Journal of Magnetic Resonance 196 (2009) 61–73

Contents lists available at ScienceDirect

Journal of Magnetic Resonance

journal homepage: www.elsevier.com/locate/jmr

A new FIR filter technique for solvent suppression in MRS signals

Jean-Baptiste Poullet [a,*], Rik Pintelon [b], Sabine Van Huffel [a]

[a] Department of Electrical Engineering, SCD-SISTA, Katholieke Universiteit Leuven, Kasteelpark Arenberg 10, 3001 Leuven, Belgium
[b] Department of Fundamental Electricity and Instrumentation, Faculty of Engineering, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium

ARTICLE INFO

Article history:
Received 22 July 2008
Revised 8 October 2008
Available online 14 October 2008

Keywords:
Magnetic resonance spectroscopy (MRS)
Frequency-selective quantitation
Finite impulse response (FIR) filter
Automated quantitation of short echo time MRS spectra (AQSES)
Non-linear least-squares (NLLS)

ABSTRACT

A large water resonance is usually present in MRS signals. Time-domain quantitation methods require a good suppression of these components in order to obtain accurate parameter estimates. In this paper, we analyze one of the most successful methods for solvent suppression, the maximum-phase finite impulse response filter (MP-FIR), and identify its drawbacks. A new filtering method is proposed to overcome the limitation of MP-FIR.

1. Introduction

The water signal intensity is often several orders of magnitude larger than the signal intensities of the metabolites in MRS signals. In order to reduce the range of amplitudes (between the water resonance and the other components), water saturation techniques are used during acquisition (see, e.g., [1–3]), yielding water-suppressed signals. However, these signals still contain a water resonance higher than the metabolite signals. Since the water components are located outside the frequency region of interest, frequency-selective quantitation methods are indicated to obtain accurate parameter estimates. Solvent suppression methods have been mainly developed for time-domain MRS signals ([4–7]), where water and metabolite components are not distinguishable. In frequency-domain quantitation methods, the residual water peak tails are often dealt with by considering them as an additional baseline (e.g., [8–12]), or the water resonance is removed before the Fourier transformation of the time-domain signal. Water-suppressed signals may show various kinds of water shapes due to water saturation and no (or little) prior knowledge can be used to model the water resonance.

In [6], a maximum-phase FIR filter (MP-FIR) has been proposed and has been slightly refined in [7]. It has been shown to be successful for water removal and outperforms most of the existing concurrent methods [4,13]. However, this filter does not use all the available prior knowledge. Indeed, although little can be said about the shape of the water resonance, the position of the water peak is usually known or can be easily determined in the Fourier transform domain. Therefore the following prior knowledge should be used:

(i) the position of the water resonance and other resonances that should be removed (outside the frequency region of interest),
(ii) the position of the metabolite resonance which is closest to the water peak in order to determine the transition band.

Similarly, the location of any nuisance component (i.e., component in the frequency region of no interest) can be used as prior knowledge. MP-FIR is not explicitly designed for Lorentzian peaks since the filter design constraints are on the unit circle (pure sinusoids). In other words, a sinusoidal signal passes through the filter unchanged (if the ripples in the magnitude response are neglected), while the amplitude of a Lorentzian passing through the filter will be reduced proportionally to the damping value of the Lorentzian (see [13] for more details). We propose a new filtering method that uses the above-mentioned prior knowledge and reduces the attenuation of the metabolite signal when passing through the filter.

2. Filter design

2.1. FIR filters and their use in time-domain quantitation methods

In time-domain quantitation methods such as the automated quantitation of short echo time MRS[1] (AQSES) [14] or the QUantitation

* Corresponding author.
E-mail addresses: jean-baptiste.poullet@esat.kuleuven.be (J.-B. Poullet), sabine.vanhuffel@esat.kuleuven.be (S. Van Huffel).

[1] Software available at http://homes.esat.kuleuven.be/~biomed/software.php.

1090-7807/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2008.10.011

based on quantum ESTimation[2] (QUEST) [15], the short echo time MRS signal y is modeled in the time domain as

$$y(n) = \sum_{k=1}^{K} \alpha_k \zeta_k^n v_k(n) + b(n) + w(n) + \varepsilon(n), \quad n = 0, \ldots, N-1, \tag{1}$$

where $\{v_k, \text{ for } k = 1, \ldots, K\}$ denotes the metabolite database, $\alpha_k \zeta_k^n$ the correction applied to each profile k in this database, $b(n)$ the baseline, $w(n)$ the water component (as well as other nuisance components), $\varepsilon(n)$ the unknown noise of zero mean and N the number of points. The complex amplitudes α_k and the complex signal poles ζ_k can be written as (with $j = \sqrt{-1}$):

$$\alpha_k = a_k \exp(j\phi_k), \qquad \zeta_k = \exp(-d_k + j2\pi f_k)\Delta t, \tag{2}$$

where a_k are the real amplitudes, ϕ_k are the phase shifts, d_k are damping corrections, f_k are frequency shifts and Δt is the sampling time. Let

$$\hat{y}_k(n) = \alpha_k \zeta_k^n v_k(n), \tag{3}$$

where $\hat{y}_k$ is the kth individually corrected metabolite profile. A similar model is used for long echo time MRS data for which quantitation methods like AMARES [16] or VARPRO [17] have been widely used:

$$y(n) = \sum_{k=1}^{K} \alpha_k \zeta_k^n + b(n) + w(n) + \varepsilon(n), \quad n = 0, \ldots, N-1, \tag{4}$$

where no metabolite profiles $v_k(n)$ is used, and $\alpha_k \zeta_k^n$ are not corrections of a metabolite profile k but the complex signal intensity of metabolite k at the discrete time point $n\Delta t$.

The result of applying a FIR filter to a time-domain signal y is defined in the time domain by the convolution

$$y_{fil}(n) = \sum_{m=0}^{M-1} h(m) y(n-m), \tag{5}$$

where $\{h(m)\}_{m=0,\ldots,M-1}$ are the constant (possibly complex) filter coefficients. The first $M-1$ points of this filtered signal are distorted since they assume zero values for the unknown $y(-1), \ldots, y(-M+1)$. It has been shown in [6] for long echo time MRS data and in [13] for short echo time MRS data that metabolite signals based on a Lorentzian model pass through a FIR filter undistorted if

$$\left|\bar{h}\bar{\zeta}_k\right| = 1 \tag{6}$$

where

$$\bar{h} = (h_{M-1}, \ldots, h_0) \tag{7}$$

and

$$\bar{\zeta}_k = \left(1\, e^{(-d_k + j2\pi f_k)\Delta t} \ldots e^{(M-1)(-d_k + j2\pi f_k)\Delta t}\right)^T. \tag{8}$$

The challenge is therefore to design a filter which matches the constraints of Eq. (6) for metabolite resonances in the frequency passband and cancels out all the components lying in the stopband. We detail below how FIR filters can be used with quantitation methods based on Eq. (1). This can be easily extended to long echo time MRS data (Eq. (4)). In the rest of the paper AQSES, as quantitation method, will be used since this method is particularly suited when used in combination with FIR filtering methods (see [14,13] for more details). However, this filter is not restricted to AQSES and can be used with other quantitation methods as explained in Section 3.3.

FIR filter embedded in the quantitation

When using a FIR filter inside a quantitation method based on the minimization of the fitting error (neglecting the baseline for sake of simplicity), the nonlinear least squares problem becomes:

$$\min_{\alpha_k, \zeta_k} \sum_{n=0}^{N-1} \left| y_{fil}(n) - \sum_{k=1}^{K} \hat{y}_{fil,k}(n) \right|^2, \tag{9}$$

where

$$\hat{y}_{fil,k}(n) = \sum_{m=0}^{M-1} h(m) \alpha_k \zeta_k^{n-m} v_k(n-m). \tag{10}$$

The first $M-1$ points of the filtered metabolite profiles $\hat{y}_{fil,k}$ are discarded to avoid distortion.

FIR filter used prior to quantitation

An FIR filter can also be used as preprocessing method. In that case, the signal is filtered prior to the quantitation. The metabolite profiles in the database can also be filtered if they contain nuisance components. The nonlinear least squares problem becomes:

$$\min_{\alpha_k, \zeta_k} \sum_{n=0}^{N-1} \left| y_{fil}(n) - \sum_{k=1}^{K} \alpha_k \zeta_k^n v_{fil,k}(n) \right|^2, \tag{11}$$

where

$$v_{fil,k}(n) = \sum_{m=0}^{M-1} h(m) v_k(n-m). \tag{12}$$

The first $M-1$ points of the filtered metabolite profiles $v_{fil,k}$ are discarded to avoid distortion.

2.2. The proposed filter

As mentioned before, the proposed filter design is based on:

- the fact that we are dealing with Lorentzian signals and not pure sinusoids,
- prior knowledge of the position of the nuisance peaks.

