ACADEMIC PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Magnetic Resonance 158 (2002) 1–14

JMR Journal of Magnetic Resonance

www.academicpress.com

Automatic frequency alignment and quantitation of single resonances in multiple magnetic resonance spectra via complex principal component analysis

Sabine Van Huffel,[a,*] Yu Wang,[a] Leentje Vanhamme,[a] and Paul Van Hecke[b]

[a] *Department of Electrical Engineering, Katholieke Universiteit Leuven, SCD-SISTA. Kasteelpark Arenberg 10, B-3001 Leuven, Belgium*
[b] *Biomedical NMR Unit, Faculty of Medicine, Katholieke Universiteit Leuven, Herestraat 49, 3000 Leuven, Belgium*

Received 29 August 2001; revised 8 July 2002

Abstract

Several algorithms for automatic frequency alignment and quantitation of single resonances in multiple magnetic resonance (MR) spectra are investigated. First, a careful comparison between the complex principal component analysis (PCA) and the Hankel total least squares-based methods for quantifying the resonances in the spectral sets of magnetic resonance spectroscopy imaging (MRSI) spectra is presented. Afterward, we discuss a method based on complex PCA plus linear regression and a method based on cross correlation of the magnitude spectra for correcting frequency shifts of resonances in sets of MR spectra. Their advantages and limitations are demonstrated on simulated MR data sets as well as on an in vivo MRSI data set of the human brain.

Keywords: In vivo MRSI; Quantitation; Principal component analysis; Total least squares; Singular value decomposition; Automatic frequency alignment; Cross correlation

1. Introduction

In vivo magnetic resonance spectroscopy (MRS) examinations may result in large sets of spectra, e.g., spectra from time series [1] or from MRS imaging (MRSI) experiments [2,3]. Separate quantitation of each spectrum is cumbersome, while automated quantitation of the complete data set may pose serious problems. Moreover, quantitation of MRS data in general, and MRSI data in particular, yielding information about metabolite concentrations, is usually hampered by low signal-to-noise ratios. Van Huffel et al. [4] applied the total least squares method to MRS data quantitation, assuming that the data were modeled as a sum of exponentially damped sinusoids. Their algorithm, called Hankel total least squares (HTLS), was shown to improve the accuracy of all MRS parameter estimates as compared to HSVD. Recently, HTLS was generalized [5] to the quantitation of sets of MRS data as in MRSI, resulting in the HTLSstack and HTLSsum algorithms.

As a widely used statistical technique, principal component analysis (PCA) was introduced in MRS data quantitation by Stoyanova et al. [6]. PCA is able to quantify simultaneously all spectra in one data set provided they are single resonances of the same lineshape. As such, the quantitation result is improved [6] since PCA exploits the common information between all spectra to suppress the effect of noise. In [6] only the real part of the MRS data is considered and transformed to the frequency domain. An interval surrounding the resonance of interest is selected and the points within this interval are stored row by row in a matrix, to which PCA is applied. To improve the accuracy and eliminate the influence of the phase difference, Elliot et al. [7] performed PCA on complex MRS data using a complex SVD. Mathematically, PCA decomposes the data set in order to extract the basic features, called principal

* Corresponding author. Fax: +32-16-321970.
E-mail addresses: sabine.vanhuffel@esat.kuleuven.ac.be (S. Van Huffel), yu.wang@esat.kuleuven.ac.be (Y. Wang), leentje.vanhamme@esat.kuleuven.ac.be (L. Vanhamme), paul.vanhecke@med.kuleuven.ac.be (P. Van Hecke).

1090-7807/02/$ - see front matter
PII: S1090-7807(02)00055-1

components (PCs). It has been shown that under some conditions PCA can successfully extract quantitative metabolite information from data sets without any prior knowledge about the lineshape [6–9].

