Computational aspects of exponential data fitting in Magnetic Resonance Spectroscopy

Diana M. Sima[1], J.B. Poullet[1], S Van Huffel[1]
[1] Department of Electrical Engineering ESAT-SCD,
Katholieke Universiteit Leuven,
Kasteelpark Arenberg 10, 3001 Leuven-Heverlee, Belgium,
diana.sima@esat.kuleuven.be

Abstract. Magnetic Resonance Spectroscopy (MRS) is one of the practical biomedical applications where exponential data fitting is an essential tool. An MRS signal is a complex-valued time-domain signal that satisfies in theory a model expressed as a sum of complex damped exponentials. The model parameters provide useful information about metabolites, *i.e.*, about the chemical content of the sample or tissue under study. *In vivo* MRS signals are characterized by the presence of noise and by deviations from the theoretical model. Along the years many fitting methods have been employed for the so-called "metabolite quantification" problem. These are, to mention just a few, subspace-based methods such as HSVD, or optimization-based methods such as VARPRO. In recent years more focus has been placed upon methods that incorporate prior knowledge into the model. Initially, prior knowledge has been included in the form of simple relationships between parameters of the exponential model (for instance, linear relations between the parameters of individual exponential components, when it is known that these components originate from the same metabolite). More advanced methods are able to include much more prior knowledge by combining whole metabolite profiles into the model; these metabolite profiles are quantum-mechanically simulated or *in vitro* measured MRS signals. Other interesting computational issues related to MRS quantification include preprocessing techniques (such as filtering out some unwanted spectral components) aimed at modifying the measured data without loss of essential information, in order to match the assumed exponential model. This chapter reviews time-domain MRS quantification methods and related preprocessing techniques, it emphasizes the computational aspects such as the use of variable projection in the optimization-based methods, and it hints towards some of the features of existing software packages for MRS.

Keywords: Magnetic Resonance Spectroscopy (MRS), metabolite quantification, Hankel Singular Value Decomposition, nonlinear least squares, variable projection

1. Introduction

1.1. Basics of the NMR phenomenon. [1]

[1]The material in this section is more extensively discussed in Chapter 1 of [**10**]. Interesting resources on the basics of NMR can also be found online; see, *e.g.*, the Wikipedia pages, or the "Virtual Textbook of Organic Chemistry".

Nuclear Magnetic Resonance (NMR) is a quantum mechanical phenomenon observed when a population of atomic nuclei is placed into an external magnetic field B_0. Nuclei such as ^1H—hydrogen (also referred to simply as proton), ^{13}C—carbon, or ^{31}P—phosphorus possess an inherent angular momentum called *spin*, capable of two equivalent states. When the nucleus is immersed in an external magnetic field, the two spin states are no longer equivalent and can be visualized as states of precession of the spin around the magnetic field direction (low energy) and around the direction opposite the field (high energy). The spin precesses about the external field axis with an angular frequency ω_0 (called Larmor frequency) equal to the product between the gyromagnetic ratio γ characteristic for the nucleus (*e.g.*, γ=42.58, 10.71, and 17.25 MHz/T, for ^{1}H, ^{13}C, and ^{31}P, respectively) and the magnetic field strength B_0:

$$\omega_0 = \gamma B_0. \tag{1}$$

Due to its two states, the nucleus can be seen as a little magnet that tends to align with the magnetic field in order to reach the state of minimum magnetic potential energy. The analogy with classical physics is limited by the fact that in quantum mechanics the energy difference between high and low energy states is quantized; this is given by

$$\Delta E = \gamma \hbar B_0 = \hbar \omega_0,\ \hbar = h/(2\pi), \tag{2}$$

where h is Planck's constant (6.626×10^{-34} J s), which means that only a photon whose frequency matches exactly the Larmor frequency ω_0 is able to be absorbed/emitted by the nucleus in order for it to jump between the two energy states.

When a population of the same nuclei is considered to be immersed into a static magnetic field B_0, its equilibrium state gives a total magnetization vector (as the sum of directions of all individual spins) that points in the direction of the magnetic field (see Figure 1). In thermal equilibrium, the ratio between the number $n-$ and $n+$ of nuclei in the high and low energy states is given by

$$n - /n+ = e^{-\Delta E/kT}, \tag{3}$$

where k is the Boltzman constant (1.381×10^{-23} J/K), T is the temperature in Kelvin and ΔE is the energy difference between the states as given in (2). The ratio $n - /n+$ is very close to 1 in normal circumstances, resulting in a tiny excess of nuclei in the low state (this excess is about four in a million in clinical NMR scanners with $B_0 = 1.5$ Tesla), but this is enough to allow sufficient signal strength.

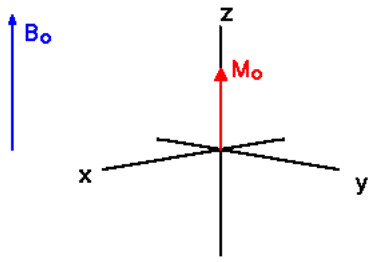


FIGURE 1. At equilibrium, the net magnetization M_0 of a large population of nuclei is aligned with the external field B_0. The explanation is that an excess of nuclei is found in the low energy state, thus these spins sum up to a nonzero vector with positive magnitude in the z-direction, $M_z = M_0$, but with zero magnitude in the x or y directions, $M_{xy} = 0$, for which the spins show no preference; *i.e.*, there is no phase coherence in the precession of the spins in the xy-plane.

The equilibrium state is non-informative and has to be perturbed in order to obtain a measurable effect. To this end, a second oscillating magnetic field B_1, created by a radio-frequency (RF) electric field in a coil of wire and oscillating at the Larmor frequency of interest in a plane perpendicular to the direction of B_0, is used to perturb the static magnetic field. The RF photons are absorbed by some of the nuclei to flip their spin from aligned to the magnetic field to anti-aligned, reducing thus the magnitude of the magnetization vector in the direction of B_0. (M_z becomes smaller than M_0, and it may reach zero magnitude in the so-called 90° RF pulses, or even $-M_0$ in the so-called 180° RF pulses). At the same time, the spins also start to precess about the second magnetic field B_1, yielding a tilt of the magnetization vector M to the plane perpendicular onto B_0, making the M_{xy} component non-zero. After the RF pulse is turned off, the spins tend to return to their low energy state, with the effect that the magnetization vector M precesses about B_0 in a spiral until it reaches the equilibrium state. This "relaxation" of the magnetization induces a current in a detector coil, which in turn is captured as an NMR signal (a free induction decay (FID) signal), see Figure 2. The time dependence of the return to equilibrium is governed by *Bloch's equations* [**5**], which are first-order differential equations describing the motion of the magnetization vector M as a function of time.

The time T_1, called the spin-lattice relaxation time, is the time to reduce the difference between the longitudinal magnetization after the

RF pulse and its equilibrium value (see Figure 1) by a factor of e,

$$M_z(t) = M_0(1 - e^{\frac{-t}{T_1}}) \tag{4}$$

where $M_z(t)$ is the longitudinal magnetization at time t.

The time constant which describes the return to equilibrium of the transverse magnetization, M_{xy}, is called the spin-spin relaxation time, T_2, and is defined such that

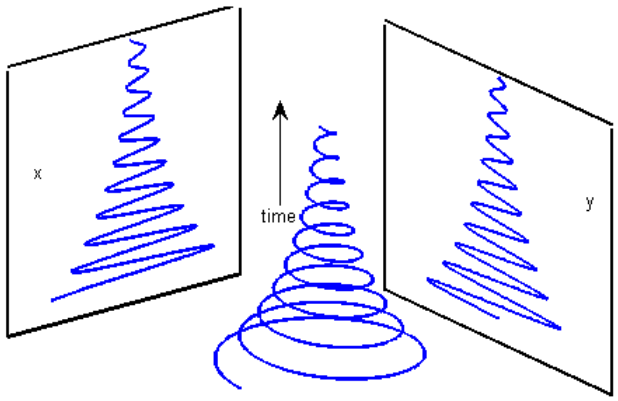
$$M_{xy}(t) = M_0 e^{\frac{-t}{T_2}}. \tag{5}$$

T_1 and T_2 depend on the chemical compound, but T_2 is always less than or equal to T_1. Note that the coils which acquire the signal are located along the x- and y-axes. This explains why T_2 is more important than T_1 in MRS. Due to Bloch's laws [**5**], the induced signal has the shape of a complex exponentially decaying sinusoid in the time domain. The amount of energy emitted by the nuclei, and thus the signal intensity (*i.e.*, its amplitude), is proportional to the number of nuclei that contribute to it, thus to the population difference between the two energy states at every given moment. Equations (2)-(3) show that a higher intensity is obtained for higher B_0 and lower temperature. For *in vivo* experiments, the temperature is fixed at body temperature and one can only vary B_0 within a secure range to prevent tissue heating.

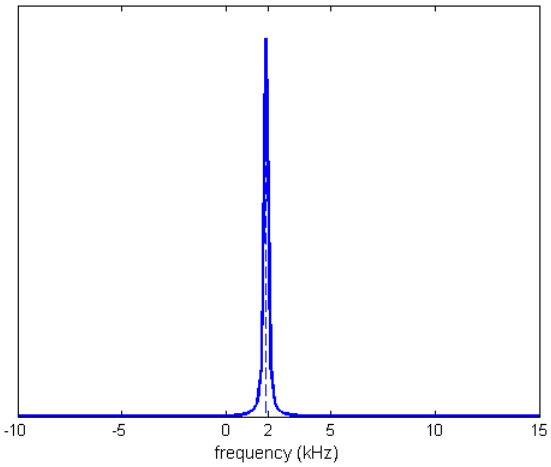
Note. In order to increase the inherently low signal-to-noise ratio (SNR), several data acquisitions are performed consecutively and the final signal is the average of all measured signals.