The general idea of the filter is to place zeros in the z-plane to cancel components of specific frequencies (corresponding to the nuisance components). The z-transform, which allows to transform the domain of the original data (i.e., the time domain for MRS data) into the z-domain (or z-plane), is a well-known concept in signal processing and is defined as

$$X(z) = \sum_{n=0}^{N-1} x(n) z^{-n}. \tag{13}$$

where $x(n)$ is a discrete signal in the time domain, $X(z)$ its transform in the z-domain and $z = e^{(-d+j2\pi f)}$. The filtered signal is given by $X_{fil}(z) = H(z)X(z)$ where $H(z)$ is the filter transfer function in the z-domain. Filtering out a component z_0 from the signal is achieved when $H(z_0) = 0$ and z_0 is a root (=a zero) of $H(z)$. In order to suppress a Lorentzian component of frequency f_r and damping d, the filter transfer function $H(z)$ should be such that $H(z) = H(e^{(-d+j2\pi f_r)}) = 0$. The largest nuisance components and their frequencies are automatically determined (see step 3 in the algorithm below) and the zeros are chosen accordingly. The algorithm starts from one pair of zeros (=zero and its conjugate) corresponding to the largest nuisance component (e.g., the water resonance), calculates the filter coefficients, filters the signal and checks whether all the components in the stopband are below the noise level. If this is not the case, it adds a new pair of zeros corresponding to the frequency of the largest nuisance component in the filtered signal, recalculates the filter coefficients and filters again the signal, and so on until the nuisance components have been sufficiently filtered out. The filter order is increased if a satisfactory attenuation is not reached, and we start back from the first pair of zeros since the position of the subsequent zeros may depend on the filter order. Consequently, the initial choice of the filter order by the

[2] Software available at http://www.mrui.uab.es/mrui/.

user is not critical. Zeros are added by pair to ensure a flat and nonleaning magnitude response of the filter. A large and wide nuisance peak may need more than one zero to be filtered out and the position of the second zero will be preferably chosen accordingly to the largest remaining nuisance component. This explains the iterative process.

The algorithm is presented in Fig. 1 and will be referred to as MP-FIR0 in the rest of the manuscript. The different blocks of the algorithm are described here and further detailed in Appendix A. The frequencies are normalized in the interval $[-0.5, 0.5]$.

Initialization

1. To use the algorithm, the user needs to define the passband, i.e., the frequency region of interest, the transition bandwidth (TBW), i.e., the frequency band located between the region of interest and the region of no interest, and the filter order. The TBW cannot be zeroed due to the numerical limitation of FIR filters with finite number of coefficients. As a rule of thumb, the TBW will be chosen sufficiently large to reduce the ripples in the passband and stopband, but not too large to suppress satisfactorily the nuisance components, which have to be located in the stopband. As previously mentioned, the choice of the filter order is not critical and 30, for instance, can be selected as starting value.

2. The noise is estimated using the last samples of the complex signal. By default, the algorithm considers the last 20 samples (or time domain points), but this number can be modified according to the signal under investigation. The principle is to take as many points as possible which only contain noise (i.e., no metabolite signal or signal of interest).

3. The frequency of the strongest nuisance peak, i.e., the one with the highest magnitude, is automatically determined. In MRS spectra, the strongest nuisance peak often corresponds to the water resonance.

4. Since we want to suppress this nuisance peak, a zero corresponding to the frequency computed in step 3 is placed in the z-plane, together with its conjugate to keep a flat and nonleaning magnitude response. The first pair of zeros is placed on the unit circle (i.e., $z_r = e^{(-d+j2\pi f_r)} = 0$ with $d = 0$), since the water resonance is rarely a pure Lorentzian and no better parameter estimates were obtained with an arbitrary damping value.

5. We construct a grid of points in the frequency-damping domain. The goal is to keep the signal of interest in the passband, $|H(z)| = |H(f_l, d)| = 1$ with f_l in the passband, while suppressing the components in the stopband $|H(f_l, d)| = 0$ with f_l in the stopband. The number of points in the grid should be large enough to ensure the desired magnitude response for all frequencies in the passband and the stopband, but not too large to avoid a high computational load and numerical problems due to an ill-conditioned matrix (which occurs when rows become too similar in the matrix G, see Appendix A, Eq. (A.7)). The damping value d is estimated from the largest peak of interest after filtering. Intuitively, the largest peak is likely to yield the best estimate of d. Since we are at the initialization step, the signal has not been filtered yet, and the damping value should be initialized by the user. The experience shows that this parameter is not critical and the user can start from $d = 0$ if he has no idea about the true damping value.

Compute h

6. A filter transfer function is typically defined as

$$H(z) = H(d, f_l) = \prod_{c=1,\ldots,C} (1 - z_c^{-1}(e^{-d+j2\pi f_l})^{-1}), \tag{14}$$

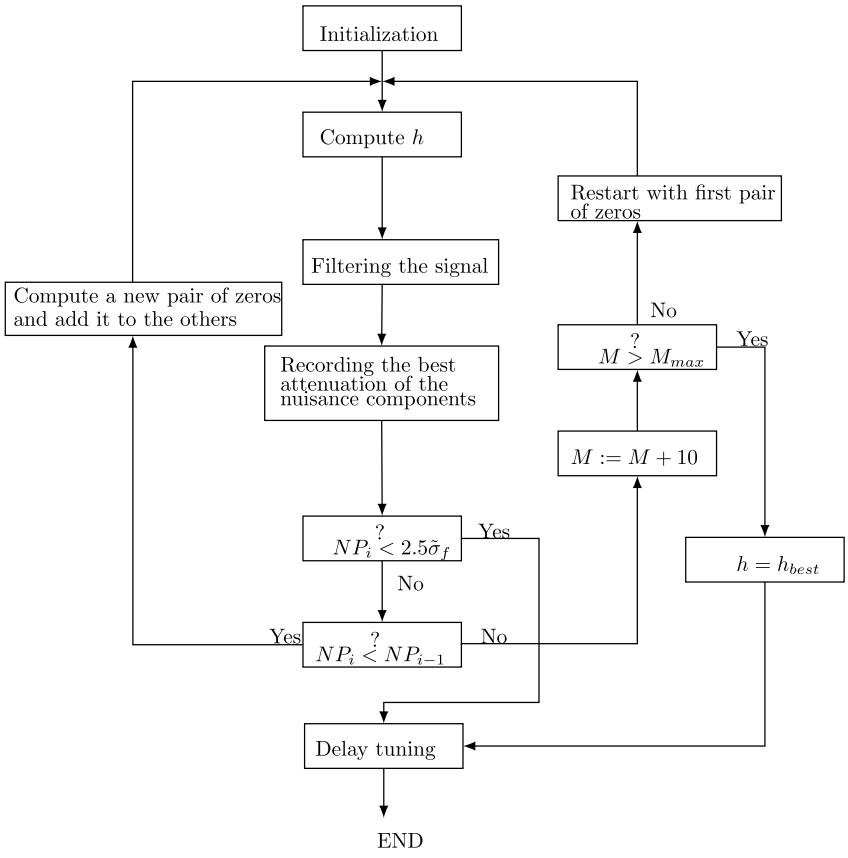

Fig. 1. Algorithm of the proposed filter, MP-FIR0. h is a vector containing the filter coefficients, NP_i is the highest magnitude in the stopband at iteration i, h_{best} are the filter coefficients which provide the smallest NP, $\tilde{\sigma}_f$ is the standard deviation of the noise in the frequency domain, M is the filter length and M_{max} the maximum filter length set to 110 by default. The mathematical expressions of these notations are given in Appendix A.

where C is the number of zeros. One pair (after initialization or after increasing the filter order) or several pairs of zeros (after step 11) have been computed and the transfer function can be expressed as $H(z) = H_1(z)H_2(z)$, where $H_1(z)$ results from the computed zeros (see Eq. (A.5)) and is fixed, and $H_2(z)$ is unknown and can be computed to ensure the desired magnitude response (see step 8). H_1 is computed at each point of the grid (see Eq. (A.5)).

7. The goal is thus to find H_2 such that $H(z) = H(f_l, d) = T(f_l)$ where $T(f_l)$ is the target vector (see Eq. (A.6)). The delay in Eq. (A.6) allows to obtain a better magnitude response since the effect of the frequency f_l is taken into account, while the magnitude of the target function is left unchanged. The value of the delay is difficult to estimate beforehand, but the experience shows that a good starting value is $0.73(M-1)$. Nevertheless, a final delay tuning (see step 13) is necessary to obtain a nonleaning magnitude response in all cases. Note that the choice of the delay value, $0.73(M-1)$, is not critical and other reasonable starting delay values will not change the results (see experiment 1 in Section 3.2).

8. $H_1(f_l, d)$ and $T(f_l)$ being computed at each grid point, we can calculate the coefficients of h_2 (see Eq. (A.7) and Appendix B). Similarly as in Eq. (13), $H_1(z) = \sum_{n=0}^{N-1} h_1(n) z^{-n}$, $H_1(z)$ is a polynomial in z^{-1} with the coefficients $h_1(n)$. Knowing the roots of the polynomial, its coefficients $h_1(n)$ can easily be calculated. The coefficients of h result from those of h_1 and h_2 since a multiplication $(H(z) = H_1(z)H_2(z))$ in the z-domain corresponds to a convolution $(h = h_1 \otimes h_2)$ in the time domain. The filter is then transformed into a maximum-phase FIR filter to minimize the energy loss of the components in the passband as explained in [13,6]. This transformation consists in projecting the zeros of the filter outside the unit circle (i.e., $z := z^{-1}$ for all zeros z inside the unit circle). Note that this operation does not change the magnitude of the filter.

