

## Understanding NMR $T_2$ spectral uncertainty

Michael Prange\*, Yi-Qiao Song

Schlumberger-Doll Research, One Hampshire Street, Cambridge, MA 02139, USA

### ARTICLE INFO

#### Article history:

Received 4 January 2010

Revised 12 February 2010

Available online 19 February 2010

#### Keywords:

Monte Carlo  
Laplace inversion  
Uncertainty  
Relaxation  
Diffusion

### ABSTRACT

NMR relaxation and diffusion data analysis commonly uses a wide range of methods from simple exponential fitting to Laplace inversions. The pros and cons of these methods are often the subject of intense debate. We show that the ill-conditioned nature of such analysis gives rise to a range of solutions for every method resulting in uncertainty in the spectral solution. Such uncertainty is in fact characteristic of the inversion method. We show a simple method of sparse spectral representation can be used to improve the statistics of multiple-exponential-based inversion schemes.

© 2010 Elsevier Inc. All rights reserved.

### 1. Introduction

Analysis of relaxation and diffusion data often uses many different methods, such as single exponential fitting for samples of single molecular species, fitting with a few exponentials (e.g., two or three), stretched exponential fitting, and Laplace inversion [1–5]. Methods other than single exponential fitting are used to analyze the multi-exponential decays resulting from heterogeneity in the underlying sample. Laplace inversion can be used to analyze arbitrarily complex decay spectra. However, there is often a debate in the literature on which criteria to use in choosing the appropriate method for a particular application, and the benefit of using more complex spectral models.

The origin of this ambiguity is the ill-conditioned nature of multi-exponential inversion [6]. Given a data set with finite noise, the solution spectrum is not unique and many spectra fit the experimental decay data within the statistics of the noise. As a result, we should expect the data to be fitted by many different spectra, both with the same functional form for the spectra and with different ones. In fact, the true solution is the entire ensemble of the solution spectra as demonstrated in our recent work [7]. In the present work, we apply a concept that is similar to these other analysis methods, such as single- and multi-exponential fitting, however an ensemble of solutions is obtained for each method. For a given data set, many different fitting methods can yield solution spectra that are consistent with the data, meaning that data misfit is not an adequate differentiator for choosing a method. We show that the Akaike information criterion [8], which weighs

both data misfit and model complexity, does provide an effective metric for choosing the best stochastic spectral representation for evaluating the uncertainty of particular spectral properties.

In addition, we demonstrate a Monte Carlo algorithm for Laplace inversion that uses a small number of variable-width boxcar functions instead of the small number of delta functions used in multi-exponential fitting. Note that in the limit of a large number of delta functions assigned to fixed spectral locations, the multi-exponential fit can be used to approximate the continuous spectral range used in traditional Laplace inversion. We show the advantage of using a sparse boxcar representation, when applied to spectral uncertainty analysis, compared to other methods. Such an approach may also have advantages when performing uncertainty analysis on higher-dimensional NMR data because it requires fewer spectral parameters than either the multi-exponential or the Laplace inversion approaches.

Once an ensemble of solution spectra have been generated, questions about the shape of the spectrum can be posed in a probabilistic form that allow the uncertainty of the answers to be quantified. In Section 4 we examine the uncertainty in functionals of the spectrum that relate to the total porosity and the bound-fluid porosity in a rock sample.

### 2. Theory

#### 2.1. General

NMR signals of  $T_2$  decay in porous materials are well modeled as a sum of decaying exponentials,

$$m(t) = \int_0^\infty f(T) \exp(-t/T) d \log(T), \quad (1)$$

\* Corresponding author.

E-mail address: [prange@slb.com](mailto:prange@slb.com) (M. Prange).

where  $m(t)$  is the signal as a function of time,  $f(T)$  is the spectrum as a function of relaxation time  $T$  satisfying  $f(T) \geq 0$ . Using a more compact operator notation,

$$m(t) = \mathcal{G}(t)f, \quad (2)$$

where  $\mathcal{G}(t)$  is the kernel operator. Our goal is to describe the uncertainty in  $f$  resulting from the inversion of a noisy measurement of  $m$ .

