

Optimisation of NMR dynamic models I. Minimisation algorithms and their performance within the model-free and Brownian rotational diffusion spaces

Edward J. d'Auvergne · Paul R. Gooley

Received: 23 February 2007 / Accepted: 6 November 2007 / Published online: 18 December 2007
© Springer Science+Business Media B.V. 2007

Abstract The key to obtaining the model-free description of the dynamics of a macromolecule is the optimisation of the model-free and Brownian rotational diffusion parameters using the collected R_1 , R_2 and steady-state NOE relaxation data. The problem of optimising the chi-squared value is often assumed to be trivial, however, the long chain of dependencies required for its calculation complicates the model-free chi-squared space. Convolutions are induced by the Lorentzian form of the spectral density functions, the linear recombinations of certain spectral density values to obtain the relaxation rates, the calculation of the NOE using the ratio of two of these rates, and finally the quadratic form of the chi-squared equation itself. Two major topological features of the model-free space complicate optimisation. The first is a long, shallow valley which commences at infinite correlation times and gradually approaches the minimum. The most severe convolution occurs for motions on two timescales in which the minimum is often located at the end of a long, deep, curved tunnel or multidimensional valley through the space. A large number of optimisation algorithms will be investigated and their performance

compared to determine which techniques are suitable for use in model-free analysis. Local optimisation algorithms will be shown to be sufficient for minimisation not only within the model-free space but also for the minimisation of the Brownian rotational diffusion tensor. In addition the performance of the programs Modelfree and Dasha are investigated. A number of model-free optimisation failures were identified: the inability to slide along the limits, the singular matrix failure of the Levenberg–Marquardt minimisation algorithm, the low precision of both programs, and a bug in Modelfree. Significantly, the singular matrix failure of the Levenberg–Marquardt algorithm occurs when internal correlation times are undefined and is greatly amplified in model-free analysis by both the grid search and constraint algorithms. The program relax (<http://www.nmr-relax.com>) is also presented as a new software package designed for the analysis of macromolecular dynamics through the use of NMR relaxation data and which alleviates all of the problems inherent within model-free analysis.

Keywords Brownian rotational diffusion · Model-free analysis · Minimisation · Newton minimisation · Optimisation · NMR relaxation

Electronic supplementary material The online version of this article (doi:10.1007/s10858-007-9214-2) contains supplementary material, which is available to authorized users.

E. J. d'Auvergne (✉)
Department of NMR-based Structural Biology, Max Planck
Institute for Biophysical Chemistry, Am Fassberg 11, D-37077
Goettingen, Germany
e-mail: edward@nmr-relax.com

P. R. Gooley
Department of Biochemistry and Molecular Biology, Bio21
Institute of Biotechnology and Molecular Science, University of
Melbourne, Parkville, VIC 3010, Australia
e-mail: prg@unimelb.edu.au

Abbreviations

AIC	Akaike's Information Criteria
BFGS	Broyden, Fletcher, Goldfarb, and Shanno quasi-Newton optimisation
CD	Coordinate descent optimisation
CG	Conjugate gradient
CSA	Chemical Shift Anisotropy
\mathcal{D}	set of diffusion tensor parameters
DMG	Double Motion Grid
\mathcal{G}	set of geometric diffusion parameters

GMW	Gill, Murray, and Wright
LM	Levenberg–Marquardt optimisation
MT	Moré and Thuente step-length selection algorithm
NR	the combined Newton–Raphson/conjugate gradient minimisation of Dasha
\mathfrak{D}	set of orientational diffusion parameters
RG	R_{ex} Grid
\mathfrak{E}	the global model, space, or universe
XH bond	heteronucleus-proton bond

Introduction

Detailed, atomic resolution information describing the dynamics of proteins or other macromolecules can be obtained experimentally through NMR by the measurement of the R_1 and R_2 heteronuclear relaxation rates and the steady-state heteronuclear NOE. These numbers by themselves are, however, difficult to interpret meaningfully and are therefore usually decomposed using the model-free analysis of Lipari and Szabo, Lipari and Szabo (1982a, b) which separates the global rotational diffusion of the macromolecule from the internal motions of the molecule's bond vectors. The relaxation values are translated into the parameters of Brownian rotational diffusion as well as those of the internal motions which consist of the generalised order parameter S^2 and the effective correlation time τ_e . The order parameter S^2 is related to the amplitude of the internal motion. An S^2 value of zero indicates large amplitude motions whereas a value of one means that the bond vector is completely rigid. The effective correlation time τ_e is correlated with the timescale of the internal motion of the bond vector albeit being dependent on the amplitude of that motion. The original Lipari and Szabo model-free analysis was extended by Clore et al. (1990) to include internal motions on two timescales. The order parameter S_f^2 and effective correlation time τ_f describe the faster of the two motions whereas S_s^2 and τ_s are associated with the slower motion. The two order parameters are related by the equation $S_f^2 \cdot S_s^2 = S^2$.

By assuming certain order parameters or correlation times to be statistically negligible, either being one or zero respectively, a number of model-free models can be constructed. An order parameter of one means that the motion is statistically insignificant whereas the correlation time of zero means that the motion is too fast for that parameter to be reliably extracted. In combining parametric restrictions, in which statistically insignificant parameters are dropped, together with the addition of a parameter accounting for chemical exchange relaxation a number of increasingly complex models of model-free motions can be constructed

(Fushman et al. 1997; Orekhov et al. 1999; Korzhnev et al. 2001; Zhuravleva et al. 2004). Extending this list the model-free models will be labelled as

$$m0 = \{\}, \quad (1.0)$$

$$m1 = \{S^2\}, \quad (1.1)$$

$$m2 = \{S^2, \tau_e\}, \quad (1.2)$$

$$m3 = \{S^2, R_{ex}\}, \quad (1.3)$$

$$m4 = \{S^2, \tau_e, R_{ex}\}, \quad (1.4)$$

$$m5 = \{S^2, S_f^2, \tau_s\}, \quad (1.5)$$

$$m6 = \{S^2, \tau_f, S_f^2, \tau_s\}, \quad (1.6)$$

$$m7 = \{S^2, S_f^2, \tau_s, R_{ex}\}, \quad (1.7)$$

$$m8 = \{S^2, \tau_f, S_f^2, \tau_s, R_{ex}\}, \quad (1.8)$$

$$m9 = \{R_{ex}\}, \quad (1.9)$$

where R_{ex} is the parameter accounting for chemical exchange relaxation which solely affects the R_2 relaxation rate. The chemical exchange is assumed to be fast and is scaled quadratically with field strength. Model $m0$ corresponds to the special situation whereby no internal motions are statistically significant.