Note. In an idealized nuclear magnetic resonance experiment, the FID decays approximately exponentially with a time constant T_2, but, in practice, small differences in the static magnetic field at different spatial locations, also called inhomogeneities, cause the Larmor frequency to vary across the body creating destructive interference which shortens the FID. The time constant for the observed decay of the FID is called T_2* ("T 2 star") relaxation time, and is always shorter than T_2.

1.2. Magnetic Resonance Spectroscopy. From the idealized description in the previous section it might seem that every type of nucleus gives rise to a single exponentially decaying sinusoid at its Larmor frequency. However this is not the case due to the influence of the electrons that are present in the atoms, which perturb the static magnetic field B_0 with their own magnetic environment. This is called a shielding effect and it is essential in separating the chemically and magnetically nonequivalent nuclei of the same isotope. For example, a hydrogen bonded to an oxygen will be shielded differently from a



a



b

FIGURE 2. **a.** Relaxation of the net magnetization M when the field B_1 is turned off can be measured as an exponentially decaying complex valued signal in the time domain. **b.** Real part of the Fourier transform of the exponential decay in **a.** exhibits a single Lorentzian peak. The area under this peak is equal to the amplitude of the time-domain signal, thus it is proportional to the concentration of the spin population.

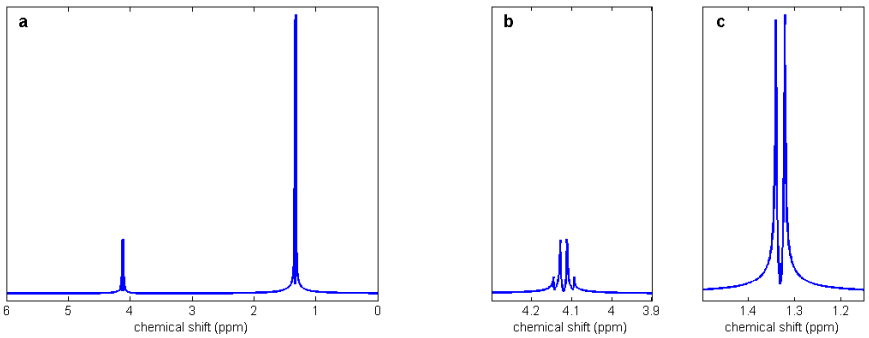


FIGURE 3. **a**. ^{1}H NMR spectrum of lactate, $CH_3CH(OH)COO$ simulated for a 400MHz spectrometer frequency. The methine proton (*i.e.*, the H in CH) resonates as a quartet at 4.11 ppm (**b**), while the three methyl protons (*i.e.*, the H in CH_3) resonate as a doublet at 1.33 ppm (**c**).

hydrogen bonded to a carbon atom. The electrons are responsible for a small magnetic field at the nucleus, which opposes the externally applied field. The amount of shielding experienced by a nucleus depends on the exact local environment, *i.e.*, on the bonding pattern between atoms in a molecule, and determines the effective magnetic field,

$$B_{\text{eff}} = B_0(1 - \sigma') \tag{6}$$

where σ' is the shielding constant, a dimensionless unit that depends on the electrical environment of a nucleus. The effective field at each nucleus will vary. This is called the *chemical shift* phenomenon and it is essential in Magnetic Resonance Spectroscopy because it allows to spectrally distinguish between different chemical compounds.

In addition to the effect of electrons, nuclei which are close to each other exert an influence on each other's effective magnetic field. This effect shows up in the NMR signal when the nuclei are nonequivalent, *i.e.*, have different chemical shifts. If the distance between nonequivalent nuclei is less than or equal to three bond lengths, the effect called *spin-spin coupling* or *J-coupling* is observable. When a nucleus is coupled to n equivalent nuclei, its signal splits into an $n+1$ multiplet with amplitude ratios following Pascal's triangle: Singlet 1; Doublet 1:1; Triplet 1:2:1; Quartet 1:3:3:1; Quintet 1:4:6:4:1; etc. Coupling to additional spins will lead to further splittings of each component of the multiplet, with the possibility to obtain, e.g., a doublet of doublets. Besides the chemical shift, J-coupling is contributing to the specific pattern of the

NMR spectrum of a chemical compound, providing thus an important feature for identifying it. This spectral *prior knowledge* (*i.e.*, the shape of an MRS signal given the chemical formula of a compound) can be used to quantum mechanically simulate *metabolite signals*, which can further be used to improve the analysis of *in vivo* measured signals, as illustrated further in Section 3.

Note. MRS signals are measured in the time domain as complex-valued decaying signals. For visualization and analysis purposes they are usually Fourier transformed to the frequency domain, yielding MR spectra.

Note. In MRS, the frequency axis is given in part per million (ppm) in order to be independent from the spectrometer frequency (that is, the RF field frequency) in the sense that, if a nucleus resonates at a certain ppm-value for a spectrometer frequency of 63.8 MHz, it will resonate at the same ppm-value for a spectrometer frequency of 500 MHz. The ppm-value of a resonance is defined by

$$\text{ppm}_M = \frac{f_{\text{Hz}} 10^6}{f_S} + \text{ppm}_{\text{Ref}} \tag{7}$$

where f_{Hz} is the resonance frequency in Hz, f_S the spectrometer frequency in Hz and ppm_{Ref} the ppm-value of a certain reference resonance.

1.3. Biomedical applications of ^{1}H NMR. The hydrogen proton (^{1}H) is the most sensitive nucleus in what concerns NMR measurements, and moreover it is abundant in *in vivo* samples. In particular, water is the largest component of the human body. For this reason the water resonance is about 10^5 times more intense compared to other resonances. One of the most widely used ^{1}H NMR techniques is Magnetic Resonance Imaging (MRI), where signals from hydrogen protons from the water molecules are used to form anatomic images that emphasize differences in tissue types. Complementary to MRI, the technique of Magnetic Resonance Spectroscopy has a much higher *spectral resolution* at the expense of lower spatial resolution compared to MRI. This information is used to determine the concentration of brain *metabolites* (*i.e.*, specific chemical compounds) such as N-acetyl aspartate (NAA), choline (Cho), creatine (Cr) and lactate (Lac) in a tissue. Specific patterns of metabolites can be helpful in the diagnosis of certain central nervous system disorders [**6**] such as Alzheimers disease or epilepsy [**18, 17**], and in the differentiation of high grade from low grade brain tumors [**12**]. *Single-voxel* MRS measures one signal from a voxel (spatial volume element), while *multi-voxel* spectroscopy

or MRSI (*MRS imaging*) is designed to combine spatial techniques from MRI with MRS, yielding a compromise between spatial and spectral resolution.

1.4. Metabolite quantification. The main computational problem associated with the analysis of MRS signals is called *metabolite quantification (or quantitation)*. This problem aims at determining the concentration of each metabolite in the voxel under investigation, by fitting an appropriate model to the measured signal and computing concentration values from the retrieved parameters. Recall that the concentration of the nucleus under investigation can be computed from the amplitude of the corresponding spectral components.

The rest of this Chapter focuses on computational techniques that perform metabolite quantification in the time-domain (the data acquisition domain), although there are some popular methods that solve this problem in the frequency domain (after applying Fourier transform to the original signal). The frequency domain methods range from very simple *peak integration*, where metabolite concentrations are estimated from the area under the peaks obtained via numerical integration of the real or absolute value of the spectrum (see Figure 2**b.**), or complex model-based optimization algorithms, such as the regularized nonlinear least squares method implemented in the widely-used commercial package LCModel [**34, 35**].

Metabolite quantification methods assume a certain model for the signals. Most certainly, *in vivo* signals do not satisfy all model assumptions, and for this reason various *preprocessing methods* have been designed with the goal of modifying the measured signals such that they better approach the desired model [**31**]. (As a simple example, phase correction is necessary if peak integration is used on the real part of a spectrum, such that all peaks point upwards.)

Although this Chapter focuses on quantification (also referred to simply as *processing*), it also illustrates a number of time-domain preprocessing techniques and explains various artifacts appearing in the signals.

2. The classical exponential model for MRS signals

Traditionally, the function used to model an MRS signal is the sum of exponentially damped complex sinusoids:

$$y(t) = \widehat{y}(t) + \varepsilon_t := \sum_{k=1}^{K} \alpha_k \, \zeta_k^t + \varepsilon_t, \quad t = t_0, \ldots, t_{m-1}, \; t_\ell = t_0 + \ell \, \Delta t, \tag{8}$$

where K is the number of exponential components, α_k is the complex amplitude of the k^{th} component, which can also be written as

$$\alpha_k = a_k \exp(j\phi_k), \tag{9}$$

with a_k the positive real amplitude, ϕ_k the phase, and ζ_k is the signal pole of the k^{th} component, which can be written as

$$\zeta_k = \exp(-d_k + j2\pi f_k), \tag{10}$$

with $d_k > 0$ the damping factor, f_k the frequency, and $j = \sqrt{-1}$.

