\frac{\mu_J}{1 + \mu_J} \right)^j, \quad j \in \{0, 1, \dots\}$$

with mean μ_J representing the expected number of peaks for a given experiment. Campa et al. (2003) found approximately fifty proteins for fractionated samples similar to the single and mean spectra described in Section 6.4, leading to perhaps seventy or so peaks due to adducts, multiply charged ions, etc. On this basis we chose $\mu_J = 100$ with a median of $J \approx 70$ peaks with symmetric 50%, 90% and 99% ranges of approximately $30 \leq J \leq 140$, $5 \leq J \leq 300$, and $0.50 \leq J \leq 532.5$, respectively. For the blank, matrix and known protein spectra, we set $\mu_J = 20$.

There is little reason to give higher prior probability to one range of TOFs than another without prior knowledge of the collection of proteins present in the samples. Thus, we take $\{\tau_j\}_{1 \leq j \leq J} \stackrel{\text{i.i.d.}}{\sim} \text{Un}(T_0, T_1)$ (independently of J and $\{\lambda_j\}_{1 \leq j \leq J}$), for some interval of length $T = T_1 - T_0$, large enough to include the TOF for all molecules of interest. To eliminate saturation by matrix molecules at the low end, and to include as wide as possible a range of the biologically relevant molecules, we chose a TOF interval corresponding to the range $[2 \text{ kDa/e} \leq m/z \leq 75 \text{ kDa/e}]$ for all experimental data. Differing sampling rates and calibration levels for different experiments lead the TOF ranges $[T_0, T_1]$ to vary across experiments (see Table 1).

We use expert opinion to construct an informed prior distribution on the resolutions $\{\rho_j\}_{1 \leq j \leq J}$ (see Section 3.2), which induces a distribution on the peak widths $\{\omega_j\}_{1 \leq j \leq J}$. It has been suggested [Siuzdak (2003), page 44] that individual peak resolutions ρ_j *should* be nearly constant across the entire TOF range, but in practice they are observed to vary [see pages 6–32 of Applied Biosystems (2001)]. To reflect this variation, we construct a hierarchical prior distribution for the resolution parameters $\{\rho_j\}_{1 \leq j \leq J}$ as follows. Independently of the number of peaks, TOFs and abundance, we take the resolutions to have log-normal distributions,

$$\begin{aligned} \log(\varrho) | \mu_\varrho, \sigma_\varrho^2 &\sim \text{N}(\log(\mu_\varrho), \sigma_\varrho^2), \\ \log(\rho_j) | \varrho, \sigma_\rho^2 &\stackrel{\text{i.i.d.}}{\sim} \text{N}(\log(\varrho), \sigma_\rho^2), \end{aligned}$$

centered at an overall “experiment” resolution ρ (which may be machine- or condition-specific). We use the manufacturer’s reported resolution ranges [Applied Biosystems (2001), Table 6-2 and Table H-6] for the Voyager workstation and set $\mu_\rho = 200$ and set $\sigma_\rho^2 = 0.49$, so a priori the distribution covers the range 50–800 with 95% probability, and 32–928 with 99% probability. The standard deviation for the individual resolutions was set to $\sigma_\rho = 0.35$, seventy percent of the population standard deviation for resolution. For a population resolution of $\mu_\rho = 50$, this leads to a prior 99% interval for individual resolutions of (20, 120), while at the upper extreme with a population resolution of $\mu_\rho = 800$, the prior 99% interval covers the range (325, 1,710). Finally, the relationship between width, TOF and resolution given by equation (3.6) induces a log-normal prior distribution on the width parameters,

$$\log(\omega_j) | \tau_j, \rho_j \stackrel{\text{ind}}{\sim} \mathbf{N}(\log(\tau_j/c\rho_j), \sigma_\rho^2), \quad j = 1, \dots, J,$$

with $c = 2$ for the Cauchy kernel and $c = 2\sqrt{\log 4}$ for the Gaussian.

For protein abundances $\{\eta_j\}$ we use the left-truncated gamma distribution $\text{Ga}(0, \lambda, \varepsilon)$ with parameters $\alpha = 0$ and λ, ε (chosen below), whose density function is given in general by

$$(4.2) \quad \text{Ga}(\eta; \alpha, \lambda, \varepsilon) \equiv \frac{\lambda^\alpha}{\Gamma(\alpha, \lambda\varepsilon)} \eta^{\alpha-1} e^{-\lambda\eta} \mathbf{1}_{\{\eta > \varepsilon\}},$$

where $\Gamma(\alpha, x) \equiv \int_x^\infty z^{\alpha-1} e^{-z} dz$ denotes the complimentary incomplete gamma function [Abramowitz and Stegun (1964), Section 6.5.3]. For $\alpha, \lambda > 0$ this is the conditional density for a gamma-distributed $\text{Ga}(\alpha, \lambda)$ random variable, given that it exceeds $\varepsilon \geq 0$; for strictly positive $\varepsilon > 0$, the distribution is well-defined for all $\alpha \in \mathbb{R}$ including the limiting case $\alpha = 0$ [Wolpert, Clyde and Tu (2011)], which we adopt. The mean is

$$\mathbb{E}[\eta] = \frac{1}{\lambda e^{\lambda\varepsilon} \mathbf{E}_1(\lambda\varepsilon)},$$

where $\mathbf{E}_1(z)$ denotes the exponential integral function [Abramowitz and Stegun (1964), page 228]. The parameter ε may be interpreted as the minimum detectable abundance. In the limit $\varepsilon \rightarrow 0$, this distribution permits an increasing number of isotopes with small abundance, while reflecting that only a few isotopes are expected to have large abundance.

