On the Use of Bayesian Probability Theory for Analysis of Exponential Decay Data: An Example Taken from Intravoxel Incoherent Motion Experiments

Jeffrey J. Neil, G. Larry Bretthorst

Traditionally, the method of nonlinear least squares (NLLS) analysis has been used to estimate the parameters obtained from exponential decay data. In this study, we evaluated the use of Bayesian probability theory to analyze such data; specifically, that resulting from intravoxel incoherent motion NMR experiments. Analysis was done both on simulated data to which different amounts of Gaussian noise had been added and on actual data derived from rat brain. On simulated data, Bayesian analysis performed substantially better than NLLS under conditions of relatively low signal-to-noise ratio. Bayesian probability theory also offers the advantages of: a) not requiring initial parameter estimates and hence not being susceptible to errors due to incorrect starting values and b) providing a much better representation of the uncertainty in the parameter estimates in the form of the probability density function. Bayesian analysis of rat brain data was used to demonstrate the shape of the probability density function from data sets of different quality.

Key words: Bayesian probability theory; nonlinear least squares; intravoxel incoherent motion.

INTRODUCTION

One of the more interesting and unusual NMR methods for detecting blood flow is the intravoxel incoherent motion (IVIM) method (1). The pulse sequence used in this method is made sensitive to the motion of blood through the application of magnetic field gradients in a manner analogous to the measurement of diffusion coefficient (2). The data obtained from application of this pulse sequence consist of a series of NMR free induction decay signals of varying amplitudes which correspond to different values of magnetic field gradient strength or timing (*vide infra*). In the simplest possible model for analysis of IVIM data, the water in the tissue of interest is described as being in either one of two compartments—intravascular or extravascular. Extravascular water moves by ordinary diffusion, and the process can be assigned a diffusion coefficient, D, with dimensions mm^2/s. Intravascular water not only moves *via* diffusion, but also moves with the bulk flow of blood. Blood flow displaces intravascular spins over much greater distances per unit time than diffusion. If the movement of these spins through vessels is modeled as random or "incoherent" motion, then this process can be assigned a pseudodiffusion coefficient, D^*, which also has dimensions mm^2/s. Because intravascular spins move greater distances per unit time, D^* is, in theory, considerably greater than D. Overall, there are three pieces of information available from this model of the IVIM experiment: f, the fraction of spins within the volume of interest that are within flowing blood (i.e., intravascular); D^*, the "pseudodiffusion coefficient" of those intravascular flowing spins; and D, the diffusion coefficient of the extravascular nonflowing spins.

The model given above is almost certainly an oversimplification. More accurate models, for example, might include: a) exchange between intravascular and extravascular compartments, b) diffusion anisotropy (different values for D depending upon the orientation of the magnetic field gradients with respect to the tissue (3)), or c) different types of movement within different types of blood vessels (e.g., artery *versus* vein).

Despite the simplicity of the model, the two parameters related to blood flow—f and D^*—are very difficult to estimate. Some authors have suggested that pseudodiffusion is not detectable (4), quite possibly because of difficulty detecting (let alone estimating) f and D^* under less than ideal conditions. The major problem in estimating these parameters is one of dynamic range. In brain, as in many organs, blood water comprises less than 5% of the total tissue water. Thus, less than 5% of the total signal arises from those water molecules which are of interest. The remaining 95% of the signal arises from extravascular water and hence is uninformative (at least theoretically) in assessing blood flow by the IVIM method. Therefore, while the total signal may be large, estimation of f and D^* depends upon measuring changes in less than 5% of the total signal. In other words, the dynamic range for the measurement is poor. There are modifications to the NMR pulse sequence used in IVIM experiments which can improve the dynamic range available for estimation of f and D^* (5, 6), and data from such an experiment will be shown.

The IVIM experiment represents one of many experiments (NMR or otherwise) whose results require analysis of exponentially varying data. Often, parameters of interested are derived or "estimated" from such data by nonlinear least squares (NLLS) analysis. Bayesian probability theory represents a powerful alternative parameter estimation method.

