Frequency-selective quantitation of short-echo time ^{1}H magnetic resonance spectra

Jean-Baptiste Poullet [a,*], Diana M. Sima [a], Sabine Van Huffel [a], Paul Van Hecke [b]

[a] Department of Electrical Engineering, SCD-SISTA, Katholieke Universiteit Leuven, Kasteelpark Arenberg 10, 3001 Leuven, Belgium
[b] Biomedical NMR Unit, Katholieke Universiteit Leuven, Gasthuisberg, 3000 Leuven, Belgium

Abstract

Accurate and efficient filtering techniques are required to suppress large nuisance components present in short-echo time magnetic resonance (MR) spectra. This paper discusses two powerful filtering techniques used in long-echo time MR spectral quantitation, the maximum-phase FIR filter (MP-FIR) and the Hankel-Lanczos Singular Value Decomposition with Partial ReOrthogonalization (HLSVD-PRO), and shows that they can be applied to their more complex short-echo time spectral counterparts. Both filters are validated and compared through extensive simulations. Their properties are discussed. In particular, the capability of MP-FIR for dealing with macromolecular components is emphasized. Although this property does not make a large difference for long-echo time MR spectra, it can be important when quantifying short-echo time spectra.
© 2007 Elsevier Inc. All rights reserved.

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)

1. Introduction

Efficient and accurate quantitation of metabolites from short-echo time *in vivo* MR spectroscopy (MRS) might be an important aid in the correct non-invasive diagnosis of human brain pathology. Quantitation of short-echo time MR spectra provides more metabolite information than long-echo time spectra, but is hampered by broad baseline signal contributions, resonance line-shape distortions, low signal-to-noise ratios (SNR) and overlapping peaks in the frequency domain (see, *e.g.* [1]). Therefore, in order to obtain accurate parameter estimates, it is recommended to disentangle the metabolite contributions from the unwanted components. Numerous methods have been developed to address this issue. In this paper, we focus on two filtering techniques widely used with long-echo time

MR spectra for solvent suppression: the maximum-phase FIR filter (MP-FIR) [2] and the Hankel-Lanczos Singular Value Decomposition with Partial ReOrthogonalization (HLSVD-PRO) [3]. Since the differences between long-echo and short-echo time MR spectra are substantial, the appropriateness of these filters for short-echo time spectra should be validated. Indeed, to fully understand the behavior of these filters it is imperative to study them in close connection with the type of signals they are applied to. Some characteristics of the filters are more important for short-echo time spectra than long-echo time spectra. Furthermore, since the final goal is to obtain accurate parameter estimates, a quantitation method is needed to validate and compare the filters. AQSES, Automated quantitation of short-echo time MRS spectra (see [4] and references therein), has been chosen as quantitation method for the reasons described below. The goal of the paper is therefore twofold: first, to validate both filtering techniques for short-echo time spectra, and second, to show the extra value of MP-FIR over HLSVD-PRO for this type of spectra.

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

The paper is organized as follows. First, we discuss briefly the choice of the filtering techniques, MP-FIR and HLSVD-PRO, and the quantitation method AQSES. Then, the properties of both filters are described and illustrated on simple examples. The influence of the filters on the estimated parameters is studied through extensive numerical studies. Finally, the advantages of MP-FIR over HLSVD-PRO are pointed out and the combined use of MP-FIR with other quantitation methods is discussed.

Why MPFIR and HLSVD-PRO? Measurement sequences fail to suppress completely the water resonance without affecting the metabolites of interest. Therefore, the so-called water suppressed signals contain a residual water contribution which cannot be depicted by an analytical formulation. Time-domain quantitation methods (see, e.g., [4–8]) cannot be used without first removing the water components, while frequency-domain methods (see, e.g., [9–13]) can be applied directly to the water suppressed signal, but should consider the water tails as part of the macromolecular baseline. Water suppression methods should not influence the parameter estimates and should have a low computational complexity. They can be categorized in two main groups: on one hand, the convolution-based methods [2,14–19] convolve the original signal with the coefficients of a filtering window, and, on the other hand, the SVD-based methods [3,20,21] make use of the singular value decomposition of a Hankel matrix. In their review, Coron et al. [22] compare five convolution-based methods: the Gabor transform based method [17], the method by Marion et al. [14], the filtering method of Sodano and Delepierre [15], the highpass butterworth filter described by Cross [16], and the maximum phase finite impulse response filter (MP-FIR) by Sundin et al. [2] which we study in this paper. Among them, MP-FIR is the most accurate and efficient method for quantifying long-echo time MRS spectra. In addition, MP-FIR allows the inclusion of prior knowledge that may be taken into account during quantitation (see [22] for more details). In parallel, numerous SVD-based methods have shown to be successful in removing the water resonances, the most common method being HLSVD developed by Pijnappel et al. [20] which reduces the computational load of the original HSVD method (Barkhuijsen et al. [23]). From HLSVD, several variants have been developed (see, e.g., [3,24]). In [3], Laudadio et al. compare HLSVD with two other proposed variants: the method based on the Lanczos algorithm with Partial ReOrthogonalization (HLSVD-PRO) and the method based on the Implicitly Restarted Lanczos Algorithm (HLSVD-IRL [24]). HLSVD-PRO and HLSVD-IRL outperform HLSVD in terms of computational efficiency and numerical reliability. Moreover, HLSVD-PRO is faster than HLSVD-IRL. MP-FIR and HLSVD-PRO have been selected since they are the most successful filters for long-echo time MR spectra and they are based on two different approaches. Although, the main objective of these filters is to remove the water resonances, we will show that MP-FIR can have a substantial effect on the baseline which

can result in better parameter estimates. Nevertheless, MP-FIR will not compete with baseline correction methods but, combined with the latter, might improve the final estimates. In this paper, we study three types of nuisance components –noise, baseline and water resonances—and the effect of the filters on these components. The choice of AQSES for validating both filtering techniques can easily be explained since

- AQSES is a time-domain quantitation method,
- both filtering techniques can easily be combined with AQSES,
- AQSES has been designed for short-echo time MR spectra, in contrast with AMARES [25], for example, typically designed for long-echo time MR spectra.

2. Theory

The goal of this section is to discuss the properties of MP-FIR and HLSVD-PRO using their analytical form in the specific context of short-echo time MR spectroscopy. The use of these filters depends on the quantitation method. For example, MP-FIR is applied to the original and estimated signals (sum of Lorentzians) in $AMARES_W$ [26], while it is applied to the original signal and each corrected metabolite profile in AQSES (see below). Note that this section is applicable to any quantitation method that is based on the model described in Eq. (1).

2.1. Model

In time-domain quantitation methods such as AQSES or QUEST [7], the short-echo time MRS signal y is modeled in the time domain as

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

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), ε_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.

The main goal of both filters is to filter out the water component w which has resonance frequencies in a known frequency interval. This interval is disjoint to the frequency region (or interval) of interest where the metabolite resonances are located. Quantitation methods based on iterative optimization procedures can reasonably combine filtering methods only if these methods are fast, especially when numerous signals have to be processed (e.g. with MRS imaging data). The computational load of SVD-based methods is known to be much larger than the one of simple convolution methods like MP-FIR (see, e.g., [2]). Moreover, the filter coefficients of MP-FIR are computed only once prior to the optimization procedure. The same filter coefficients are then used all along the iterative optimization procedure such that MP-FIR boils down to a simple matrix multiplication, which is considerably faster than calculating the filter coefficients. To be used in the optimization procedure, HLSVD-PRO would require an SVD decomposition of a large matrix [3] at each iteration. A matrix multiplication is much faster than an SVD decomposition, making HLSVD-PRO much less attractive for being used in the iterative optimization procedure.