The simplest model of Brownian rotational diffusion is the tumbling of a sphere and is parameterised either by global correlation time τ_m or the isotropic diffusion rate \mathfrak{D}_{iso} . For spherical diffusion all three eigenvalues of the diffusion tensor are equal, $\mathfrak{D}_x = \mathfrak{D}_y = \mathfrak{D}_z$. When only two of the eigenvalues are equal the molecule diffuses as a spheroid. This model, which is synonymous with axially symmetric anisotropic diffusion, is parameterised by $\mathfrak{D} = \{\mathfrak{D}_\perp, \mathfrak{D}_\parallel, \theta, \phi\}$ where the geometric parameters \mathfrak{G} are defined as $\mathfrak{D}_\perp \equiv \mathfrak{D}_x = \mathfrak{D}_y$ and $\mathfrak{D}_\parallel \equiv \mathfrak{D}_z$ (d'Auvergne and Gooley 2007). The orientational parameters $\mathfrak{D} = \{\theta, \phi\}$ are the azimuthal angle and the polar angle respectively. Alternatively spheroidal diffusion can be parameterised by $\mathfrak{D} = \{\mathfrak{D}_{iso}, \mathfrak{D}_a, \theta, \phi\}$ where \mathfrak{D}_a is the anisotropic diffusion parameter which when positive denotes a prolate spheroid and when negative denotes an oblate spheroid (d'Auvergne 2006). The most complex model of rotational diffusion is when all three eigenvalues are different and is the diffusion of an ellipsoid (Einstein 1905; Perrin 1934, 1936; Favro 1960; Woessner 1962). The ellipsoid can be parameterised by $\mathfrak{D} = \{\mathfrak{D}_x, \mathfrak{D}_y, \mathfrak{D}_z, \alpha, \beta, \gamma\}$ where the geometric parameters \mathfrak{G} are the eigenvalues of the tensor and the orientational parameters \mathfrak{D} are the Euler angles defined using the z - y - z notation. Alternatively the ellipsoid can be characterised by the parameters $\mathfrak{D} = \{\mathfrak{D}_{iso}, \mathfrak{D}_a, \mathfrak{D}_r, \alpha, \beta, \gamma\}$ where \mathfrak{D}_r is a measure of the rhombicity (d'Auvergne 2006).

Not only can the diffusion tensor be optimised as a global model affecting all residues of the protein but a set

of model-free models can be constructed in which each residue is assumed to diffuse independently. In these models a single local τ_m parameter approximates the true, multiexponential description of the Brownian rotational diffusion of the molecule. Each residue of the macromolecule is treated independently. Another set of model-free models which add the local τ_m parameter to Models $m0$ to $m9$ can be created and will be labelled $tm0$ to $tm9$. These models are an extension of the ideas introduced in Barbato et al. (1992) and Schurr et al. (1994) whereby the model $m2$, the original Lipari and Szabo model-free equation, is extended with the local τ_m parameter to avoid the issues associated with inaccurate diffusion tensor approximations.

A number of computer programs have been written to minimise the χ^2 value in search of optimal parameter values of the model-free models. The majority of these programs implement solely local minimisers as it is assumed that only a single minimum exists. In certain situations two minima within a single model-free space have been identified (Orekhov et al. 1995b). Nevertheless, global minimisers such as simulated annealing or genetic algorithms are generally considered superfluous for model-free analysis. The most commonly used program in the literature is the Modelfree program (Palmer et al. 1991; Mandel et al. 1995), the current version of which will be labelled as Modelfree4. Another widely used model-free optimisation program is Dasha (Orekhov et al. 1995a) in which two local optimisation algorithms are available. Less widely used are the programs DYNAMICS (Fushman et al. 1997) and Tensor 2 (Blackledge et al. 1998; Cordier et al. 1998; Dosset et al. 2000; Tsan et al. 2000).

This article also presents a new software package named ‘relax’ which is designed for the analysis of NMR relaxation data. The program is open source and can be freely downloaded from <http://www.nmr-relax.com>. All aspects of data analysis starting from the initial peak intensities and ending with the visualisation of the final results are implemented within the program. The software is highly flexible as relax has been written as a modular collection of data analysis tools. The advantages of relax include a high precision of optimisation; an advanced and efficient constraint algorithm; the ability to avoid the Levenberg–Marquardt minimisation algorithm which tends to be problematic for model-free analysis; built-in support for both model elimination (d’Auvergne and Gooley 2006) and numerous model selection techniques (d’Auvergne and Gooley 2003); support for all types of rotational diffusion tensor; the use of the numerically stabilised model-free spectral density functions; simplified diffusion equations (d’Auvergne 2006); and a documented collapse of the diffusion tensor symmetries (d’Auvergne 2006). The program can also create the input files for both Modelfree (Palmer et al. 1991; Mandel et al. 1995) and Dasha

(Orekhov et al. 1995a), execute both programs in-line, and then read their output. This means that both programs can be used as optimisation engines replacing the minimisation algorithms built into relax. A number of other data analysis techniques involving NMR relaxation data are also supported by relax including the calculation of the steady-state NOE and its error, exponential curve-fitting to determine the R_1 and R_2 relaxation rates through optimisation and error propagation through Monte Carlo (MC) simulation, and reduced spectral density mapping (Lefevre et al. 1996) using MC simulation to translate the uncertainties.

Theory and methods

Model-free theory

The original model-free spectral density functions presented in Lipari and Szabo (1982a) and Clore et al. (1990) are not the most numerically stable form of these equations. An important problem encountered in optimisation is round-off error in which machine precision influences the result of mathematical operations. The double reciprocal $\tau^{-1} = \tau_m^{-1} + \tau_e^{-1}$ used in the equations are operations which are particularly susceptible to round-off error, especially when $\tau_e \ll \tau_m$. By incorporating these reciprocals into the model-free spectral density functions and then simplifying the equations this source of round-off error can be eliminated. The new, numerically stabilised form of the model-free equations can be expressed as

$$J(\omega) = \frac{2}{5} \tau_m \left(\frac{S^2}{1 + (\omega\tau_m)^2} + \frac{(1 - S^2)(\tau_e + \tau_m)\tau_e}{(\tau_e + \tau_m)^2 + (\omega\tau_e\tau_m)^2} \right), \quad (2)$$

$$J(\omega) = \frac{2}{5} \tau_m \left(\frac{S^2}{1 + (\omega\tau_m)^2} + \frac{(1 - S_f^2)(\tau_f + \tau_m)\tau_f}{(\tau_f + \tau_m)^2 + (\omega\tau_f\tau_m)^2} + \frac{(S_f^2 - S^2)(\tau_s + \tau_m)\tau_s}{(\tau_s + \tau_m)^2 + (\omega\tau_s\tau_m)^2} \right). \quad (3)$$

The perfect, noise-free grids

To thoroughly test the optimisation algorithms across a large range of model-free motions, two grids covering the entire model-free space of models $m4$ (1.4) and $m5$ (1.5) were constructed. The parameter values at each grid point were used to create a synthetic data set consisting of the R_1 , R_2 and NOE at both 600 and 500 MHz. Each grid point was treated as a separate, independent residue of a large protein. These two grids are the same as presented in d’Auvergne and Gooley (2003, 2006) and will be labelled

the R_{ex} Grid (RG) for the data of model $m4$ and the Double Motion Grid (DMG) for the data of model $m5$.