MRM 29:642–647 (1993)

From the Department of Pediatrics, Division of Pediatric Neurology, Washington University School of Medicine, (J.J.N.), and the Department of Chemistry, Washington University (G.L.B.), St. Louis, Missouri.

Address correspondence to: Jeffrey J. Neil, Ph.D., Division of Neurology, Department of Pediatrics, Campus Box 8116, St. Louis' Children's Hospital, 400 S. Kingshighway, St. Louis, MO 63110.

Received July 21, 1992; revised November 10, 1992; accepted November 10, 1992.

This work was supported by Grants NS01453, GM30331, and RR 02469 from the National Institutes of Health.

0740-3194/93 $3.00

In its broadest sense, Bayesian probability theory is a quantitative theory of inference. It is applicable to virtually any problem in which logical inferences must be made. In practice, it has been applied mainly to situations in which there is uncertainty in the problem (i.e., noise in the data) such that no unique solution exists. All of the information relevant to estimating a parameter is summarized in a probability density function. The argument of the probability density function is a hypothesis. For example, one could hypothesize that the diffusion coefficient had a particular value, x. The value of the probability density function for $D = x$ expresses how much one should believe, in the colloquial sense, that x represents the correct value for D (*vide infra*). The rules used for manipulating probabilities are the standard rules of statistics plus Bayes' theorem. Thus, Bayesian probability theory reinterprets the rules for manipulating probabilities. This method of data analysis has been used with a number of systems, from corn crop yields (7) to neutron scattering molecular tunneling data (8).

In addition to improved accuracy of the parameter estimates, Bayesian probability theory offers two other advantages over nonlinear least squares (NLLS). First, Bayesian probability theory does not require starting values for parameter estimation, while NLLS does. Using incorrect starting values for NLLS parameter estimation may cause the searching algorithm to find a local minimum of χ^2, not the global minimum, giving incorrect estimates. On the other hand, prior information about the model is easily incorporated into the Bayesian theory. Incorporation of such information improves the accuracy of the estimates. Second, Bayesian analysis provides a much better representation of the uncertainty in the parameter estimates in the form of the probability density function. As will be discussed below, the probability density function for a single data set is in some ways analogous to the sampling distribution derived from multiple data sets. But the probability density function offers the unique advantage of estimating the uncertainty in the parameter estimate from a *single* data set.

In this study we compare NLLS and probability theory by analyzing both simulated IVIM data (which offers the advantage that the "correct" parameter values are known) and real data from rat brain. The analysis presented here is not restricted to IVIM experiments, and may be generalized to *any* physical process for which a biexponential function is used as the model.

THEORY

The simple, two-compartment model described above can be expresses as:

$$S(b_i)/S_o = (1 - f)e^{-b_iD} + fe^{-b_iD^*},\ D < D^*,\ f \geq 0 \quad [1]$$

where $S(b_i)$ is the echo amplitude with diffusion gradients on at amplitude b_i, and S_o is the echo amplitude with diffusion gradients off. The parameters f, D, and D^*, are as defined above. The convention $D < D^*$ enforces the condition that the larger of the two decay rates corresponds to the pseudodiffusion rate. Given that $S_o > 0$, the requirement that $f \geq 0$ enforces the condition that both exponents must be positive, a requirement that is known from prior information. The variable b_i is given by the expression

$$b_i = \gamma^2 G_i^2 \delta^2 \left(\Delta - \frac{1}{3}\delta\right) \quad [2]$$

where G_i is the gradient strength in gauss/mm for the ith data point (assuming one is varying gradient strength to vary b; one could just as easily vary δ or even Δ), γ is the magnetogyric ratio, δ is the length of time for which each gradient pulse was applied, and Δ is the time interval between gradient pulses. The parameter b_i is under experimental control and has the units s/mm^2. For a drawing of the pulse sequence timing used to determine b_i (see ref. 5).