2.2. MP-FIR

MP-FIR is a maximum-phase FIR filter based on the constrained least squares design of FIR filters proposed by Selesnick et al. [27], who does not specifically consider the case of complex damped exponentials. Nevertheless, the use of a maximum-phase FIR filter instead of a general FIR filter is motivated by the assumption of such a model (sum of complex damped exponentials) as detailed below. Moreover, the filter length is optimized with respect to the noise and the water resonances of the signal under processing. All components in the frequency region of no interest are attenuated such that the largest magnitude in that region is smaller than twice the standard deviation of the noise. Although the filter coefficients are optimized before starting the iterative optimization for quantitation, MP-FIR is applied to the MRS signal and each *corrected* metabolite profile $\{y_k\}_{k=1,\dots,K}$ in each iteration of the optimization procedure. In order to avoid signal distortion, the first $M-1$ points are discarded from the filtered signal [2], $M-1$ being the filter order.

Analytical validation of MP-FIR for short-echo time MRS quantitation. MP-FIR was initially designed for long-echo time MRS quantitation [2]. This section shows that this filter can be applied to short-echo time MR spectra as well.

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

$$y_{\text{fil}}(n) = \sum_{m=0}^{M-1} h_m y(n-m), \tag{4}$$

where $\{h_m\}_{m=0,\dots,M-1}$ are the constant (possibly complex) filter coefficients. By truncating the distorted first

$M-1$ data points[1] of this filtered signal, with $n = 0,\dots,N-M$, $\hat{y}_{\text{fil}}(n)$ can be expressed as

$$
\begin{aligned}
\hat{y}_{\text{fil}}(n) &= \sum_{m=0}^{M-1} h_m \sum_{k=1}^{K} \hat{y}_k(n-m+M-1) \\
&= \sum_{m=0}^{M-1} h_m \sum_{k=1}^{K} \alpha_k \zeta_k^{n-m+M-1} \sum_{p=1}^{P_k} \alpha_{k,p} \zeta_{k,p}^{n-m+M-1} \\
&= \sum_{m=0}^{M-1} h_m \sum_{k=1}^{K} \sum_{p=1}^{P_k} \alpha_k \zeta_k^n v_k(n) \zeta_k^{-m+M-1} \zeta_{k,p}^{-m+M-1} \\
&= \sum_{k=1}^{K} \hat{y}_k(n) \sum_{m=0}^{M-1} \sum_{p=1}^{P_k} h_m \zeta_k^{-m+M-1} \zeta_{k,p}^{-m+M-1},
\end{aligned}
\tag{5}
$$

where we assume that each metabolite profile $v_k(n)$ can be modeled by a sum of Lorentzians, i.e., $v_k(n) = \sum_{p=1}^{P_k} \alpha_{k,p} \zeta_{k,p}^n$, P_k being the number of Lorentzians used to model the metabolite profile k. The linear part, $\alpha_{k,p}$, and the non-linear part, $\zeta_{k,p}$ are defined similarly as in Eq. (2). Sundin et al. [2] showed that each individual Lorentzian (or by extension each corrected Lorentzian) lying in the frequency region of interest is not distorted by the filter if the condition of no distortion of Lorentzian (p,k) is written as

$$|\bar{h}\bar{\zeta}_{k,p}| = 1 \tag{6}$$

where

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

and

$$\bar{\zeta}_{k,p} = (1\ e^{(-d_{k,p}+j2\pi f_{k,p})\Delta t} \dots e^{(M-1)(-d_{k,p}+j2\pi f_{k,p})\Delta t})^{\text{T}}. \tag{8}$$

Selesnick et al. [27] provided tools for calculating the coefficients h_m, which ensure that the magnitude response of $\bar{h}$ (i.e., $|\bar{h}(1\ e^{j2\pi f_{k,p}\Delta t} \dots e^{(M-1)(j2\pi f_{k,p})\Delta t})^{\text{T}}|$) is approximately equal to 1 in the frequency region of interest and equal to zero elsewhere. If we use the linear FIR filter by Selesnick, the components with higher damping factors $d_{k,p}$ will undergo a smaller gain (i.e., smaller $|\bar{h}\bar{\zeta}_{k,p}|$) than the other ones resulting in signal distortion as described in [2]. In order to reduce the effect of this distortion vector, Sundin et al. [2] proposed to transform Selesnick's FIR filter into a maximum-phase FIR filter, moving most of the energy towards the first coefficients of $\bar{h}$ (i.e., h_{M-1}, h_{M-2},…). The same principle is applicable for short-echo time MR spectra.

To understand the effect of the filter, we can compare the filtered signal to the original signal in terms of signal-to-baseline and signal-to-noise ratios. The baseline components are assumed to have large damping values compared to the metabolite peaks. We neglect the fact that y_{fil} is limited to $N-M+1$ points and not N. Each Lorentzian (k,p) is thus multiplied by the complex number $\bar{h}\bar{\zeta}_{k,p}$. The term $|\bar{h}\bar{\zeta}_{k,p}|$ decreases exponentially (or linearly in a logarithmic scale) with the damping values as illustrated in Fig. 1,

---

[1] They are distorted since they assume zero values for the unknown $y(-1),\dots,y(-M+1)$.

resulting in an *increased signal-to-baseline ratio*. The broader the passband, the more stable $|\bar{h}\bar{\zeta}_{k,p}|$ with respect to the damping coefficients. The influence of the damping is thus mainly observable for narrow-passband filters. The noise will not be attenuated (or negligibly) by the filter in the region of interest, resulting in a *decreased signal-to-noise ratio*. Choosing the broadest possible passband seems to be reasonable. However, the signal-to-baseline ratio will be higher for narrow-passband filters (see Example 1 below). The choice of the filter region will result in a trade-off between the increase of the signal-to-baseline ratio and the decrease of signal-to-noise ratio. In other words, narrow-passband filters will be preferred for signals with large baseline components, while broad-passband filters will be more indicated for signals with small baseline components like long-echo time MR spectra.