Optimisation

The model-free parameters

Prior to optimisation using relax a coarse grid search of 11 increments per dimension was carried out to find the initial parameter vector, with $0 \leq S^2 \leq 1$, $0 \leq \tau_i \leq 10$ ns and $0 \leq R_{ex} \leq 20$ s⁻¹. This grid search, which is a ubiquitous optimisation technique, is in no way related to the RG and DMG described above which were simply used to sample different model-free motions. A grid search over the same range was carried out for Modelfree4 but with 20 increments per dimension, the granularity of which reflects the values presented in the examples section of the Modelfree4 manual and hence would be similar to that used by much of the model-free literature. Although the steps of 526.316 ps in the correlation time dimension are coarse compared to the values extracted in an analysis, this is finer than the 1,000 ps steps used with relax for the comparison. If optimisation is unsuccessful, then the granularity of the grid search will have a large effect on the final results. The numerous optimisation algorithms tested, the precision of optimisation and issue of parameter deconvolution are all addressed in the supplementary material.

The diffusion tensor parameters

All optimisations of the diffusion parameters presented within this paper employed Newton minimisation in conjunction with the backtracking line search algorithm (Nocedal and Wright 1999) and the Gill, Murray and Wright modified Cholesky Hessian modification algorithm (Gill et al. 1981). For the initial grid search of the diffusion tensor parameters all searches were carried out between the values $1 \text{ ns} \leq \tau_m \leq 12 \text{ ns}$, $0 \leq \mathfrak{D}_r \leq 1$, and $0 \leq \theta, \phi, \alpha, \beta, \gamma \leq \pi$. For the prolate spheroid and the ellipsoid the parameter \mathfrak{D}_a ranged from 0 to $1e^7$. For the oblate spheroid the parameter ranged from $-1e^7$ to 0.

Mapping of the model-free space

A third type of grid was used for the mapping of the model-free and diffusion tensor χ^2 spaces. By laying out a grid over select parameter values the chi-squared value was calculated at each position. These values were then fed into the OpenDX program (Open Visualization Data Explorer 4.3.2, <http://www.opendx.org>) to create images of the model-free spaces. The spaces were mapped to a resolution of 100 points per dimension.

Analysis of the relaxation data of cytochrome c_2

The R_1 and R_2 relaxation rates for the 44 helical residues of cytochrome c_2 from *Rhodobacter capsulatus* were obtained from the supplementary material of Blackledge et al. (1998). For consistency and to enable a direct comparison with the original publication a PDB file was created from the amide NH coordinates from the same supplementary material. For optimisation model-free model $m1$ (1.1) was selected for all residues. Rather than using a grid search for finding the optimal parameter values to use as an initial position for minimisation, this position was chosen to be where the S^2 values of all residues are 0.8. The diffusion tensor parameters were set approximately to the values in the supplementary material of Blackledge et al. (1998). For the constants used in minimisation the values from Cordier et al. (1998) were used. The NH bond length was set to 1.01 Å, the CSA value set to -170 ppm, and the field strength of the relaxation data was assumed to be exactly 600 MHz.

Results and discussion

The R_{ex} and DMG

All grid points of both the R_{ex} Grid and DMG were optimised using three different programs: Modelfree4, Dasha and relax. The program relax (<http://www.nmr-relax.com>) implements all of the optimisation algorithms utilised in this paper. The exact versions of the programs used are version 4.15 of Modelfree, version 3.48c of Dasha, and various versions of relax (from 1.0 to 1.2). To visualise the results the difference between each optimised parameter and the true value was calculated. For each parameter the difference for every grid point was joined to create a surface using OpenDX. If the minimum has been successfully located the parameter difference should be zero for each point and hence the surface should be perfectly flat with a height of zero. The figures for Modelfree4, Dasha and relax of the S^2 , τ_e and R_{ex} difference surfaces of the RG and S^2 , S_f^2 and τ_s of the DMG are presented in the supplementary material (Figs. S1–S18). A select few of these surfaces are presented in Fig. 1.

Comparison of optimisation algorithms

The curvature of the single motion model-free space

For a large majority of the single timescale motions of models $m1$ (1.1) to $m4$ (1.4) represented by the R_{ex} Grid,

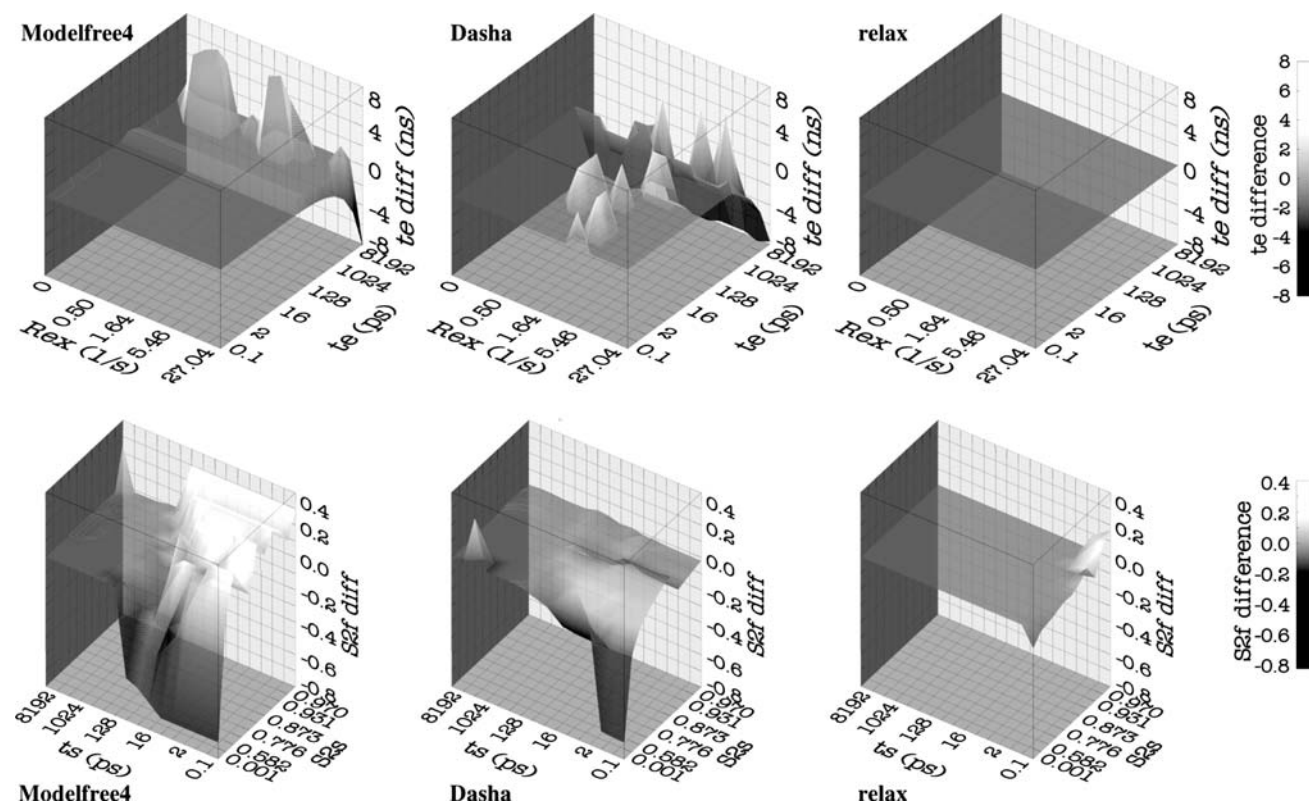


Fig. 1 The results of the optimisation of both the R_{ex} Grid (RG) and the DMG using Modelfree4, Dasha and relax presented as difference surfaces. For each grid point the programs were used to optimise either model m_4 or m_5 for the RG and DMG, respectively. If the minimum has been found for each point the difference between the optimised and true parameter values should be zero. Positive and negative differences correspond to over and underestimation respectively. As a surface has been draped over the discrete differences, perfect optimisation should result in a flat surface of zero height. The