In Bayesian probability theory, all the information in the data relevant to estimating each of the three parameters is given by three probability density functions. For the fraction of total spins within the intravascular space, f, this is given symbolically by:

$$P(f \mid X,I) = \int\int\int d\sigma\, dD^*\, dD\, P(f,D^*,D,\sigma \mid X,I) \quad [3]$$

where $P(f \mid X,I)$ is the probability for the fraction of total spins, f, given the data, X, and the prior information, I; and $P(f,D^*,D,\sigma \mid X,I)$ is the joint posterior probability for f,D^*,D, and σ (the noise standard deviation) given the data and the prior information. Using uninformative prior probabilities, $P(f,D^*,D,\sigma \mid X,I)$ is given by:

$$P(f,D^*,D,\sigma \mid X,I) \propto \sigma^{-(N+1)} D^{-1} (D^*)^{-1} \exp\left\{-\frac{1}{2\sigma^2}\sum_{i=1}^{N}(X_i - S(b_i)/S_o)^2\right\} \quad [4]$$

where N is the number of data points and X_i represents the ith data point. Similarly, the probabilities for the other two parameters are given by:

$$P(D^* \mid X,I) = \int\int\int d\sigma\, dD\, df\, P(f,D^*,D \mid X,I) \quad [5]$$

and

$$P(D \mid X,I) = \int\int\int d\sigma\, dD^*\, df\, P(f,D^*,D \mid X,I). \quad [6]$$

The integrals were performed numerically. The limits of the integrals were set so that outside the integration region the joint distribution contributed a negligible amount to the integral. The parameter estimates used in this study were the value of the parameter which maximized each of the three probability density functions. For more detailed descriptions of the application of Bayes' theorem, including its application to complex free induction decay parameter estimation (i.e., amplitude, frequency, decay rate, and phase) (see refs. 7, 9–17).

For estimating the parameters in Eq. [1] using NLLS, the Levenberg-Marquardt algorithm (18) was used with the *correct values* for the parameter taken as the initial estimates when simulated data were analyzed. This, of course, is not possible with actual data.

MATERIALS AND METHODS

Both real and computer simulated amplitude data were analyzed. The simulated data were generated from b values of 10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 200, 300, 400, 500, 600, and 700 s/mm^2 using Eq. [1]. The values for f, D^*, and D were 0.05, 0.010 mm^2/s, and 0.001 mm^2/s, respectively. These values for b, f, D^*, and D correspond to those used in a study of NLLS analysis of IVIM data reported by Pekar *et al.* (19). The ratio of D^* to D for the computer simulated data is one theoretical estimate of 10 (1), although in rat experiments we have found the ratio to be closer to 200 (5). The maximum time domain signal amplitude (corresponding to $b = 10$) of the simulated data was normalized to 100. Gaussian noise was then added at root-mean-square (RMS) amplitudes of 0.25 and 2.5. In other words, the maximum signal-to-RMS-noise ratio was either 400 or 40. Note that throughout this manuscript, we refer to the signal-to-noise ratio (SNR) of the time domain amplitude data, as opposed to the SNR of the frequency domain NMR signal. Seventeen thousand data sets were generated at each of the two noise levels. Each data set was analyzed using both NLLS and Bayesian probability theory.

The parameter estimates obtained from each method of analysis were plotted as sampling distributions (i.e., parameter estimate on the abscissa, number of occurrences of a particular estimate on the ordinate). When preparing the sampling distributions, estimates outside the plot ranges were discarded. At SNR=40, 1308 and 1513 estimates were discarded and at SNR=400, 26 and 14 estimates were discarded from the Bayesian and NLLS estimates, respectively. The mean and standard deviation from each sampling distribution was then calculated.

Real NMR data obtained from rat brain (from a study published previously (5)) were also analyzed. For these data, an intravascular relaxation agent and modified pulse sequence were used to suppress signal from extravascular (i.e., unwanted) spins, thereby improving dynamic range. It was possible to modify the pulse sequence such that the suppression of unwanted signal was more or less effective (by varying the inversion time TI, see (5)). This had the effect of producing data sets from which more or less of the total signal was derived from intravascular spins, and hence pertinent to the estimation of D^* (see Discussion). Two of these data sets were analyzed to demonstrate the form of the probability density function obtained through Bayesian analysis.