Example 1. `[referring to Fig. 1] Considering a sampling frequency of 1 kHz, an individual Lorentzian with` $d_k = 0.016$ `will be attenuated by ±0.9 dB (≈10% loss in magnitude) by a narrow-passband filter ([0.25,4.2] ppm), while it will be attenuated by ±0.25 dB (≈3% loss in magnitude) by a broad-passband filter ([−3,7] ppm). Note that this Lorentzian would be much more attenuated by a linear-phase filter: 5.6 dB (≈48% loss in magnitude) for [0.25,4.2] and 4.9 dB (≈43% loss in magnitude) for [−3,7]. This shows why a linear-phase filter is not suitable for Lorentzians.`

2.3. HLSVD-PRO

HLSVD-PRO is a subspace-based method for modeling a sum of exponentially damped sinusoids (Lorentzians), which can be used as a frequency-selective filter.
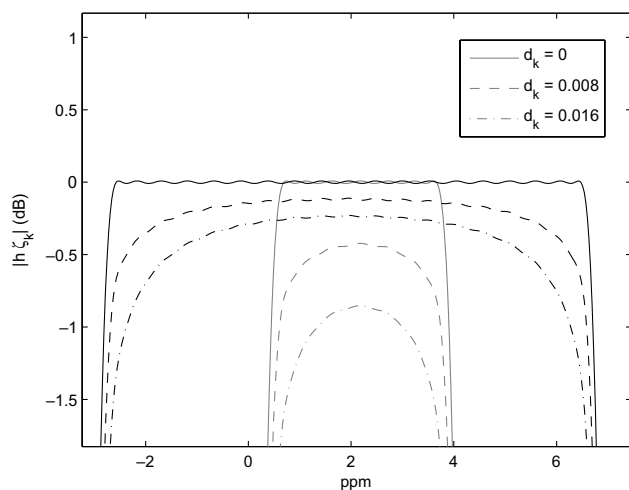

Fig. 1. Magnitude of $|\bar{h}\bar{\zeta}_{k,p}|$ as a function of the frequency for different damping coefficients d_k (in kHz*rad) and different filtering regions (in black for the frequency interval [0.25,4.2] ppm and in gray for [−3,7] ppm). The filter order is 70.

HLSVD-PRO is applied, before entering the iterative quantitation procedure, to the MRS signal, as well as to each metabolite profile included in the database. HLSVD-PRO models the original signal and the metabolite profiles such that

$$\hat{y}(n) = \sum_{i=1}^{L_0} \alpha_{0,i}\zeta_{0,i}^n = \sum_{i=1}^{W_0} \alpha_{0,j(i)}\zeta_{0,j(i)}^n + \sum_{i=W_0+1}^{L_0} \alpha_{0,j(i)}\zeta_{0,j(i)}^n$$
$$\hat{v}_k(n) = \sum_{i=1}^{L_k} \alpha_{k,i}\zeta_{k,i}^n = \sum_{i=1}^{W_k} \alpha_{k,j(i)}\zeta_{k,j(i)}^n + \sum_{i=W_k+1}^{L_k} \alpha_{k,j(i)}\zeta_{k,j(i)}^n \tag{9}$$

where $j(i)$, $i = 1,\ldots,L_0$ (respectively, L_k) is a permutation of $i = 1,\ldots,L_0$ (respectively, L_k) such that the indices of Lorentzians with central frequency inside the water resonance region are the first W_0 (respectively, W_k) ones. L_0 (respectively, L_k) is the model order of HLSVD-PRO applied to the original signal (respectively, metabolite profile k), W_0 (respectively, W_k) is the number of Lorentzians falling within the water resonance region when applying HLSVD-PRO to the original signal (respectively, metabolite profile k). The Lorentzians located in the frequency region of no interest are then subtracted from the original signal y, as well as from the initial metabolite profiles v_k, $k = 1,\ldots,K$, such that

$$y_{\text{fil}}(n) = y(n) - \sum_{i=1}^{W_0} \alpha_{0,j(i)}\zeta_{0,j(i)}^n,$$
$$\hat{y}_{\text{fil}}(n) = \sum_{k=1}^{K} \alpha_k \zeta_k^n \left(v_k(n) - \sum_{i=1}^{W_k} \alpha_{k,j(i)}\zeta_{k,j(i)}^n \right), \tag{10}$$

where $\hat{y}_{\text{fil}}$ denotes the estimate of the filtered MRS signal. The filtered model $\hat{y}_{\text{fil}}$ is used to fit the filtered signal y_{fil}.

Besides being relatively slow compared to MP-FIR, the major problem with HLSVD-PRO is that it assumes a Lorentzian model which is optimistic given the random nature of the remaining water resonances. This results in a *decreased signal-to-baseline ratio*, which can be illustrated by Example 2. The *signal-to-noise ratio* is however kept *constant*.

Example 2. `[referring to Fig. 2] Suppose a signal containing three peaks, 2 Lorentzians at 2 and 3.5 ppm for the metabolites and 1 Gaussian at 4.7 ppm for the water peak, as plotted in Fig. 2a. We apply HLSVD-PRO on the frequency region [0.25 4] ppm. The model order was set at 4 since at least 2 Lorentzians are used for modeling a Gaussian. In Fig. 2b, we reconstruct the water components that have been removed by the filter. A bump is clearly visible in the transition region between the water and the metabolites. This bump is due to the nature of the model used in HLSVD-PRO. Increasing the model order would help to reduce the water component but would not help to substantially reduce this "bump effect". Consequently, the residual water tails will`
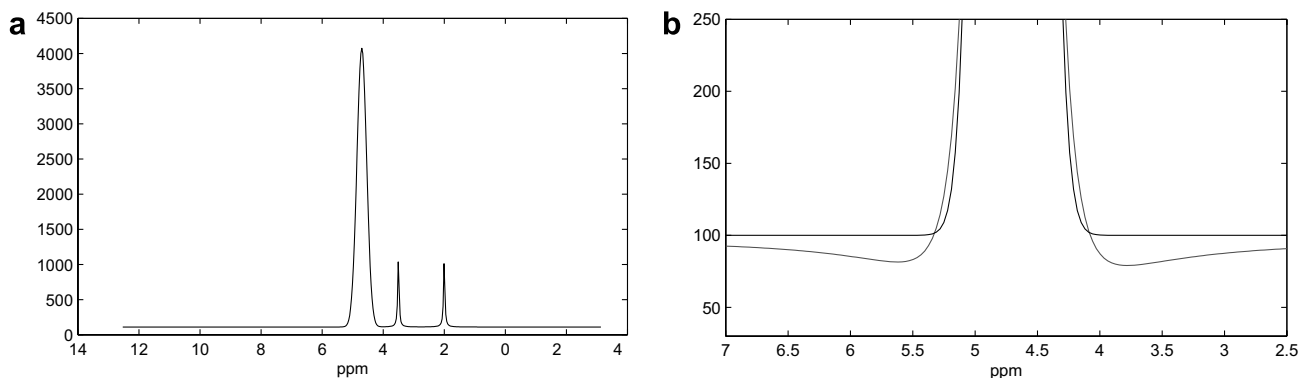
Fig. 2. (a) Real part of the spectrum. (b) Real part of the water components in the frequency domain (true components in black and estimated ones in gray). Bump effect illustrated by applying HLSVD-PRO on a signal with a Gaussian water peak.

Table 1
Example 2: Amplitude estimates obtained with AQSES in combination with HLSVD-PRO or MP-FIR (last column)

Lorentzian at	True	HLSVD-PRO	MP-FIR
2 ppm	10	10.27	10.04
3.5 ppm	10	10.69	9.96

The true amplitudes are 10 for each Lorentzian.

```
be seen as baseline components overlapping
with the frequency region of interest, result-
ing in a decreased signal-to-baseline ratio.
    To get an idea of the influence of this bump
effect on the amplitude estimates, we apply
AQSES to the filtered signal. The results are
shown in Table 1. For comparison, the results of
MP-FIR are also shown. The residual water tail
affects the parameter estimates, resulting in
a 2.7% error for the peak at 2 ppm and 6.9% error
for the peak at 3.5 ppm, while the errors are
less than 0.5% when using MP-FIR.
```

3. Materials and methods

The properties of the filters are illustrated on simple examples in the previous section. The goal of the following sections is to compare these filters on more complex and more realistic signals. Their robustness is studied with respect to the choice of a variety of nuisance components and filtering regions. An *in vivo* experiment shows the influence of these filters on the sensitivity of the quantitation algorithm.