The model $\widehat{y}$ can be seen as the discretized solution of a linear time-invariant (LTI) system. The theory of LTI system identification opens a wide range of possible methods for going from data to a model and then to a particular parametrization. In the MRS community, state-space and linear prediction methods have become popular as black-box identification methods [**2**]. However, optimization methods based on the nonlinear least squares fit of the measured y with the model $\widehat{y}$ are even more appreciated due to the easy incorporation of extra parameter constraints coming from biochemical prior knowledge and to the possibility of modifying the model function to account for various distortions.

2.1. Subspace-based methods. HSVD (Hankel singular value decomposition) [**27**] is among the most well-known subspace-based methods applied to the exponential data fitting problem in the MRS application. It is very popular due to its full automation (no essential user interaction is required), but it has also some potential drawbacks. In particular, deviations of the data from the assumed model may lead to bad parameter estimates. Moreover, the original HSVD gives full freedom to the model parameters, while MRS quantification might benefit from imposing some prior knowledge. As further described in Section 4, HSVD is most successfully used as a black-box modeling technique in the preprocessing step of *removal of the water component*.

HSVD requires computing the truncated singular value decomposition (SVD) of a Hankel matrix H, containing the noisy signal in its first column and row. If the signal is a sum of K complex damped exponentials (giving K Lorentzian peaks, when the signal is Fourier transformed to the frequency domain), then the data matrix H has rank exactly equal to K. Noise in the signal or a signal not satisfying the exponentially decaying model yields a full rank matrix H. By truncating the SVD to lower rank, one can approximately model the signal as a sum of complex damped exponentials. A recent review of the basic derivation of HSVD and numerical implementation issues (such as

the application of the Lanczos method in the computation of a truncated SVD, with a partial reorthogonalization procedure for accuracy improvement [**23**]) can be found in [**22**]. See also the discussion in the first chapter of this book. Moreover, HSVD can be improved by considering a total least squares solution (yielding the HTLS method [**49**]) instead of a least squares solution in approximately solving an incompatible linear system. Special developments allowing to impose prior knowledge include:

- frequencies and damping factors of some exponentials are known [**8, 9**];
- frequencies and phases of some exponentials are known [**7**];
- frequencies, dampings and phases of some exponentials are known [**7**];
- phases of some exponentials are known [**7**];
- equal phases, equal dampings, linearly related amplitudes are imposed, and frequency shifts are known in signals with multiplet structure [**24**].

These methods can be more computationally demanding compared to the original HSVD. Furthermore, they require user interaction for choosing the correct prior knowledge or for selecting the type and the frequency region of a multiplet structure.

2.2. Optimization-based methods: VARPRO and AMARES. These methods explicitly minimize the function $\|\mathbf{y} - \widehat{\mathbf{y}}\|^2$ (where $\mathbf{y}$, $\widehat{\mathbf{y}}$ stand for vectors obtained from the signals $y(t)$, $\widehat{y}(t)$ at time instants $t = t_0, \ldots, t_{m-1}$) over the model parameters a_k, d_k, f_k, ϕ_k. The variable projection (VARPRO) [**14**] method has been used since the 1980's [**48**] in the context of MRS quantification. The reason of this success story is that the sum of complex damped exponential model suits the separability condition of VARPRO, since its variables can be split into linear—the complex amplitudes α_k, and nonlinear—the complex signal poles ζ_k. Let Φ_k denote the column vector $[(\zeta_k)^{t_0}, \ldots, (\zeta_k)^{t_{m-1}}]^T$ and $\Phi(\zeta)$ the matrix $[\Phi_1, \ldots, \Phi_K]$. With this notation, the model (8) can be written as $\widehat{\mathbf{y}} = \Phi(\boldsymbol{\zeta})\boldsymbol{\alpha}$. Then for each fixed instance of the nonlinear parameter vector $\boldsymbol{\zeta} = [\zeta_1, \ldots, \zeta_K]$, the linear variables $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_K]$ can be computed in the least squares sense as $\boldsymbol{\alpha}(\boldsymbol{\zeta}) = \Phi(\boldsymbol{\zeta})^\dagger \mathbf{y}$, and the nonlinear least squares objective becomes $\|\mathbf{y} - \Phi(\boldsymbol{\zeta})\Phi(\boldsymbol{\zeta})^\dagger \mathbf{y}\|$. The advantage of this formulation is that no starting values are needed for the linear parameters, and that the number of parameters is halved. Moreover, it is well known and proven by theory [**39**] and practice that variable projection always converges with less iterations than the full functional approach. This includes convergence in cases when the full

functional approach diverges. A historical note on the application of VARPRO to MRS data quantification can be read in Section 17 of the review paper [**15**].

In [**48, 47**] imposing prior knowledge about the model parameters is also discussed. The main advantage that can be used is eliminating redundant parameters from a multiplet structure. For example, consider a simple doublet (*e.g.*, the Lactate doublet in Figure 3**c.**). For this doublet we know that

$$a_1 = a_2, \quad \phi_1 = \phi_2, \quad d_1 = d_2, \quad f_1 = f_2 + \delta f, \quad \delta f \text{ known.}$$

Thus the number of variables can be effectively halved and the VARPRO method can still be used. The admissible prior knowledge (linear relations between amplitudes, equal phases, equal dampings, constant frequency shifts) should be in such a way that a common factor can be extracted from a group of multiplet peaks.

Extensive comparisons between minimization of the full and the variable projection functional based upon the software implementations for nonlinear optimization available at the time have been presented for this MRS model in [**52, 50**] (see also the comments related to this comparison in Section 17 of the review paper [**15**]). In general, the two approaches are equally efficient, but for some of the optimization algorithms the implementation used for VARPRO was less robust, *i.e.*, it had a slightly larger number of convergence failures. The tested implementations differ, however, in some other respects that could have in fact led to this conclusion. For instance, ensuring positivity of the damping factors d_k was done in the VARPRO implementation by using the square of d_k as optimization variable, while the inequality bound $d_k > 0$ was explicitly imposed in the minimization of the full functional. Moreover, the tested VARPRO implementation did not include Kaufman's simplification for Jacobian evaluation [**20**]. However, the main reason for recommending the minimization of the full functional as implemented in AMARES (advanced method for accurate, robust and efficient spectral fitting) [**52**] is the possibility of allowing more flexibility to the prior knowledge (*e.g.*, equal dampings between 2 peaks, but different amplitudes and phases). The implementation of AMARES into jMRUI [**26**] allows for an interactive peak-picking procedure where the user identifies the peaks of the spectrum by mouse-clicking, and starting values for the nonlinear model parameters d_k and f_k are thus computed. The linear parameters are initialized to the linear least squares solution when these nonlinear parameter values are fixed, thus improving considerably the efficiency and robustness of the method.

Note. Optimization-based methods allow for more variation of the model function, by considering other lineshapes apart from the Lorentzian function. Gaussian or Voigt lineshapes (the latter being a combination of Lorentzian and Gaussian functions) are often used. In this case

$$\widehat{y}(t) := \sum_{k=1}^{K} \alpha_k \, (\zeta_k)^t (\eta_k)^{t^2}, \quad t = t_0, \ldots, t_{m-1}, \tag{11}$$

where

(12)

$$\alpha_k = a_k \exp(j\phi_k),$$

$$\zeta_k = \begin{cases} \exp(-d_k + j2\pi f_k), & \text{for Lorentzian and Voigt lineshapes} \\ \exp(j2\pi f_k), & \text{for Gaussian lineshapes} \end{cases}$$

$$\eta_k = \begin{cases} \exp(j2\pi e_k), & \text{for Lorentzian lineshapes} \\ \exp(-g_k + j2\pi e_k), & \text{for Gaussian and Voigt lineshapes} \end{cases}$$

where a_k are the real amplitudes, ϕ_k are the phases, d_k are Lorentzian dampings, f_k are frequencies, g_k are Gaussian dampings, and e_k are eddy current correction terms.

3. The model for short echo time MRS signals

The model and methods described in the previous section are appropriate for fitting the so-called long echo-time MRS signals. Nowadays, the nuclear magnetic resonance scanners record short echo-time signals that are richer in information because they display the responses of significantly more metabolites. As we saw in Section 1, due to chemical shift and J-coupling, the response of each metabolite can be spread out over the whole spectrum. The methods described above might break down if peaks that are in the same frequency region but are unwanted (nuisance peaks) have large amplitudes or are close, in frequency, to the peaks of interest [**38, 51**]. Although methods such as AMARES have been applied quite successfully to short echo time MR spectra [**19**], the nuisance peaks and the more intensive user interaction tend to encourage methods based on the use of *metabolite profiles*, where more prior knowledge can be implicitly included in the model, especially information related to experimental conditions of acquisition.