The computer code for Bayesian analysis was written in FORTRAN and run on a Sun SPARC 2 workstation. Bayesian analysis of one data set, including generation of the files required to display the probability density functions for the three estimated parameters, required approximately one CPU minute.

RESULTS

Simulated Data

The means and standard deviations for the estimates of all three parameters (f, D^*, and D) at both noise levels (SNR=40 and SNR=400) are shown in Table 1. At SNR=40, neither NLLS nor Bayesian analysis provides good estimates for f or D^*; although with knowledge of the true answers it is clear that the Bayesian estimates of f are better both in standard deviation and mean value. Bayesian analysis performs much better than NLLS in estimating D, as demonstrated by the factor of two difference between the standard deviations of the parameter estimates (0.07 *versus* 0.14) using the two methods. This difference can also be seen in the sampling distributions shown in Fig. 1A. Note that the curve for Bayesian analysis (solid line) is more sharply peaked than the corresponding curve for NLLS (dotted line), with a higher proportion of estimates for D near the true value of 1.0×10^{-3} mm^2/s.

For the parameter estimates at SNR=400, both Bayesian analysis and NLLS provide good estimates of f and D (Table 1). However, Bayesian analysis does a much better job of estimating D^*—the standard deviations of the estimates differ by a factor of ten.

Figure 1B shows a typical probability density function for D derived by applying Bayesian analysis to a *single data set* at SNR=40. Note the similarity in overall shape of Fig. 1B and the Bayesian curve in Fig. 1A.

Rat Brain Data

Figure 2 shows the probability density functions for two determinations of D^* on the same animal. Suppression of extravascular signal was used for each determination, but the experimental conditions were set such that the suppression was optimized for the determination represented by the solid line, but not optimized for that represented by the dashed line. The values of D^* at the peak of the probability density function for each data set are virtually identical: 0.148 mm^2/s for optimized and 0.155 mm^2/s for nonoptimized. As can be seen from the plot, the probability density function for the optimized data is much more sharply peaked than the one for the nonoptimized data. When the same data were analyzed using

TABLE 1
A Comparison of Parameter Estimates from Bayesian Probability Theory and NLLS

Parameter[a]	Signal-to-noise ratio	Bayesian analysis		NLLS analysis	
		Mean	SD[b]	Mean	SD
f	40	5.9	4.5	8.5	7.4
D^*	40	21	38	20	30
D	40	1.03	0.07	0.96	0.14
f	400	5.2	1.3	5.3	1.2
D^*	400	9.1	0.9	10.2	9.6
D	400	1.00	0.02	0.99	0.02

[a] Units for f are percent. Units for D^* and D are $\times 10^{-3}$ mm^2/s. Correct values are: $f = 5.0\%$, $D = 1.0 \times 10^{-3}$ mm^2/s, $D^* = 10.0 \times 10^{-3}$ mm^2/s.
[b] Standard deviation.

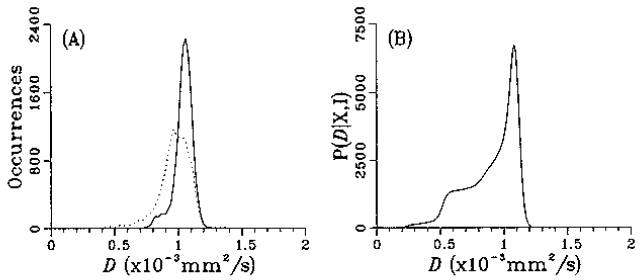

FIG. 1. (A) The sampling distribution of parameter estimates using NLLS (dotted line) and Bayesian analysis (solid line). The estimates shown are derived from approximately 17,000 data sets with SNR=40. The correct value of D is 1.0×10^{-3} mm²/s. The abscissa represents estimates for the value of D. The ordinate represents the number of occurrences of each particular estimate among the 17,000 data sets. (B) The probability density function generated through Bayesian analysis for a single simulated data set at SNR=40. The abscissa represents possible values for D. The ordinate represents the probability of each particular value given the data and prior information. The area under the curve is normalized to 1.0. The correct value for D is 1.0×10^{-3} mm²/s.