Filtering the signal

9. The signal is filtered by h using Eq. (5), and the first $M-1$ points are discarded to avoid any distortion due to the fact that $y(M-1), \ldots, y(-1)$ are not defined (setting these coefficients to 0 generates distortion).

Checking the attenuation of the nuisance components

10. In this step, we first record the best attenuation of the nuisance components obtained until here. Then, we verify whether the attenuation criterion is reached, i.e., the signal of the nuisance components is below the noise signal (see Eq. (A.11) for more details). If so, the last step is to ensure a nonleaning passband by refining the delay value (step 13). Otherwise, we check whether the attenuation is improved compared with the last iteration. If it is the case, we add a new pair of zeros (step 11) corresponding to the frequency of the highest residual nuisance peak in order to remove it when filtering. Otherwise, adding new zeros will not improve the attenuation and increasing the filter order is necessary (step 12).

Increasing the number of zeros

11. A new pair of zeros is added as explained in Appendix A. Adding more zeros directly is not appropriate for two reasons: the number of nuisance peaks is often low and fixing zeros reduces the flexibility of the algorithm, which may result in suboptimal solutions.

Increasing the filter order

12. The order of the filter is increased by 10. The number 10 results from a tradeoff between a too small step, which would not yield significant changes in terms of attenuation of the nuisance components from one iteration to another, and a too large step, which would result in a too long filter with the risk of numerical instabilities, especially when transforming to a maximum-phase FIR filter which requires a root finding algorithm. We recommend to use $M_{max} = 100$ or =110 to avoid numerical problems.

Delay tuning

13. The goal of the delay tuning is to ensure a flat (nonleaning) magnitude response in the frequency passband. One can use a brute-force search, testing all delay values or start from an initial delay value and search around it until no flatter magnitude response is found. Based on the second approach, we propose the simple search detailed in Appendix A, step 13. To ensure a nonleaning passband, we want the magnitude response on the left-hand side of the center of the passband to be equivalent to the one on the right-hand side. For this, we calculate the difference defined in Eq. (A.13). Then, we try to find a delay value which minimizes this difference. Our simple search algorithm defines two delays around the previous delay value (the interval between the delay values is defined by D_{step}, see Appendix A, step 13.a.) and recomputes the difference between the sum of the magnitude responses for both delays. The delay value with the smallest corresponding difference is selected for the next iteration. The interval between the delay values is reduced at each iteration in the algorithm, which is stopped when no better delay value is found.

3. Experimental results using simulations and real-life examples

In this section the proposed FIR filter-based suppression technique is evaluated to determine the sensitivity of the final parameter estimates with respect to the choice of filter design parameters. The method is compared with the successful MP-FIR method and tests on *in vivo*, *ex vivo* and *in vitro* signals are also performed.

3.1. Material

3.1.1. Simulated signals

The simulated signals are constructed as described in Tables 1 and 2 in [6]. The parameters are repeated here in Tables 1 and 2 for clarity of exposition.

The simulated signals are the sum of the 5 metabolite signals, the nuisance peak (7 components in Table 1) and some noise:

$$y(n) = \sum_{k=1}^{5} v_k(n) + w(n) + \varepsilon(n), \quad n = 0, \ldots, N-1, \tag{15}$$

and

$$v_k(n) = a_{mk} e^{j\phi_{mk}} e^{(-d_{mk} + j2\pi f_{mk})n\Delta t} \tag{16}$$

where the values for a_{mk}, ϕ_{mk}, d_{mk} and f_{mk} are in Table 2, ε is circular complex distributed white Gaussian noise. The nuisance components in the simulated signals are constructed from Table 1:

$$w(n) = \sum_{wk=1}^{7} a_{wk} e^{j\phi_{wk}} e^{(-d_{wk} + j2\pi f_{wk})n\Delta t}. \tag{17}$$

No baseline $b(n)$ is added to the simulated signals.

In experiment 4, the value of d_{m2} in Table 2 has been replaced by 0.021 kHz (=21 Hz) to study the effect of MPFIR0 on wider peaks, the rest was kept unchanged. The SNR is defined with respect to the largest metabolite peak (peak 2) and is measured in decibels (dB):

Table 1
Estimated water signal parameters used in the reconstruction of the water peak.

f_{wk} (Hz)	d_{wk} (Hz)	ϕ_{wk} (°)	a_{wk} (a.u.)
−8.48	5.1	−90.88	15.01
−5.25	8.28	−45.19	64.74
−2.16	10.51	−2.95	321.25
−0.18	12.45	179.97	1142.3
−0.17	4.24	−170.39	251.92
3.09	6.79	36.77	201.11
6.31	4	81.11	12.3

Table 2
Metabolite parameters used in the simulated signals.

Peak k	f_{mk} (Hz)	d_{mk} (Hz)	ϕ_{mk} (°)	a_{mk} (a.u.)
1	61	7	0	20
2	118	7	0	30
3	189	7	0	20
4	231	7	0	20
5	311	7	0	20

$$\mathrm{SNR} = 20\log\left(\frac{a_{m2}}{\sigma}\right). \tag{18}$$

The metabolite profiles in the database used for quantitation are constructed from the 5 components in Table 2 ($\nu_k, k = 1,\ldots,5$), using exactly the same parameter values to reduce the effects of the quantitation method on the parameter estimates except for experiment 4 where f_{m1} = 0.056 kHz instead of 0.061 kHz.

Examples of simulated spectra are given in Fig. 2. Fig. 2(d) shows the metabolite signals constructed with a Gaussian model with the same parameter values for a_{mk}, f_{mk} and ϕ_{mk} as in Table 2 and $d_{G_k} = 0.0002$ for $k = 1,\ldots,5$:

$$m_G(n) = \sum_{k=1}^{K} a_{mk} e^{j\phi_{mk}} e^{-d_{G_k}(n\Delta t)^2 + j2\pi f_{mk} n\Delta t}, \quad n = 0,\ldots,N-1. \tag{19}$$

3.1.2. In vivo, ex vivo and in vitro signals

We analyze 4 *in vivo* data chosen randomly from the short echo time INTERPRET[3] database, coming from different centers: Institut de Diagnostic per la Imatge (IDI), Fundacion para la Lucha contra las Enfermedades Neurologicas de la Infancia (FLENI), Uniwersytet Medyczny w Lodzi (MUL). Detailed information about the acquisition parameters for each center are provided in Table 3. Quality control criteria have been applied to MR system performance and all data in this study passed a strict validation process [18]. These water-suppressed data were acquired at 1.5 T.

The *ex vivo* high resolution magic angle spinning (HRMAS) signal was acquired with an 1D PRESAT (pulse-and-acquire) pulse sequence at 11.7 T (500 MHz for ^{1}H), at 0–4 °C, and 4000 Hz spinning rate using a BRUKER Analytik GmbH spectrometer. The biopsy was extracted from brain tumor tissue (grade II astrocytoma) and frozen at −80 °C until use (TE/TR = 31/2000 ms).

The *in vitro* signal, *N*-acetyl-aspartate (NAA), was acquired on a 1.5 T Philips spectrometer using a PRESS sequence with an echo time of 23 ms and a repetition time of 6000 ms, and a volume box of $2 \times 2 \times 2$ cm^3.

3.2. Description of the experiments and results

Two experiments have been set up to study the influence of the filter design parameters, which are the initial delay and the transition bandwidth. In two other experiments, MP-FIR (method proposed in [7]) and MP-FIR0 (the proposed filter) are compared. The quantitation method AQSES is used for all the experiments. However, the baseline is not considered and therefore not included in the model, the Lorentzian lineshape is assumed. AMARES could have been chosen as well as AQSES for the simulated signals and AQSES is appropriate for long echo time signals as long as the metabolite profiles are properly chosen. AQSES can be used with any FIR filter (see [14]). The following parameters are used for the simulated signals: sampling frequency = 1 kHz, spectrometer frequency = 63.83 MHz. A final experiment shows the effects of MP-FIR0 on *in vivo*, *ex vivo* and *in vitro* signals.

3.2.1. Comparison criterion

In each experiment on simulated data, the quality of the amplitude estimates is measured with the relative root mean squared error (RRMSE) in percent,

$$\mathrm{RRMSE}_k = 100\sqrt{\frac{1}{C}\sum_{c=1}^{C}\frac{\left(a_k - \tilde{a}_k^c\right)^2}{a_k^2}}, \quad k = 1,\ldots,5 \tag{20}$$

where C is the number of simulation runs (C = 100 in each experiment) and $\tilde{a}_k^c$ denotes the estimate of a_k obtained in simulation run c. Note that a_k (for k = 1,...,5) is equal to 1 in all the experiments based on simulated data.

3.2.2. Experiment 1: influence of the initial delay on the amplitude estimates

MP-FIR0 is used as a preprocessing method with the following filter design parameters:

- the transition bandwidth TBW = 0.02 kHz,
- the length of the filter $M = 71$,
- the passband PB = [5.5, 11] ppm,
- the delays $D = [0.6(M-1), 0.67(M-1), 0.73(M-1), 0.8(M-1), 0.87(M-1)]$.

AQSES is used without filtering method and without constraints on the phase.

The results of *experiment 1* are given in Fig. 3. The results are stable for delay factors between 0.6 and 0.73, but too large delay factors make the algorithm unstable in some cases, resulting in higher RRMSEs.