We start with noisy samples of  $m$  measured at discrete times and put in the measurement vector  $m = \{m(t_i), i = 1, \dots, M\}$ . Parker and Song [9] have demonstrated that the noise in such a measurement can often be modeled as an independent identically distributed (IID) normal random process. In this case, with some simplifying assumptions described below, the uncertainty of  $f$  is given by the probability density function [7]

$$\pi(f) \propto \exp \left[ -\frac{1}{2\sigma^2} \|m - \mathcal{G}(t)f\|^2 \right], \quad f(T) \geq 0, \quad (3)$$

where  $\sigma^2$  is the independently determined noise variance,  $t$  is the vector of measurement times and  $\mathcal{G}(t)f$  is the vector of predicted signal associated with the spectrum  $f$ . Note that in the parlance of Bayesian inference [10], the distribution given by Eq. (3) is actually a likelihood function, and is usually denoted by  $\pi(m|f)$ . The uncertainty in  $f$  is described by the posterior distribution,  $\pi(f|m)$ . The posterior is related to the likelihood through Bayes' Rule,  $\pi(f|m) \propto \pi(m|f)\pi(f)$ , where  $\pi(f)$  is called the prior on  $f$ . Here we make the simplifying assumption that  $\pi(f)$  is constant, in which case the posterior is proportional to the prior, yielding Eq. (3).

The uncertainty in  $f$  is fully described by Eq. (3). One way to visualize this uncertainty is by randomly sampling this distribution and using those samples to represent the uncertainty in  $f$ . This is called the Monte Carlo method. From these samples histograms can be generated and statistics such as mean and variance computed.

## 2.2. Sparse basis representations

Although the uncertainty in  $f$  is formally described by Eq. (3), the numerical computation of random samples from Eq. (3) usually begins by expressing  $f$  in a discrete form with a finite number of parameters. In [7] the spectrum was discretely represented by 100 spectral bins. Here we express  $f$  as a weighted sum of  $N$  basis functions:

$$f(T) = \sum_{i=1}^N \alpha_i \psi_i(T; \lambda_i), \quad (4)$$

where the basis functions  $\psi_i(T; \lambda_i)$  are non-negative compact functions of  $T$  whose shapes are controlled by the parameter vectors  $\lambda_i$ . These basis functions are normalized such that

$$\int_0^\infty \psi_i(T; \lambda_i) d \log(T) = 1. \quad (5)$$

Substituting Eq. (4) into Eq. (1) yields

$$m(t) = \sum_{i=1}^N \alpha_i \phi_i(t; \lambda_i), \quad \alpha_i \geq 0, \quad (6)$$

where

$$\phi_i(t; \lambda_i) = \int_0^\infty \psi_i(T; \lambda_i) \exp(-t/T) d \log(T). \quad (7)$$

In the case of using Dirac delta functions as the basis functions, the parameter vectors  $\lambda_i$  contain the spectral position of the delta function, namely  $T_j$ . Thus  $\psi_i(T; T_i) = \delta(T - T_i)$  and  $\phi_i(t, T_i) = \exp(-t/T_i)$ .