3.1. Database and simulated signals

The database used in AQSES is identical to the one used in [4] with 6 metabolites + 2 lipids: Myo-inositol (Myo), Phosphorylcholine (PCh), Creatine (Cr), Glutamate (Glu), N-acetylaspartate (NAA), Lactate (Lac), Lipid at 1.3 ppm (Lip1), Lipid at 0.9 ppm (Lip2). Simulated data were generated to compare the two filtering methods, HLSVD-PRO

and MP-FIR. A signal free from nuisance components (except for the reference peaks at 0 and 8.44 ppm which are 3-trimethylsilyl-1-propane-sulfonic acid and formate, respectively) has been chosen from set 1 in [4]. This signal was quantified perfectly (*i.e.*, no error in amplitude estimation) with AQSES. This guarantees that all estimation errors are due to nuisance components. Although AQSES was applied on the unphased signal, we display the phased signal in Fig. 3 for a better signal assessment. For sake of clarity and space, only one signal (*i.e.*, one set of parameters a_k, ϕ_k, d_k and f_k) have been analyzed. The true amplitude values are given in Table 2.

Although different signal parameter values will lead to other results, one can expect that the general trends (*i.e.*, limitations and potentials of each filtering technique) are preserved. The nuisance components have been added to this signal to generate four different sets of signals as follows:

- set 1 = signal + low noise
- set 2 = signal + low noise + high baseline + water
- set 3 = signal + high noise
- set 4 = signal + high noise + high baseline + water

The baseline distortion was based on information from Table 1 in [28]; the baseline is the sum of Gaussians referred to as lip3, lip4, lip5, mm2, mm3 and mm4 in that
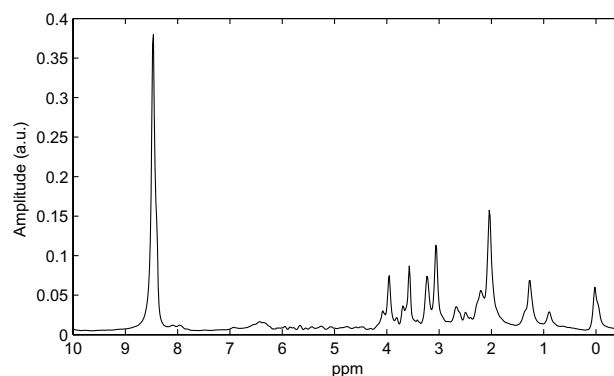

Fig. 3. Real part of the phased spectrum chosen from set 1 in [4].

Table 2
True amplitudes in arbitrary units used to construct the signal in Fig. 3

NAA	Myo	Cr	Pch	Glu	Lac	Lip1	Lip2
40.60	13.40	27.12	9.64	29.51	16.65	0.21	0.10

paper. This baseline b overlaps with the frequency region of the metabolites. The water profile has been extracted from a water suppressed *in vivo* spectrum by means of HLSVD-PRO. Low and high noise levels correspond to SNR=300 and SNR=75, respectively. The SNR is defined as the ratio of the reference peak height at 8.44 ppm and the standard deviation of the circular white gaussian noise, both in the frequency domain. Each set contains 256 simulated spectra. As illustration, one signal from set 4 is plotted in Fig. 4 (the signal phase is corrected).

3.2. Methodology

In order to test the sensitivity of each method with respect to the frequency bounds of the nuisance region, we defined two filtering regions: [0.25,4.2], [0.25,4.5] in ppm. The frequency region close to the water resonance at 4.7 ppm being more sensitive than the other one, we limited our study to the variations of the bounds for that region.

The effects of the nuisance components are tested with 6 model settings, which differ from each other regarding three options: choice between HLSVD-PRO or MP-FIR, the use of a baseline in the model or not, and the bounds of the region to be filtered. Note that these options can be easily defined in AQSES-GUI, the graphical user interface of AQSES [4]. The regularization parameter λ, which controls the degree of smoothness of the baseline, has been defined manually. This parameter is used only when the baseline is included in the model (as a linear combination of spline basis functions). For each set, the best value of λ (*i.e.*, providing the amplitude estimates closest to the true amplitudes) was chosen for each filtering technique, resulting in 2 * 4 = 8 values of λ for all sets.
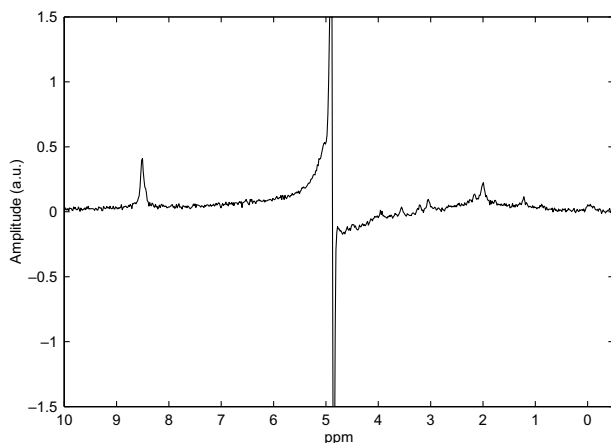
The model settings were defined as follows:

- Model setting 1: Use of the **baseline** in the model and **MP-FIR** filtering in [0.25,**4.2**] ppm
- Model setting 2: Use of the **baseline** in the model and **MP-FIR** filtering in [0.25,**4.5**] ppm
- Model setting 3: **No** use of **baseline** in the model and **MP-FIR** filtering in [0.25,**4.2**] ppm
- Model setting 4: Use of the **baseline** in the model and **HLSVD-PRO** filtering in [0.25,**4.2**] ppm
- Model setting 5: Use of the **baseline** in the model and **HLSVD-PRO** filtering in [0.25,**4.5**] ppm
- Model setting 6: **No** use of **baseline** in the model and **HLSVD-PRO** filtering in [0.25,**4.2**] ppm

Each set defined in Section 3.1 was quantified using each model setting. In order to compare the results obtained with the different model settings, we use the relative root mean square error (RRMSE), defined as

$$RRMSE_k = 100 \sqrt{\frac{1}{L} \sum_{l=1}^{L} \frac{(a_k - \tilde{a}_{k,l})^2}{a_k^2}}, \tag{11}$$

where a_k (respectively, $\tilde{a}_{k,l}$) is the true (respectively, estimated) amplitude for metabolite profile k, l refers to the lth simulation and L is the total number of simulations within each set, *i.e.*, 256. The model order used in HLSVD-PRO was fixed at 25 as recommended in [29].