top three plots correspond to the RG τ_e difference surfaces for the subset of the grid whereby $S^2 = 0.952$. The bottom three plots correspond to the DMG S_f^2 difference surfaces for the subset of the grid whereby $S_f^2 = 0.698$. The optimisation methods used are: Modelfree4, the Levenberg–Marquardt algorithm; Dasha, the combined Newton–Raphson/conjugate gradient algorithm; relax, Newton optimisation together with the backtracking line search, GMW Hessian modification, and the Method of Multipliers constraint algorithm

optimisation is relatively simple requiring only a small number of iterations for most algorithms to find the minimum. Nevertheless, minimisation can be complicated by being drawn towards and then being forced to navigate along shallow, long, curved valleys which originate at infinite correlation times but slowly head to the minimum (for example see Fig. 2). This is the same topological feature which can induce model failure due to distortion by experimental noise—the true local minimum is lost while a new minimum appears at infinite τ_e values (d’Auvergne and Gooley 2006). Despite this topology, in which most algorithms can navigate without problem, the curvature in the vicinity of the minimum is relatively quadratic.

The tunnels of the double motion model-free space

For the models representing motions on two timescales, models m_5 to m_8 (1.5–1.8), the curvature of the space can be much more convoluted. In certain situations to reach the

single local minimum within the model-free space the optimisation algorithms must pass through a long, curved and deep tunnel or multidimensional valley in which minimisation can be quite complex. The closer the time-scales of the two motions, τ_f and τ_s , the more drawn out, crescent-shaped and finer the tunnel becomes and therefore the more difficult it is to reach the local minimum. Examples of this type of curvature are shown in Fig. 3 (as well as Figs. S19, S21 and S22 of the supplementary material).

The best optimisation techniques

As discussed in the supplementary material, of the numerous optimisation algorithms tested the Newton line search algorithm combined with the backtracking line search (Nocedal and Wright 1999) and GMW Hessian modification (Gill et al. 1981) was found to be the best combination for model-free analysis. The use of parameter




Fig. 2 The failure of the constraints algorithm in Modelfree4 demonstrated by minimisation terminating at the upper limit of the τ_e parameter of 10 ns. The figure is a map of the chi-squared space of model *m4* which is composed of the parameters $\{S^2, \tau_e, R_{ex}\}$. Different χ^2 values are demonstrated by the four isosurfaces which, from outermost to innermost, possess the values of 1, 0.5, 0.3 and 0.05. The relaxation data was generated by back calculation from the model-free parameter values of $S^2 = 0.970$, $\tau_e = 2048$ ps and

$R_{ex} = 0.149 \text{ s}^{-1}$. As no noise was added the minimum for this model is located at this position. The four spheres in the plot correspond to the results as found by Modelfree4 using the Levenberg–Marquardt algorithm (black sphere), Dasha using either Levenberg–Marquardt or Newton–Raphson–CG minimisation (grey spheres), and relax using Newton minimisation together with the backtracking line search and the GMW Hessian modification (white sphere). The exact coordinates of the spheres are listed in Table S6 of the supplementary material

deconvolution in the double motion models, by optimising the parameters $\{S^2, S_f^2, \tau_s\}$ rather than $\{S_f^2, S_s^2, \tau_s\}$, is also important for simplifying the model-free optimisation problem.

Failure of optimisation


By investigating the results of minimisation of Modelfree4, Dasha and relax across the grids spanning the model-free space a number of points of failure have been identified. These are described in detail below and include: the failure of the limits algorithm within Modelfree4; the singular matrix failure of the Levenberg–Marquardt optimisation algorithm in all programs; the low precision of both Modelfree4 and Dasha causing incomplete optimisation;

and a bug within the Levenberg–Marquardt algorithm of Modelfree4 causing early termination of minimisation.

Failure of the limits algorithm

The type of curvature which complicates model-free optimisation with boundaries is illustrated in Fig. 2 by the grid point $S^2 = 0.970$, $\tau_e = 2048$ ps and $R_{ex} = 0.149 \text{ s}^{-1}$ where the shallow valley parallel to the correlation time dimension is clearly visible. The results of optimisation using Modelfree4, Dasha and relax are presented in Table S6 in the supplementary material, the corresponding positions in the model-free space being indicated by spheres in the figure. This example was chosen because not only does it illustrate the failure of the constraints

Fig. 3 The effect of optimisation precision on the final model-free parameter values. The deep and curved tunnel of the model-free space is illustrated by the four isosurfaces which correspond, from outermost to innermost, to chi-squared values of 50, 20, 5 and 0.5, respectively. The χ^2 space belongs to the grid point $S_f^2 = 0.952$, $S_s^2 = 0.582$, and $\tau_s = 32$ ps. The white sphere corresponds to the true parameter values as well as the results of the program relax, the black sphere corresponds to the results of Modelfree4, and the two grey spheres are the results from Dasha using the two available optimisation algorithms. The parameter and chi-squared values of these positions are given in Table S7 of the supplementary material



algorithm in Modelfree4 but it also demonstrates a failure of the Levenberg–Marquardt algorithm in Dasha which is due either to the singular matrix failure of the algorithm or to Dasha’s constraint algorithm. The results of Dasha using the combined Newton–Raphson/conjugate gradient optimisation are close to the truth yet low precision prevents the exact values from being found. relax is able to find the exact values to within machine precision.

Within Modelfree4 when a trajectory step of the Levenberg–Marquardt algorithm exceeds a constraint, that step terminates flush at the boundary. If the negative of the gradient at that position points back into the box then optimisation will continue, otherwise optimisation terminates. In Fig. 2 the downhill gradient at the point where Modelfree4 terminates points outside of the box, hence the minimum is not found. In Dasha constraints are implemented by imposing a penalty on steps outside the limit. The penalty is added to the chi-squared value as the distance outside the constraint to the sixth power (Orekhov et al. 1995a). The program relax implements the Method of Multipliers algorithm (Nocedal and Wright 1999), also known as the Augmented Lagrangian, which allows optimisation to be successful in this case. As a penalty is added

to the chi-squared value when a constraint is violated the optimisation trajectory can exit the box in which minimisation is constrained. During the progression of the Method of Multipliers algorithm the penalty is increased. The end result is that optimisation will slide along the constraint face, albeit just outside the box, allowing the local minimum to be found.

