The Open-Access Journal for the Basic Principles of Diffusion Theory, Experiment and Application

A Fast Monte Carlo Sampler for NMR T_2 Inversion

Michael Prange,[1] *Yi-Qiao Song*[1]

[1] Schlumberger-Doll Research, Cambridge, MA, USA

Corresponding author: Michael Prange, Schlumberger-Doll Research, One Hampshire St., Cambridge, MA 02139, E-Mail: prange@slb.com

(received 16 July 2008, accepted 26 January 2009)

Abstract

The inversion of noisy NMR T_2 echo data into a T_2 spectrum is widely recognized as an inherently non-unique process [1]. One approach to quantifying this uncertainty is to use Monte Carlo sampling. Uncorrelated measurement noise combine with the non-negativity constraint on T_2 spectral values to yield spectra following a non-negative normal distribution. There are two published samplers for truncated normal distributions [2], of which non-negative normal samples are a subset, but we show that these converge too slowly to be practical for the T_2 spectral inversion problem. This is because they are based on Gibbs' samplers that update the spectral estimate just one T_2 component at a time. When all of the spectral elements are fixed but one, that one has little room for change without violating the noise constraints on the data. Thus each spectral sample can only be slightly different from the preceding sample, indicating a high degree of statistical correlation and slow convergence. Our solution is to simultaneously update two neighboring spectral components at a time, allowing changes due to one spectral component to be offset by changes in its neighbor. Central to this improvement is a fast 2D slice sampler for non-negative normal distributions. This improves convergence by more than two orders of magnitude. Such speedup allows routine Monte Carlo inversion of 1D NMR spectra, and opens the door for the inversion of 2D NMR spectra.

Keywords

NMR Laplace Inversion, Monte Carlo Inversion, Gibbs' Sampler

1. Introduction

NMR spin relaxation spectra and diffusion constants are often used as a finger-print of the molecular species, structure and dynamics. For example, water and crude oils present in oil reservoirs can be distinguished by diffusion and relaxation experiments [3]. Data analysis often involves Laplace inversion to obtain a spectrum of relaxation times or diffusion constants. Such an inversion is ill-conditioned in the sense that for a given set of data with finite noise, many solutions will fit the data within the statistics of the noise. The well-established methods for spectral inversion use, for example, Tikonov regularization [4] and the maximum entropy method [5] to find one solution that fits the data and satisfies some other simultaneous constraint. This type of approach effectively makes a choice of the class of solution based on independent criteria. In the case of the regularization solution, smoother spectra are preferred over more spiky spectra. Different algorithms essentially use different preferences and thus result in different "best" solutions. However, it is difficult to justify these choices.

Instead of focusing interpretation efforts on a single "best" spectrum, we submit that a better approach is to present the multiplicity of spectra that match the data, allowing interpreters to provide an error estimate along with their interpretations. In the following we present a Monte Carlo inversion for NMR spectra that is capable of generating 10000 spectral samples in about 4 seconds on a desktop computer. Using a synthetic example, we demonstrate the method by estimating the total porosity and unbound fluid porosity in a sandstone rock.

2. Monte Carlo Inversion

Using 1D NMR T_2 inversion as an example, the continuum T_2 equation,

$$m(t) = \int_0^\infty f(T)\exp(-t/T)\,\mathrm{d}\log(T),\ f(T) \geq 0 \tag{1}$$

where $m(t)$ is the NMR echo measurement as a function of time and $f(T)$ is the spectrum as a function of T_2, is put into the vector form **m = Gf**, where **m** is the vector of measurements, **G** is the kernel matrix and **f** is the spectral vector. Since measurement noise is often in the form of an uncorrelated normal distribution [1], the probability density function (*pdf*) describing the variability in **f** can be expressed as the likelihood function

$$\pi(\mathsf{f}) \propto \exp\left[-\frac{1}{2}(\mathsf{m}-\mathsf{G}\mathsf{f})^T\Lambda^{-1}(\mathsf{m}-\mathsf{G}\mathsf{f})\right],\ \mathsf{f} \geq 0 \tag{2}$$

where Λ is a covariance matrix describing the measurement noise. This is a truncated multinormal distribution. All spectral samples drawn from $\pi(\mathbf{f})$ are valid spectra for the measured T_2 signals.