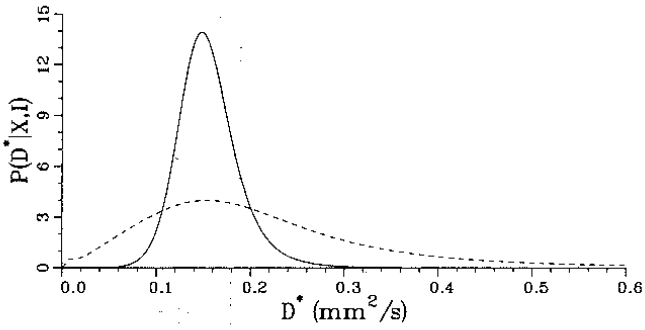

FIG. 2. A plot of the probability density functions for two determinations of the pseudodiffusion coefficient, D^*, in rat brain (see text for details). The solid line represents a data set for which suppression of extravascular signal was optimized ($f = 0.25$); the dashed line represents a data set for which the suppression was not optimized ($f = 0.17$). Both data sets were obtained from the same animal.

least squares, the estimates for D^* from these data were 0.165 mm²/s for the optimized and 0.204 mm²/s for the nonoptimized data sets.

DISCUSSION

The quality of parameter estimates in IVIM experiments depends upon a number of factors, including: the number of data points obtained, their b values, errors in calculation of the b values (such as due to nonlinear amplifier performance or local magnetic field gradients), the "apparent" value of f (*vide infra*), and the applicability of the model to the data. All else being equal, at high SNR, both Bayesian analysis and NLLS provide good parameter estimates. For the conditions of Table 1, both methods perform well in estimating f and D at SNR=400. Conversely, at low enough SNR, neither method provides good parameter estimates, as in estimation of D^* at SNR=40. However, there is a window in which Bayesian analysis provides good parameter estimates and NLLS does not. In Table 1, there are two instances which clearly fall within this window—estimation of D at SNR=40 and estimation of D^* at SNR=400. In both instances, the standard deviation of the Bayesian analysis estimates is substantially less than those for the corresponding NLLS estimates. For estimating f at SNR=40, Bayesian analysis gives estimates of f that are closer to the true value (18% *versus* 70% error). Though the standard deviations for the estimates from both methods are rather large in this particular instance, that from Bayesian analysis is about half that of NLLS. In short, Bayesian analysis provides a marked advantage over NLLS in parameter estimation under certain, but not all, circumstances.

It is important to note that the two methods were not on exactly equal footing in this comparison. The NLLS analysis was operating under ideal circumstances because the starting values employed in the Levenberg-Marquardt search algorithm were the correct ones. The Bayesian analysis, in contrast, did not require initial estimates. On the other hand, it is possible to incorporate information about the model into Bayesian analysis *prior to* analyzing the data. For the analysis used in this study, the conventions $D^* > D$ and $f \geq 0$ were used. Forcing $D^* > D$ ensured that the larger of the two estimated decay rates would be assigned to D^*. Not allowing f to be negative ensured that both exponential components were positive and was taken as a part of the model itself.

Bayesian analysis offers another advantage in that the probability density function is obtained as an inherent part of the process of estimating each parameter and provides information regarding the uncertainty of the parameter estimate. There is no associated accuracy estimate when NLLS is used to estimate the parameters. Typically, one of two procedures is used to express the uncertainty in a parameter estimate from NLLS—either a Monte-Carlo simulation or a confidence interval (18). Both of these procedures express the uncertainty in the form of a mean ± standard deviation. Assuming a normally distributed population, an estimated value falls within one standard deviation of the mean 68% of the time if the parameter is estimated *multiple times*. In the case of Monte-Carlo, one simulates the parameter estimate multiple times and then computes a sampling distribution (as was done here). The accuracy estimates are then given as the width of this sampling distribution. The sharper the peak in such a sampling distribution, the better the parameter estimate and the smaller the standard deviation. As can be seen in Fig. 1A, the sampling distribution obtained from the Bayesian estimates (solid line) is sharper than that for NLLS (dotted line).