Incomplete minimisation due to low precision

A number of internal program settings influence the precision of optimisation. These include the maximum number of iterations of the minimisation algorithm, the function tolerance, the gradient tolerance and a step length tolerance. For a simple space low optimisation precision will have a negligible effect, however the model-free space is highly convoluted. Both the curved, shallow valley parallel to the correlation time dimensions as well as the curved, deep tunnels through the double motion spaces are particularly sensitive to optimisation precision.

The grid point $S_f^2 = 0.952$, $S_s^2 = 0.582$, and $\tau_s = 32$ ps perfectly illustrates the results of the different optimisation

precisions of Modelfree4, Dasha and relax. It was chosen due to its optimisation complexity and hence its ability to thoroughly test the precision of optimisation. The space close to the minimum demonstrates the classic curved and deep tunnel of model-free double motion spaces (Fig. 3). Despite both the chi-squared difference between iterations and the length of each individual step being quite small the addition of numerous small steps results in a large change in parameter values and a large overall decrease in the chi-squared value.

The optimisation precisions of Modelfree4, Dasha and relax are compared in Table 1. From the low number of iterations as well as the relatively high cutoff tolerances, of all programs, Modelfree4 by far exhibits the lowest precision. This is reflected in the optimisation results where Modelfree4 terminates optimisation at the very start of the tunnel, the result being that the S_f^2 and S_s^2 parameters appear to be swapped and the correlation time is overestimated by an order of magnitude. The chi-squared difference between the results and the true minimum is 3.47, a statistically significant value which strongly biases model selection away from model $m5$. The relax cutoffs allow the true parameter values to be found regardless of the curvature of the model-free space thereby facilitating unbiased model selection. The precision of both Modelfree4 and Dasha can be increased to the high default values of relax by modification and recompilation of the source code. The high precision model-free results of both programs are presented in Table S8 of the supplementary material.

Singular matrix failure of Levenberg–Marquardt minimisation

The Levenberg–Marquardt optimisation algorithm is a widely used technique which, for most problems, performs extremely well. However this method is known to, at times, experience outright failure. Its Achilles heel is when the Levenberg–Marquardt matrix is singular which, when

Table 1 The optimisation precision of various model-free minimisation programs

Program	Max iter ^a	f tol ^b	g tol ^b	s tol ^b
Modelfree4	50	$1e^{-4}$		
Dasha (LM) ^c	$80n$	$1e^{-14}$	Off	$1e^{-14}$
Dasha (NR) ^d	Off	$1e^{-5}$	$1e^{-5}$	$1e^{-4}$
relax	$1e^7$	$1e^{-25}$	Off	

^a The maximum number of iterations

^b The function, gradient and step length tolerances

^c Levenberg–Marquardt minimisation. n is the number of parameters being optimised

^d Combined Newton–Raphson/conjugate gradient minimisation

encountered, causes the algorithm necessarily terminate at the current position. The failure is linked to the curvature of the space. The model-free space is particularly prone to the zero pivot or singular matrix failure of the Levenberg–Marquardt technique. This is due to the Lorentzian form of the model-free spectral density functions and occurs whenever one of the internal correlation times is undefined due to the corresponding order parameter being equal to one. When both order parameters S_f^2 and S_s^2 of the extended model-free spectral density function are equal to zero the τ_s parameter is additionally undefined.

Two features of model-free optimisation tend to amplify the problem, the initial grid search prior to optimisation and the constraints used during optimisation. When searching through the order parameter dimensions the grid searches have upper and lower limits of one and zero respectively. This results in a large number of initial starting positions whereby the first iteration of the Levenberg–Marquardt algorithm is preconditioned to fail. If constraints are used during optimisation the order parameters are limited to being between zero and one, again accentuating the problem. The number of failures encountered by Modelfree4 are given in Table 2. The percentages from both relax and Dasha should be lower due to the constraint algorithm. The result of these failures is that the model-free parameter values are meaningless. The convolution of the space ensures that proximity to the minimum cannot be guaranteed and the chi-squared value is significantly overestimated.

The Modelfree4 bug

One final issue affecting the results of model-free analysis was identified in the Modelfree program and has been traced back to a bug. The result is that optimisation is terminated prematurely. In many cases optimisation stops on the very first step—the final model-free parameters are those found by the initial grid search. To demonstrate the failure the grid point $S^2 = 0.388$, $\tau_e = 128$ ps and $R_{ex} = 0.223$ s⁻¹ was chosen, the simple curvature of which is

Table 2 The number of grid points of the RG and DMG whereby optimisation has failed due to the singular matrix failure of the Levenberg–Marquardt algorithm in Modelfree4

Model	RG (2,640 points)		DMG (1,940 points)	
	No.	%	No.	%
$m1$	0	0.00	0	0.00
$m2$	772	29.24	0	0.00
$m3$	0	0.00	0	0.00
$m4$	112	4.24	0	0.00
$m5$	834	31.59	232	11.96

demonstrated in Fig. 4. The gradient at the black sphere, which is the position of the Modelfree4 results, is perpendicular to the isosurface passing through it hence failure is clearly not due to optimisation complexity. The results of optimisation using Modelfree4, Dasha and relax are presented in Table S9 of the supplementary material. The number of failures for both grids are presented in Table 3. This bug has been fixed in Modelfree versions 4.20 and above.

The double minima of the model-free space

Local versus global optimisation

In the model-free optimisation of models m_0 to m_9 (1.0–1.9), models tm_0 to tm_9 , the Brownian rotational diffusion tensor parameters, or the totality of all model-free and diffusion parameters one fundamental assumption about the chi-squared space is made. The programs Modelfree4, Dasha, DYNAMICS and relax all use local rather than global optimisation algorithms. Their success is reliant upon the presence of only one minimum within the space

Table 3 The number of grid points of the RG and DMG whereby optimisation has failed due to the Levenberg–Marquardt bug in Modelfree4


Model	RG (2,640 points)		DMG (1,940 points)	
	No.	%	No.	%
m_1	545	20.64	461	23.76
m_2	1,093	41.40	529	27.27
m_3	345	13.07	597	30.77
m_4	715	27.08	735	37.89
m_5	1,176	44.55	696	35.88

being searched. If this is indeed true, finding the local minimum is synonymous with finding the global minimum. However it has been noted that there are cases where this basic premise collapses.