During an *in vivo* NMR experiment, the measured time-domain signal consists of responses from all metabolites (including macromolecules), noise and (partially suppressed or unsuppressed) water. Spectra of metabolites that are visible during *in vivo* spectroscopy can be measured *in vitro* or they can also be simulated using quantum mechanical knowledge [**16, 45**], and these signals can be grouped in a basis set of

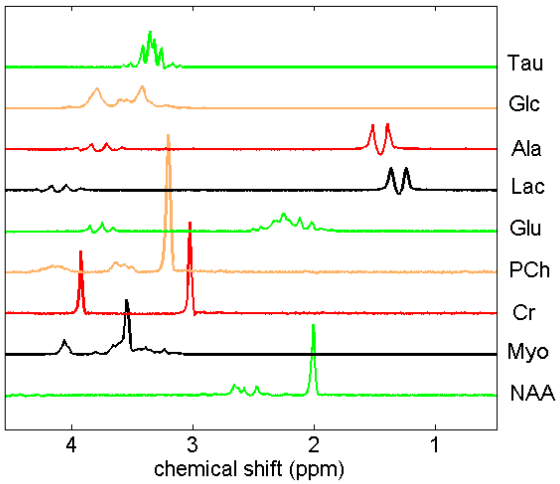


FIGURE 4. Examples of *in vitro* metabolite spectra (NAA, Myo, Cr, Pch, Glu, Lac, Ala, Glc, Tau) acquired on a 1.5 T Philips NT Gyroscan using a PRESS sequence. They can be used for fitting *in vivo* signals acquired under the same conditions.

metabolite profiles. As discussed in the Introduction, an *in vivo* time-domain signal has, in theory, the shape of a sum of complex damped exponentials. Instead of modeling these individual components, the *in vivo* short echo time MRS signal can be modeled using the profiles in the basis set, such that the prior knowledge that relates individual peaks in the *in vivo* spectrum is implicitly imposed. The quantities of interest, the metabolite *concentrations*, can be estimated from the weighting coefficients (amplitudes) of the linearly combined *in vitro* profiles. The linear combination should allow small corrections in spectral parameters like frequency shifts, damping corrections and phase shifts, since these parameters may vary from measurement to measurement [**37**], *e.g.*, from *in vitro*, to *in vivo*.

We consider that we are given a "metabolite database" as a set $\{v_k, \text{ for } k = 1, \ldots, K\}$ of complex-valued time series of length N, representing *in vitro* measured NMR responses. Examples of *in vitro* metabolite spectra are given in Figure 4.

When lineshape distortions are ignored, we assume that an *in vivo* measured MRS signal y satisfies the model

$$y(t) = \widehat{y}(t) + \varepsilon_t, \quad t = t_0, \ldots, t_{m-1}, \qquad \text{with} \tag{13}$$

$$\widehat{y}(t) = \sum_{k=1}^{K} \alpha_k \, (\zeta_k)^t \, v_k(t) + b(t) + w(t), \tag{14}$$

where $\alpha_k, \zeta_k \in \mathbb{C}$ are unknown parameters that account for amplitudes of the metabolites in the database and for the necessary corrections of the database signals. The complex amplitudes α_k and the complex values ζ_k can be expanded as in (9-10), or Gaussian/Voigt correction terms can be considered as in (12). In (13), the term ε_t denotes an unknown noise perturbation with zero mean, and in (14), $b(t)$ denotes the "macromolecular baseline", and $w(t)$ denotes the water component.

The identification of complex amplitudes α_k and complex correction factors ζ_k, for $k = 1, \ldots, K$, can be accomplished by minimizing the least squares criterion: $\sum_{t=t_0,\ldots,t_{m-1}} |y(t) - \widehat{y}(t)|^2$. We ignore for the moment the nuisance signals of water $w(t)$ and macromolecular baseline $b(t)$, which will be dealt with in Section 4. This nonlinear least squares problem formulation constitutes the kernel of the algorithms and software implementations such as QUEST [**37**] (included in jMRUI) and AQSES [**30**] (included in the Java interface AqsesGUI [**11**] and in the Matlab package SPID [**28**]). This problem is a separable problem, where linear parameters α_k can be projected out of the least squares problem in closed-form expression. An iterative minimization algorithm of the Levenberg-Marquardt type [**25**] can be used for the numerical minimization of the variable projection functional. The solution of this local optimization method depends on the starting values. In AQSES, for instance, we set initial values to zero, which means that we start the optimization with no spectral corrections to the signals in the database, which is appropriate for reasonably frequency-aligned signals. Moreover, we use linear bounds on each nonlinear parameter (soft constraints), since, for instance, frequency shifts should be restricted only to small variations in order to remain relevant.

As explained in [**41**], a variable projection approach [**14**] with approximate Jacobian evaluation is used in AQSES, modified in order to include simple box constraints on the linear and nonlinear variables. According to the authors' experience, this is more robust than optimizing the nonlinear least squares problem directly over all linear and nonlinear parameters without separation, and more computationally efficient, especially in the case when a baseline is also modeled (see

Section 4 and the Appendix). More importantly, in the variable projection method nearly zero amplitudes a_k do not yield a rank-deficient Jacobian, as it is the case with the full functional. The trust region or Levenberg-Marquardt methods used with either of the functionals are naturally guarded against rank-deficient Jacobians, but it seems preferable to avoid reaching this situation altogether.

Note. Handling (nearly) zero amplitudes without damaging the method's convergence is very important in the fitting of the metabolite-based model. The reason is that we can include more metabolite profiles in the basis set than what is actually present in the *in vivo* signal. Recall that some metabolites are present or not depending of the tissue type, *e.g.*, normal versus tumor. Obtaining nearly zero amplitudes plays the role of "de-selecting" the metabolites that are not contained in the signal.

4. Model distortions and preprocessing methods

The aim of the preprocessing steps is to remove irrelevant information, while enhancing the key features (*i.e.*, the metabolic components of interest in the MRS signal) in order to extract parameters (*e.g.*, concentrations) more accurately. MRS signals are affected by the presence of artifacts, instrumental errors, noise and other unwanted components. The spectral quality can be considerably improved by appropriate manipulation of the data. Typical preprocessing steps are enumerated and illustrated below, in their logical sequential order.

Frequency alignment. Frequency alignment is needed when a batch of signals need to be automatically processed, as is the case for MRSI data (*i.e.*, signals in a 2D or 3D grid of voxels). The reason why signals might be spectrally misaligned is the fact that the resonating frequencies are proportional to the static magnetic field B_0, and any spatial inhomogeneities of B_0 are directly translated in frequency misalignment between signals from different regions. This step is performed before other preprocessing steps since most of the subsequent methods require correctly aligned spectra. For the task of aligning two signals, a frequency shift Δf has to be estimated, and the signals have to be shifted by point-wise multiplication in the time-domain with $\exp(2\pi\Delta f t)$. In *in vivo* MR data at 1.5T, large peaks such as NAA, Cho or Cr are generally used as reference for alignment. Alternatively, the spectra can be manually shifted in the frequency domain (for instance, jMRUI offers a drag-and-drop utility for this task) and transformed back into the time domain by the inverse Fourier transformation. For *in vitro* signals, a substance such as TSPS (3-trimethylsilyl-1-propane-sulfonic

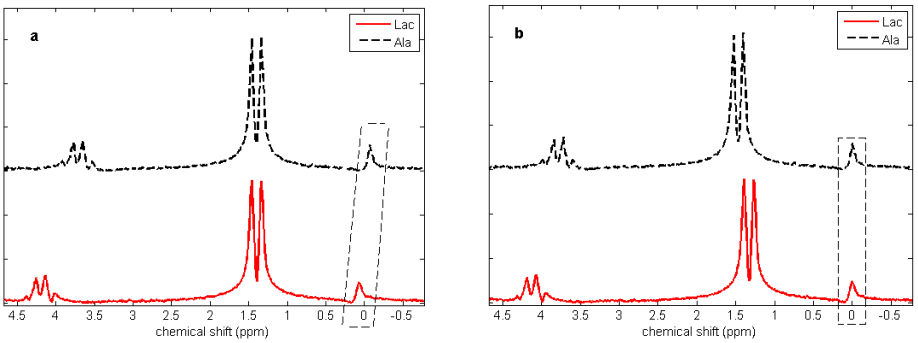


FIGURE 5. Frequency alignment is essential for these two similarly looking metabolite spectra measured *in vitro*, Alanine and Lactate. **a**) The two signals are not aligned according to the reference peak of TSPS at 0 ppm, thus the doublets at around 1.4 ppm are indistinguishable from each other. **b**) The two signals are aligned according to the reference peak at 0 ppm, making the doublets and quartets of Alanine and Lactate more easily distinguishable.

acid) can be added to each sample and later be used as reference for frequency alignment (see Figure 5).

Phase and lineshape correction. The zero-order phase is the overall phase of the signal, *i.e.*, the common phase of all spectral components at time 0. The signal can be trivially corrected if an estimate of the phase at time 0 is available. First-order or higher-order phase distortions, which are due to hardware imperfections or timing errors, may depend upon the frequency, thus their correction is more challenging. Time-domain quantification methods account for higher-order phase distortions by incorporating different phases ϕ_k for each spectral component into the model.