In using the confidence interval method, constant contours of χ^2 are used to determine the standard deviation. This requires knowledge of the standard deviation of the noise, which is usually not known. But if the errors in the data are normally distributed, the expected value of χ^2 is the number of degrees of freedom. Making this assumption, the noise standard deviation is set to ensure that the expected value of χ^2 is the number of degrees of freedom. This then allows calculation of confidence intervals. As with the Monte-Carlo method, these confidence intervals are defined and derived in terms of a theoretical situation in which the parameter has been estimated multiple

times and make no statement about the accuracy of the parameters estimated from *one particular* data set.

When parameter estimates are made using Bayesian probability theory, the probability density function expresses what is known about the parameter. Parameter values for which the probability density function returns a low probability are values of the parameter for which one should have low "belief." Those for which the probability density function returns a high probability are those for which one should have high "belief." The value corresponding to the peak or maximum of the probability density function is the value for which one should have the strongest "belief." In Figs. 1A and 1B the sampling distribution of the Bayesian estimates of D at SNR=40 (from roughly 17,000 data sets, solid line) and the probability density function (from a single data set) are side by side. The overall shape of the probability density function in Fig. 1B is similar to the corresponding sampling distribution in Fig. 1A, though the probability density function comes from a single data set. While this probability density function is typical, each data set has a *unique* probability density function for each estimated parameter. For the simulated data, this depends on the noise that was added to the data set. The added noise either widens or narrows the shape of the probability density function depending upon how the noise affects the fit. As can be seen in the plots, the probability density function in this case is broader than the corresponding sampling distribution. This indicates that for this particular data set, D was estimated less well than what one would have assumed from the sampling distribution. Thus, sampling distribution, while useful, can be misleading when applied to any particular data set.

It is possible to generalize the mean ± 1 standard deviation estimate using Bayesian probability theory. Bayesian probability density functions are normalized such that their total integrated area is one, and one could determine the region around the peak of the probability density function that would encompass some fraction of the total probability. For example, in Fig. 2, one could draw horizontal lines at successively lower probabilities until the area enclosed by the horizontal line and a probability density function was 0.68, or encompassed 68% of the total probability. As the probability density functions are somewhat asymmetric, the values for D^* at which the line intersects the probability density function may not be symmetrically placed about the peak. Using this procedure on the probability density functions given in Fig. 2, the values are $0.148 + 0.030$ and -0.026 mm^2/s and $0.155 + 0.110$ and -0.088 mm^2/s for the dynamic range optimized and nonoptimized data sets, respectively. With this procedure, one can say that there is a 68% probability that the correct value of D^* was within the estimated range, and this conclusion is based on the actual, observed data set. In contrast, no statement can be made about a particular data set when using estimates derived through NLLS.

The probability density functions derived through Bayesian analysis are sometimes asymmetric, as can be seen in Fig. 2. This asymmetry is seen in both simulated and real data, and is more pronounced at low SNRs. It is a consequence of the model used. For data with low SNR, the probability density function for D^* is broad and its low end has a tendency to overlap the estimated values for D. But the model assumes that D^* is greater than D. As a result, the lower portion of the probability density function falls off more steeply from the peak than the upper portion and the curve is asymmetric. For data with high SNR, the probability density function for D^* is sharply peaked and does not significantly overlap the estimated values for D, so these probability density functions are nearly symmetric.

The situation is quite different for data sets of very low SNR (e.g., less than 5, data not shown). Under these circumstances, Bayesian analysis returns values for D and D^* which are virtually identical. The probability density function for D^* is quite asymmetric, with a sharp peak at the first allowed values of D^* greater than D, and a long "tail" toward higher values of D^*. Bayesian analysis is reporting, in effect, that there is not enough information in the data set to estimate the second component. But the model assumes that this component, or parameter, is present and greater than D. The shape of $P(D^* \mid X,I)$ indicates that the value must be bounded by D and some maximum. So while probability theory can not really tell you the value of D^* under these conditions, it can tell you what values of D^* have been ruled out by the data.