The two minima

In Orekhov et al. (1995b), and subsequently discussed in Korzhnev et al. (1997), it was demonstrated that there are

Fig. 4 An example of optimisation failure in Modelfree4 caused by the bug in the Levenberg–Marquardt algorithm. The chi-squared space belongs to the grid point where the true parameter values are $S^2 = 0.388$, $\tau_e = 128$ ps and $R_{ex} = 0.223$ s⁻¹. From outermost to innermost, the five isosurfaces illustrating the curvature of the space correspond to chi-squared values of 1371.79, 500, 100, 20 and 7. The true parameter values which were found by both relax and Dasha are indicated by the white sphere whereas the black sphere corresponds to the final parameter values found by Modelfree4 ($S^2 = 0.263$, $\tau_e = 526.316$ ps and $R_{ex} = 1.053$ s⁻¹). The failure occurred in the first iteration hence the final parameter values are those of the grid point with the lowest chi-squared value (1371.79)



situations in which two minima exist within a single space. This occurred when motions were occurring on two timescales but the fitted model only represented a single internal motion. The example presented was composed of synthetic data representing *m5* (Model 5) where the parameter values were $S_f^2 = 0.9$, $S_s^2 = 0.8$, and $\tau_s = 0.8$ ns (ibid.). The global correlation time was set to 7 ns and, as there is only a single residue, this model is identical to *tm5*. To approximately replicate the example the CSA and bond length were set to -170 ppm and 1.02 Å, respectively. In Fig. 5a the two minima of the space are clearly distinguishable. When using noise-free synthetic data, as in Orekhov et al. (1995b), the first and broadest of the minima is located at the position $S^2 = 0.749$, $\tau_e = 0.89$ ns and $\tau_m = 7.1$ ns, its chi-squared value being 8.27. The second minimum, which is narrower but is the global minimum, is at $S^2 = 0.790$, $\tau_e = 0.10$ ns and $\tau_m = 6.5$ ns. Its chi-squared value is $1e^{-29}$ and, if found by the local minimiser instead of the broad minimum, AIC model selection (Akaike 1973; d’Auvergne and Gooley 2003) will pick *tm2* over the true *tm5* model.

Avoidance of the problem

The selection of the simplistic and incorrect double minima model can be avoided in a number of indirect ways. The first is implicit within experimental data as the addition of noise to the synthetic test case of Orekhov et al. (1995b) differentially shifts the chi-squared values of all the models and, in most cases, tilts the balance towards model *tm5* (Fig. 5b). For the first and second minima of the *tm2* space

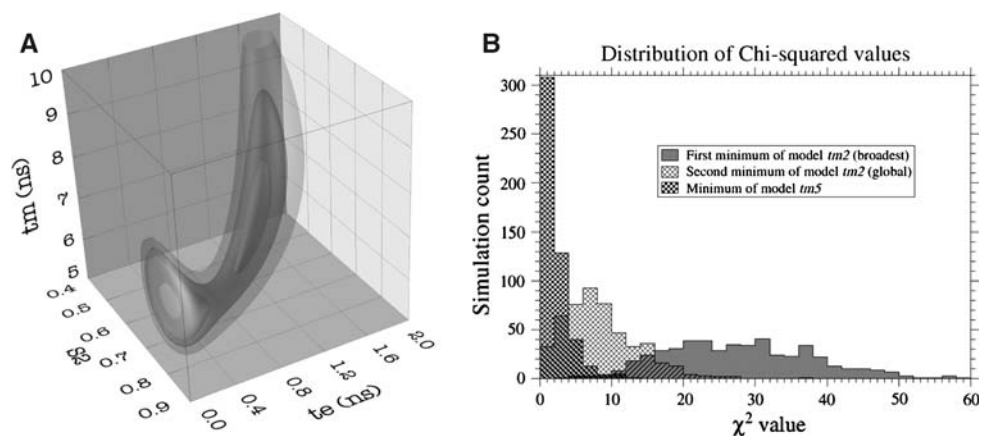


Fig. 5 A rare example of the double minima phenomenon existent in a space of a simplistic approximation to the true model. In this case the true model is *tm5* where $\tau_m = 7$ ns, $S_f^2 = 0.9$, $S_s^2 = 0.8$, and $\tau_s = 0.8$ ns whereas the space in (a) is a map of model *tm2*. The back calculated relaxation data used to generate this plot consisted of data at 500 and 600 MHz. The four isosurfaces shown correspond, from outermost to innermost, to χ^2 values of 30, 22, 17 and 5 respectively. The resolution of the plot is 100 data points per dimension. The three

the average chi-squared values shifted to 28.0 ± 10.5 and 8.5 ± 5.2 , respectively. The chi-squared value of model *tm5* shifted from zero to an average of 2.0 ± 2.0 and, in using AIC model selection, this model is selected 84% of the time over the global minimum of model *tm2* (assuming that this minimum is found 100% of the time). Therefore for most experimental data AIC model selection will discard the double minima model.

Second, the double minima phenomenon is only a tiny subset of a much bigger problem—the hiding of nanosecond motions when the whole protein exhibits motions on these timescales. This larger problem will be discussed in detail in the next paper of this series, Paper II, and can be avoided by using the new model-free optimisation protocol which is also presented in that paper. Finally, if all else fails, collection of data at another field strength will solve the problem (Orekhov et al. 1995b). The addition of 800 MHz data to the above synthetic, noise-free example causes model *tm5*, which now has a minimised chi-squared value of $3.4e^{-10}$, to be chosen by AIC model selection over model *tm2* which has a new global minimum of 6.82. The chi-squared value of the broader minimum is 16.98.

Cytochrome *c*₂: optimisation of the diffusion tensor

When solving the model-free puzzle in order to obtain the dynamic description of the system, the optimisation of models *m0* to *m9* is only one part of a multifaceted enigma. A second minimisation problem needs to be resolved and that is the optimisation of the Brownian rotational diffusion tensor.

chi-squared distributions shown in (b), corresponding to the two minima of model *tm2* and the single minimum of model *tm5*, were generated by 500 randomisations of the original noise-free data assuming Gaussian noise. The chi-squared values were binned to increments of 2. These distributions demonstrate that experimental noise will cause most instances of model *tm2* in which two minima are present to be eliminated by AIC model selection as model *tm5* is selected instead

Two different approaches: global versus local optimisation

An ideal, real protein test case for the optimisation of the rotational diffusion tensor is that of the ferrocycytochrome c_2 molecule from the photosynthetic alpha-proteobacterium *Rhodobacter capsulatus* (Blackledge et al. 1998; Cordier et al. 1998). This system, a proven ellipsoid, is perfect for testing the local optimisation of the diffusion tensor parameters using the new model-free optimisation protocol (Paper II). In the analysis of Blackledge et al. (1998) the optimisation algorithm utilised was not a local but a global minimiser—simulated annealing combined with a fuzzy logic algorithm utilising the Levenberg–Marquardt local minimiser as a sub-algorithm. The R_1 and R_2 relaxation rates collected at 600 MHz were used to characterise the diffusion tensor of cytochrome c_2 using the program Tensor. This approach employed the original ideas of Kay et al. (1989) whereby the R_2/R_1 ratio is used to find the diffusion tensor by chi-squared optimisation. To replicate as closely as possible the results of Blackledge et al. (1998) and Cordier et al. (1998) and to allow for comparison of the global and local optimisation techniques model-free model $m1$ (Model 1) was chosen for all residues. The use of this model is important as the chi-squared equation used in Tensor and that used in relax are defined differently.

Although the global and local optimisation algorithms are operating in χ^2 spaces with different dimensions and dimensionalities the final diffusion tensors should nevertheless be the same. This is demonstrated in Table S10 of the supplementary material where the original results from the program Tensor (Blackledge et al. 1998; Cordier et al. 1998; Dosset et al. 2000; Tsan et al. 2000) are contrasted with the results from the program relax. Comparing the eigenvalues of the various diffusion tensors both methodologies can be seen to return the same geometric picture.