Lineshape distortions in the MRS signals are usually caused by eddy currents and magnetic field inhomogeneities. Eddy currents are in turn appearing when a conductor encounters a varying magnetic field, which is the case for the coils of the NMR scanner. Good hardware techniques are nowadays available for compensating against eddy current effects [**46**], but signal processing can also provide an effective solution. For instance, in the case a water-unsuppressed signal is available together with the water-suppressed signal, Klose's method [**21**] can be carried out for phase and eddy current correction. This corrects for eddy currents by dividing the water-suppressed signal to the phase term of the

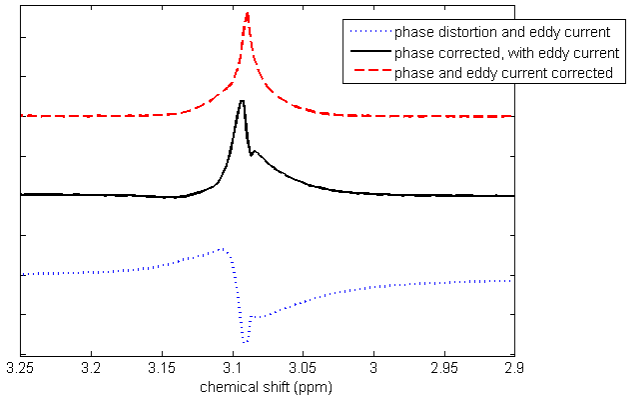
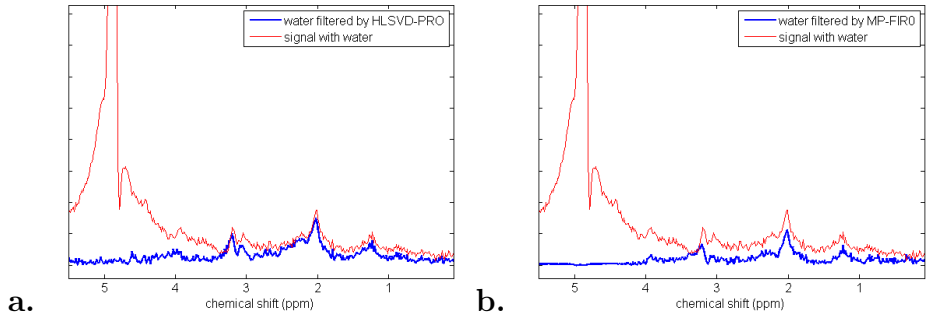


FIGURE 6. Phase and eddy current correction (illustrated here on a singlet) can be accomplished in the scanner or by Klose's method.

water-unsuppressed signal. (See Figure 6.) Lineshape correction methods like QUALITY [**10**] and QUECC [**4**] can also correct for distortions in the magnitude of the signal, while other phase artifacts are dealt with by phase filtering in [**53**]. The goal of these methods is to transform the original signal such that it follows a Lorentzian model (*i.e.*, sum of complex damped exponentials), which is often used in quantification methods.

Water filtering. As mentioned before, the water component in ^{1}H MRS signals is several orders of magnitude higher than the other metabolite contributions. For this reason many efforts have been put into hardware-based suppression techniques of the water signal through the design of special pulse sequences [**33, 43, 44**]. It is nowadays common practice to acquire so-called "water-suppressed" signals besides/instead of the "water-unsuppressed" signals that are normally obtained by MRS. However, these signals still contain water resonances higher than the metabolite signals. Signal processing methods must and can in fact deal with these two types of water signals, in order to extract only the metabolite information. Frequency selective filtering is the key tool, since the metabolites of interest and the water components are resonating at sufficiently distant frequencies.

Residual water removal can be performed by a black-box modeling approach such as HSVD [**3**], based on the idea of first fitting the whole signal as a sum of complex damped exponentials and then subtracting all the components whose frequencies were estimated in the frequency



a. **b.**

FIGURE 7. Filtering water components at around 4.7 ppm from a realistically simulated noisy MRS signal. **a.** HLSVD-PRO [**23**], **b.** MP-FIR0 [**29**].

region of the water resonance. In Figure 7**a.** we illustrate this procedure using the HLSVD-PRO implementation [**23**]. Wavelets and the Gabor transform [**1**] have also been successful for this task. Finally, pass-band filtering can also be performed by FIR filtering, after appropriate filter design. Here we cite the MP-FIR (maximum-phase finite impulse response filter) [**51**] and its recent improvement MP-FIR0 [**29**]. These methods are based on finding the coefficients of an FIR filter that will be convolved with the MRS signal. The filters are optimized to suppress the resonances that are outside the frequency region of interest and to keep the region of interest (where the metabolites are visible) with minimal distortions [**28**]. (See Figure 7**b.**)

Note. AQSES [**30**] does not use MP-FIR in a preprocessing step, but includes it *during* the nonlinear least squares optimization. Therefore, the quantification algorithm fits a filtered model based on metabolite profiles to a filtered *in vivo* signal. This approach leaves model parameters (in particular, the amplitudes) intact, since the FIR filtering is a linear operation, *i.e.*, it can be expressed as a matrix-vector product. Results in [**32**] indicate that this approach is superior in terms of accuracy of estimated amplitudes compared to quantifying preprocessed signals by water filtering with HLSVD-PRO.

Baseline correction. Baselines in the NMR spectra are caused by distortion of the first few data points in the time domain signal. Two different phenomena may contribute to this distortion: 1) the amplitudes of the initial data points are distorted due to instrumental imperfections (baseline distortion), 2) actual resonances of weakly mobile *macromolecules* are present in the signal (macromolecular background). Macromolecular signals, coming from the macromolecules present in the tissue under investigation, are characterized by broad

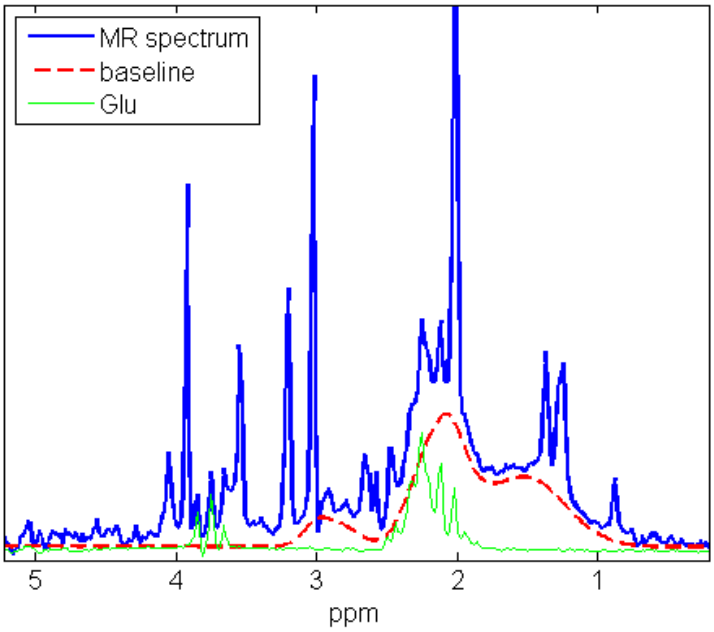


FIGURE 8. Real part of a realistically simulated MRS signal (thick blue line) with an added macromolecular baseline (interrupted red line). The metabolite Glutamate (Glu), highly overlapping with the simulated baseline, is also shown as a thin green line.

spectral lines (short T_2), which often overlap in the frequency domain with metabolite components. Short echo time ^{1}H-MR spectra of human brain contain a significant macromolecular background that complicates drastically the quantification process.

We illustrate in Figure 8 a realistically simulated signal (using 8 metabolites from the basis set in Figure 4 at relative amplitudes corresponding to values found in healthy human brain), and we artificially add a simulated macromolecular baseline (created by adding 5 Gaussian components as suggested in [**40**]).

There exist several types of baseline correction methods.

- Multiplication of the FID by an exponentially decaying function. The baseline is fitted as the product of the original FID and an exponentially decreasing *apodization* function $\exp(-d_0 t)$ [**42**], which requires only the decay as parameter to be selected. As the baseline consists of broad fast decaying components, their main contribution is in the first part of the FID. After

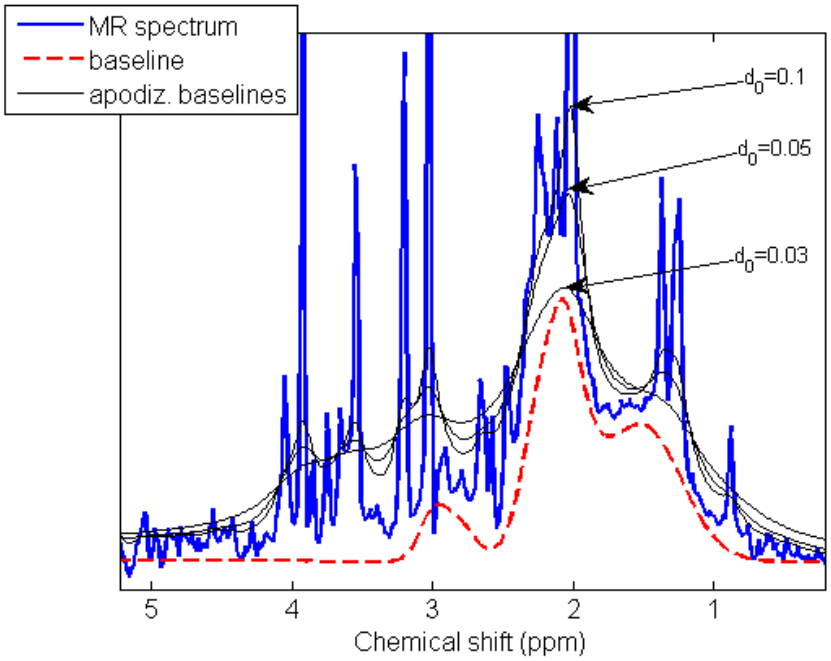


FIGURE 9. Real part of a simulated MRS signal with added macromolecular baseline, and three baseline estimates obtained by apodization with decay parameters $d_0 = 0.1, 0.05, 0.03$. For this example, the baseline is systematically overestimated.

multiplication with a decaying function, the resulting function will mainly contain the broad baseline components. As illustrated in Figure 9 for a range of decay parameters d_0, the obtained baseline also takes a significant part of the metabolite components and strongly overestimates the baseline contribution.