Discussions with spectrometrists suggest that the smallest peak that can possibly be distinguished from noise is about 5–10% of the average signal. Using the midpoint of this interval, we take $\varepsilon/\mathbb{E}[\eta] = \lambda\varepsilon e^{\lambda\varepsilon} \mathbf{E}_1(\lambda\varepsilon) = 0.075$. For t well away from the boundary of $[T_0, T_1]$, $\int_{T_0}^{T_1} k(t; \tau_j, \omega_j) d\tau_j \approx 1$ (the kernels are density functions), and

$$(4.3) \quad \mathbb{E}[f(t)] \approx \frac{\mu J}{T \lambda e^{\lambda\varepsilon} \mathbf{E}_1(\lambda\varepsilon)}, \quad T_0 \ll t \ll T_1.$$

Because $f(t)$ is the normalized signal, we take $\mathbf{E}[f(t)] = 1$ a priori, leading to the solution $\varepsilon = 0.075T/\mu_J$. Inverting the function $\lambda\varepsilon e^{\lambda\varepsilon}\mathbf{E}_1(\lambda\varepsilon) = 0.075$, we obtain $\lambda\varepsilon = 0.0227$, which determines λ for any specified ε . Values for the different experiments are provided in Table 1.

4.3. *Prior distribution for matrix background.* As with other peaks, we use a log-normal distribution for the initial peak width ω_0 and a left-truncated gamma model for the abundance η_0 . Because the exponential decay varies greatly from experiment to experiment, we utilize modestly informative priors based on the data. To construct these, we first fit a linear regression with mean function $\log(\beta_0(t))$ from equation (3.7) to the log intensities in an initial segment of the spectrum ($2 \text{ kDa/e} < m/z < 3.5 \text{ kDa/e}$). We use the slope and intercept from this fit to construct point estimates $\widehat{\omega}_0$ and $\widehat{\eta}_0$ to center the prior distributions

$$\log(\omega_0) \sim \mathbf{N}(\log(\widehat{\omega}_0), \sigma_{\omega_0}^2)$$

with $\sigma_{\omega_0}^2 = 0.25$ and

$$\eta_0 \sim \text{Ga}(0, \widehat{\lambda}_0, \varepsilon),$$

where $\widehat{\lambda}_0$ is the solution to $(\widehat{\lambda}_0 e^{\varepsilon \widehat{\lambda}_0} \mathbf{E}_1(\varepsilon \widehat{\lambda}_0))^{-1} = \widehat{\eta}_0$.

Finally, for the signal fraction S we use a beta prior distribution

$$S \sim \text{Be}(a_S, b_S),$$

with data-based mean $a_S/(a_S + b_S) = 1 - \bar{Y}^N/\bar{Y}$ and $b_S = 1$ (here \bar{Y}^N is the observed mean intensity in a “noise” region of the spectrum with low intensity and no apparent peaks, while \bar{Y} is the overall mean intensity). With this choice the mode of the prior density for S is one whenever $\bar{Y} > 2\bar{Y}^N$, suggesting that the signal dominates, and zero when the mean in the noise region exceeds half the overall mean, favoring the thermal noise component over the nonparametric signal model.

5. Posterior analysis. To support inference about protein location and abundance, and about other model parameters, we construct an ergodic Markov chain on the space Θ of possible parameter vectors $\boldsymbol{\theta} = \{\zeta, J, \{\tau_j, \omega_j, \eta_j\}_{1 \leq j \leq J}, S, (\omega_0, \eta_0), \lambda_J, \rho\}$ with the posterior distribution as its stationary distribution. At each Markov chain step we select one of the components of $\boldsymbol{\theta}$ and either update it via a Gibbs step (replace the current value with a draw from its complete conditional posterior distribution given the other components) or, if this is not feasible, a random-walk Metropolis–Hastings (M–H) step by proposing a small change in that component which is then accepted or rejected according to the Hastings probabilities. Note that each proposed change in J (which we always take to be a random-walk step of size

one) changes the *dimension* of θ (by three). Such dimension changing M–H algorithms, introduced in Green (1995), are called reversible jump MCMC (RJ-MCMC) algorithms. Our approach is modeled after that of Wolpert and Ickstadt (2004), where a general RJ-MCMC procedure for Lévy random field models is presented. For updating the varying dimensional parameters $\{\tau_j, \omega_j, \eta_j\}_{1 \leq j \leq J}$ we consider the possible moves of three basic types: peak *birth* [incrementing J by one and introducing a new triplet $(\tau_*, \omega_*, \eta_*)$]; peak *death* [decrementing J by one and removing a randomly-chosen triplet $(\tau_j, \omega_j, \eta_j)$]; and peak *update* [moving a randomly-chosen triplet $(\tau_j, \omega_j, \eta_j)$ within \mathbb{R}^3]. We also incorporate two additional move types, peak *splitting*, in which a single peak is replaced by a pair of smaller ones, and the reverse move, peak *merging*, in which two nearby peaks are replaced with a single larger one. These lead to a vast improvement in algorithmic efficiency over RJ-MCMC algorithms using only birth/death and update steps.

For sufficiently large spectra or complex protein mixtures, convergence to the posterior distribution from random starting values may require upward of a million iterations. To reduce computation time, we begin the Markov chain close to a mode, located using an EM algorithm [Dempster, Laird and Rubin (1977)] for a simple Gaussian approximation to our LARK model; see House (2006) for details of this and the RJ-MCMC algorithm.