We also consider basis functions which have finite spectral extent in order to evaluate the effect of spectral smoothness in the evaluation of spectral uncertainty. We choose to use as basis func-

tions the boxcar functions so that Eq. (7) can be analytically evaluated. Boxcar functions are constant within a given range of  $T$  and zero otherwise. Our boxcar functions are defined by

$$B(T; T^{\text{Lo}}, T^{\text{Hi}}) = \begin{cases} \frac{1}{\log T^{\text{Hi}} - \log T^{\text{Lo}}}, & T^{\text{Lo}} < T \leq T^{\text{Hi}} \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where the normalization of Eq. (5) has been applied. Note that this definition allow the basis functions to overlap, allowing, for example, two overlapping basis functions to form a three-stepped function akin to a winners' podium at a sports arena. In a limiting case, this basis function can approach zero width, resulting in a properly-normalized delta function. Substituting this basis function into Eq. (7) yields

$$\phi_i(t; T_i^{\text{Lo}}, T_i^{\text{Hi}}) = \frac{\Gamma_0\left(\frac{t}{T_i^{\text{Hi}}}\right) - \Gamma_0\left(\frac{t}{T_i^{\text{Lo}}}\right)}{\log(T_i^{\text{Hi}}) - \log(T_i^{\text{Lo}})}, \quad (9)$$

where  $\Gamma_0$  is the incomplete gamma function defined by  $\Gamma_0(x) = \int_x^\infty e^{-\xi}/\xi d\xi$ . The incomplete gamma function is commonly available in scientific computing libraries.

## 2.3. Monte Carlo sampler

With this sparse-basis representation of the spectrum, the uncertainty in spectral parameters  $\alpha$  and  $\lambda$  is found by substituting Eq. (6) into Eq. (3) to get

$$\pi(f) \propto \exp \left[ -\frac{1}{2\sigma^2} \left\| m - \sum_{i=1}^N \alpha_i \phi_i(t; \lambda_i) \right\|^2 \right], \quad \alpha_i \geq 0, \quad (10)$$

We implement the Monte Carlo sampling in two steps: a non-linear step in which the  $\lambda$  parameters are sampled, followed by a quasi-linear step in which the  $\alpha$  parameters are sampled for a given sample of  $\lambda$  parameters. The sample of  $\alpha$  is computed using the non-negative multi-normal sampler presented in [7].


In the first step a refinement of a Gibbs' sampler, herein called the directed Gibbs' sampler, is used to sample the  $\lambda$  parameters for one sampling iteration (see Appendix A for a description of our directed Gibbs' sampler). A standard Gibbs' sampler, such as Algorithm A.39 in [11], could not be used because it was found to converge too slowly from the strong correlation between the parameters in  $\lambda$ . During Gibbs' sampling, as  $\lambda$  is scanned over a range of values,  $\alpha$  is set to its maximum likelihood value found by using the non-negative least squares (NNLS) algorithm of Lawson and Hanson [12]. We are essentially generating a random sample of  $\lambda$  while making the  $\alpha$  parameter dependent on  $\lambda$ .

## 3. Numerical experiments

### 3.1. Monte Carlo with different bases

Here we apply the above two sparse-basis approaches to examining the spectral uncertainty for a synthetic data example. To facilitate comparison we use the same synthetic example, with identical noise realization, as was used in [7]. This example is shown in Fig. 1.

An example of the diversity of spectral solutions for delta-function bases is shown in Fig. 2a for  $N = 3$  and in Fig. 2c for  $N = 4$ , where the spectra for 20 samples from the total of  $10^4$  generated Monte Carlo samples are superimposed on each plot. In Fig. 2b and d are plotted the statistics of the time-domain measurement predictions associated with the  $10^4$  spectral samples associated with Fig. 2a and c, respectively. The bold curve in Fig. 2b and d provide the median of the  $10^4$  time-domain predictions, and the four



**Fig. 1.** This example  $T_2$  spectrum was used to synthesize the  $T_2$  echo data shown in the inset. The 8192 echos are sampled at an echo spacing of 0.0002 s starting at 0.0002 s. The noise standard deviation is 0.025 of the maximum echo value.

dashed curves on each plot provide the minimum and maximum values at each time along with the 10% and 90% quantiles, with these last two curves being barely visible at this plot scale. The noisy echo data is also shown for comparison. Note that three delta functions provide a good fit to the measurements. Using four delta functions occasionally overfits the measurements, meaning that the noise on the first few measurement points is sometimes being fit. This is clearly seen in the upper dashed line in Fig. 2d and in the increase in spectral values with diminishing  $T_2$ .