- Truncation of the initial points in the FID proposed by Ratiney *et al.* [**36**] in the method QUEST. The baseline, being a broad component in the spectrum, decays very rapidly (T_2-value of the macromolecules is small compared to the T_2 of the metabolites of interest). Hence, removing the initial part of the FID should suppress the baseline components. However, this operation also suppresses part of the metabolite contributions and might reduce the accuracy of the estimates since the SNR is decreased. QUEST fits a basis set of truncated metabolite signals to the truncated signal, then extrapolates the model

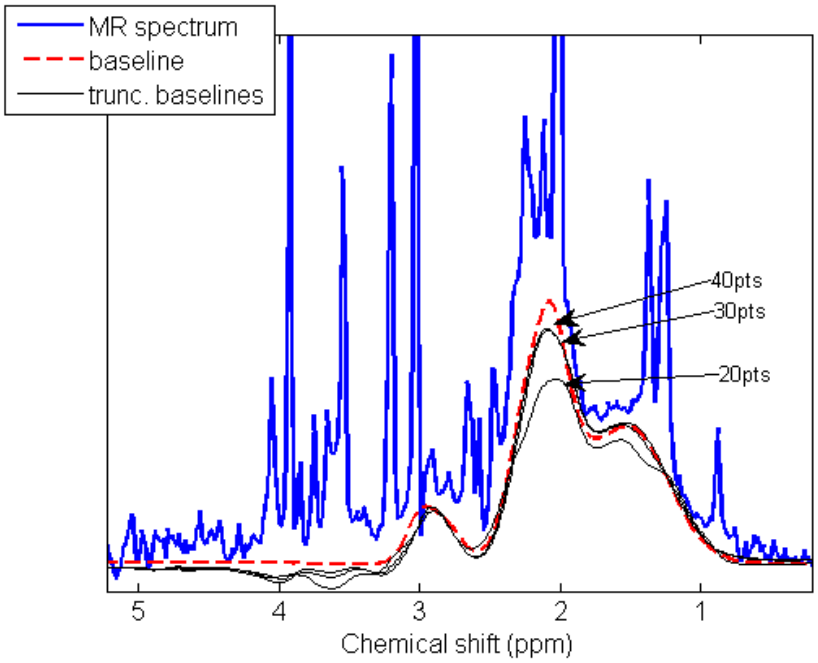


FIGURE 10. Real part of a simulated MRS signal with added macromolecular baseline, and three baseline estimates obtained by truncation, back-extrapolation, subtraction and denoising—a QUEST-like procedure—with number of truncation points equal to 20, 30, and 40. For this example, the baseline is systematically underestimated.

backwards to time 0, and subtracts it from the original signal. What is obtained is a (noisy) estimation of the baseline. Figure 10 illustrates several baseline estimates obtained for different number of truncation points and denoised using smoothing spline.

- Nonparametric modeling of the baseline, characterized by a smooth function in the frequency domain, by penalized splines [**13**] in AQSES [**30, 41**]. In order to fit the model and the smooth baseline at the same time, a regularized nonlinear least squares criterion is considered

$$\min \ \|\mathbf{y} - \widehat{\mathbf{y}}\|^2 + \lambda \|D\mathbf{c}\|^2, \tag{15}$$

where $\widehat{\mathbf{y}}$ is given in (14), but with the baseline b replaced by the formula $b = \mathcal{A} \cdot \mathbf{c}$, which is the inverse Fourier transformation of the frequency-domain baseline linearly modeled as a

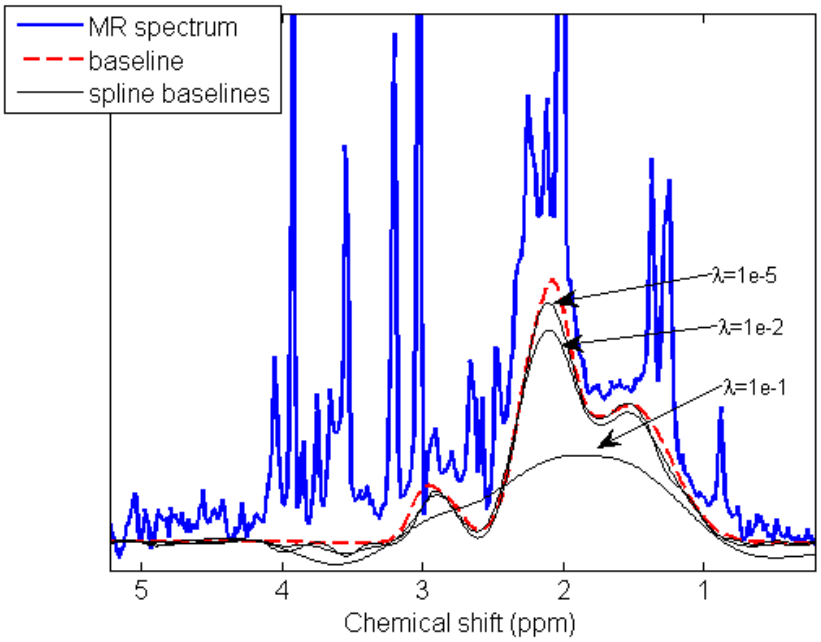


FIGURE 11. Real part of a simulated MRS signal with added macromolecular baseline, and three baseline estimates obtained by simultaneous fitting the metabolites and the spline baseline in the time-domain—the AQSES procedure—with regularization parameters λ equal to 1e−5, 1e−2, and 1e−1. For this example, the baseline is systematically underestimated.

sum of spline functions, $A \cdot \mathbf{c}$. In (15), the whole term $\lambda \|D\mathbf{c}\|^2$ is responsible for ensuring a certain degree of smoothness to the baseline b. The regularization matrix D is a (combination of) discrete derivative operator(s) [**13**]. The regularization parameter λ is a positive scalar that controls the degree of smoothness. Figure 11 illustrates the baseline fit obtained for the same example as above; here the regularization parameter was modified to vary the degree of smoothness.

None of the methods above is perfect, partly because the main assumption made is only about the smoothness of the baseline in the frequency domain (equivalent to saying that the time-domain counterpart decays in the first few time points), and this smoothness has to be controlled through tuning certain hyper-parameters. However, there are other limitations not taken into account, as for example the fact

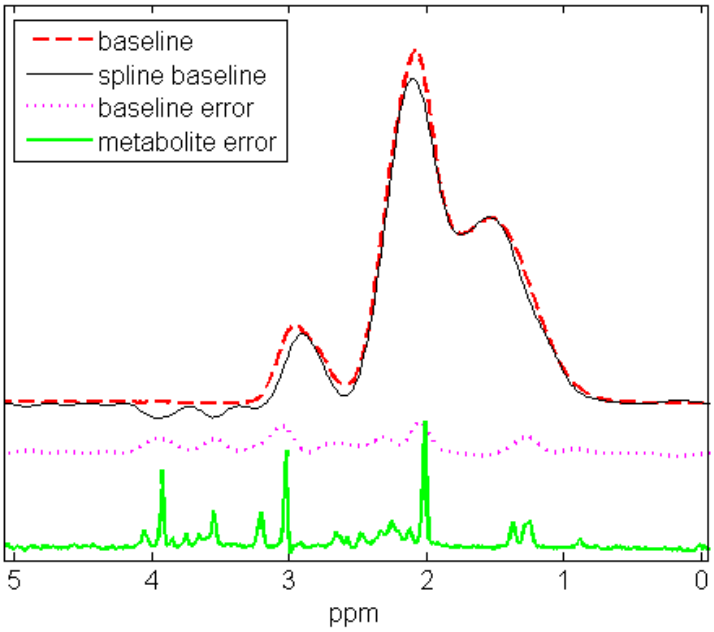


FIGURE 12. Real part of a simulated MRS macromolecular baseline (interrupted red line), and a spline baseline estimate (thin black line). The error in the estimation of the metabolites, *i.e.,* an overestimation of about 7% compared to the true values, is shown as a thick green line. This error is highly correlated with the baseline estimation error (dotted magenta line).

that the baseline is highly overlapping to some metabolites of interest, thus their contributions to the whole signal cannot always be disentagled. We illustrate this fact in Figure 12, where the metabolite error (*i.e.,* the difference between simulated and reconstructed signals based only on the metabolite profiles) is shown together with the simulated baseline, a reconstructed baseline (with splines in AQSES), as well as the baseline residual fit. The estimated amplitudes of all metabolites are overestimated by around 7% compared to the true values, while the baseline is underestimated. The high spectral overlap between metabolites and the baseline thus leads to systematic bias in the estimation.

Note. When a baseline is modeled nonparametrically with splines, a large number of additional linear variables (the spline coefficients) need also to be estimated. In this case it is essential to make use of the variable projection approach, as illustrated in the appendix.

5. Conclusions

In this Chapter we discussed Magnetic Resonance Spectroscopy from basic theoretical concepts to typical signal analysis techniques. The concepts of chemical shift and J-coupling have been instrumental for the definition of biochemical prior knowledge, which makes it possible to differentiate between different metabolites present in the voxel under investigation. The problem of metabolite quantification can be formulated mathematically as a nonlinear least squares fit between the data and a model. Traditional models based on a sum of complex damped exponential, although very successful for signals with few spectral components and no artifacts, have given way in more recent years to methods that can incorporate more prior knowledge, *e.g.,* models based on a weighted sum of metabolite signals. The numerical methods used for metabolite quantification range from black-box subspace-based methods that can only incorporate very limited prior knowledge to optimization-based methods that are allowing more flexibility to the model. The separation of the model parameters into linear and nonlinear motivates the applicability of the variable projection technique in a number of MRS models.