5.1. *Peak identification.* Features of the configuration $\{(\tau_j, \omega_j, \eta_j)\}_{1 \leq j \leq J}$ are updated at each iteration of the RJ-MCMC sampler, with the number of peaks J and associated parameters changing. While the posterior mean of J , J^{PM} , and the posterior mean function $\mathbb{E}[\mu(t)|Y]$ [see equation (4.1)] are well-defined quantities that may be used to summarize the RJ-MCMC output, the well known label switching problem complicates peak identification. We have two ways of identifying peaks in an RJ-MCMC run. The first is to use the J^{HP} peak locations $\{\tau_j^{\text{HP}}\}$ in the single RC-MCMC iteration with highest posterior (HP) density, which is proportional to the product of the likelihood function and prior density evaluated at the parameter vector for that iteration. Alternatively, we may use model averaging to identify local maxima in the denoised signal by identifying the J^∇ down-crossings of the derivative $(d/dt)\mathbb{E}[\mu(t)|Y] = \mathbb{E}[\mu'(t)|Y]$; these local modes $\{\tau_j^\nabla\}$ are the collection of J^∇ solutions of

$$(5.1) \quad \frac{d}{dt}\mathbb{E}[\mu(t-)|Y] > 0 > \frac{d}{dt}\mathbb{E}[\mu(t+)|Y].$$

Because derivatives of densities may be used to construct wavelets, the implied LARK model for the derivative process actually uses a continuous wavelet dictionary based on the “Mexican Hat” (or Marr’s) family [Vidakovic (1999), pages 48–49]. This method is closely related to the recent paper of Nguyen et al. (2010) who use zero-crossings of the derivatives of

Gaussian wavelets. Typically J^∇ is smaller than either J^{PM} or J^{HP} [one reason is that some pairs of peaks with small inter-peak distance $|\tau_j - \tau_k|$ will combine to generate a single local maximum of $\mu(t)$, but all the major peaks will be represented among the $\{\tau_j^\nabla\}$].

6. Examples. In this section we illustrate our method using five data sets: a blank spectrum, a matrix spectrum, a “spiked” spectrum from a sample of known protein composition, and two spectra (one single and one mean) for a serum sample from a patient diagnosed with lung cancer. All data were generated using the Voyager DE spectrometer [Applied Biosystems (2001)] at Duke University. Prior hyperparameters were chosen as described in Section 4 and are given in Table 1 for the five data sets. All examples in this section use the Cauchy kernel, selected because of its better fit to similar data in a preliminary investigation. All RJ-MCMC were run for 500,000 iterations to ensure convergence (burn-in), with an additional 500,000 iterations used for posterior inference.

6.1. *Blank spectrum.* Figure 4(a) shows the recorded spectrum from the average of ten blank-plate spectra each based on 32 laser shots, with the posterior mean $\mathbb{E}[\mu(t)|Y]$ shown as a solid curve. The two rows of tick marks on the horizontal axis represent peak locations identified using the highest posterior realization $\{\tau_j^{\text{HP}}\}$ (top row) and local maxima under model averaging $\{\tau_j^\nabla\}$ (bottom). The highest-posterior realization included $J^{\text{HP}} = 38$ peaks, the local modes under model averaging included $J^\nabla = 22$ peaks, while the posterior mean (and standard deviation) were $J^{\text{PM}} = \mathbb{E}[J|Y] = 46.92$ (4.22). With no solution or matrix on the metal plate, there can be no protein signature, but nevertheless the spectrum shows numerous low-resolution peaks. The posterior expected resolution in the blank spectrum was $\mathbb{E}[\varrho|Y] = 16.76$ (2.92), lower than the informative prior mean and significantly lower than the typical resolutions for protein peaks in the other examples (see Table 2). These apparent peaks may reflect laser fluctuations or resonances in the detector. The number of periods and intensities may vary unpredictably across spectroscopic samples, so rather than use a harmonic function as in Harezlak et al. (2008), we instead allow the adaptive LARK model to identify and fit these low-resolution peaks as part of $f(\cdot)$. They may be discriminated from protein peaks in post-processing by their lower resolution.

6.2. *Matrix spectrum.* Figure 4(b) shows the average of ten spectra, each based on 32 laser shots from a sinapinic matrix solution containing no protein serum sample. The posterior mean under the LARK model shows the characteristic near-exponential spectral fall-off arising from the very low-molecular-weight sinapinic acid matrix ions, as well as several low-resolution peaks [posterior mean for $\varrho = 15.03$ (1.70)] comparable to those in the blank spectrum.