3.2.3. Experiment 2: influence of the transition bandwidth on the amplitude estimates

MP-FIR0 is used as a preprocessing method with the following filter design parameters:

- TBW = [0.01, 0.02, 0.03, 0.04, 0.05] kHz,
- $M = 71$,
- PB = [5.5, 11] ppm,
- $D = 0.73(M-1)$.

AQSES is used without filtering method and without constraints on the phase. The transition bandwidths are illustrated in Fig. 4.

The results of *experiment 2* are given in Fig. 5. The algorithm is in general robust with respect to the choice of the transition bandwidth. However, large errors are observed in some simulation runs for TBW = 0.01 and SNR = 30. This can be explained by too severe filter design constraints since the algorithm aims to attenuate the water peak under the noise level, which is very low, while allowing a too narrow transition band.

3.2.4. Experiment 3: comparison MP-FIR0 vs. MP-FIR

In this experiment, MP-FIR0 and MP-FIR are compared. In that respect, 4 simulations have been carried out:

[3] http://azizu.uab.es/INTERPRET/.

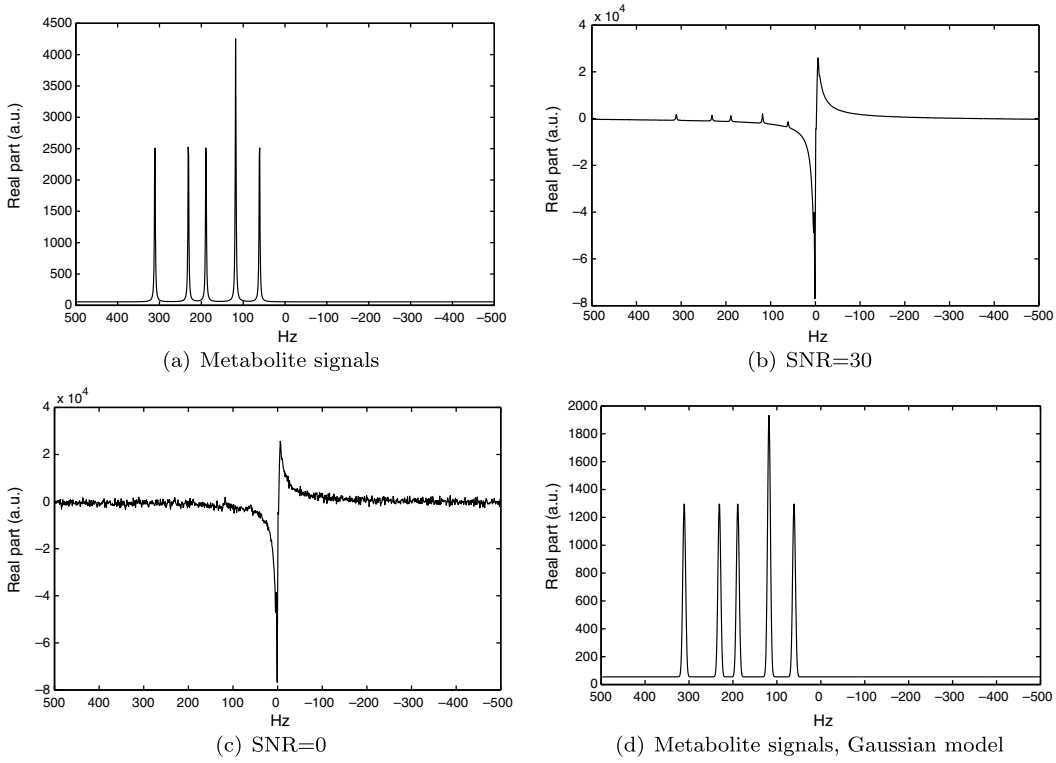

Fig. 2. Examples of simulated signals. (a) Metabolite signals with a Lorentzian model. (b) Simulated signal with SNR = 30 (Eq. (15)). (c) Simulated signal with SNR = 0 (Eq. (15)). (d) Metabolite signals with a Gaussian model (Eq. (19)).

Table 3
Overview of the acquisition parameters for the single voxel MRS data. For each center the manufacturer of the scanner, the sequence, the echo time, the repetition time, the spectral width (SW) and the number of points in the original FID (N) are displayed.

Center	Manufacturer	Sequence	TE (ms)	TR (ms)	SW (Hz)	N
FLENI	GE	PRESS	30	2000	2500	2048
IDI	Philips	PRESS	30	2000	1000	512
MUL	Siemens	STEAM	20	2000	1000	512

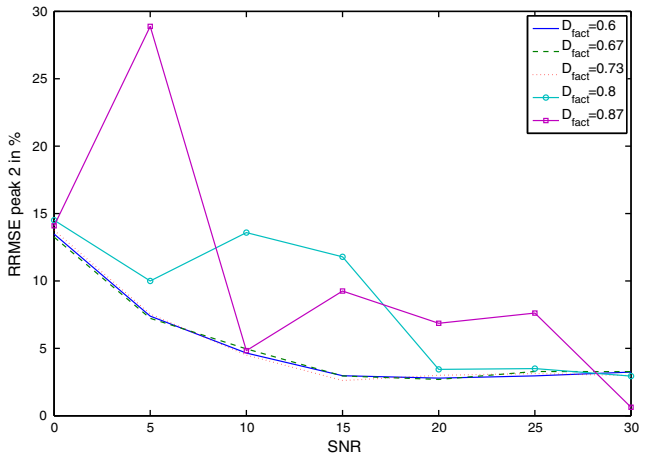

Fig. 3. Experiment 1: RRMSE as a function of SNR obtained from 100 simulation runs using different initial delay factors.

a. MP-FIR0 used as a preprocessing method with the following filter design parameters:
   - TBW = 0.02 kHz,
   - $M = 71$,
   - PB = [5.5, 11] ppm,
   - $D = 0.73(M-1)$.

   AQSES is used without filtering method and without constraints on the phase.
b. AQSES is used with MP-FIR0 as filtering method and with constraints on the phase (equal phases). The filter design parameters are identical to the ones in simulation a.
c. MP-FIR used as a preprocessing method with the following filter design parameters:
   - ripple = 0.01,
   - $M = 71$,
   - PB = [5.5, 11] ppm.

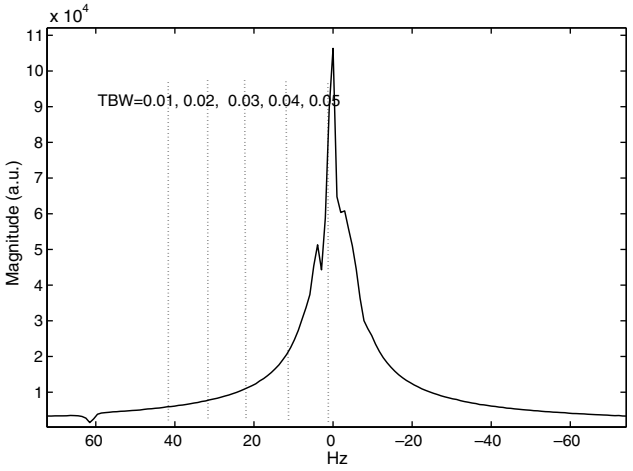

Fig. 4. Example of simulated spectrum (SNR = 30) in absolute values. The stopband is delimited by the dotted lines for the different transition bandwidths (TBW = 0.01, 0.02, 0.03, 0.04, 0.05).

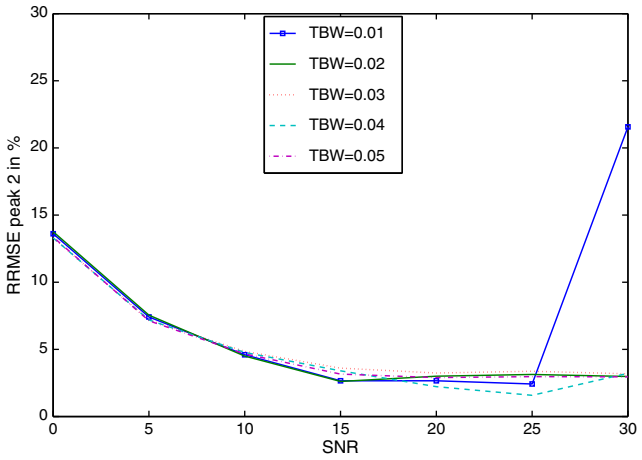

Fig. 5. Experiment 2: RRMSE as a function of SNR obtained from 100 simulation runs using different transition bandwidths.

AQSES is used without filtering method and without constraints on the phase.

d. AQSES is used with MP-FIR as filtering method and with constraints on the phase (equal phases). The filter design parameters are identical to the ones in simulation c.