3.3. In vivo experiment

The performances of MP-FIR and HLSVD-PRO with respect to the regularization parameter λ are tested on an *in vivo* MR proton spectrum of normal brain tissue, acquired on a 1.5 T Philips NT Gyroscan using a PRESS sequence with an echo time of 23 ms, and a PRESS box of $4 \times 3 \times 3$ cm^3. This water suppressed signal was eddy current corrected using Klose's method [30] (see Fig. 5). Note that other methods for eddy current correction such as QUECC [31] could also be used.

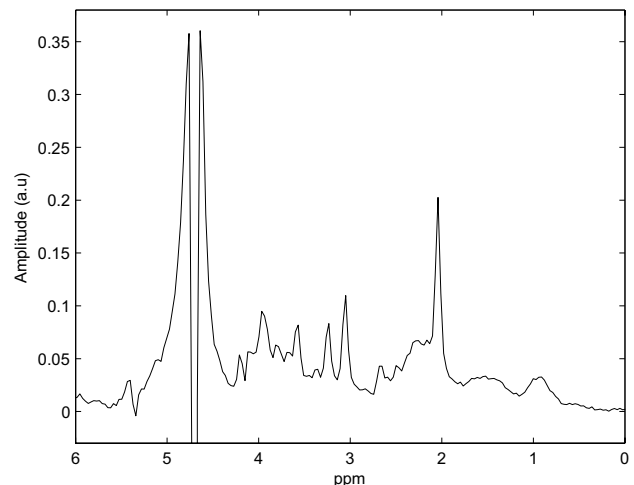Fig. 4. Real part of the phased spectrum from set 4.

Fig. 5. Real part of the *in vivo* spectrum corrected for eddy current effects.

The database used to fit the signal is identical to the one used for the simulated data. Glu is used to estimate the contribution of glutamate and glutamine and will be denoted Glx and PCh for the contribution of choline compounds.

4. Results

Representative fits of sets 1 and 4, obtained by using model settings 1 and 4, are given as illustration in Fig. 6.

The results are shown in Fig. 7. Each subfigure corresponds to one specific set. Model setting 1, where the baseline is modeled and MP-FIR is used, provides the best results with a RRMSE that remains under 50% for all metabolites and all sets. In most of the cases, the RRMSEs are larger when using HLSVD-PRO instead of MP-FIR with some values larger than 100%, especially for lower concentration metabolites such as Lip1 and Lip2.

Including the *baseline* into the model is especially interesting when the signal baseline is large even if we observe relatively large errors for sets 2 and 4. Cr and Myo are less affected by the addition of the baseline than the other metabolites such as Glu. Cr is known to be less correlated with the baseline [7] and Myo shows up in a frequency

range where the baseline is less present (see, *e.g.*, [28]). The profile of Glu is relatively widely spread in the frequency domain. Therefore, in the presence of a strong baseline, Glu tends to fit a part of the baseline resulting in large errors for this metabolite. The differences between HLSVD-PRO and MP-FIR are larger when the signal contains a strong baseline. Indeed, strong differences between HLSVD-PRO and MP-FIR appear in set 1 only for Lac, Lip1 and Lip2, while in sets 2 and 4 with high baseline there are in addition also visible differences between HLSVD-PRO and MP-FIR for Cr, Glu and NAA. Comparing the bottom plots in Fig. 6 corresponding to set 4, one can see that a large part of the baseline has been filtered out by MP-FIR while it is still present in the filtered signal (in light gray) when using HLSVD-PRO. The signal-to-baseline ratio is increased as predicted in Section 2.2, resulting in better amplitude estimates. Modeling the baseline seems to be useful when using HLSVD-PRO even if no baseline is added to the signal (see Myo, PCh, Cr, Glu and NAA in Fig. 7a). This might be due to the difficulty of HLSVD-PRO to model the water peak with pure Lorentzians. Another explanation can be the frequency misalignments between the original signal and the metabolite profiles as described in the discussion section.

The *noise* mainly affects the metabolites of lower concentration such as Lip1 and Lip2. If we only look at set 3 (*i.e.*, signals with high noise and without water nor baseline), all model settings give similar results, thus no model setting seems to be preferable over the others. The RRMSEs of metabolites of larger concentration (NAA, Cr, Glu) turn out to be similar between sets 1 and 3 for model setting 6 (error mainly explained by this model setting), but much better for other model settings. When comparing sets 2 and 4 (Fig. 7b and Fig. 7d), the effect of adding a larger noise level is clearly visible. The tested model settings seem to be robust against noise in the sense that if a small RRMSE is obtained for a certain metabolite in set 1 (respectively, set 2), a small RRMSE is also obtained in set 3 for the same metabolite (respectively, set 4). For each metabolite, the shapes depicted by the 6 model settings (6 bars) are quite similar for sets 2 and 4, meaning that the effect of the added noise on the RRMSE does not depend on the model setting.

The magnitude of the *water resonance* has a very small impact on the RRMSE for both MP-FIR and HLSVD-PRO. In contrast, the shape of the water resonance has more influence, especially for HLSVD-PRO as described in Section 2.3. More important for MP-FIR is the closeness between the region of interest and the water resonance. Regarding the choice of the filtering region, HLSVD-PRO seems to be less sensitive than MP-FIR, which exhibits substantially larger errors when the bound of the filtering region gets closer to the water resonance. This effect is especially noticeable for low noise signals (compare model settings 1 and 2 in sets 1 and 2, Fig. 7). Comparable or even better results are obtained with HLSVD-PRO when the filtering region is set to [0.25,4.5] ppm.
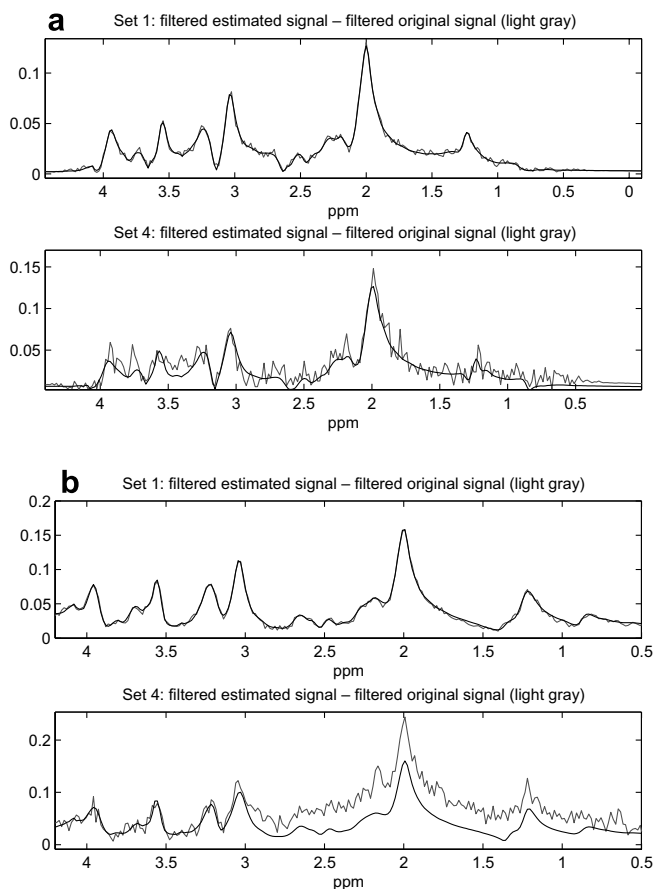