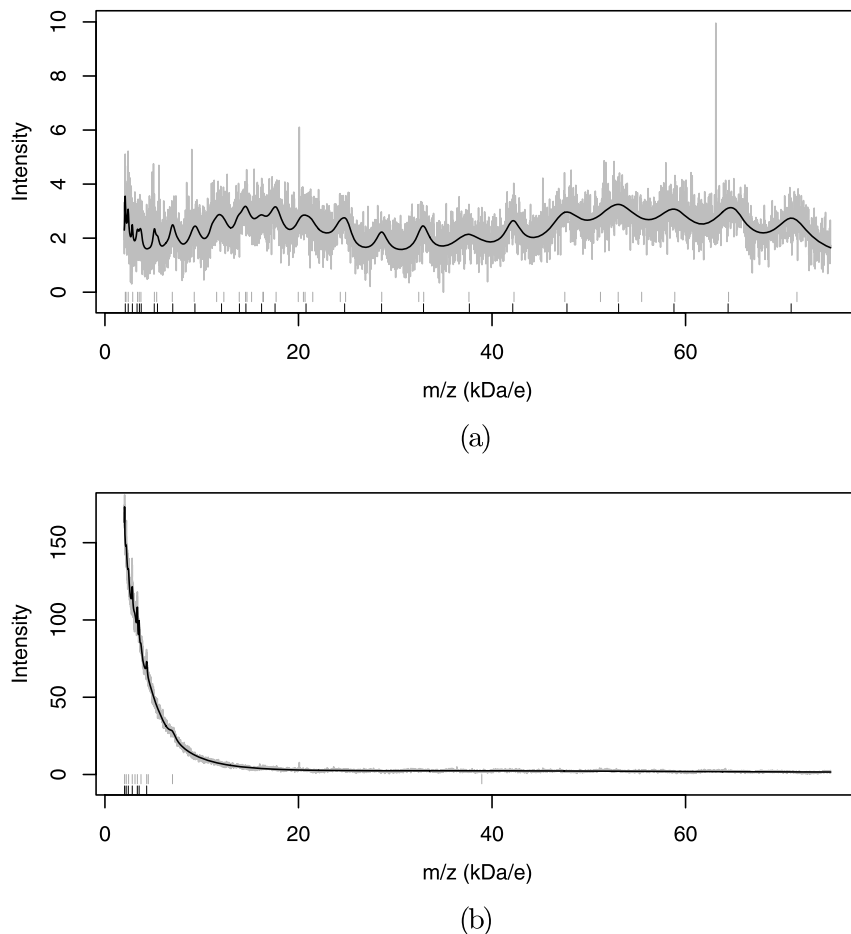


FIG. 4. The mean spectra based on ten replicates for (a) an empty plate illustrating thermal noise and (b) the sinapinic acid matrix with no proteins. The posterior means from LARK are shown as solid curves with identified peak locations indicated as two rows of tick marks: highest posterior realization (top, gray) and local modes under model averaging (bottom, black). Note the difference in scale of Y axes.

6.3. *Known protein spectrum.* Figure 5 shows the average of ten spectra (each with 32 shots) from a preparation of five proteins with known masses provided by Professor M. Fitzgerald in the Department of Chemistry at Duke. The five known masses of singly-charged molecules are indicated by solid triangles and the five peaks for doubly-charged molecules are indicated by open triangles; each open triangle for a doubly-charged peak lies at one-half the m/z value of the singly-charged peak for the same molecule. Finally, one triply-charged peak at 22.1 kDa/e (one-third the singly-charged value) is indicated by an inverted triangle. Peaks identified by our procedure are

TABLE 2
Posterior means and (standard deviations) for model parameters for the five experimental data sets

Data set	S	φ	ϱ	η_0	ω_0	J^{PM}	J^{HP}	J^{∇}
Blank	0.5660 (0.0537)	7.385 (0.126)	16.76 (2.92)	38.59 (11.14)	233.70 (150.80)	46.92 (4.22)	38	22
Matrix	0.9569 (0.0063)	6.770 (0.121)	15.03 (1.70)	163.70 (2.63)	17.41 (0.21)	12.27 (0.45)	11	7
Known	0.9016 (0.0021)	2.609 (0.093)	96.43 (5.54)	78.26 (2.87)	22.80 (0.95)	42.24 (0.96)	42	28
Single	0.9996 (0.0004)	0.310 (0.005)	49.46 (3.56)	182.30 (1.77)	16.76 (0.15)	52.68 (2.65)	45	45
Mean	0.9846 (0.0031)	3.952 (0.066)	69.35 (3.55)	188.40 (0.97)	15.77 (0.07)	76.91 (1.47)	74	54

indicated by vertical tick marks; these include all eleven “true” peaks, plus several additional peaks. These may reflect contaminants, differential isotopic compositions or thermal noise. Several of these identified peaks have resolutions in the range of the median resolution for the blank spectrum (Table 2), suggesting that the model is capturing the thermal noise component. The peak at 3.903 kDa (just below the smallest “true” peak) has higher resolution than typical thermal peaks, and also higher abundance. It is clearly present in all ten replicates, suggesting a potential contaminant in the mixture. Features of the LARK model, such as the resolution and abundance parameters, may aid in reducing false positives and prioritizing masses for further study, beyond using estimated mass alone.

6.4. *Lung cancer protein spectrum.* Figure 6 displays posterior reconstructions from LARK for different segments of the single and mean spectra for the complete data depicted in Figure 1. The noise reduction from averaging several spectra results in higher estimated precision φ , just as one would expect (Table 2), approximately ten times higher than that for a single spectrum. The resolution is also higher in the mean spectrum, leading to the identification of a larger number of peaks. Posterior means for other fixed dimensional summaries are comparable for the single and mean spectra. Figure 6(d)–(f) illustrates the ability of the LARK model to deconvolve a single large peak into several peaks. While there is a local mode at 33.5 kDa, corresponding to the doubly charged peak for albumin (mass 67 kDa), the relationship between resolution and mass suggests that there are other molecules present that lead to the wider than expected peak and its asymmetric shape. The highest posterior realization from LARK provides