MRS signals measured in the human brain may contain nuisance components originating from inhomogeneous magnetic field, patient movement, interference from the skull, etc, which must be removed during preprocessing or taken into account at the modeling stage. Although water suppression or eddy current corrections are well studied at both hardware and signal processing stages, other problems with artifacts are still under study, such as the best model for the macromolecular background, or the automatic estimation of non-Lorentzian lineshape models.

Appendix

In this appendix we illustrate the convergence of the quantification method based on a metabolite basis set (AQSES [**30**]) in the cases when variable projection is implemented or not in the nonlinear least squares objective, and when a baseline modeled with splines is used into the model or not. Results are obtained in Matlab® (The Mathworks) using the Optimization Toolbox. For the variable projection version we use the AQSES implementation in SPID, while the minimization of the full functional is available in an older AQSES version.

The considered example is a noise-free signal of length 1024 containing 8 metabolite profiles, each perturbed by a frequency shift and

a damping variation, but without other artifacts. Note that this zero-residual example is not typical, as there are always artifacts and noise in MRS signals, but these would impair the illustrative effect of this example, since convergence would not be so easily visualized.

When no baseline is estimated, the model has an equal number of linear and nonlinear variables (*i.e.*, 16 (non)linear variables for the 8 metabolites). In this case, the evaluation of the projected functional and Jacobian (with Kaufman's simplification [**20**]) is more expensive than of the full functional. The number of iterations needed for convergence is similar for the two approaches because good starting values for the linear parameters are provided in the minimization of the full functional by a linear least squares fit with fixed initial nonlinear parameters values, which makes the variable projection slower for this example (9 versus 4 seconds on a 2.4GHz processor; without counting plotting time, the effective time becomes 4.7s versus 0.7s). (See Figure 13 **a.-b.** and the corresponding videos.)

We also illustrate the case when a baseline is considered in the model, which is nonparametrically fitted with splines. In this case there are additional linear parameters (the spline coefficients, where we considered a spline function per 10% of the data points, leading to 102 additional linear parameters). Without doubt, the variable projection is faster and more stable in this case. The number of iterations needed for convergence is lower for the variable projection approach, which makes the variable projection faster for this example (28 versus 113 seconds on a 2.4GHz processor; without counting plotting time, the effective time becomes 24s versus 109s). (See Figure 13 **c.-d.** and the corresponding videos.)

Acknowledgments

Dr. Diana M. Sima is a postdoctoral fellow of the Fund for Scientific Research-Flanders. Dr. Sabine Van Huffel is a full professor at the Katholieke Universiteit Leuven, Belgium. Research supported by

- Research Council KUL: GOA-AMBioRICS, CoE EF/05/006 Optimization in Engineering (OPTEC), IDO 05/010 EEG-fMRI, IOF-KP06/11, several PhD/postdoc and fellow grants;
- Flemish Government:
  - FWO: PhD/postdoc grants, projects, G.0407.02 (support vector machines), G.0360.05 (EEG, Epileptic), G.0519.06 (Noninvasive brain oxygenation), FWO-G.0321.06 (Tensors/Spectral Analysis), G.0302.07 (SVM), G.0341.07 (Data fusion), research communities (ICCoS, ANMMM);
  - IWT: PhD Grants;

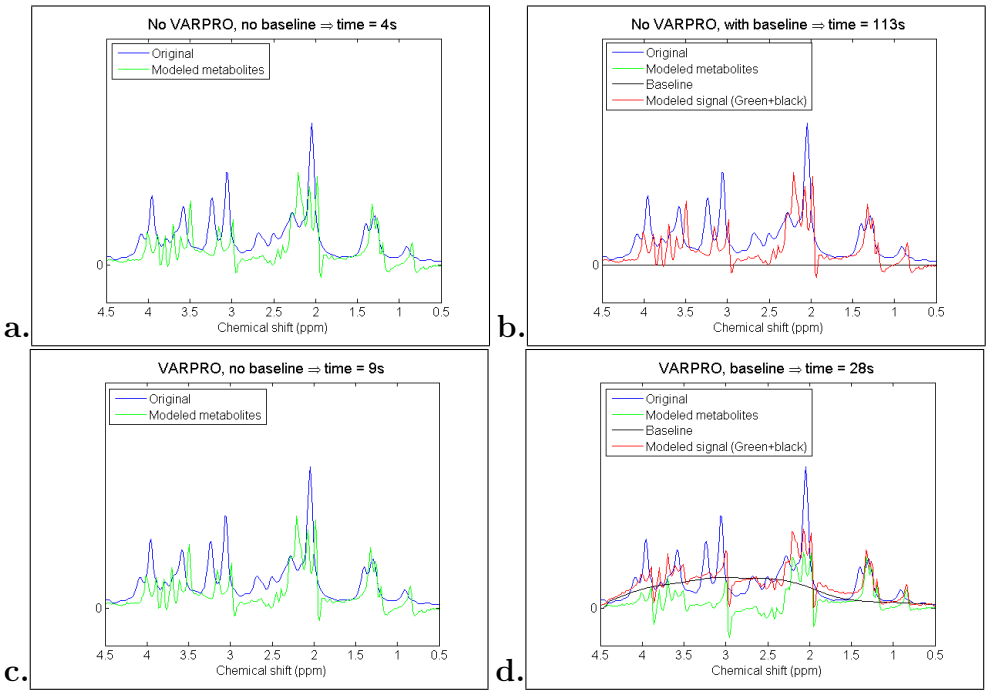


FIGURE 13. Real part of a simulated MRS signal (blue) and the initial fit with a combination of 8 metabolite signals (green in **a.** and **c.**, red in **b.** and **d.**, where a spline baseline in black is also added). Recall that the initialization is set to no corrections in the metabolite profiles. **a.** Convergence occurs in 12 iterations and takes 4 seconds when the full functional is used and no baseline is included. **b.** Convergence occurs in 17 iterations and takes 113 seconds when the full functional is used and baseline is included. **c.** Convergence occurs in 15 iterations and takes 9 seconds when variable projection is used and no baseline is included. **d.** Convergence occurs in 12 iterations and takes 28 seconds when variable projection is used and baseline is included.

- Belgian Federal Science Policy Office IUAP P6/04 (DYSCO, 'Dynamical systems, control and optimization', 2007-2011);
- EU: BIOPATTERN (FP6-2002-IST 508803), ETUMOUR (FP6-2002-LIFESCIHEALTH 503094), Healthagents (IST-2004-27214), FAST (FP6-MC-RTN-035801)
- ESA: Cardiovascular Control (Prodex-8 C90242)

Bibliography

1. J. P. Antoine, A. Coron, and J. M. Dereppe, *Water peak suppression: time-frequency vs time-scale approach*, J. Magn. Reson. **144** (2000), no. 2, 189–194.
2. H. Barkhuijsen, R. de Beer, W. M. Bovee, J. H. Creyghton, and D. van Ormondt, *Application of linear prediction and singular value decomposition (LPSVD) to determine NMR frequencies and intensities from the FID*, Magn. Reson. Med. **2** (1985), no. 1, 86–9.
3. H. Barkhuijsen, R. de Beer, and D. van Ormondt, *Improved algorithm for non-iterative time-domain model fitting to exponentially damped magnetic resonance signals*, J. Magn. Reson. **73** (1987), no. 3, 553–57.
4. R. Bartha, D. J. Drost, R. S. Menon, and P. C. Williamson, *Spectroscopic lineshape correction by QUECC: combined QUALITY deconvolution and eddy current correction*, Magn. Reson. Med. **44** (2000), no. 4, 641–5.
5. F. Bloch, *Nuclear induction*, Physics Review **70** (1946), 460–473.
6. I.M. Burtscher and S. Holtøas, *Proton MR spectroscopy in clinical routine*, Journal of Magnetic Resonance Imaging **13** (2001), no. 3, 560–567.
7. H. Chen, S. Van Huffel, A. van den Boom, and P. van den Bosch, *Subspace-based parameter estimation of exponentially damped sinusoids using prior knowledge of frequency and phase*, Signal Process. **59** (1997), no. 1, 129–36.
8. H. Chen, S. Van Huffel, D. van Ormondt, and R. de Beer, *Parameter estimation with prior knowledge of known signal poles for the quantification of NMR spectroscopy data in the time domain*, J. Magn. Reson. A **119** (1996), no. 2, 225–34.
9. H. Chen, S. Van Huffel, and J. Vandewalle, *Improved methods for exponential parameter estimation in the presence of known poles and noise*, IEEE T. Signal Proces. **45** (1997), no. 5, 1390–1393.
10. A. A. de Graaf, *QUALITY: quantification improvement by converting lineshapes to the Lorentzian type*, Magn. Reson. Med. **13** (1990), 343–57.
11. B. De Neuter, J. Luts, L. Vanhamme, P. Lemmerling, and S. Van Huffel, *Java-based framework for processing and displaying short-echo-time magnetic resonance spectroscopy signals*, Computer Methods and Programs in Biomedicine **85** (2007), 129–137.
12. A. Devos, L. Lukas, J. A. K. Suykens, L. Vanhamme, A. R. Tate, F. A. Howe, C. Majos, A. Moreno Torres, C. van der Graaf, M. Arus, and S. Van Huffel, *Classification of brain tumours using short echo time ^1H MR spectra*, J. Magn. Reson. **170** (2004), no. 1, 164–75.
13. P. H. C. Eilers and B. D. Marx, *Flexible smoothing with B-splines and penalties*, Stat. Sci. **11** (1997), 89–121.