Fig. 6. (a) Model setting 1 (MP-FIR). (b) Model setting 4 (HLSVD-PRO). Magnitude of the filtered original spectra *vs.* filtered estimated spectra for sets 1 and 4 using model settings 1 and 4.
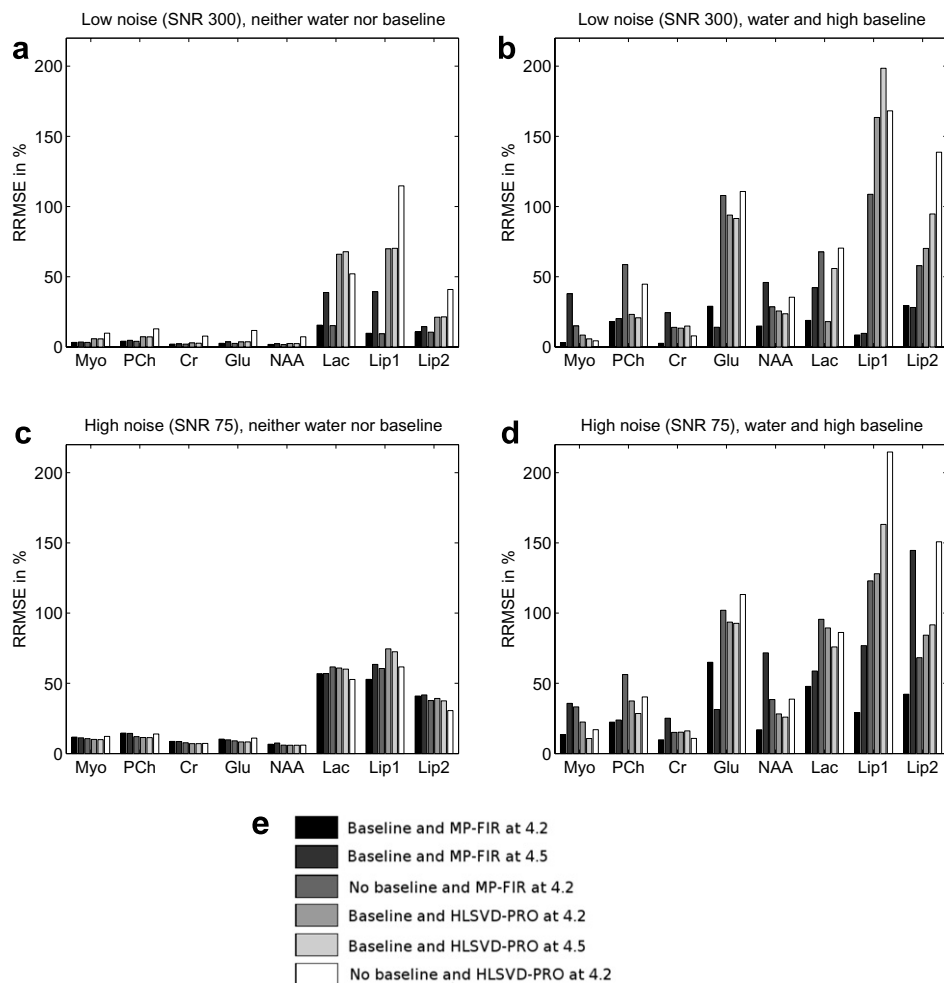
Fig. 7. (a) Set 1. (b) Set 2. (c) Set 3. (d) Set 4. (e) Model settings. Sensitivity of AQSES on the amplitude estimates with respect to the filtering techniques, the nuisance components and the choice of the filtering region.

The *computation times* needed per spectrum for quantitation have been reported in Table 3 for model settings 1, 3, 4 and 6. AQSES has been run on a Windows XP platform with a Pentium 4 (3 GHz CPU, 1 Gb RAM). These computation times have been averaged out over sets 2 and 4 when modeling the baseline and over sets 1 and 3 when no baseline was modeled. Although modeling the baseline significantly increases the computation time (see [4] for more details), both filtering techniques used as described above are comparable in terms of efficiency.

4.1. In vivo experiment

The bounds of the filters are [0.25 4.1] ppm. They are chosen in such a way that the unwanted components outside the frequency range of interest are completely removed by filtering. The model order of HLSVD-PRO was fixed at 30 to guarantee a complete suppression of the water resonances. The amplitude estimates for the main 5 metabolites are shown in Fig. 8 for different **smoothness levels** of the modeled baseline. The results obtained with MP-FIR or HLSVD-PRO are sensitive to these levels of smoothness, lower levels (levels 1–3 in Fig. 8) providing more variability in the amplitude estimates. Nevertheless, less variability in the results was observed with MP-FIR for higher levels of smoothness. Similar observations were made when estimating the best λ for the simulated spectra (the results are not reported for sake of space). On the contrary, a not sufficiently smooth baseline alters the amplitude estimates when filtering with MP-FIR. In that case, the modeled baseline tends to fit the metabolites widely spread in the frequency domain like Myo and Glx. Note that the amplitude

Table 3
Averaged computation times (in seconds) per spectrum for model setting 1 (MP-FIR and baseline), 3 (MP-FIR and no baseline), 4 (HLSVD-PRO and baseline) and 6 (HLSVD-PRO and no baseline)

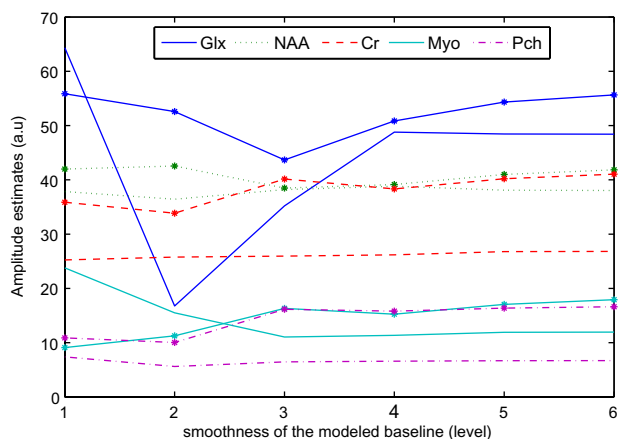|	MP-FIR	HLSVD-PRO
Baseline (sets 2 and 4)	2.13	2.52
No baseline (sets 1 and 3)	0.61	0.72

Fig. 8. *In vivo* signal of Fig. 5: amplitude estimates of the 5 main metabolites (Myo, Cr, PCh, Glx and NAA) versus the smoothness of the modeled baseline ('-' for MP-FIR and '-*-' for HLSVD-PRO): level 1 of smoothness means "low smoothness" and level 6 means "high smoothness".

estimates of Cr, PCh and NAA are constant for the different smoothness levels of the baseline.

Relatively large differences appear in the amplitude estimates when comparing the results obtained with MP-FIR and HLSVD-PRO for a specific smoothness level of the modeled baseline. Most often, the frequency correction reaches its maximum fixed value (set at 20 Hz in AQSES) for most of the metabolites when using HLSVD-PRO, suggesting that the metabolites either try to fit the baseline or try to fit other metabolites that are not contained in the database. We also notice larger residue when using HLSVD-PRO. Although conclusions cannot be drawn, one can think that the amplitude obtained with HLSVD-PRO are overestimated. The simulations above confirm this trend to overestimate the metabolite amplitudes when the original signal contains a non-smooth baseline. The results in terms of amplitude ratios are in agreement with the ratios of metabolite concentrations reported by Govindaraju et al. (Table 2 in [32]).