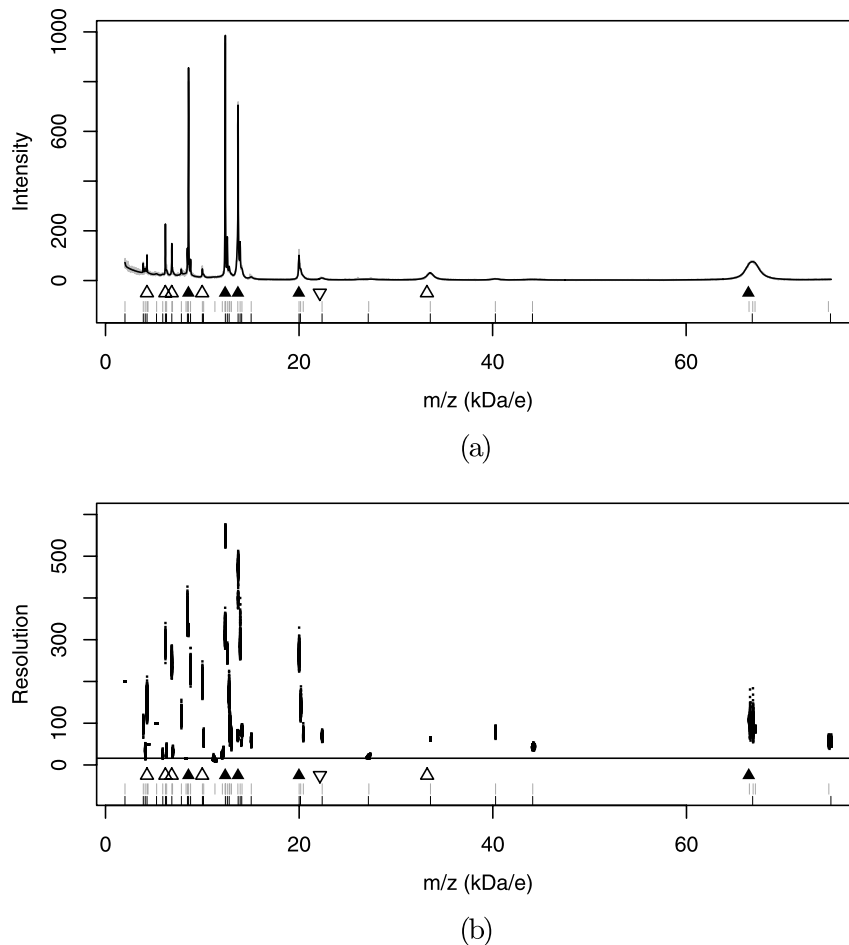


FIG. 5. *LARK* posterior mean (solid line) and data from a mixture of five known proteins (a) and associated posterior distribution of resolution (b). Solid triangles represent singly-charged molecules, open triangles represent doubly-charged molecules and the inverted triangle represents a triply-charged protein. The rows of tick marks represent identified peak locations using the highest posterior draw (top, gray) and local modes under model averaging (bottom, black). The horizontal line in (b) corresponds to the median resolution from the blank spectrum analysis.

a way to estimate these masses (in contrast, local modes of the estimated signal would suggest just a single protein, while methods that identify peaks by regions where the estimated signal is greater than some specified threshold would report almost the entire range between 31–39 kDa).

7. Simulation study. Coombes et al. (2005a) construct a mathematical model for MALDI-TOF mass spectrometry based on the physics of the process, providing a virtual mass spectrometer. Morris et al. (2005) use this

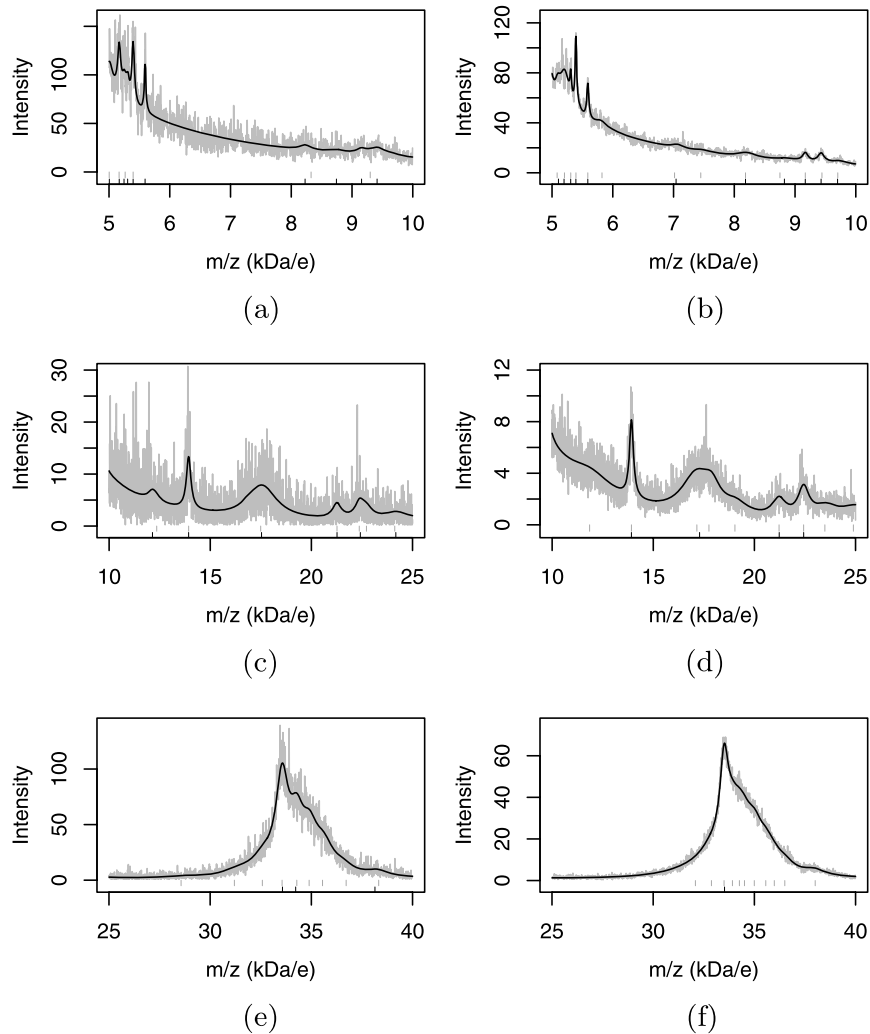


FIG. 6. Peak reconstruction using the LARK posterior mean (solid line) for the lung cancer patient using the single spectrum (left) and mean spectrum (right) for segments of the respective spectra. The two rows of tick marks on the horizontal axis correspond to identified peak locations using the local maxima from highest probability draw (top, gray) and model averaging (bottom, black).

simulator to generate spectra to explore operating characteristics of their mean-spectrum undecimated wavelet threshold (MUDWT) peak detection method. They generated 100 data sets, each comprised of 100 simulated spectra generated with 150 true peaks and additive i.i.d. Gaussian errors (sd $\sigma = 66$) [for details, see Section 4.3 of Morris et al. (2005); raw data and code are available on the Cromwell (2004) website]. Although in theory the simulated spectra should not need alignment before calculating mean spec-

tra, our exploratory analysis revealed boundary effects in the initial segment (below approximately 3 kDa) due to averaging, and visual peaks outside of the true peak mass $\pm 0.3\%$ tolerance window used by Morris et al. (2005) to classify matches of identified peaks to true peaks.⁵ Thus, for this analysis we use mean spectra from each of the 100 data sets for all masses greater than 3 kDa. In addition to the MUDWT approach of Morris et al. and LARK, we identified peaks using the R package `PRocess` [Li (2005)], which provides automatic baseline subtraction and peak identification.