The corresponding plots for the boxcar-basis function (for  $N = 2$  and  $N = 3$ ) are shown in Fig. 3a–d. Comparing the dashed curves in Fig. 3d, that mark the maximum and minimum of the time-domain predictions, with those in Fig. 2d, it is clear that the extreme samples for the delta-function basis are overfitting the data for small times, but the same is not true for the boxcar-basis. A comparison of Fig. 3c with Fig. 2c shows that the delta-function basis contains spectral energy at small  $T_2$  values, while these are absent for the

boxcar-basis. This demonstrates that the boxcar-basis is much less prone to overfitting. The reason of this is that the low probability of basis functions of both small width and small  $T_2$  values makes it unlikely to occur in this region that is largely unconstrained by the measurements. The impact of this will be discussed later.

### 3.2. Akaike information criterion


When fitting data (e.g.,  $m$ ) with different models, models with more parameters are expected to fit the data better than less complex models. However, the problem of overfitting must be considered, since the additional model complexity may be fitting noise instead of true signal. In signal processing, the Akaike information criterion (AIC) is often used to balance the improvement in the  $\chi^2$  fit of the model to the data with increasing  $k$  (the number of model parameters) against the commensurate increase in model complexity [13]. This is an application of Occam's razor to the model fitting problem. The AIC is defined by

$$\text{AIC} = \chi_{\min}^2 + 2k, \quad (11)$$


where  $\chi^2 = \sigma^{-2} \|m - \sum_{i=1}^N \alpha_i \phi_i(t; \lambda_i)\|^2$ ,  $\chi_{\min}^2$  is the minimum value of  $\chi^2$  over the range of model parameters and  $k$  is the number of model parameters used to fit the data. For the delta-function basis  $k = 2N$  and for the boxcar-function basis  $k = 3N$ . The conventional use of the AIC is to choose the model that minimizes the AIC value.

We compute  $10^4$  Monte Carlo samples (after burn-in) for each of our spectral models and approximate  $\chi_{\min}^2$  by the minimum value of  $\chi^2$  over all the samples. The resulting values of the AIC versus  $k$  for delta function and boxcar-function bases is plotted in Fig. 4. AIC values contain an unknown additive constant that is dependent only on the data, so here we plot  $\Delta \text{AIC} = \text{AIC} - \text{AIC}_{\min}$ .

The AIC values in Fig. 4 indicate that six parameters are ideal for both the delta and boxcar bases. This supports our conclusion from Fig. 2 that  $N = 3$  is superior to  $N = 4$  for the delta function basis, and also shows that  $N = 2$  results in a very poor fit for this basis. Furthermore, the AIC suggests that the delta function basis is pre-



**Fig. 2.** The spectral uncertainty for the synthetic data shown in Fig. 1 is illustrated here for delta-function bases with (a)  $N = 3$  and (c)  $N = 4$ . The spectral values are superimposed as gray bars for 20 of the  $10^4$  sampled solutions. The true spectrum is overlain as a dashed curve. The statistics of the time-domain predictions associated with the  $10^4$  spectral samples are shown in (b) for  $N = 3$  and (d) for  $N = 4$ , with the median curve shown in bold, and the four dashed curves showing minimum values, the 10% and 90% quantiles, and the maximum values. Only the minimum and maximum curves are readily apparent in (b), as the tightly-bounded quantile curves are hidden by the thickness of the median curve.