5. Discussion

SVD-based methods such as HLSVD-PRO are often chosen for removing the water components in MRS spectra. As preprocessing method, they exhibit several advantages: the phase is not distorted, the water resonance seems visually removed, the number of parameters to be chosen by the user is restricted to 2 (the model order and the filtering region). Nevertheless, we showed that, in presence of non-Lorentzian water peaks (see Section 2), this algorithm fails to completely remove the residual water tails in the frequency region of interest. The metabolite contributions are normally kept unchanged if there is no misalignment between the metabolite profiles and the original signal. Consequently, the signal-to-baseline ratio (SBR) is

decreased. Fortunately, this decrease is usually not dramatic since the residual water tails are often much smaller than the baseline itself. On the contrary, MP-FIR will increase the signal-to-baseline ratio as described in Section 2. The signal-to-noise ratio decreases when using MP-FIR while it is stable in the case of HLSVD-PRO. MP-FIR will not remove the baseline but will reduce it more than it reduces the metabolite contributions. In this sense, MP-FIR can be useful for better disentangling the metabolites from overlapping macromolecules. HLSVD-PRO encounters indeed more difficulties when dealing with signals affected by large baselines (see Section 4). MP-FIR also reduces the sensitivity of the regularization parameter λ in AQSES. The results show that the amplitude estimates remain constant over a relatively large range of λ values. We can intuitively think that similar conclusions would be observed for other quantitation methods based on a semiparametric model.

The results show that MP-FIR outperforms HLSVD-PRO whatever the nuisance components. A major part of the explanations is due to the increase of signal-to-baseline ratio when using MP-FIR. The fact that HLSVD-PRO is used as a preprocessing step can also generate errors in the amplitude estimation. This will be the case if the metabolite profiles and the original signal are not correctly aligned and if the bounds of the filtering region are too close to the metabolite peaks. Since MP-FIR is applied during quantitation at each iteration to the corrected metabolite profiles (see Eq. (5)), the transition band of this signal and the metabolite profiles of the basis set will match after a couple of iterations thereby avoiding border problems. As described above, HLSVD-PRO could be applied after correction of the metabolite profiles, but would result in an inefficient method barely usable when large sets of signals have to be processed. The efficiency depends also on the model order used in HLSVD-PRO. However, we noticed that the choice of the model order was not crucial as long as it remains around 25 and is kept sufficiently high. A similar observation was made by Sundin et al. [2] for long-echo time MR spectra with a model order of 12 instead of 25. For a typical model order of 25, HLSVD-PRO can encounter technical problems due to shortage of memory when processing very long signals. This is rarely the case for short-echo time signals that are usually shorter or equal to 2048 points.

As for long-echo time spectra, the choice of the MP-FIR filter order should result from a tradeoff between fulfilling the filter requirements (stop band suppression, passband ripple, *etc.*) and preserving most of the useful information from the original signal. For example, a higher filter suppression and therefore a higher filter order is required if some of the nuisance peaks in the stop band have a fast decay [26]. Nevertheless, the filter order should be bounded to a maximum value to keep a good SNR and to avoid numerical problems caused by the computation of polynomial roots and coefficients when transforming the FIR filter into a maximum-phase FIR filter. Numerical

problems were encountered in some cases for M larger than 80. We noticed experimentally that a satisfactory transition band and a satisfactory suppression in the stop band are obtained for $M = 60$. A typical stop band suppression of such a filter is -65 dB, which suffices for *in vivo* MR spectra with presaturated water. The experiments show that the filter design algorithm proposed by Vanhamme et al. [26] converges in most of the cases to suitable filter parameters for short-echo time MR spectra. It optimizes the filter order and the ripple magnitudes in the frequency region of stop band and passband and finds the best corresponding transition bands. The results show that the amplitude estimates can be affected when the cutoff frequency is too close to the water resonances, but good amplitude estimates were obtained in a reasonable range around 4.2 ppm ([4.0 4.3] ppm). In summary, the filter order of MP-FIR should lie between 30 and 80 to ensure a good filtering without facing numerical issues. The region of interest should be kept sufficiently far away from the water resonances (*e.g.*, the frequency region [0.25 4.2] ppm). The baseline should be included in the model in AQSES when a baseline is expected in the original signal. The regularization parameter should be chosen within a range where the amplitude estimates are constant.

As explained before, MP-FIR gives an added value to AQSES compared to an SVD-based method like HLSVD-PRO. First, MP-FIR improves the performance of AQSES when dealing with signals affected by a baseline; and, second, MP-FIR can be used in the optimization procedure without increasing significantly the computation times (simple matrix multiplication).

Increasing the signal-to-baseline ratio improves the initial estimates and thereby reduces the risk of converging to wrong parameter estimates (local minimum). Methods based on the Levenberg–Marquardt least-squares minimization method (see, *e.g.*, [4,8,33–35]) are indeed sensitive to the initial parameter estimates. For example, Ratiney et al. indicated in [7] that the uncertainty on the final estimates is sensitive to the truncated initial data points which influence directly the initial parameter estimates. Soher et al. studied in [36] four different optimization procedures: the iterative procedures based on Young's method with wavelet or spline baseline characterization, and the single-pass optimization with wavelet or spline baseline function (semiparametric model). He concluded that the iterative methods were more dependent on having good starting values for the baseline parameters, without dramatic differences in performance regarding the choice of the baseline model (wavelets or splines). This tends to show that the baseline should be included in the model, as confirmed by our results, at least if an initial estimate of the baseline is rather uncertain (*i.e.*, with high risk of wrong initial guess).

Although MP-FIR has been designed for processing time-domain signals, it might be helpful for frequency-domain methods like LCModel [34]. Used as preprocessing method prior to LCModel, it can reduce the contribution of the baseline while suppressing completely the water resonances which might have some contributions in the frequency region of interest. Using LCModel, Seeger et al. [28] showed that extending the standard basis set of metabolites by inclusion of parameterized components for macromolecules and lipids can improve the parameter estimates and especially estimates of metabolites such as NAA, Glx or Lac. Our results have shown that those metabolites are indeed sensitive to the baseline. Nevertheless, adding more components (Seeger et al. added five components to take into account the lipids and the macromolecules) complicates the model, increasing the degrees of freedom and so the risk of converging to a local minimum. MP-FIR might limit this risk since it reduces the contribution of large damping components (like macromolecules and lipids). Furthermore, additional prior knowledge might help in identifying appropriate initial parameter estimates. For example, Hofmann et al. [37] showed that the macromolecule baseline is different in composition and concentration between white and gray matter volumes, and depends on the age but not on the gender of the subject. MP-FIR allows to incorporate this prior knowledge in the model.