The PCA methods discussed above only apply if corresponding single resonances in all spectra have the same frequency, lineshape (and phase in case of real-valued data sets). However, Brown and Stoyanova [9] developed a PCA-based method for aligning the spectra in frequency and phase using the real-valued part of the data set. This method was further improved in [12] by combining PCA with linear regression. In a broader sense, frequency-shift alignment is always needed as a preprocessing technique for in vivo applications in which frequency shifts are introduced due to instrumental imperfections, field inhomogeneity, patient movement, etc. Very recently, Stoyanova and Brown [13,14] further improved and generalized these methods to correct simultaneously for amplitude variations, as well as frequency, phase, and linewidth variations. In essence, they apply the same linear regression approach as presented in [12] but extend it with linewidth corrections and use, instead of the original spectral data matrix, a rank-four approximation.

This paper analyzes the advantages and limitations of the use of complex PCA for quantitation of single resonances, as well as for frequency alignment. Here, the entire complex-valued data set (real and imaginary part) is used, in contrast to [6,9,12–14] where only the real part is considered. As already shown in [7,10] for quantitation and proven here for frequency alignment, working with the entire complex-valued data set has the large advantage that the spectra no longer need to be aligned in-phase for quantifying the areas under each resonance using PCA. First of all, cPCA is compared via an extensive Monte Carlo simulation study with HTLS, HTLSsum, and HTLSstack. Next, we extend the frequency alignment procedures described in [9,12] to the complex domain by combining cPCA and linear regression, and we also discuss further improvements. The performance of this newly developed method, called cPCA-LR(f), is compared in detail to the existing methods in the literature, in particular its real-valued counterpart (rPCA-LR(f,p)) and cross correlation (cCross(f)). It should be noticed that so far, no extensive evaluation of performances of PCA-based methods compared to the other state-of-the-art methods has been published. These are the main contributions of this paper.

The structure of this paper is as follows: we start with the theory section, where the algorithms for quantitation are discussed. Thereafter, we present frequency alignment methods. In the next section, the advantages and limitations of all these methods are demonstrated on simulated MR data sets of single resonances and an in vivo MRSI data set of the human brain.

2. Theory

2.1. MRS data model function

It is assumed that, basically, MRS data can be modeled as a sum of exponentially damped complex-valued sinusoids (Lorentzians),

$$y_n = \breve{y}_n + e_n = \sum_{k=1}^{K} a_k e^{j\phi_k} e^{(-d_k + j2\pi f_k)t_n} + e_n, \tag{1}$$

where $n = 0, 1, \ldots, N-1$; y_n is the nth measured data point of an MRS signal; $\breve{y}_n$ represents the nth value of the model function, $j = \sqrt{-1}$; a_k is the amplitude, ϕ_k the phase, d_k the damping factor, and f_k the frequency of the kth sinusoid $k = 1, 2, \ldots, K$; K is the number of the sinusoids: $t_n = n\Delta t + t_0$ with Δt the sampling interval and t_0 the time between the effective time origin and the first data point to be included in the analysis; and e_n is complex white Gaussian noise. The amplitude a_k is directly related to the concentration of a certain metabolite and should be estimated as precisely as possible. Alternatively, the data can be modeled as a sum of Gaussians,

$$y_n = \breve{y}_n + e_n = \sum_{k=1}^{K} a_k e^{j\phi_k} e^{(-g_k t_n + j2\pi f_k)t_n} + e_n, \tag{2}$$

where g_k is the damping factor and the other parameters are as defined above.

2.2. Hankel total least squares

Arranging the data points y_n, $n = 0, 1, \ldots, N-1$ of one MRS signal $\mathbf{y}$ in a Hankel matrix H of dimensions $L \times M$, $L \geqslant K$, $M \geqslant K$, $N = L + M - 1$, yields:

$$H = \begin{bmatrix} y_0 & y_1 & \cdots & y_{M-1} \\ y_1 & y_2 & \cdots & y_M \\ \vdots & \vdots & \vdots & \vdots \\ y_{L-1} & y_{L-2} & \cdots & y_{N-1} \end{bmatrix}. \tag{3}$$

Computing the SVD of the Hankel matrix H, we obtain

$$H_{L\times M} = U_{L\times L} \Sigma_{L\times M} V^H_{M\times M}, \tag{4}$$

where $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_q)$, $\sigma_1 \geqslant \sigma_2 \geqslant \cdots \geqslant \sigma_q$, $q = \min(M, L)$, and the superscript H denotes the hermitian conjugate. According to [4], H should be chosen as square as possible in order to get the best parameter accuracy.