**Fig. 3.** The spectral uncertainty for the synthetic data shown in Fig. 1 is illustrated here for boxcar-function bases with (a)  $N = 2$  and (c)  $N = 3$ . The spectral values are superimposed as gray curves for 20 of the  $10^4$  sampled solutions. The true spectrum is overlain as a dashed curve. The statistics of the time-domain predictions of these spectral samples are shown in (b) for  $N = 2$  and (d) for  $N = 3$ , with the median curve shown in bold, and the four dashed curves showing minimum values, the 10% and 90% quantiles, and the maximum values. Only the minimum and maximum curves are readily apparent, as the tightly-bounded quantile curves are mostly hidden by the thickness of the median curve.

ferred over the boxcar-basis. A more quantitative comparison can be made by translating  $\Delta AIC$  into an odds ratio, which is the probability of one model over another [14]. This odds ratio is given by  $\exp(\Delta AIC)$ . Using an odds ratio, the best Delta function model ( $N = 3$ ) is preferred to the best boxcar function model ( $N = 2$ ) by about 2.3 to 1. This is a weak preference [14]. By contrast, the best model is preferred over the next best models with  $k < 6$  by more than  $10^8$  to 1. This is decisive evidence that at least six parameters are needed to model the noisy time-domain measurements in this example. As we will see later, the choice of delta-function versus boxcar-function model depends also on how the spectrum is being used, with smooth spectral models being preferred in some applications.

#### 4. Uncertainty comparisons using functionals

Our objective here is to compute the uncertainty in one-dimensional  $T_2$  spectra and to compare these uncertainty estimates with



**Fig. 4.** Values of the Akaike information criterion (AIC) are plotted versus the number of model parameters for delta-function bases (solid) and boxcar-function bases (dashed). Lower values of the AIC indicate better models.

those made by an alternative algorithm presented in [7]. We choose to examine the uncertainty in  $T_2$  spectra by examining the impact of this uncertainty on functionals of the spectra, as was done in [7]. In this comparison we use the functionals for total porosity ( $\rho_T$ ) and bound porosity ( $\rho_B$ ) defined by

$$\begin{aligned} \rho_T &= \int_0^\infty f(T) d \log(T) \quad \text{and} \\ \rho_B &= \int_0^{T_c} f(T) d \log(T), \end{aligned} \tag{12}$$


where we use  $T_c = 0.033$  s, a typical value used for sandstone [15]. Histograms of  $\rho_T$  and  $\rho_B$  were generated from Monte Carlo samples of  $f(T)$ .

One might suppose that spectral uncertainty would be better examined by directly displaying the uncertainty in  $f$ , for example via a plot of the mean spectrum along with 95% confidence limits. Unfortunately such displays are misleading because they ignore the strong correlations that exist between parameters [9,7]. For example, an increase in the spectrum for one value of  $T$  can be compensated for by decreasing the spectrum at a neighboring value of  $T$ . This results in overly large confidence limits since such a plot only shows the marginal uncertainty in the spectrum for any given value of  $T$ . We sidestep this complexity through the use of functionals – functionals map the spectrum into a single random variable that can be easily histogrammed.

##### 4.1. Uncertainty in total porosity

Estimates of the uncertainty of  $\rho_T$  were computed from  $10^4$  Monte Carlo samples for both types of basis function. These results are plotted versus  $k$  in Fig. 5. The spectrum used to synthesize our synthetic data was chosen such that  $\rho_T$  is precisely 0.1.

First consider the case when the model is inadequate to represent the data, namely for  $k < 6$ . Here the misfit is large and the variance is underestimated, leaving the true value of  $\rho_T$  well outside of the error bounds of the estimate. This is hardly surprising, because



**Fig. 5.** Estimates of  $\rho_T$ , along with their uncertainty, are plotted versus the number of parameters ( $k$ ) as mean curves and the mean  $\pm$  one standard deviation confidence limits for delta-function bases (solid line and circles) and boxcar-function bases (dashed lines and open circles). The true value of 0.1 is indicated by a dotted line. The “uncorrected” and “corrected” values (open squares), taken from [7], are computed for the same synthetic data using 100 delta functions at fixed spectral locations ( $k = 100$ ), where the correction is a post-processing step designed to remove the bias.

the trade-off between goodness of fit and variance is the foundation of how statisticians frame the principle of parsimony. Parsimony seeks to find the model with “the smallest number of parameters for adequate representation of the data” [13, p. 31]. We use the AIC to manage this trade-off.