The results of *experiment 3* are given in Fig. 6. It shows in Fig. 6(a) that the use of MP-FIR or MP-FIR0 inside AQSES is preferable, at least if the SNR is high. When using MPFIR or MP-FIR0 outside AQSES, peak 1 is attenuated and the RRMSE remains above a certain level since the database profiles remain unchanged. Using MP-FIR0 reduces this attenuation: the RRMSE level reached by MP-FIR0 is twice as small as the one reached by MP-FIR. The attenuation decreases for peaks like peak 2 that are further away from the passband borders. This is illustrated in Fig. 7. In Fig. 7(c) and (d), one can see that MP-FIR0 yields a smaller attenuation than MP-FIR with a flatter magnitude response curve in the passband. A damping of 0.021 has been chosen instead of 0.007 for a better contrast between the magnitude response of MP-FIR and MP-

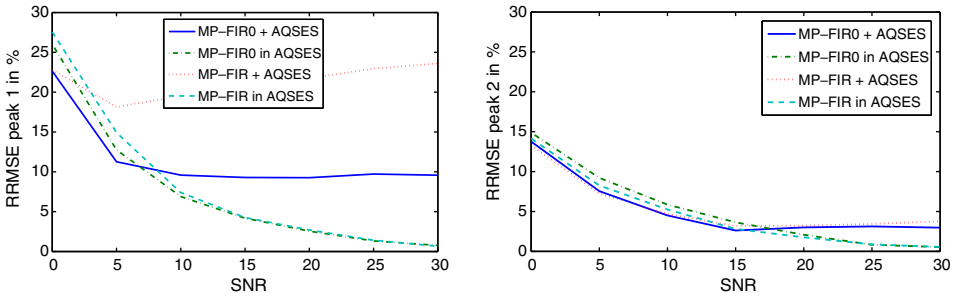

(a) Peak 1 (b) Peak 2

Fig. 6. Experiment 3: RRMSE as a function of SNR obtained from 100 simulation runs using different quantitation approaches.

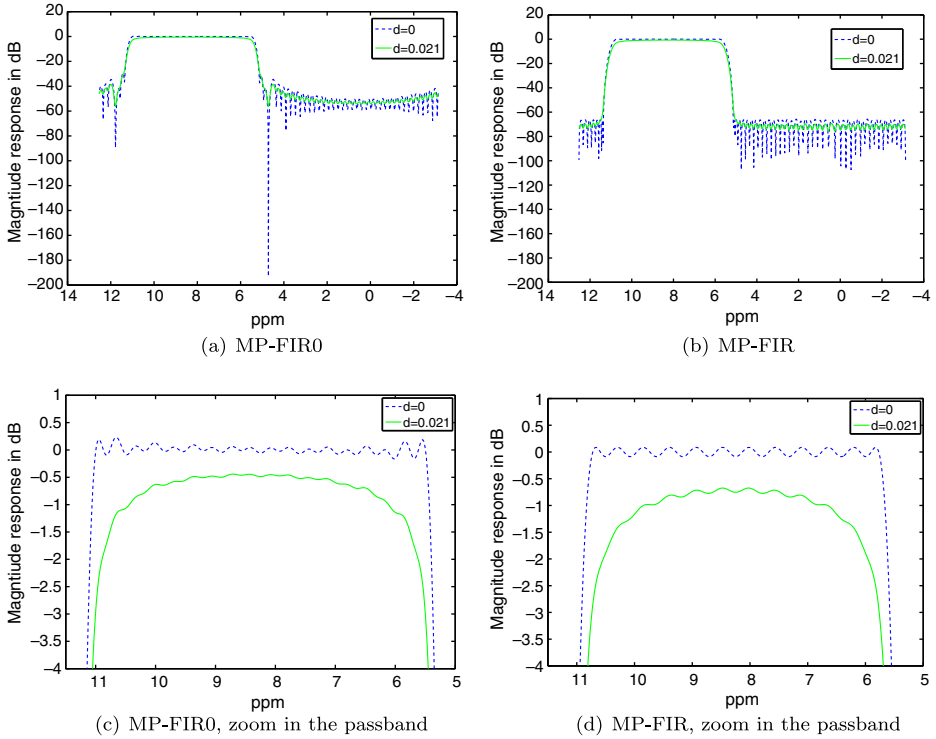

(a) MP-FIR0 (b) MP-FIR

(c) MP-FIR0, zoom in the passband (d) MP-FIR, zoom in the passband

Fig. 7. Experiment 3: examples of magnitude response of MP-FIR0 and MP-FIR in dB for components with different damping values: $d = 0$ and $d = 0.021$.

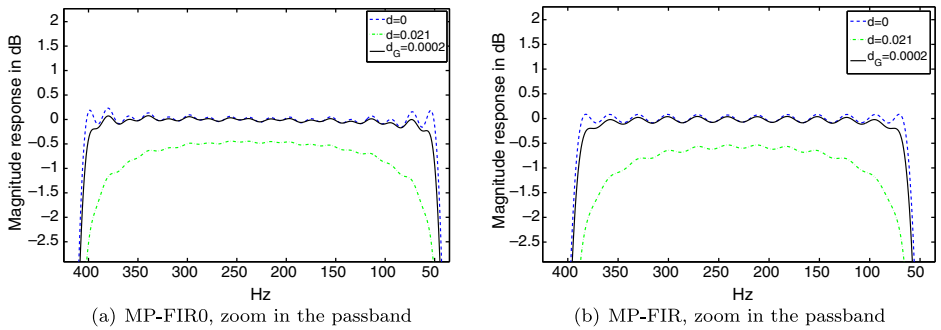

(a) MP-FIR0, zoom in the passband (b) MP-FIR, zoom in the passband

Fig. 8. Experiment 3: examples of magnitude response of MP-FIR0 and MP-FIR in dB for Gaussian (d_G) and Lorentzian (d) components.

FIR0. The location of the zeros in MP-FIR0 ensures a better suppression of the highest nuisance components as it is illustrated in Fig. 7(a) and (b). The suppression in the rest of the stopband is better when using MP-FIR. However, there is no peak to be suppressed in the rest of the passband, which illustrates that MP-FIR is not optimal in the sense that it does not take into account the position of the high nuisance peaks, their number being limited.

Fig. 8 shows what is the effect of MP-FIR and MP-FIR0 on a Gaussian peak. An example of Gaussian peaks with $d_{G_k} = 0.0002$ are given in Fig. 2(d). The passband attenuation is smaller for a Gaussian peak than for a Lorentzian peak.

3.2.5. Experiment 4: influence of deviations from the database in terms of frequency and damping

Experiment 4 only differs from experiment 3 in the fact that AQSES is always used without constraints on the phase, and that the frequency of the first metabolite in the *database* is 0.056 kHz instead of 0.061 kHz and the damping of the second metabolite in the *signal* is 0.021 kHz instead of 0.007 kHz. The goal of this experiment is to show that more complex problems can be handled with MP-FIR0.

The results of *experiment 4* are given in Fig. 9. Peak 1 (resonance frequency f_{m1} = 56 Hz) in the database is located at the border of the passband (passband = [51.06,402.13] in Hz). A better RRMSE is observed when using MP-FIR0 instead of MP-FIR outside AQSES. The differences are less obvious when using MP-FIR or MP-FIR0 inside AQSES.

3.2.6. Experiment 5: effects of MP-FIR0 on in vivo, ex vivo and in vitro data

MP-FIR0 is used as a preprocessing method with the following filter design parameters:

- TBW = [0.02] kHz,
- $M = 61$,
- PB = [0.5,4] ppm,
- $D = 0.73(M - 1)$.

MP-FIR is used as a preprocessing method with the following filter design parameters:

- ripple = 0.01,
- $M = 61$,
- PB = [0.5,4] ppm.

AQSES is not used here.

The *in vivo* results of *experiment 5* are given in Fig. 10. The magnitude of the signals is plotted for two reasons. First, in order to display the signal in the usual way, a phase correction of the filtered signal is needed. An *approximate* phase correction can easily be calculated from the filter phase response, but we prefer to plot the magnitude spectrum to circumvent the user-dependent phasing of the signal. Second, by displaying the magnitude spectra, all information present in the signal is visualized. This does not mean that the FIR filtering methods are not sensitive to the phase. The different water signals are reduced under the noise level by MP-FIR0.

The *ex vivo* results are given in Fig. 11. The water resonance is perfectly removed. The lipids at 0.9 ppm (comprised in the passband) are more attenuated with MP-FIR than MP-FIR0 as mentioned above for simulated signals. This is also true for large damping components, even if they are not especially close to passband borders (see, e.g., the fatty acids under the lactate doublet at 1.3 ppm).

The *in vitro* results are given in Fig. 12. Fig. 12(a) shows that all the nuisance components (components outside the passband [0.5,4] ppm) are removed. The algorithm automatically places the zeros to remove the largest nuisance components as it is illustrated by the magnitude response of the filter (for $d = 0$) in Fig. 12(b).

3.3. Discussion

The advantages of a method based on a FIR filter compared to a singular value decomposition method such as HLSVD-PRO [5] have been given in [13]. In particular,

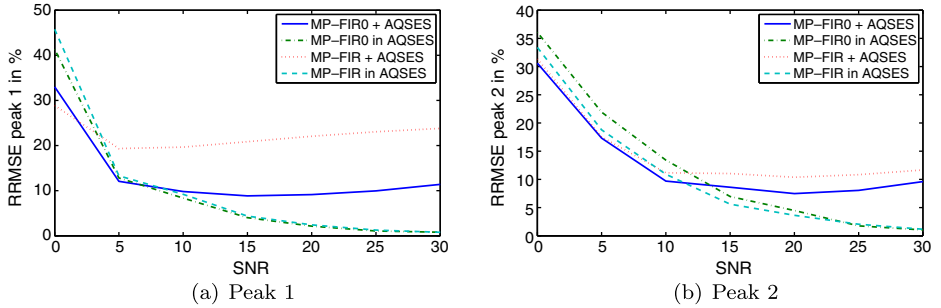

(a) Peak 1 (b) Peak 2

Fig. 9. Experiment 4: RRMSE as a function of SNR obtained from 100 simulation runs using different quantitation approaches.