Truncate H to a matrix H_K of rank K (K is the estimated number of signal poles):

$$H_K = U_K \Sigma_K V_K^H. \tag{5}$$

U_K and V_K are the first K columns of U and V, respectively. Σ_K is the $K \times K$ upper-left submatrix of Σ.

Now compute the total least squares (TLS) [15] solution E of the following (incompatible) set:

$$V_K^{(t)} E^H \approx V_K^{(b)}. \tag{6}$$

$V_K^{(b)}$ and $V_K^{(t)}$ are derived from V_K by omitting its first and last row, respectively. The K eigenvalues of E yield the signal pole estimates:

$$\hat{z}_k = \mathrm{e}^{(-\hat{d}_k + \mathrm{j}2\pi\hat{f}_k)\Delta t} \quad k = 1, \ldots, K. \tag{7}$$

Filling in the estimated frequencies $\hat{f}_k$ and damping factors $\hat{d}_k$ into the model Eq. (1) yields the set:

$$y_n \approx \sum_{k=1}^{K} c_k \mathrm{e}^{(-\hat{d}_k + \mathrm{j}2\pi\hat{f}_k)t_n} \quad n = 0, 1, \ldots, N-1.$$

From its least squares solution $\hat{c}_k = \hat{a}_k \mathrm{e}^{\mathrm{j}\hat{\phi}k}$, we find $\hat{a}_k$ and $\hat{\phi}_k$ the estimated amplitudes and phases.

A more detailed description of the HTLS algorithm and underlying principles is given in [15].

2.3. HTLSsum and HTLSstack

Two extensions of the HTLS algorithms quantifying simultaneously sets of MRS signals are presented now. A thorough discussion of these algorithms can be found in [5].

Suppose we have P MRS signals $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_\mathbf{P}$ of N data points each. Let

$$\bar{\mathbf{y}} = \mathbf{y}_1 + \mathbf{y}_2 + \cdots + \mathbf{y}_\mathbf{P}.$$

Now we can apply the HTLS algorithm to the sum $\bar{\mathbf{y}}$ of these signals to first obtain the signal poles and afterward calculate the corresponding amplitudes and phases for each signal separately. How to assign the signal poles to corresponding signals is not trivial. A procedure is outlined in [5]. This algorithm is called HTLSsum.

The HTLSstack algorithm is as follows:

Arrange N data points of each MRS signal $\mathbf{y}_p$ into a Hankel matrix as in Eq. (3):

$$H_p = \begin{bmatrix} y_{0p} & y_{1p} & \cdots & y_{(M-1)p} \\ y_{1p} & y_{2p} & \cdots & y_{Mp} \\ \vdots & \vdots & \vdots & \vdots \\ y_{(L-1)p} & y_{(L-2)p} & \cdots & y_{(N-1)p} \end{bmatrix}, \tag{8}$$

where $p = 1, 2, \ldots, P$. Vertically stack these P Hankel matrices to obtain:

$$H_{\text{stack}} = \begin{bmatrix} H_1 \\ H_2 \\ \vdots \\ H_p \end{bmatrix}.$$

Equations (4) and (6) can be solved using H_{stack} in order to obtain frequencies and damping factors for all MRS signals together. Note that here signal poles need to be assigned to corresponding signals. The same procedure is used in HTLSsum [5].