The AIC identifies  $k = 6$  as the optimum model complexity for both the delta-function and boxcar-function bases. The  $k = 6$  uncertainty estimates for both basis types result in reasonable estimates of  $\rho_T$ , meaning that the true value is a probable sample from the uncertainty distribution.

For  $k > 6$  the delta-function models still provide reasonable estimates of  $\rho_T$ , but the means are upwardly biased and the variance is increased to compensate for this bias. The reason for this positive bias is discussed in depth in [7]. In short, the addition of extra parameters that are not needed to match the data allows these parameters to be assigned to small  $T_2$  values whose spectra are not strongly constrained by the data. Since the spectral values are constrained to be positive, this produces an upward bias in the estimate of  $\rho_T$ . This effect can be seen in Fig. 2c, where the addition of one extra delta-function has added substantial spectral energy at small  $T_2$  values.


From this reasoning an upward bias is also expected with the boxcar basis models for  $k > 6$ , but the results plotted in Fig. 5 do not possess this bias. The absence of this bias is not clearly understood. Since each basis function represents an interval of the spectrum instead of just a point value as in the delta-function basis, perhaps this extended nature makes it more difficult to match just the noise at small  $T_2$  values.

Our results in Fig. 5 are now compared with the uncorrected and bias-corrected estimates presented in [7]. These are plotted at  $k = 100$  in Fig. 5 because that uncertainty quantification approach is equivalent to using our delta-function basis approach with 100 delta functions of fixed position and varying amplitude. The bias in that uncorrected result is much larger than that of our results. However, the bias-corrected result agrees well in mean with our boxcar-basis result, but has a standard error that is about twice the magnitude of our result. In this example we have demonstrated that our sparse boxcar-basis results are much improved over an approach that represents the spectrum by many delta functions of fixed position and varying amplitude. This is probably due to the increase in variance resulting from the bias correction required by that approach.


#### 4.2. Uncertainty in bound porosity

In saturated porous media, water in small pores is bound by capillary force, and its amount can be determined by integrating the  $T_2$  spectrum up to a cutoff value (typically 0.033 s). Uncertainty in the estimate of such bound porosity,  $\rho_B$ , is plotted versus  $k$  in Fig. 6 for both delta- and boxcar-function bases. Since  $\rho_B$  focuses on the lower  $T_2$  part of the spectrum, the effect of a positive bias in that zone will have a proportionally larger impact than with  $\rho_T$ . The true value in this case is  $\rho_B \approx 0.014$ . As expected, for  $k < 6$  both bases provide poor estimates of the uncertainty. For  $k \geq 6$ , the boxcar-function basis yields good estimates, while the delta-function basis once again has a positive bias. Thus we have shown that in this example the delta-function basis provides a poor estimator of  $\rho_B$  even when the optimum  $k$  is chosen using the AIC.

A key difference between the delta-function and boxcar-basis estimators is highlighted in Fig. 7, where the two estimators are compared for  $k = 6$ . The histogram for the boxcar-basis estimator is unimodal and is peaked near the true value. In contrast, the estimator using the delta-function basis is bimodal, with low probability at the true value. The reason for this bimodality is apparent in Fig. 2a. When the leftmost of the three spikes is to the right of the cutoff time,  $T_c = 0.033$  s, the value of  $\rho_B$  is zero, while it is non-zero otherwise. As this spike dances around from sample to sample, an



**Fig. 6.** Estimates of  $\rho_B$ , along with their uncertainty, are plotted versus the number of parameters ( $k$ ) as mean curves and the mean  $\pm$  one standard deviation confidence limits for delta-function bases (solid lines and circles) and boxcar-function bases (dashed lines and open circles). The true value of 0.014 is indicated by a dotted line.