14. G. H. Golub and V. Pereyra, *The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate*, SIAM J. Numer. Anal. **10** (1973), 413–32.
15. ———, *Separable nonlinear least squares: the variable projection method and its applications*, Inverse Probl. **19** (2003), no. 2, 1–26.
16. D. Graveron Demilly, A. Diop, A. Briguet, and B. Fenet, *Product-operator algebra for strongly coupled spin systems*, J. Magn. Reson. A **101** (1993), no. 3, 233–39.
17. Y.Y. Hsu, A.T. Du, N Schuff, and Weiner M.W., *Magnetic resonance imaging and magnetic resonance spectroscopy in dementias*, J. Geriatr. Psychiatry Neurol. **14** (2001), no. 3, 145–166.
18. W. Huang, G. E. Alexander, L. Chang, H. U. Shetty, J. S. Krasuski, S. I. Rapoport, and M. B. Schapiro, *Brain metabolite concentration and dementia severity in Alzheimer's disease: a ^1H MRS study.*, Neurology **57** (2001), no. 4, 626–632.
19. M. Kanowski, J. Kaufmann, J. Braun, J. Bernarding, and C. Tempelmann, *Quantitation of simulated short echo time ^1H human brain spectra by LCModel and AMARES*, Magn. Reson. Med. **51** (2004), no. 5, 904–912.
20. L. Kaufman, *A variable projection method for solving separable nonlinear least squares problems*, BIT **15** (1975), 49–57.
21. U. Klose, *In vivo proton spectroscopy in presence of eddy currents*, Magn. Reson. Med. **14** (1990), 26–30.
22. T. Laudadio, *Subspace-based quantification of magnetic resonance spectroscopy data using biochemical prior knowledge*, Ph.D. thesis, Faculty of Engineering, K.U.Leuven, Leuven, Belgium, 2005.
23. T. Laudadio, N. Mastronardi, L. Vanhamme, P. Van Hecke, and S. Van Huffel, *Improved Lanczos algorithms for blackbox MRS data quantitation*, J. Magn. Reson. **157** (2002), no. 2, 292–7.
24. T. Laudadio, Y. Selën, L. Vanhamme, P. Stoica, P. Van Hecke, and S. Van Huffel, *Subspace-based MRS data quantitation of multiplets using prior knowledge*, J. Magn. Reson. **168** (2004), no. 1, 53–65.
25. J. J. Moré, *The Levenberg-Marquardt algorithm: Implementation and theory*, Numerical Analysis: Proceedings of the Biennial Conference held at Dundee, June 28-July 1, 1977 (Lecture Notes in Mathematics #630) (1978), 104–16.
26. A. Naressi, C. Couturier, I. Castang, R. de Beer, and D. Graveron-Demilly, *Java-based graphical user interface for MRUI, a software package for quantitation of in vivo/medical magnetic resonance spectroscopy signals.*, Comput Biol Med **31** (2001), no. 4, 269–286.
27. W. W. F. Pijnappel, A. van den Boogaart, R. de Beer, and D. van Ormondt, *SVD-based quantification of magnetic resonance signals*, J. Magn. Reson. **97** (1992), no. 1, 122–34.
28. J. Poullet, *Quantification and classification of magnetic resonance spectroscopic data for brain tumor diagnosis*, Ph.D. thesis, Faculty of Engineering, K.U.Leuven, Leuven, Belgium, 2008.
29. J. Poullet, R. Pintelon, and S. Van Huffel, *A new fir filter technique for solvent suppression in mrs signals*, Journal of Magnetic Resonance **196** (2009), 61–73.
30. J. Poullet, D. M. Sima, A. W. Simonetti, B. De Neuter, L. Vanhamme, P. Lemmerling, and S. Van Huffel, *An automated quantitation of short echo time MRS*

spectra in an open source software environment: AQSES, NMR Biomed. **20** (2007), no. 5, 493–504.
31. J. Poullet, D. M. Sima, and S. Van Huffel, *MRS signal quantitation: a review of time- and frequency-domain methods*, J. Magn. Reson. **195** (2008), 134–144.
32. J. Poullet, Diana M Sima, S. Van Huffel, and P. Van Hecke, *Frequency-selective quantitation of short-echo time ^1H magnetic resonance spectra*, J. Magn. Reson. **186** (2007), no. 2, 293–304.
33. E. Prost, P. Sizun, M. Piotto, and J. M. Nuzillard, *A simple scheme for the design of solvent-suppression pulses*, J. Magn. Reson. **159** (2002), no. 1, 76–81.
34. S. W. Provencher, *Estimation of metabolite concentrations from localized in vivo proton NMR spectra*, Magn. Reson. Med. **30** (1993), no. 6, 672–79.
35. ———, *Automatic quantitation of localized in vivo ^1H spectra with LCModel*, NMR Biomed. **14** (2001), no. 4, 260–4.
36. H. Ratiney, Y. Coenradie, S. Cavassila, D. van Ormondt, and D. Graveron-Demilly, *Time-domain quantitation of ^1H short echo-time signals: background accommodation*, MAGMA **16** (2004), no. 6, 284–96.
37. H. Ratiney, M. Sdika, Y. Coenradie, S. Cavassila, D. van Ormondt, and D. Graveron-Demilly, *Time-domain semi-parametric estimation based on a metabolite basis set*, NMR Biomed. **18** (2005), no. 1, 1–13.
38. R. Romano, A. Motta, S. Camassa, C. Pagano, M. T. Santini, and P. L. Indovina, *A new time-domain frequency-selective quantification algorithm*, J. Magn. Reson. **155** (2002), no. 2, 226–35.
39. A. Ruhe and P.A. Wedin, *Algorithms for Separable Nonlinear Least Squares Problems*, SIAM Rev. **22** (1980), no. 3, 318–37.
40. U. Seeger and U. Klose, *Parameterized evaluation of macromolecules and lipids in proton MR spectroscopy of brain diseases*, Magn. Reson. Med. **49** (2003), 19–28.
41. D. M. Sima and S. Van Huffel, *Regularized semiparametric model identification with application to NMR signal quantification with unknown macromolecular baseline*, J. Roy. Stat. Soc. B **68** (2006), no. 3, 383–409.
42. A. W. Simonetti, W. Melsen, M. van der Graaf, G. Postma, A. Heerschap, and L. M. C. Buydens, *A chemometric approach for brain tumor classsification using magnetic resonance imaging and spectroscopy*, Anal. Chem. **75(20)** (2003), 29–36.
43. A. J. Simpson and S. A. Brown, *Purge NMR: effective and easy solvent suppression*, J. Magn. Reson. **175** (2005), no. 2, 340–6.
44. M. A. Smith, J. Gillen, M. T. McMahon, P. B. Barker, and X. Golay, *Simultaneous water and lipid suppression for in vivo brain spectroscopy in humans*, Magn. Reson. Med. **54** (2005), no. 3, 691–6.
45. S. A. Smith, T. O. Levante, B. H. Meier, and R. R. Ernst, *Computer simulations in magnetic resonance. An object-oriented programming approach*, J. Magn. Reson. A **106** (1994), no. 1, 75–105.
46. M. Terpstra, P. M. Andersen, and R. Gruetter, *Localized eddy current compensation using quantitative field mapping.*, J Magn Reson **131** (1998), no. 1, 139–143.
47. A. van den Boogaart, F. A. Howe, L. M. Rodrigues, M. Stubbs, and J. R. Griffiths, *In vivo ^{31}P MRS: absolute concentrations, signal-to-noise and prior knowledge*, NMR Biomed. **8** (1995), no. 2, 87–93.

48. J. W. van der Veen, R. de Beer, P. R. Luyten, and D. van Ormondt, *Accurate quantification of in vivo ^{31}P NMR signals using the variable projection method and prior knowledge*, Magn. Reson. Med. **6** (1988), no. 1, 92–8.
49. S. Van Huffel, H. Chen, C. Decanniere, and P. Van Hecke, *Algorithm for time-domain NMR data fitting based on total least squares*, J. Magn. Reson. A **110** (1994), 228–237.
50. L. Vanhamme, S. Van Huffel, P. Van Hecke, and D. van Ormondt, *Time-domain quantification of series of biomedical magnetic resonance spectroscopy signals.*, J. Magn. Reson. **140** (1999), no. 1, 120–130.
51. L. Vanhamme, T. Sundin, P. Van Hecke, S. Van Huffel, and R. Pintelon, *Frequency-selective quantification of biomedical magnetic resonance spectroscopy data.*, J. Magn. Reson. **143** (2000), no. 1, 1–16.
52. L. Vanhamme, A. van den Boogaart, and S. Van Huffel, *Improved method for accurate and efficient quantification of MRS data with use of prior knowledge*, J. Magn. Reson. **129** (1997), 35–43.
53. J. M. Wild, *Artifacts introduced by zero order phase correction in proton NMR spectroscopy and a method of elimination by phase filtering*, Journal of Magnetic Resonance **137** (1999), 430–436.