The quality of parameter estimation is mainly dependent upon the SNR. The study by Pekar *et al.* (19) gave estimates of what SNR is necessary to obtain reasonable parameter estimates. The NLLS data presented here are consistent with their conclusions, though Bayesian analysis allows reasonable estimates at lower SNR than NLLS. It is interesting to note that for estimating the parameters D and D^*, the "signal" in "signal-to-noise ratio" is essentially the SNR for each of the individual decay components. For a value of $f = 0.05$, 95% of the signal is derived from "extravascular" spins, with only the remaining 5% coming from "intravascular" spins. This makes the determination of D (for "extravascular" spins) roughly 20 times better than that of D^* (for "intravascular" spins). This difference is reflected in the observation that D^* could not be determined at SNR=40 while D could. In experiments in which this ratio (the "apparent" value of f) can be altered, it is possible to obtain better estimates of D^* (5, 6). This is what was done in the animal experiments used for Fig. 2. The data sets were obtained under identical conditions except that the pulse sequence was modified to alter the amount of suppression of extravascular signal (or the apparent values of f). For the optimized curve, the apparent value of f was 0.25; whereas for the nonoptimized curve, the apparent value of f was 0.17.

Overall, Bayesian analysis offers substantially better parameter estimates than NLLS for IVIM data. In addition, Bayesian analysis provides a very powerful way of assessing the uncertainty in a parameter estimate by means of the probability density function.

ACKNOWLEDGMENT

We express our sincere thanks to Dr. Joseph J. H. Ackerman for useful discussion and criticism.

REFERENCES

1. D. LeBihan, E. Breton, D. Lallemand, M. L. Aubin, J. Vignaud, M. Laval-Jeantet, *Radiology* **168,** 497 (1988).
2. E. O. Stejskal, J. E. Tanner, *J. Chem. Phys.* **42,** 288 (1965).
3. D. LeBihan, *Magn. Reson. Q.* **7,** 1 (1991).
4. K. K. Kwong, P. C. McKinstry, D. Chien, A. P. Crawley, J. D. Pearlman, B. R. Rosen, *Magn. Reson. Med.* **21,** 157 (1991).
5. J. J. Neil, L. A. Scherrer, J. J. H. Ackerman, *J. Magn. Reson.* **95,** 607 (1991).
6. J. J. Neil, C. S. Bosch, J. J. H. Ackerman, *J. Magn. Reson.* **98,** 436 (1992).
7. G. L. Bretthorst, *in* "Bayesian Spectrum Analysis and Parameter Estimation," vol. 48, Springer-Verlag, New York, 1988.
8. D. S. Sivia, C. J. Carlile, *J. Chem. Phys.* **96,** 170 (1992).
9. G. L. Bretthorst, C.-C. Hung, D. A. D'Avignon, J. J. H. Ackerman, *J. Magn. Reson.* **79,** 369 (1988).
10. G. L. Bretthorst, J. J. Kotyk, J. J. H. Ackerman, *Magn. Reson. Med.* **9,** 282 (1989).
11. G. L. Bretthorst, *Magn. Reson. Med.* **88,** 533 (1990).
12. G. L. Bretthorst, *Magn. Reson. Med.* **88,** 552 (1990).
13. G. L. Bretthorst, *Magn. Reson. Med.* **88,** 571 (1990).
14. G. L. Bretthorst, *J. Magn. Reson.* **93,** 369 (1991).
15. G. L. Bretthorst, *J. Magn. Reson.* **98,** 501 (1992).
16. J. J. Kotyk, N. G. Hoffman, W. C. Hutton, G. L. Bretthorst, *J. Magn. Reson.* **98,** 483 (1992).
17. P. W. Anderson, *Phys. Today* **45,** 9 (1992).
18. W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, *in* "Numerical Recipes: The Art of Scientific Computing," Cambridge University Press, Cambridge, 1986.
19. J. Pekar, C. T. W. Moonen, P. C. M. van Zijl, *Magn. Reson. Med.* **23,** 122 (1992).