**Fig. 7.** Histograms for  $\rho_B$  estimated using the delta-function (wide bins) and boxcar-function (narrow bins) bases for  $k = 6$ . The arrow indicates the true value of  $\rho_B$ .



artificial transition is created that is simply due to the infinitesimal width of the spike. This problem does not occur in the estimation of  $\rho_T$  because there is no cutoff time in that case. It was this observation led us to postulate that a more robust estimator might be created by allowing the spikes to have width (the boxcar-basis), thereby minimizing the effect of this transition.

## 5. Discussion

The examples in Figs. 2 and 3 show clearly that models with dramatically different functional forms (and shapes of the relaxation spectra) can all fit the data within the noise characteristics. All solutions within each functional form comprise a portion of the entire ensemble of spectral solutions that satisfy the data. Since the spectral shapes are substantially different, we expect the solutions of different functional forms occupy somewhat different regions of the solution space. Within each functional form, when the number of degrees of freedom increases, there is more spectral variability and the size of the solution region increases. Thus it is natural to observe an increase in uncertainty with increasing  $k$ .

When  $k$  is much less than the optimal value by the AIC criterion, the fit is generally poor and the significantly reduced uncertainty does not reflect the uncertainty of the data. When  $k$  is larger than optimal, the uncertainty can vary significantly for different bases and  $k$ . This is a direct consequence of the different solution regions obtained through different basis functions and methods. In fact, the choice of specific functional form can be a very strong regulator for reducing model variation and uncertainty provided that the assumption of the functional form is properly justified. Often, such justification comes from understanding the physical, chemical and symmetry properties of the material and the experimental techniques.

The boxcar-basis is an interesting example of the effect of basis function choice. Since there is no limit on the width of the boxes and the Monte Carlo procedure is free to choose very narrow box widths (reducing to delta functions in the limit of zero box width), the phase space for the boxcar-basis expansion contains and is much larger than that for the delta-function basis. Thus it is rare to see a very narrow box in the spectral solution. This preference is likely the cause of the improved uncertainty estimates compared to those using the delta-function basis. The concept of using more restrictive basis functions is similar in spirit to the use of a regularization method (e.g., Ref. [5]).

Although we demonstrate our approach with one-dimensional data, all aspects of the approach generalize straightforwardly to higher dimensions. One of the motivations of this work was to speed up the Monte Carlo method for two- and higher-dimensional data by reducing the number of free parameters by using sparse-basis representations. Although the boxcar-function is shown to be efficient, it is possible that other functions may also be used to achieve similar results.

## 6. Conclusions

This paper discusses the uncertainty of Laplace inversion results obtained from several different methods. Although all of these methods fit the data well, their uncertainty behavior is markedly different. The choice of different basis functions can be considered as applying a prior knowledge of the system, since they restrict the range of spectral solutions as in a regularization approach. The commonly-used delta-function basis often possesses a positive bias and a large variance in the estimate of  $\rho_T$  and  $\rho_B$ . This positive bias issue was also demonstrated in [7], where the spectrum is represented by a large number of spectral bins. On the other hand, the

boxcar-basis function with variable width does not exhibit such deficiencies and possesses no positive bias. Moreover, the boxcar basis provides a sparse representation of higher-dimensional spectra, making it a good candidate for use in an efficient algorithm for two- and higher-dimensional Laplace inversion.

## Appendix A. Directed Gibbs' sampler

A standard Gibbs' sampler, such as Algorithm A.39 in [11], evaluates the  $\pi(f)$  function (Eq. 10) along each parameter in  $\lambda$  in turn, forming a probability distribution along that coordinate direction. The parameter value is then updated with a value sampled from that distribution. Once all the parameters in  $\lambda$  have been updated, the resulting value of  $\lambda$  is the new sample drawn from  $\pi(f)$ .