Because the simulated data were generated using Gaussian errors with constant variance, we replaced the gamma likelihood of the LARK model with a Gaussian likelihood, with mean $\mu(t)$ and constant precision $\phi = 1/\sigma^2$. Otherwise, all model parameters have the same interpretation as before. Preliminary investigation for a region of the spectra with a single known peak indicated that the Gaussian kernel provided a better fit than the Cauchy kernel we used for the experimental data. This also suggested that the simulated spectra featured higher resolution than those of the Voyager machine, leading us to select $\mu_\varrho = 300$, keeping $\sigma_\varrho = 0.49$ as before. Following Morris et al., we use a robust estimate of σ from the wavelet decomposition as an alternative to the data-based prior described in Section 4 for the non-constant variance model. For all remaining parameters, we used the default methods described in Section 4 to specify hyperparameters (see Table 1). We ran the RJ-MCMC algorithm for one million iterations, half for burn-in and half for posterior inference. Peaks were identified using the local mode under model averaging and the highest posterior realization.

For all methods, true peaks were classified as a *true positive* or *true discovery* if the mass of any identified peak was within $\pm 0.3\%$ of a true mass [as in Morris et al. (2005)]. Identified peaks outside of the $\pm 0.3\%$ tolerance window of true peaks were regarded as *false positives*. The True Positive Rate (TPR), the proportion of true discoveries, and False Discovery Rate (FDR), the proportion of false positives out of all identified peaks, were calculated for each of the 100 simulated mean spectra and are summarized in Table 3. Overall, `PRocess` has the lowest average FDR of all methods—however, its TPR is much worse than either of the LARK or MUDWT methods, which estimate the baseline and denoise simultaneously. Peak identification using down-crossings under model averaging (LARK-MA) has the better FDR of the two LARK methods, but has a lower TPR because local modes may miss overlapping peaks. Peak identification using the LARK HP realization and the wavelet method provide comparable TPR and FDR performance, with the MUDWT method having a slightly better TPR, while the LARK-HP method has a slightly improved FDR. The absolute differences of TPR

⁵Jeffrey S. Morris (M. D. Anderson), personal communication.

TABLE 3

Summaries of true positive rates (TPR) and false discovery rates (FDR) for 100 simulated mean spectra using the mean-spectrum undecimated wavelet transform (MUDWT) [Morris et al. (2005)], LARK with highest posterior realization (LARK-HP), LARK with local modes under model averaging (LARK-MA) and PROcess [Li (2005)]

Summary	Method	Average	95% range
TPR	MUDWT	0.83	0.77–0.89
	LARK-HP	0.82	0.75–0.89
	LARK-MA	0.73	0.61–0.81
	PROcess	0.31	0.10–0.43
FDR	MUDWT	0.06	0.003–0.208
	LARK-HP	0.05	0.000–0.132
	LARK-MA	0.01	0.000–0.050
	PROcess	0.0033	0.000–0.032

and FDR rates for the two methods are both about 0.01, “statistically significant” using a paired t -test but not practically significant, leading to the discovery of an extra 1–2 proteins by MUDWT (on average) at the cost of a comparable number of additional false positives. A further breakdown of the TPR for LARK by prevalence and abundance [as in Morris et al. (2005), Table 4] is provided in the supplemental materials [House, Clyde and Wolpert (2011)], with LARK-HP having substantially higher TPR than MUDWT for peaks in the higher prevalence groups across all abundance categories, but poorer performance than MUDWT for the two lowest prevalence groups.

8. Discussion. The Gaussian LARK model leads to true positive rates and false discovery rates comparable to adaptive nonparametric wavelet methods for simulated Gaussian intensities data, while the gamma LARK model is able to capture the mean/variance relationship that is observed in experimental data. Exploratory data analysis may be used to decide which model is more appropriate by examining the mean/variance relationship or residual analysis, with other error models easily substituted to define alternative likelihood functions given the mean function.

A key feature of the LARK methodology is the ability to deconvolve large peaks into mixtures of protein signatures. The model is able to identify all masses for the laboratory experiment with known protein mixture, even though many of the doubly charged proteins have low abundance. To capture the multiply charged nature of proteins in MALDI-TOF even more effectively, LARK models may be constructed using a kernel tailored to this purpose as a mixture of two peaks centered at the singly and doubly charged

masses, with a mixing weight to control the relative abundance of the two charges. This constraint reduces the number of free parameters and may lead to a more efficient algorithm. The method described in this paper is intended for use with either a single spectrum or a mean spectrum. Hierarchical versions of the LARK model for MALDI-TOF data are under development for modeling multiple spectra, which provide automatic calibration of multiple spectra and permit classification of subjects into groups.