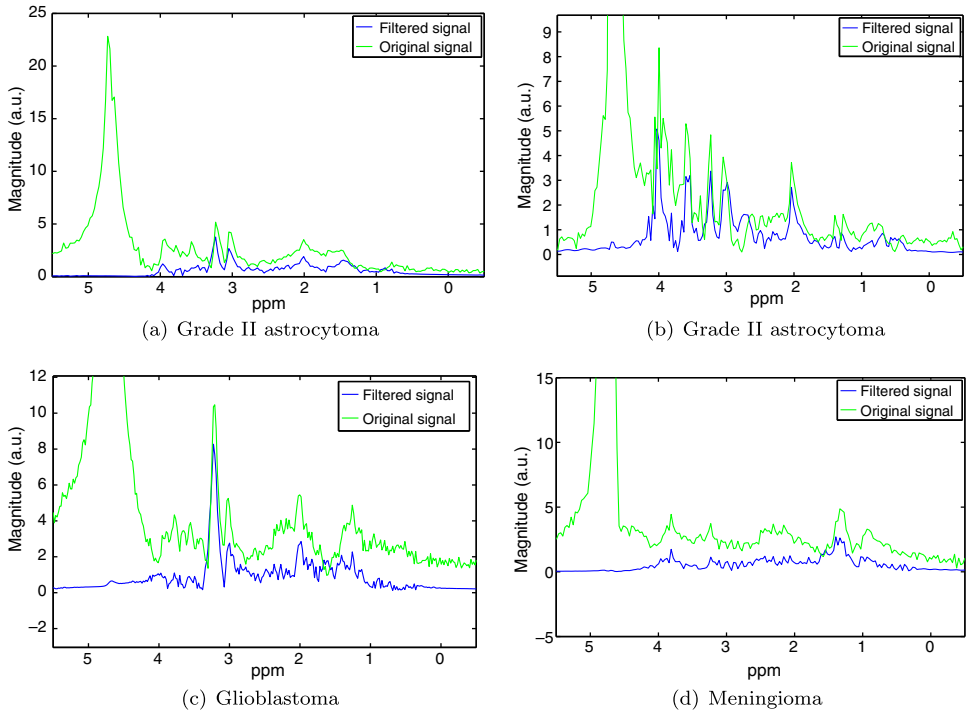

(a) Grade II astrocytoma (b) Grade II astrocytoma

(c) Glioblastoma (d) Meningioma

Fig. 10. Experiment 5: effect of MP-FIR0 on *in vivo* signals: magnitude spectrum of eddy current corrected signal (top signal) and magnitude spectrum of the same signal after FIR filtering (bottom signal). (a and b) Water-suppressed proton magnitude spectrum from the brain (tumor: grade II astrocytoma – center: IDI). (c) Water-suppressed proton magnitude spectrum from the brain (tumor: glioblastoma – center: FLENI). (d) Water-suppressed proton magnitude spectrum from the brain (tumor: meningioma – center: MUL).

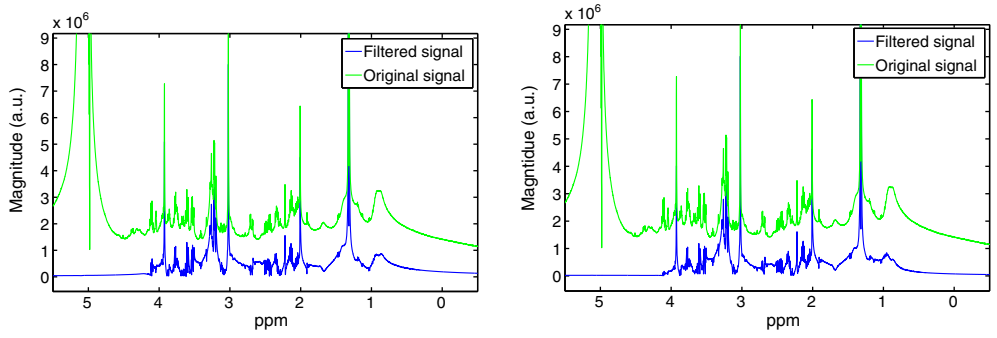

(a) MP-FIR0 (b) MP-FIR

Fig. 11. Experiment 5: effect of (a) MP-FIR0 and (b) MP-FIR on an *ex vivo* signal (tumor: grade II astrocytoma): magnitude spectrum of the signal (top signal) and magnitude spectrum of the same signal after FIR filtering (bottom signal).

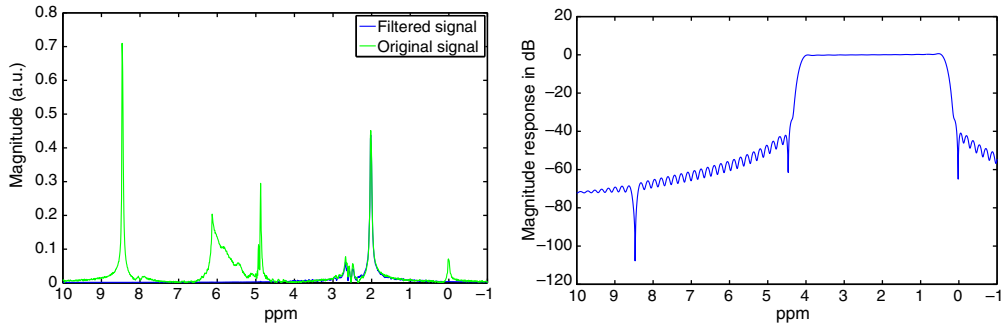

(a) NAA signal (b) Magnitude response of MP-FIR0

Fig. 12. Experiment 5: effect of MP-FIR0 on an *in vitro* signal (NAA). (a) Magnitude spectrum of the signal (top signal) and magnitude spectrum of the same signal after FIR filtering (bottom signal), (b) magnitude response of the filter in dB. The minima in the magnitude response correspond to the frequencies of the largest nuisance components.

- it is fast since it boils down to a simple matrix computation once the filter coefficients are known, and is thereby appropriate in an optimization procedure, whereas HLSVD-PRO requires a singular value decomposition at each iteration of the optimization procedure;
- it is linear and thus can be switched with any other linear operator: in AQSES [14], for example, the filtered signal is fitted with a linear combination of metabolite signals or profiles, and thus, filtering the linear combination of metabolite signals is equal to taking the linear combination of the filtered metabolite signals.

Moreover, MP-FIR0

- ensures that the water components and other unwanted components located outside the frequency region of interest are reduced under the noise level; the order of the filter $M-1$ (i.e., its length minus 1) and the number of iterations are optimized in that perspective;
- uses the information that we deal with nonpure sinusoidal peaks in the frequency region of interest, while MP-FIR is designed for sinusoidal signals;
- uses the water/unwanted peak positions as prior knowledge;
- is capable of dealing with non-Lorentzian water components as shown in the results above.

The constrained least squares design used by Sundin et al. for MP-FIR ensures equiripples in both the stopband and passband. However, the amplitudes in the stopband are far from being equal (only one or a few peaks are generally above the noise level) and imposing equal attenuation of all frequency components is usually not optimal. The results show that peaks close to the transition band are more attenuated by MP-FIR than MP-FIR0, yielding less accurate parameter estimates of these peaks.

Some remarks about the design of MP-FIR0:

- The filter design constraints are imposed outside the unit circle since we assume that the lineshape of a metabolite resonance is closer to a Lorentzian than to a pure sinusoid (i.e., an undamped Lorentzian, $d = 0$). The lineshape is also assumed to be similar for all metabolite resonances, and therefore only one value of d is considered when designing the filter. The reason is twofold: it reduces the risk of numerical issues and it decreases the computational load.
- Only one pair of zeros is added at each iteration (see step 11 in the algorithm) to keep the required filter order as low as possible, while still reducing the nuisance components under the noise level. Moreover, one pair of zeros is usually sufficient when there is only one large narrow nuisance component.
- The zeros of the large nuisance components are chosen on the unit circle since their shapes are usually unknown.