This standard approach becomes inefficient when the distribution is highly correlated, as it is in the spectral inversion problem. Imaging a two-dimension problem in which the uncertainty ellipse is long, narrow and tilted at an angle to the coordinate axes. Since the standard Gibbs' sampler samples along the coordinate axes, the narrow width of the uncertainty ellipse forces the distributions along each coordinate axis to be commensurately narrow, thereby preventing long jumps along the long axis of the ellipse. This means that traversing the long axis of the ellipse will require many steps of the sampler, i.e., the samples are highly correlated, making an inefficient sampler.

This problem is avoided by choosing the Gibbs' sampling directions to align with the axes of the uncertainty ellipse. This makes large jumps possible. The inverse covariance matrix is periodically estimated by computing the Hessian of  $-\log \pi(f)$ , and the Gibbs' sampling directions are set to the eigenvectors of this matrix. Since the number of  $\lambda$  parameters is small and the Hessian need not be frequently updated, the computation expense of this operation is insignificant.

## References

- [1] W.J. Wiscombe, J.W. Evans, Exponential-sum fitting of radiative transmission functions, *Journal of Computational Physics* 24 (4) (1977) 416–444.
- [2] J.W. Evans, W.B. Gragg, R.J. LeVeque, On least squares exponential sum approximation with positive coefficients, *Mathematics of Computation* 34 (149) (1980) 203–211.
- [3] K. Tittelbach-Helmrich, An integration method for the analysis of multiexponential transient signals, *Measurement Science and Technology* 4 (1993) 1323–1329.
- [4] Å. Ukkelberg, G.H. Sørlandb, E.W. Hansen, H.C. Widerøe, ANAHESS, a new second order inverse Laplace transform algorithm, in: Presentation at the Ninth International Bologna Conference on Magnetic Resonance in Porous Media (MRPM9), Cambridge MA, USA, 2008.
- [5] Y.-Q. Song, L. Venkataramanan, M.D. Hurlimann, M. Flaum, P. Frulla, C. Straley,  $t_1 - t_2$  correlation spectra obtained using a fast two-dimensional laplace inversion, *Journal of Magnetic Resonance* 154 (2002) 261–268.
- [6] C.L. Epstein, J. Schotland, The bad truth about Laplace's transform, *SIAM Review* 50 (3) (2008) 504–520.
- [7] M. Prange, Y.-Q. Song, Quantifying uncertainty in NMR T2 spectra using Monte Carlo inversion, *Journal of Magnetic Resonance* 196 (2009) 54–60.
- [8] H. Akaike, A new look at the statistical model identification, *IEEE Transactions on Automatic Control* 19 (6) (1974) 716–723.
- [9] R.L. Parker, Y.-Q. Song, Assigning uncertainties in the inversion of NMR relaxation data, *Journal of Magnetic Resonance* 174 (2005) 314–324.
- [10] D.S. Sivia, *Data Analysis: A Bayesian Tutorial*, second ed., Oxford University Press, Oxford, UK, 2007.
- [11] C.P. Robert, G. Casella, *Monte Carlo Statistical Methods*, Springer, New York, 2004.
- [12] C.L. Lawson, R.J. Hanson, *Solving Least Squares Problems*, Prentice-Hall, Englewood Cliffs, NJ, 1974.
- [13] K.P. Burnham, D.R. Anderson, *Model Selection and Multimodel Inference*, second ed., Springer, New York, 2002.
- [14] A.R. Liddle, Information criteria for astrophysical model selection, *Monthly Notices of the Royal Astronomical Society: Letters* 377 (1) (2007) L74–L78.
- [15] R.L. Kleinberg, *Well logging*, *Encyclopedia of Nuclear Magnetic Resonance*, vol. 8, Wiley, New York, 1996, pp. 4960–4969.