The LARK models are implemented in an R [R Development Core Team (2010)] package, with a shared library written in C and FORTRAN for the RJ-MCMC algorithm. Although the LARK model for peak identification is more computationally intensive than the wavelet method of Morris et al. (2005), with 10,000 iterations taking 10 minutes on a dual 3 GHz Quad Core Xeon Mac Pro for the simulation study (running on a single processor), its running times increase only linearly with the number of peaks and volume of data, since no matrix inversion is required. Despite the computational overhead of RJ-MCMC, Clyde and Wolpert (2007) and Wolpert, Clyde and Tu (2011) have demonstrated that LARK models can provide significant reductions in mean squared error in comparison with some of the best wavelet methods such as the nondecimated wavelet approach of Johnstone and Silverman (2005) and the continuous wavelets of Chu, Clyde and Liang (2009). Future work will incorporate advances in adaptive MCMC methods which may accelerate convergence for random-walk update steps or lead to improved proposal distributions for peak birth based on residuals or peak death utilizing abundance. The software is available from the first author's website.

Acknowledgments. The authors would like to thank Michael J. Campa, Michael C. Fitzgerald, Nayela Khan, Sheila Lee, Edward Patz, Jr. and Petra L. Roulhac at Duke University for providing the experimental data used in our examples and Jeff Morris from M.D. Anderson for providing the simulated data. We have benefited greatly from many fruitful discussions with these individuals. We also wish to thank the anonymous referees, Associate Editors and the Editor for their helpful comments on earlier versions of this manuscript.

SUPPLEMENTARY MATERIAL

Additional results for the simulation study

(DOI: [10.1214/10-AOAS450SUPP](https://doi.org/10.1214/10-AOAS450SUPP); .pdf). True positive rates for LARK estimates from the simulation study broken down by peak prevalence and average intensity of peaks across samples.

REFERENCES

- ABRAMOWITZ, M. and STEGUN, I. A., eds. (1964). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. National Bureau of Standards Applied Mathematics Series 55*. For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, DC. [MR0167642](#)
- APPLIED BIOSYSTEMS (2001). *Voyager Biospectrometry Workstation with Delayed Extraction Technology User Guide Version 5.1*. Applied Biosystems, Foster City, CA.
- BAGGERLY, K. A., COOMBES, K. R. and MORRIS, J. S. (2006). An introduction to high-throughput bioinformatics data. In *Bayesian Inference for Gene Expression and Proteomics* (K.-A. DO, P. MÜLLER and M. VANNUCCI, eds.) Chapter 1, 1–39. Cambridge Univ. Press, Cambridge.
- BAGGERLY, K. A., MORRIS, J. S. and COOMBES, K. R. (2004). Reproducibility of SELDI-TOF protein patterns in serum: Comparing datasets from different experiments. *Bioinformatics* **20** 777–785.
- CAMPA, M. J., WANG, M. Z., HOWARD, B. A., FITZGERALD, M. C. and PATZ, E. F. JR. (2003). Protein expression profiling identifies MIF and Cyclophilin A as potential molecular targets in non-small cell lung cancer. *Cancer Research* **63** 1652–1656.
- CHU, J.-H., CLYDE, M. A. and LIANG, F. (2009). Bayesian function estimation using continuous wavelet dictionaries. *Statist. Sinica* **19** 1419–1438. [MR2589190](#)
- CLYDE, M. A., HOUSE, L. L. and WOLPERT, R. L. (2006). Nonparametric models for proteomic peak identification and quantification. In *Bayesian Inference for Gene Expression and Proteomics* (K.-A. DO, P. MÜLLER and M. VANNUCCI, eds.) Chapter 15, 293–308. Cambridge Univ. Press, Cambridge.
- CLYDE, M. A. and WOLPERT, R. L. (2007). Nonparametric function estimation using overcomplete dictionaries. In *Bayesian Statistics 8* (J. M. BERNARDO, M. J. BAYARRI, J. O. BERGER, A. P. DAWID, D. HECKERMAN, A. F. M. SMITH and M. WEST, eds.) 91–114. Oxford Univ. Press, Oxford. [MR2433190](#)
- COOMBES, K. R., KOOMEN, J. M., BAGGERLY, K. A., MORRIS, J. S. and KOBAYASHI, R. (2005a). Understanding the characteristics of mass spectrometry data through the use of simulation. *Cancer Informatics* **1** 41–52.
- COOMBES, K. R., TSAVACHIDIS, S., MORRIS, J. S., BAGGERLY, K. A., HUNG, M. C. and KUERER, H. M. (2005b). Improved peak detection and quantification of mass spectrometry data acquired from surface-enhanced laser desorption and ionization by denoising spectra with the undecimated discrete wavelet transform. *Proteomics* **5** 4107–4117.
- CROMWELL (2004). Cromwell MatLab package. M. D. Anderson Cancer Center, Houston, TX. Available at <http://bioinformatics.mdanderson.org/cromwell.html>.
- DASS, C. (2001). *Principles and Practice of Biological Mass Spectrometry*. Wiley, New York.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B* **39** 1–38. With discussion. [MR0501537](#)
- FRANZEN, J. (1997). Improved resolution for MALDI-TOF mass spectrometers: A mathematical study. *International Journal of Mass Spectrometry and Ion Processes* **164** 19–34.
- GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82** 711–732. [MR1380810](#)
- GUINDANI, M., DO, K. A., MÜLLER, P. and MORRIS, J. S. (2006). Bayesian mixture models for gene expression and protein profiles. (K.-A. DO, P. MÜLLER and M. VANNUCCI, eds.) Chapter 12, 238–253. Cambridge Univ. Press, Cambridge.