There are two Monte Carlo samplers in the literature for sampling truncated multinormal distributions [2]. These are Gibbs' samplers that sample an N-dimensional *pdf* as a sequence of one-dimensional sampling problems. Unfortunately, neither works well for NMR inversion because the *pdf* is both nearly singular and has a mean that is deep within the infeasible region. In our example below, both samplers generated samples that are still strongly correlated after 2000 samples, indicating severely inefficient sampling.

Our solution is to use a Gibbs' sampler that samples in two dimensions at a time. Each 2D sample is computed using a slice sampler [6]. This sampler was used to generate the 10000 samples used in the results below, requiring only about 4 seconds for the computation. These samples decorrelate in less than ten samples, indicating an efficient sampler.

3. Demonstration

Functionals of the spectra can be linked to important petrophysical properties. We use one such functional, unbound porosity defined by

$$\rho_U = \int_{\log T_c}^\infty f(T)\,\mathrm{d}\log T \ , \tag{3}$$

to demonstrate the value of Monte Carlo spectral inversion in NMR interpretation. $T_c = 0.033$ s is a typical cutoff value for sandstones [3]. We created an example T_2 spectrum and synthesized a noisy measurement vector from it using (1). Using our new sampler we generated 10000 Monte Carlo samples and computed unbound porosity using (3) for each of these samples. The results are histogramed in Fig 1, which also shows the true value as a black dot. The true solution is well within the acceptable error range for the Monte Carlo estimate. Unlike the single-solution methods for spectral inversion, Monte Carlo inversion allows an error estimate to be presented along with the "best" estimate.

3. Conclusions

Although the regularized approach to NMR T_2 spectral inversion yields a solution that is stable in the presence of noise, it fails to capture the considerable uncertainty present in the spectral inversion problem. We demonstrated that the spectral inversion problem can be expressed as a Monte Carlo sampling problem with a truncated multinormal distribution for which samplers exist in the literature. We show that these one-dimensional Gibbs' samplers are ineffective for the T_2 spectral inversion problem. We propose a modest extension of these samplers in which the Gibbs' sampling is done in two dimensions instead of one, and demonstrate that it is an efficient sampler for the T_2 spectral inversion problem when applied to the estimation of unbound porosity.

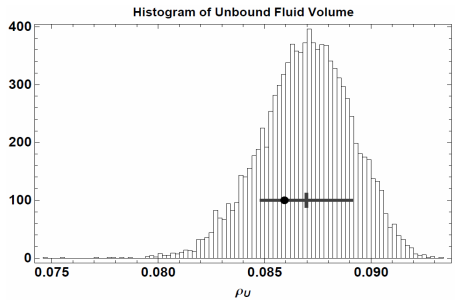

Fig 1: Histogram of the estimate of unbound porosity**.** The mean and standard deviation are indicated by gray bars, while the true value is shown as a black dot. The vertical axis indicates the number of samples in each histogram bin.

References

[1] R. Parker, Y-Q Song, Assigning uncertainties in the inversion of NMR relaxation data, J. Mag. Res. 174 (2005) 314-324.

[2] G. Rodriguez-Yam et al., Efficient Gibbs sampling of truncated multivariate normal with application to constrained linear regression, Technical Report, Colorado State University (2004).

[3] R. Kleinberg, Well logging, in: Encyclopedia of nuclear magnetic resonance, Vol. 8, Wiley, New York, 1996, pp. 4960–4969.

[4] A. Tikhonov, V. Arsenin, Solution of Ill-posed Problems, Winston and Sons, 1977.

[5] J. Skilling, Classic maximum entropy, in: J. Skilling (Ed.), Maximum Entropy and Bayesian Methods, Cambridge, Kluwer, 1989, pp. 45–52.

[6] C. Robert, G. Casella, Monte Carlo Statistical Methods, Springer, New York, 2004.