Similarly to MP-FIR, MP-FIR0 can be used in several quantitation methods although we have limited our analysis to the use of AQSES for sake of space. In [7], Vanhamme et al. show that MP-FIR can be used inside AMARES. Since the design of the filter is independent from the quantitation method (true for MP-FIR and MP-FIR0), it is obvious that MP-FIR0 can also be used inside AMARES. AQSES is based on the variable projection algorithm like VARPRO, therefore MP-FIR0 can similarly be used in VARPRO. The use of MP-FIR0 can also be extended to quantitation methods such as the one proposed by Elster et al. [19] or QUEST [15]. The original algorithm of QUEST can be summarized as follows: (1) filter out the nuisance components with an SVD-based method such as HLSVD [20], (2) truncate the initial points of both the signal under investigation and the metabolite profiles, and quantitate with QUEST the metabolites, (3) estimate the baseline from the metabolite-free signal by an SVD-based method or AMARES, (4) subtract the parameterized baseline from the raw signal, and quantitate the metabolites. With MP-FIR0, the algorithm would become: (1) truncate the initial points of both the filtered signal and the filtered metabolite profiles, and quantitate with QUEST, (2) estimate the baseline from the metabolite-free and water-free signal by an SVD-based method or AMARES, (3) subtract the parameterized baseline from the raw signal (with nuisance components), and quantitate the metabolites using MP-FIR0 inside QUEST. The number of truncated points in this case is smaller since the baseline is already reduced by the filter as it has been shown in the *in vivo* and *ex vivo* experiments. MP-FIR0 is not restricted to time-domain quantitation methods. For instance, in the time-domain frequency-domain quantitation method TDFDFit [21], it can be used just before transforming numerically the time-domain points into the frequency domain. The minimization problem is still solved in the frequency domain, but based on the filtered data. Note that filtering would occur at each iteration of the optimization procedure (i.e., inside the quantitation algorithm). In order to be used with the frequency-domain quantitation method LCModel [11], MP-FIR0 should be applied prior to the quantitation algorithm (MP-FIR0 uses the complex data while LC-Model only uses the real part of the data). As shown in experiments 3 and 4, MP-FIR0 is preferably used inside a quantitation method, and, in that sense, would be less appropriate for such quantitation methods.

4. Conclusions

The filter design of the proposed filter exploits more prior knowledge than MP-FIR. The following main advantages have been emphasized:

- a stronger attenuation of the large nuisance components since the choice of the zeros is directly related to the position of these peaks;
- a smaller attenuation of the signals of interest, especially with large damping values at the passband borders.

MP-FIR0 also shows to be successful with different types of signals (simulated, *in vivo*, *in vitro*, *ex vivo*) acquired from different spectrometers (Philips, Siemens, GE, Bruker) from different centers (IDI, FLENI, etc). MP-FIR0 can be used in several widely used quantitation methods.

Acknowledgments

We thank Bernardo Celda and M Carmen Martinez-Bisbal of University of Valencia, Spain, for providing us with the HRMAS data. We also thank Carles Arus and other partners from the INTERPRET project for the use of *in vivo* data. 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 Ph.D./postdoc and fellow grants.
- Flemish Government:
- FWO: Ph.D./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: Ph.D. Grants.
- 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).
– Grants from the Ministerio de Educación y Ciencia del Gobierno de España (SAF2007-6547).

Appendix A. Algorithm of MPFIR0

Initialization

1. The user defines the passband $PB = [f_{pl}, f_{pu}]$, the transition bandwidth (TBW) in the frequency domain and a starting value for the filter order.

2. An estimate of the noise standard deviation $\tilde{\sigma}$ is calculated as the standard deviation of the complex numbers of the last samples of the signal $y(n)$.

3. The frequency f_{NP} of the strongest nuisance peak is automatically determined. The strongest nuisance peak is defined by the peak in the stopband with the highest magnitude, NP_0, in the frequency domain,

$$NP_0 = \max \left| \frac{Y(f_l)}{\sqrt{N}} \right|, \tag{A.1}$$

where $f_l \in [f_{sl}, f_{su}]$ (stopband), N is the number of points in the signal, Y is the signal in the frequency domain (spectrum) after Fourier transformation of the time-domain signal:

$$Y(f_l) = \sum_{n=0}^{N-1} y(n) e^{\frac{-2\pi j n l}{N}}. \tag{A.2}$$

Note that the standard deviation of the noise in the frequency domain is given by

$$\tilde{\sigma}_f = \frac{\tilde{\sigma}}{\sqrt{N}}. \tag{A.3}$$

4. A pair of zeros is placed at $e^{(j2\pi f_{NP})}$ and its conjugate in the z-plane.

5. A grid of frequency-damping domain points is generated such that the distance between the newly computed zeros and the points is sufficiently large to avoid any numerical issue (division by zero): the couples

$$(f_l, d) \text{ such that } f_l \in [-0.5, 0.5], |f_l - f_r| > 0.001, \tag{A.4}$$

f_r are the frequencies of the zeros, $r = 1, 2$ here (in the initialization step). The points in the interval $[-0.5, 0.5]$ are equally spaced (a step of 0.001 is used in the experiments). An initial value for the damping is given in kHz (e.g., $d = 0.008$). This damping value is refined all along the filter design process based on the estimated width of the largest peak of interest after filtering.

Compute h

6. The transfer function of the filter made from the zeros is calculated at each point of the grid

$$H_1(f_l, d) = \prod_{r=1,\ldots,R} \left(1 - z_r^{-1} (e^{-d+j2\pi f_l})^{-1}\right) \tag{A.5}$$

where R is the number of zeros (2 at the first iteration), $z_r = e^{(-d+j2\pi f_r)}$ are the zeros.

7. The target vector is defined as

$$T(f_l) = \begin{cases} e^{-j2\pi f_l D} & \text{if } f_l \in [f_{pl}, f_{pu}] \\ 0 & \text{else} \end{cases} \tag{A.6}$$

where D is the delay. The default value for the delay is $0.73(M-1)$. This value is based on experiments. Note that this delay is refined at the end of the algorithm (see *Delay tuning*, below).

8. The filter coefficients are calculated following the steps:

– The coefficients of the vector h_2 are computed (more details in Appendix B),

$$h_2 = (G^H W G)^{-1} G^H W T_{H_1} \tag{A.7}$$

where

$$G = \begin{bmatrix} 1 & e^{(d-j2\pi f_1)} & \ldots & e^{(d-j2\pi f_1)(M-1)} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & e^{(d-j2\pi f_L)} & \ldots & e^{(d-j2\pi f_L)(M-1)} \end{bmatrix},$$

and L is the number of points on the grid. W is a diagonal matrix with $W_{l,l} = H_1^2(f_l, d)$ for $l = 1, \ldots, L$, and $T_{H_1} = \left[\frac{T(f_1)}{H_1(f_1,d)} \cdots \frac{T(f_L)}{H_1(f_L,d)}\right]'$. The superscript H indicates the conjugate transpose and the superscript $'$ the transpose.

– The coefficients of h_1 are calculated from its roots (zeros).
– The filter is computed:

$$h = h_1 \otimes h_2, \tag{A.8}$$

where $\otimes$ is the convolution operator, and h_1 is the impulse response of the filter in Eq. (A.5).

– The filter is transformed into a maximum-phase FIR filter (see, e.g., [6] for more details).

Filtering the signal

9. The signal y is filtered

$$y_{fil}(n) = \sum_{m=0}^{M-1} h(m) y(n-m), \tag{A.9}$$

and the first $M-1$ points of the filter are discarded to avoid any distortion (see [13] for more details).

Checking the attenuation of the nuisance components

10. The attenuation of the nuisance components is checked:

– Recording the best attenuation:

$$\begin{cases} \text{if } NP_i < NP_{best} & NP_{\text{best}} = NP_i,\ h_{\text{best}} = h \\ \text{else} & \text{nothing} \end{cases} \tag{A.10}$$

where NP_i is the maximal magnitude of the nuisance components in the spectrum at iteration i, h_{best} are the coefficients of the filter which yields the best attenuation of the nuisance components and NP_{best} is the corresponding maximal magnitude of the nuisance components in the spectrum.

– Checking whether the attenuation is under the noise level:

$$\{ \text{if } NP_i < 2.5\tilde{\sigma}_f \quad \text{go to 13.} \tag{A.11}$$

where the value 2.5 results from the fact that the magnitude of the noise $|N(n)|$ over its standard deviation σ follows a chi distribution ($X = |N(n)|/\sigma \sim \chi_2$). $P(X \leqslant x) = 95\%$ gives $x = 2.45$ which is rounded to 2.5 in the algorithm.

– Checking whether the attenuation is improved compared to the last iteration:

$$\begin{cases} \text{if } NP_i < NP_{i-1} & \text{go to 11} \\ \text{else} & \text{go to 12} \end{cases} \tag{A.12}$$

where NP_{i-1} is the maximal magnitude of the nuisance components in the spectrum at iterations $i-1$.

Increasing the number of zeros

11. A new pair of zeros is placed at $e^{(j2\pi f_{NP_i})}$ and its conjugate in the z-plane, where f_{NP_i} is the frequency of the strongest nuisance peak in the filtered spectrum Y_{fil}. Go to 6.

Increasing the filter order

12. The order of the filter is increased by 10 ($M := M + 10$).

- If the filter length M is larger than a maximum value M_{max}, $h = h_{best}$ and go to 13,
- else, keep only the first pair of zeros calculated in the initialization step and go to 6.