The glide reflection and translational symmetries

Figure 6 shows a subset of the chi-squared space which is searched in the optimisation of both the prolate and oblate spheroid diffusion tensors. Four distinct minima are clearly visible—two correspond to the prolate spheroid, two to the oblate spheroid. The duplication of minima within each of the spheroid types is due to the glide reflection symmetry of the orientational parameter space \mathfrak{D} (d’Auvergne 2006). The paired minima are identical and are related by $\{\theta, \phi\} = \{\pi - \theta, \phi - \pi\}$. This figure is very similar to Fig. 1 of Blackledge et al. (1998) except $\mathfrak{D}_{\text{ratio}}$ has been swapped for \mathfrak{D}_a and the glide reflection has been shown.

The map of the subspace of ellipsoid orientational parameters $\mathfrak{D} \in \mathfrak{S}$ (d’Auvergne and Gooley 2007) shown in Fig. S24 of the supplementary material was constructed

similarly to that of the spheroid subspace. The geometric diffusion parameters \mathfrak{G} and all S^2 values were fixed to those of the minimised diffusion tensor while the Euler angles of \mathfrak{D} were varied. The figure clearly demonstrates the four identical minima caused by the glide reflection and translational symmetries of the space \mathfrak{D} .

The adequacy of local minimisers


For each of the diffusion models the diffusion parameter space \mathfrak{D} contains only a single local minimum because of three restrictions: the eigenvalue permutation restriction of the geometric parameter subspace \mathfrak{G} by $\mathfrak{D}_x \leq \mathfrak{D}_y \leq \mathfrak{D}_z$; the collapse of the glide reflection and translational symmetries of the subspace \mathfrak{D} by $\alpha = \alpha - \pi, \{\alpha, \beta, \gamma\} = \{\pi - \alpha, \pi - \beta, \gamma - \pi\}, \{\alpha, \beta\} = \{\pi - \alpha, \beta - \pi\}$ and $\{\theta, \phi\} = \{\pi - \theta, \phi - \pi\}$; and the isolation of the prolate and oblate subspaces by the restrictions $\mathfrak{D}_a \geq 0$ and $\mathfrak{D}_a \leq 0$ respectively (d’Auvergne 2006). The smoothness of the space and the presence of only a single local and, therefore, global minimum in either the spheroid subspaces or full ellipsoid space is demonstrated in both Figs. 6 and S24. Because of these two factors Tensor’s global minimiser, which is based on simulated annealing combined with a fuzzy logic algorithm, is not essential for optimising the diffusion tensor.

Conclusion

Two topological features convolute the model-free space causing optimisation to be a non-trivial problem. The best local optimisation algorithm within the model-free space was found to be Newton minimisation in conjunction with the backtracking line search and the Gill, Murray and Wright (GMW) Hessian modification. In using the cytochrome c_2 relaxation data (Blackledge et al. 1998; Cordier et al. 1998), local optimisation algorithms were shown to be sufficient for optimising the diffusion tensor parameters.

Four problems have been identified in the currently used model-free optimisation software. The singular matrix failure of the Levenberg–Marquardt minimisation algorithm is a problem which occurs when correlation times are undefined, a phenomenon associated with order parameters of one. Model-free analysis is highly prone to this failure as both the initial coarse grid search and constraint algorithms can place the order parameter at one. This common algorithm should be avoided for model-free analysis. The second problem is low precision causing the final values to be distant from the local minimum, affecting both Model-free and Dasha. To find the minimum in the twisted model-free space high precision optimisation with fine cutoffs and a large number of iterations is essential. The

Fig. 6 Map of the chi-squared space of the spheroid diffusion tensor parameters $\{\mathcal{D}_a, \theta, \phi\}$ of cytochrome c_2 . This figure demonstrates the multiple minima of the space due to both the prolate and oblate approximations to the true ellipsoid diffusion tensor together with the glide reflection symmetry of the space. The glide reflection is most evident in the top right image where the subspace between $\pi \leq \phi \leq 2\pi$ is a duplication and mirror image reflection about θ of the subspace between $0 \leq \phi \leq \pi$. To map this space all S^2 parameters and τ_m were fixed to the values of the minimised prolate spheroid. The χ^2 values of the four isosurfaces from outermost to innermost are 300, 200, 100 and 70, respectively (dark grey to white). The resolution of the plot is 100 data points per dimension



third failure is due to Modelfree4's constraint algorithm in which sliding along the limits is never allowed. Finally, Modelfree4 is heavily affected by a bug which causes early termination of the Levenberg–Marquardt algorithm. The result of these four failures of optimisation is that minimisation is terminated early, often far from the true model-free parameter values.

All the steps of model-free analysis from the initial spectral peak intensities to the final model-free description of the dynamics of the system are implemented by the program relax. The model selection methodologies described in d'Auvergne and Gooley (2003), relaxation curve-fitting, NOE calculation, reduced spectral density mapping and model-free analysis are all available within relax. This software solves all of the optimisation issues described herein.

Acknowledgements We thank Andrew J. Perry for discussions on both the science and programming aspects behind relax. A number of relax's design features, usability and core layout ideas originated from those discussions. This includes the concepts of the prompt-based interface with the help system and tab completion, the automatic generation of the documentation on the user functions by parsing their docstrings, and the powerful Python scripting available to the user. Haydyn D. T. Mertens and James D. Swarbrick are acknowledged for helpful discussions. We are indebted to Chris A. MacRaild for many

code contributions quashing a large number of bugs, to Gary S. Thompson for both improvements to the exponential relaxation curve-fitting C modules and uncovering problems in relax, and to James Horne and Horst J. Schirra for help in tracking down bugs throughout the program.

References

- Akaike H (1973) Information theory and an extension of the maximum likelihood principle. In: Petrov BN, Csaki F (eds) Proceedings of the second international symposium on information theory. Budapest, Akademia Kiado, pp 267–281
- Barbato G, Ikura M, Kay LE, Pastor RW, Bax A (1992) Backbone dynamics of calmodulin studied by ^{15}N relaxation using inverse detected two-dimensional NMR spectroscopy: the central helix is flexible. *Biochemistry* 31(23):5269–5278
- Blackledge M., Cordier F, Dosset P, Marion D (1998) Precision and uncertainty in the characterization of anisotropic rotational diffusion by ^{15}N relaxation. *J Am Chem Soc* 120(18): 4538–4539
- Broyden CG (1970) The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations. *J Inst Math Appl* 6(1):76–90
- Clore GM, Szabo A, Bax A, Kay LE, Driscoll PC, Gronenborn AM (1990) Deviations from the simple 2-parameter model-free approach to the interpretation of ^{15}N nuclear magnetic-relaxation of proteins. *J Am Chem Soc* 112(12):4989–4991
- Cordier F, Caffrey M, Brutscher B, Cusanovich MA, Marion D, Blackledge M (1998) Solution structure, rotational diffusion