The fact that MP-FIR can be used in the optimization procedure allows to avoid frequency misalignments between the metabolite profiles and the original signal. Therefore, procedures like QUEST [7] using solvent suppression methods (like HLSVD) prior to quantitation will probably require a more stringent alignment of the original signal and the metabolite profiles to avoid that some peaks are (partially) filtered out while they should not. MP-FIR could be used prior to QUEST or inside the iterative procedure in QUEST with the In-Base method [35]. To fully benefit from it, MP-FIR should be used prior to the baseline extraction, and therefore, it is not recommended inside the iterative procedure in QUEST with the Subtract method [35].

MP-FIR, already successful for water removal for long-echo time MR spectra, might even be more powerful for short-echo time MR spectra due to its ability to deal with the baseline, thereby increasing substantially (compared to HLSVD-PRO) the accuracy level of the parameter estimates obtained with AQSES. It can be used within the optimization procedure of most of the time-domain methods based on the Levenberg–Marquardt algorithm (Elster et al.'s method [8], $AMARES_W$ [26], VARPRO [38], *etc.*).

6. Conclusions

This paper shows that HLSVD-PRO and MP-FIR, two filtering techniques used in long-echo MRS data quantitation, can be successfully applied to short-echo time spectra. The potentials and limitations of both techniques have been studied and some recommendations regarding their use have been given. Simulations show that MP-FIR outperforms HLSVD-PRO in terms of accuracy, computational efficiency and robustness against the nuisance

components. This is explained by two main reasons: MP-FIR increases the signal-to-baseline ratio while it slightly decreases with HLSVD-PRO, and MP-FIR is used inside the optimization procedure while HLSVD-PRO is used prior to quantitation. MP-FIR turns out to be especially powerful for time-domain quantitation methods based on a semiparametric model.

Acknowledgments

References

[1] 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, S. Van Huffel, Classification of brain tumours using short echo time ^{1}H MR spectra, J. Magn. Reson. 170 (1) (2004) 164–175.

[2] 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.

[3] 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.

[4] J.-B. 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. in press.

[5] L. Vanhamme, T. Sundin, P. Hecke, S. Van Huffel, MR spectroscopy quantitation: a review of time-domain methods, NMR Biomed. 14 (4) (2001) 233–246.

[6] R. Romano, A. Motta, S. Camassa, C. Pagano, M.T. Santini, P.L. Indovina, A new time-domain frequency-selective quantification algorithm, J. Magn. Reson. 155 (2) (2002) 226–235.

[7] 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.

[8] 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.

[9] V.A. Mandelshtam, H.S. Taylor, A.J. Shaka, Application of the filter diagonalization method to one- and two-dimensional NMR spectra, J. Magn. Reson. 133 (2) (1998) 304–312.

[10] S. Mierisova, M. Ala-Korpela, MR spectroscopy quantitation: a review of frequency domain methods, NMR Biomed. 14 (4) (2001) 247–259.

[11] Y. Hiltunen, J. Kaartinen, J. Pulkkinen, A.M. Hakkinen, N. Lundbom, R.A. Kauppinen, Quantification of human brain metabolites from in vivo ^{1}H NMR magnitude spectra using automated artificial neural network analysis, J. Magn. Reson. 154 (1) (2002) 1–5.

[12] P. Stoica, N. Sandgren, Y. Selen, L. Vanhamme, S. Van Huffel, Frequency-domain method based on the singular value decomposition for frequency-selective NMR spectroscopy, J. Magn. Reson. 165 (1) (2003) 80–88.

[13] R.E. Gabr, R. Ouwerkerk, P.A. Bottomley, Quantifying in vivo MR spectra with circles, J. Magn. Reson. 179 (1) (2006) 152–163.

[14] D. Marion, M. Ikura, A. Bax, Improved solvent suppression in one- and two-dimensional NMR spectra by convolution of time-domain data, J. Magn. Reson. 84 (2) (1989) 425–430.

[15] P. Sodano, M. Delepierre, Binomial frequency response to non-binomial pulse sequences for efficient water suppression, J. Biomol. NMR 3 (4) (1993) 471–477.

[16] K.J. Cross, Improved digital filtering technique for solvent suppression, J. Magn. Reson. Ser. A 101 (2) (1993) 220–224.

[17] J.P. Antoine, A. Coron, J.M. Dereppe, Water peak suppression: time-frequency vs time-scale approach, J. Magn. Reson. 144 (2) (2000) 189–194.

[18] J.P. Antoine, C. Chauvin, A. Coron, Wavelets and related time-frequency techniques in magnetic resonance spectroscopy, NMR Biomed. 14 (4) (2001) 265–270.

[19] U.L. Gunther, C. Ludwig, H. Ruterjans, WAVEWAT-improved solvent suppression in NMR spectra employing wavelet transforms, J. Magn. Reson. 156 (1) (2002) 19–25.

[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] H. Chen, S. Van Huffel, J. Vandewalle, Bandpass prefiltering for exponential data fitting with known frequency region of interest, Signal Process 48 (2) (1996) 135–154.

[22] A. Coron, L. Vanhamme, J.P. 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.

[23] H. Barkhuijsen, R. de Beer, D. van Ormondt, Improved algorithm for noniterative time-domain model fitting to exponentially damped magnetic resonance signals, J. Magn. Reson. 73 (3) (1987) 553–557.

[24] D. Calvetti, L. Reichel, D.C. Sorensen, An implicitly restarted Lanczos method for large symmetric eigenvalue problems, Electron. Trans. Numer. Anal. 2 (1994) 1–21.

[25] 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.

[26] 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.

[27] I.W. Selesnick, M. Lang, S.C. Burrus, Constrained least square design of FIR filters without specified transition bands, IEEE Trans. Signal Process 44 (1996) 1879–1892.

[28] U. Seeger, U. Klose, Parameterized evaluation of macromolecules and lipids in proton MR spectroscopy of brain diseases, Magn. Reson. Med. 49 (2003) 19–28.

[29] E. Cabanes, S. Confort Gouny, Y. Le Fur, G. Simond, P.J. Cozzone, Optimization of residual water signal removal by HLSVD on simulated short echo time proton MR spectra of the human brain, J. Magn. Reson. 150 (2001) 116–125.

[30] U. Klose, In vivo proton spectroscopy in presence of eddy currents, Magn. Reson. Med. 14 (1990) 26–30.

[31] R. Bartha, D.J. Drost, R.S. Menon, P.C. Williamson, Spectroscopic lineshape correction by QUECC: combined QUALITY deconvolution and eddy current correction, Magn. Reson. Med. 44 (4) (2000) 641–645.

[32] V. Govindaraju, K. Young, A.A. Maudsley, Proton NMR chemical shifts and coupling constants for brain metabolites, NMR Biomed. 13 (3) (2000) 129–153.

[33] 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.

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

[35] H. Ratiney, Y. Coenradie, S. Cavassila, D. van Ormondt, D. Graveron-Demilly, Time-domain quantitation of ^{1}H short echo-time signals: background accommodation, Magn. Reson. Mater. Phys. 16 (6) (2004) 284–296.

[36] B.J. Soher, K. Young, A.A. Maudsley, Representation of strong baseline contributions in ^{1}H MR spectra, Magn. Reson. Med. 45 (6) (2001) 966–972.

[37] L. Hofmann, J. Slotboom, C. Boesch, R. Kreis, Characterization of the macromolecule baseline in localized ^{1}H-MR spectra of human brain, Magn. Reson. Med. 46 (5) (2001) 855–863.

[38] 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.