- HAREZLAK, J., WU, M., WANG, M., SCHWARTZMAN, A., CHRISTIAN, D. and LIN, X. (2008). Biomarker discovery for Arsenic exposure using functional data analysis and feature learning of mass spectrometry proteomic data. *Journal of Proteome Research* **7** 217–224.
- HOUSE, L. L. (2006). Nonparametric Bayesian models in expression proteomic applications. Ph.D. dissertation. Dept. Statist. Sci., Duke Univ., Durham, NC.
- HOUSE, L. L., CLYDE, M. A. and WOLPERT, R. L. (2011). Supplement to “Bayesian nonparametric models for peak identification in MALDI-TOF mass spectroscopy.” DOI:10.1214/10-AOAS450SUPP.
- JOHNSTONE, I. M. and SILVERMAN, B. W. (2005). Empirical Bayes selection of wavelet thresholds. *Ann. Statist.* **33** 1700–1752. MR2166560
- KEMPKA, M., SÖDAHL, J., BJÖRK, A. and ROERAADE, J. (2004). Improved method for peak picking in matrix-assisted laser desorption/ionization time-of-flight mass spectrometry. *Rapid Communications in Mass Spectrometry* **18** 1208–1212.
- LI, X. (2005). PROcess: Ciphergen SELDI-TOF Processing. R Package Version 1.24.0. Available at <http://www.bioconductor.org/help/bioc-views/2.6/bioc/html/PROcess.html>.
- MALYARENKO, D. I., COOKE, W. E., ADAM, B.-L., MALIK, G., CHEN, H., TRACY, E. R., TROSSET, M. W., SASINOWSKI, M., SEMMES, O. J. and MANOS, D. M. (2005). Enhancement of sensitivity and resolution of surface-enhanced laser desorption/ionization time-of-flight mass spectrometric records for serum peptides using time-series analysis techniques. *Clin. Chem.* **51** 65–74.
- MORRIS, J. S., COOMBES, K. R., KOOMEN, J., BAGGERLY, K. A. and KOBAYASHI, R. (2005). Feature extraction and quantification for mass spectrometry in biomedical applications using the mean spectrum. *Bioinformatics* **21** 1764–1775.
- MORRIS, J. S., BROWN, P. J., BAGGERLY, K. A. and COOMBES, K. R. (2006). Analysis of mass spectrometry data using Bayesian wavelet-based functional mixed models. In *Bayesian Inference for Gene Expression and Proteomics* (K.-A. DO, P. MÜLLER and M. VANNUCCI, eds.) Chapter 14, 269–292. Cambridge Univ. Press, Cambridge.
- MORRIS, J. S., BROWN, P. J., HERRICK, R. C., BAGGERLY, K. A. and COOMBES, K. R. (2008). Bayesian analysis of mass spectrometry proteomic data using wavelet-based functional mixed models. *Biometrics* **64** 479–489, 667. MR2432418
- MÜLLER, P., BAGGERLY, K. A., DO, K.-A. and BANDYOPADHYAY, R. (2010). A Bayesian mixture model for protein biomarker discovery. In *Bayesian Modeling in Bioinformatics* (D. K. DEY, S. GHOSH and B. K. MALLICK, eds.). Chapman & Hall/CRC Press, Boca Raton, FL.
- NGUYEN, N., HUANG, H., ORAINTARA, S. and VO, A. (2010). Mass spectrometry data processing using zero-crossing lines in multi-scale of Gaussian derivative wavelet. *Bioinformatics* **26** i659–i665.
- R DEVELOPMENT CORE TEAM (2010). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna.
- SUZDAK, G. (2003). *The Expanding Role of Mass Spectrometry in Biotechnology*. MCC Press, San Diego, CA.
- TIBSHIRANI, R., HASTIE, T., NARASIMHAN, B., SOLTYS, S., SHI, G., KOONG, A. and LE, Q.-T. (2004). Sample classification from protein mass spectrometry, by ‘peak probability contrasts’. *Bioinformatics* **20** 3034–3044.
- VIDAKOVIC, B. (1999). *Statistical Modeling by Wavelets*. Wiley, New York. MR1681904
- WAND, M. P. and JONES, M. C. (1995). *Kernel Smoothing. Monographs on Statistics and Applied Probability* **60**. Chapman & Hall, London. MR1319818

- WANG, X., RAY, S. and MALLICK, B. K. (2007). Bayesian curve classification using wavelets. *J. Amer. Statist. Assoc.* **102** 962–973. [MR2354408](#)
- WANG, M. Z., HOWARD, B. A., CAMPA, M. J., PATZ, E. F. JR. and FITZGERALD, M. C. (2003). Analysis of human serum proteins by liquid phase iso-electric focusing and matrix-assisted laser desorption/ionization mass spectrometry. *Proteomics* **3** 1661–1666.
- WOLPERT, R. L., CLYDE, M. A. and TU, C. (2011). Stochastic expansions using continuous dictionaries: Lévy adaptive regression kernels. *Ann. Statist.* To appear.
- WOLPERT, R. L. and ICKSTADT, K. (2004). Reflecting uncertainty in inverse problems: A Bayesian solution using Lévy processes. *Inverse Problems* **20** 1759–1771. [MR2107235](#)
- YASUI, Y., MCLERRAN, D., ADAM, B.-L., WINGET, M., THORNQUIST, M. and FENG, Z. (2003). An automated peak identification/calibration procedure for high-dimensional protein measures from mass spectrometers. *J. Biomed. Biotechnol.* **2003** 242–248.
- ZHIGILEI, L. V. and GARRISON, B. J. (1998). Velocity distributions of analyte molecules in matrix assisted laser desorption from computer simulations. *Rapid Communications in Mass Spectrometry* **12** 1273–1277.

L. L. HOUSE
DEPARTMENT OF STATISTICS
VIRGINIA TECH
BLACKSBURG, VIRGINIA 24061-0439
USA
E-MAIL: llhouse@vt.edu

M. A. CLYDE
R. L. WOLPERT
DEPARTMENT OF STATISTICAL SCIENCE
DUKE UNIVERSITY
DURHAM, NORTH CAROLINA 27708-0251
USA
E-MAIL: clyde@stat.duke.edu
rlw@stat.duke.edu