Delay tuning

13. Simple search:

a. Initialization: delay step $D_{step} = \min(1 - D_{fact}, D_{fact} - 0)/5$, where the delay factor is the one fixed from the beginning, e.g., $D_{fact} = 0.73$.

b. The passband is divided into 4 disjoint equal-size intervals. The intervals I_2 and I_3 around the center of the passband are used to calculate the difference:

$$Df_0 = \left| \sum_{f_{l_2}} H(f_{l_2}, d) - \sum_{f_{l_3}} H(f_{l_3}, d) \right| \tag{A.13}$$

where H is the magnitude response of the filter, and f_{l_2} and f_{l_3} are the frequencies in I_2 and I_3.

c. Define 2 new delays around D_{i-1} ($D_0 = 0.73(M-1)$, for example), the delay at iteration $i-1$: $D_{i_1} = D_{i-1} - D_{step}$ and $D_{i_2} = D_{i-1} + D_{step}$, and recompute H for each delay.

d. Compute Df such as in Eq. (A.13) for each delay (Df_1 and Df_2),

$$Df_{best_i} = \min(Df_1, Df_2, Df_{best_{i-1}}), \tag{A.14}$$

$$\begin{cases} \text{if } Df_{best_i} = Df_1, & D_i = D_{i_1} \\ \quad \text{if } D_{i-1} = D_{(i-1)_1}, & D_{step} := D_{step}/1.1, \text{ else } D_{step} := D_{step}/2, \\ \quad \text{go to } c. & \\ \text{elseif } Df_{best_i} = Df_2, & D_i = D_{i_2} \\ \quad \text{if } D_{i-1} = D_{(i-1)_2}, & D_{step} := D_{step}/1.1, \text{ else } D_{step} := D_{step}/2, \\ \quad \text{go to } c. & \\ \text{else} & \text{go to END}, \end{cases} \tag{A.15}$$

where $Df_{best_0} = Df_0$, $D_{(i-1)_1}$ and $D_{(i-1)_2}$ are equivalent notations to D_{i_1} and D_{i_2} but at iteration $i-1$ instead of iteration i.

Appendix B. More details on Eq. (A.7)

We want to minimize the least squares cost function:

$$\min_{H} \sum_{f_l=f_1}^{f_L} (H(f_l, d) - T(f_l))^2 \tag{B.1}$$

$$\min_{H_2} \sum_{f_l=f_1}^{f_L} (H_1(f_l, d) H_2(f_l, d) - T(f_l))^2 \tag{B.2}$$

$$\min_{H_2} \sum_{f_l=f_1}^{f_L} H_1^2(f_l, d) \left(H_2(f_l, d) - \frac{T(f_l)}{H_1(f_l, d)} \right)^2 \tag{B.3}$$

$$\min_{h_2} \sum_{f_l=f_1}^{f_L} H_1^2(f_l, d) \left(\sum_{m=0}^{M-1} h_2(m) e^{(d-j2\pi f_l)m} - \frac{T(f_l)}{H_1(f_l, d)} \right)^2 \tag{B.4}$$

This yields in matrix notation,

$$G^H W G h_2 = G^H W T_{H_1}, \tag{B.5}$$

and finally gives Eq. (A.7). The inverse of the matrix $G^H WG$ may generate numerical problems depending on the step between two adjacent points of the grid, (f_l, d) and (f_{l+1}, d), $l = 1, \ldots, L-1$. $A = G^H WG$ is an Hermitian matrix, which means that $A^H = A$. Indeed, all the elements $W_{i,i}$ of W are real positive since $H_1(f_l, d)$ is real for all f_l due to the fact that each zero has its conjugate (see step 4 in the algorithm), and the diagonal elements of A are real positive. With a step of 0.001, the experiments showed that A was positive definite (all eigenvalues are positive) and therefore invertible. In that case, the Cholesky decomposition is recommended to solve the least squares problem. In case of very small steps, A becomes ill-conditioned and orthogonalization techniques may be needed. One satisfactory method is the row-oriented modified gram-schmidt method with reorthogonalization proposed by Dax [22]. The condition number of A can be computed in order to select the right method to solve the least squares problem.

Other numerical problems can occur when the distance between the selected zeros z_r and the points of the grid (f_l, d), $l = 1, \ldots, L$ is too small since in that case $H_1(f_l, d)$ will be close to 0 and $\frac{T(f_l)}{H_1(f_l, d)}$ close to ∞. Choosing grid points too far away from the zeros z_r will generate large ripples in the region where no constraints are imposed (between z_r and the grid points).

References

[1] E. Prost, P. Sizun, M. Piotto, J.M. Nuzillard, A simple scheme for the design of solvent-suppression pulses, J. Magn. Reson. 159 (1) (2002) 76–81.

[2] A.J. Simpson, S.A. Brown, Purge NMR: effective and easy solvent suppression, J. Magn. Reson. 175 (2) (2005) 340–346.

[3] M.A. Smith, J. Gillen, M.T. McMahon, P.B. Barker, X. Golay, Simultaneous water and lipid suppression for in vivo brain spectroscopy in humans, Magn. Reson. Med. 54 (3) (2005) 691–696.

[4] A. Coron, L. Vanhamme, J. Antoine, P. Van Hecke, S. Van Huffel, The filtering approach to solvent peak suppression in MRS: a critical review, J. Magn. Reson. 152 (1) (2001) 26–40. http://dx.doi.org/10.1006/jmre.2001.2385.

[5] T. Laudadio, N. Mastronardi, L. Vanhamme, P. Van Hecke, S. Van Huffel, Improved Lanczos algorithms for blackbox MRS data quantitation, J. Magn. Reson. 157 (2) (2002) 292–297.

[6] T. Sundin, L. Vanhamme, P. Van Hecke, I. Dologlou, S. Van Huffel, Accurate quantification of ^{1}H spectra: from finite impulse response filter design for solvent suppression to parameter estimation, J. Magn. Reson. 139 (2) (1999) 189–204.

[7] L. Vanhamme, T. Sundin, P. Van Hecke, S. Van Huffel, R. Pintelon, Frequency-selective quantification of biomedical magnetic resonance spectroscopy data, J. Magn. Reson. 143 (1) (2000) 1–16. http://dx.doi.org/10.1006/jmre.1999.1960.

[8] F. Abilgaard, H. Gesmar, J. Led, Quantitative analysis of complicated nonideal Fourier transform NMR spectra, J. Magn. Reson. A 79 (1988) 78–89.

[9] Y. Hiltunen, M. Ala-Korpela, J. Jokisaari, S. Eskelinen, K. Kiviniitty, M. Savolainen, Y.A. Kesniemi, A lineshape fitting model for ^{1}H NMR spectra of human blood plasma, Magn. Reson. Med. 21 (2) (1991) 222–232.

[10] A.A. de Graaf, W.M. Bove, Improved quantification of in vivo ^{1}H NMR spectra by optimization of signal acquisition and processing and by incorporation of prior knowledge into the spectral fitting, Magn. Reson. Med. 15 (2) (1990) 305–319.

[11] S.W. Provencher, Estimation of metabolite concentrations from localized in vivo proton NMR spectra, Magn. Reson. Med. 30 (6) (1993) 672–679.

[12] K. Young, B.J. Soher, A.A. Maudsley, Automated spectral analysis II: application of wavelet shrinkage for characterization of non-parameterized signals, Magn. Reson. Med. 40 (6) (1998) 816–821.

[13] J. Poullet, D.M. Sima, S. Van Huffel, P. Van Hecke, Frequency-selective quantitation of short-echo time ^{1}H magnetic resonance spectra, J. Magn. Reson. 186 (2) (2007) 293–304. http://dx.doi.org/10.1016/j.jmr.2007.03.015.

[14] J. Poullet, D.M. Sima, A.W. Simonetti, B. De Neuter, L. Vanhamme, P. Lemmerling, S. Van Huffel, An automated quantitation of short echo time MRS spectra in an open source software environment: AQSES, NMR Biomed. 20 (5) (2007) 493–504.

[15] H. Ratiney, M. Sdika, Y. Coenradie, S. Cavassila, D. van Ormondt, D. Graveron-Demilly, Time-domain semi-parametric estimation based on a metabolite basis set, NMR Biomed. 18 (1) (2005) 1–13.

[16] L. Vanhamme, A. van den Boogaart, 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.

[17] J.W. van der Veen, R. de Beer, P.R. Luyten, D. van Ormondt, Accurate quantification of in vivo ^{31}P NMR signals using the variable projection method and prior knowledge, Magn. Reson. Med. 6 (1) (1988) 92–98.

[18] M. van der Graaf, M. Julià-Sapé, F.A. Howe, A. Ziegler, C. Majós, A. Moreno-Torres, M. Rijpkema, D. Acosta, K.S. Opstad, Y.M. van der Meulen, C. Arús, A. Heerschap, MRS quality assessment in a multicentre study on MRS-based classification of brain tumours, NMR Biomed. 21 (2) (2008) 148–158. http://dx.doi.org/10.1002/nbm.117.
[19] C. Elster, F. Schubert, A. Link, M. Walzel, F. Seifert, H. Rinneberg, Quantitative magnetic resonance spectroscopy: Semi-parametric modeling and determination of uncertainties, Magn. Reson. Med. 53 (2005) 1288–1296.
[20] W.W.F. Pijnappel, A. van den Boogaart, R. de Beer, D. van Ormondt, SVD-based quantification of magnetic resonance signals, J. Magn. Reson. 97 (1) (1992) 122–134.
[21] J. Slotboom, C. Boesch, R. Kreis, Versatile frequency domain fitting using time domain models and prior knowledge, Magn. Reson. Med. 39 (6) (1998) 899–911.
[22] A. Dax, A modified Gram–Schmidt algorithm with iterative orthogonalization and column pivoting, Linear Algebra Appl. 310 (2000) 25–42.