- anisotropy and local backbone dynamics of *Rhodobacter capsulatus* cytochrome c2. *J Mol Biol* 281(2):341–361
- d'Auvergne EJ (2006) Protein dynamics: a study of the model-free analysis of NMR relaxation data. Ph.D. thesis, Biochemistry and Molecular Biology, University of Melbourne. <http://eprints.infodiv.unimelb.edu.au/archive/00002799/>
- d'Auvergne EJ, Gooley PR (2003) The use of model selection in the model-free analysis of protein dynamics. *J Biomol NMR* 25(1):25–39
- d'Auvergne EJ, Gooley PR (2006) Model-free model elimination: a new step in the model-free dynamic analysis of NMR relaxation data. *J Biomol NMR* 35(2):117–135
- d'Auvergne EJ, Gooley PR (2007) Set theory formulation of the model-free problem and the diffusion seeded model-free paradigm. *Mol Biosyst* 3(7):483–494
- Dosset P, Hus JC, Blackledge M, Marion D (2000) Efficient analysis of macromolecular rotational diffusion from heteronuclear relaxation data. *J Biomol NMR* 16(1):23–28
- Einstein A (1905) Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen (The motion of elements suspended in static liquids as claimed in the molecular kinetic theory of heat). *Ann Physik* 17(8):549–560
- Favro LD (1960) Theory of the rotational brownian motion of a free rigid body. *Phys Rev* 119(1):53–62
- Fletcher R (1970) A new approach to variable metric algorithms. *Comput J* 13(3):317–322
- Fletcher R, Reeves CM (1964) Function minimization by conjugate gradients. *Comp J* 7(2):149–154
- Fushman D, Cahill S, Cowburn D (1997) The main-chain dynamics of the dynamin pleckstrin homology (PH) domain in solution: analysis of 15N relaxation with monomer/dimer equilibration. *J Mol Biol* 266(1):173–194
- Gill PE, Murray W, Wright MH (1981) Practical optimization. Academic Press
- Goldfarb D (1970) A family of variable-metric methods derived by variational means. *Math Comp* 24(109):23–26
- Hestenes MR, Stiefel E (1952) Methods of conjugate gradients for solving linear systems. *J Res Nat Bur Stand* 49(6):409–436
- Kay LE, Torchia DA, Bax A (1989) Backbone dynamics of proteins as studied by 15N inverse detected heteronuclear NMR spectroscopy: application to staphylococcal nuclease. *Biochemistry* 28(23):8972–8979
- Korzhnev DM, Billeter M, Arseniev AS, Orekhov VY (2001) NMR studies of Brownian tumbling and internal motions in proteins. *Prog NMR Spectrosc* 38(3):197–266
- Korzhnev DM, Orekhov VY, Arseniev AS (1997) Model-free approach beyond the borders of its applicability. *J Magn Reson* 127(2):184–191
- Lefevre JF, Dayie KT, Peng JW, Wagner G (1996) Internal mobility in the partially folded DNA binding and dimerization domains of GAL4: NMR analysis of the N–H spectral density functions. *Biochemistry* 35(8):2674–2686
- Levenberg K (1944) A method for the solution of certain non-linear problems in least squares. *Quart Appl Math* 2:164–168
- Lipari G, Szabo A (1982a) Model-free approach to the interpretation of nuclear magnetic-resonance relaxation in macromolecules I. Theory and range of validity. *J Am Chem Soc* 104(17):4546–4559
- Lipari G, Szabo A (1982b) Model-free approach to the interpretation of nuclear magnetic-resonance relaxation in macromolecules II. Analysis of experimental results. *J Am Chem Soc* 104(17):4559–4570
- Mandel AM, Akke M, Palmer AG 3rd (1995) Backbone dynamics of *Escherichia coli* ribonuclease HI: correlations with structure and function in an active enzyme. *J Mol Biol* 246(1):144–163
- Marquardt DW (1963) An algorithm for least squares estimation of non-linear parameters. *SIAM J* 11:431–441
- More' JJ, Thuente DJ (1994) Line search algorithms with guaranteed sufficient decrease. *ACM Trans Maths Software* 20(3):286–307
- Nocedal J, Wright SJ (1999) Numerical optimization, Springer series in operations research. New York, Springer-Verlag
- Orekhov VY, Korzhnev DM, Diercks T, Kessler H, Arseniev AS (1999) H-1-N-15 NMR dynamic study of an isolated alpha-helical peptide (1-36)- bacteriorhodopsin reveals the equilibrium helix-coil transitions. *J Biomol NMR* 14(4):345–356
- Orekhov VY, Nolde DE, Golovanov AP, Korzhnev DM, Arseniev AS (1995a) Processing of heteronuclear NMR relaxation data with the new software DASHA. *Appl Magn Reson* 9(4):581–588
- Orekhov VY, Pervushin KV, Korzhnev DM, Arseniev AS (1995b) Backbone dynamics of (1-71)Bacterioopsin and (1-36)Bacterioopsin studied by 2-dimensional H-1-N-15 NMR-spectroscopy. *J Biomol NMR* 6(2):113–122
- Palmer AG, Rance M, Wright PE (1991) Intramolecular motions of a zinc finger DNA-binding domain from Xfin characterized by proton-detected natural abundance C-12 heteronuclear NMR-spectroscopy. *J Am Chem Soc* 113(12):4371–4380
- Perrin F (1934) Mouvement Brownien d'un ellipsoïde (I). Dispersion diélectrique pour des molécules ellipsoïdales. *J Phys Radium* 5:497–511
- Perrin F (1936) Mouvement Brownien d'un ellipsoïde (II). Rotation libre et dépolarisation des fluorescences. Translation et diffusion de molécules ellipsoïdales. *J Phys Radium* 7:1–11
- Polak E, Ribière G (1969) Note sur la convergence de méthodes de directions conjuguées. *Revue Française d'Informatique et de Recherche Opérationnelle* 16:35–43
- Schurr JM, Babcock HP, Fujimoto BS (1994) A test of the model-free formulas. Effects of anisotropic rotational diffusion and dimerization. *J Magn Reson B* 105(3):211–224
- Shanno DF (1970) Conditioning of quasi-Newton methods for function minimization. *Math Comp* 24(111):647–656
- Steihaug T (1983) The conjugate gradient method and trust regions in large scale optimization. *SIAM J Numer Anal* 20(3):626–637
- Tsan P., Hus JC, Caffrey M, Marion D, Blackledge M (2000) Rotational diffusion anisotropy and local backbone dynamics of carbon monoxide-bound *Rhodobacter capsulatus* cytochrome c2. *J Am Chem Soc* 122(23):5603–5612
- Woessner DE (1962) Nuclear spin relaxation in ellipsoids undergoing rotational brownian motion. *J Chem Phys* 37(3):647–654
- Zhuravleva AV, Korzhnev DM, Kupce E, Arseniev AS, Billeter M, Orekhov VY (2004) Gated electron transfers and electron pathways in azurin: a NMR dynamic study at multiple fields and temperatures. *J Mol Biol* 342(5):1599–1611