2.4. Principal component analysis

PCA is usually performed by means of an eigenvalue decomposition. Here we use the complex SVD. Although PCA can be applied to the time-domain data, the corresponding MR spectra are generally considered and are obtained from the original time-domain signal in each voxel by applying the discrete Fourier transformation (DFT). Assume P voxels and denote the MR spectrum in the ith voxel by $\mathbf{y_i}$. Arrange the P spectra $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_\mathbf{P}$ as rows in a matrix D and compute the complex SVD,

$$D_{P\times N} = \begin{bmatrix} \mathbf{y}_1^T \\ \mathbf{y}_2^T \\ \vdots \\ \mathbf{y}_\mathbf{P}^T \end{bmatrix} = (U_{P\times P}\Sigma_{P\times N})V_{N\times N}^H = S_{P\times N}V_{N\times N}^H, \tag{9}$$

where $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_q)$, $q = \min(P, N)$. The columns of V contain the basic lineshapes of the signals, called PCs. The elements in $S_{P\times N} = U_{P\times P}\Sigma_{P\times N}$ are called the scores of each principal component.

Now, suppose our set of P MRS signals is composed of a single resonance, which means that the matrix D has rank one. In this case, only $\sigma_1 \neq 0$, $\sigma_2 = \sigma_3 = \cdots = \sigma_q = 0$. Now the score matrix S contains only one column

$$S_1 = [\, s_1 \quad s_2 \quad \cdots \quad s_P \,]^{\mathrm{T}},$$

which is different from zero. The first PC, i.e., the first column $\mathbf{v}_1$ of V, represents the basic lineshape, which, in our case, contains only one peak.

When noise is present, $\sigma_2, \sigma_3, \ldots, \sigma_q$ are not equal to zero. But they are normally quite small compared to σ_1.

Each signal $\mathbf{y}_p$ in the set can be expressed in the following way:

$$\mathbf{y}_p = s_p\mathbf{v}_1 \quad p = 1, \ldots, P. \tag{10}$$

Using frequency-domain data, obtained by applying the DFT to the rows of MRS signals, we need to normalize $\mathbf{v}_1$ so that it has unit area. Denoting the normalized PC by $\tilde{\mathbf{v}}_1$, this implies:

$$\sum_{i=1}^{N} \tilde{v}_{i1} = 1. \tag{11}$$

In this case, Eqs. (9) and (10) become (in matrix notation)

$$D = \tilde{S}_1\tilde{\mathbf{v}}_1^H = \left(S_1 \sum_{i=1}^{N} v_{i1}\right) \frac{\mathbf{v}_1^H}{\sum_{i=1}^{N} v_{i1}}. \tag{12}$$

The magnitude values of the normalized scores $\tilde{S}_1 = S_1 \sum_{i=1}^{N} v_{i1}$ then represent the areas under the resonance in each spectrum, from which the corresponding metabolite concentrations can be obtained. This algorithm is called cPCA, where "c" represents complex.

Using the properties of the DFT, it is easy to prove that PCA performed in the frequency domain is equivalent to its counterpart in the time domain except for the normalization, which agrees with our experimental results.

Using the original time-domain data as rows in the matrix D, the first PC $\mathbf{v}_1$ needs to be normalized in such a way that it represents a unit amplitude signal. To normalize $\mathbf{v}_1$, we need to divide $\mathbf{v}_1$ by its first element v_{11} so that Eq. (12) becomes (in matrix notation)

$$D = \tilde{S}_1 \tilde{\mathbf{v}}_1^H = (S_1 v_{11}) \frac{\mathbf{v}_1^H}{v_{11}}. \tag{13}$$

The normalized scores $\tilde{S}_1 = S_1 v_{11}$ then represent the amplitudes of the original signals, from which the corresponding metabolite concentrations can be obtained.

Stoyanova and Brown applied PCA to quantify a spectral data set which contains only real parts of spectra. For a detailed presentation of Stoyanova et al.'s method, we refer to [6]. This method is called rPCA since it is only applied to the real part of the spectrum. Elliott et al. proved in their paper [7] that applying PCA to the whole complex spectrum can increase the accuracy of resonance area estimation by $\sqrt{2}$.

Quantitation of MRS signals containing more than one resonance using PCA is much more complex. Since PCA does not enable automated quantitation of each single resonance separately (each PC always contains a mixture of both resonances), we need to filter out each resonance before applying PCA, thereby introducing a bias. The best results with minimal bias are obtained using a maximum phase time domain Finite Impulse Response (FIR) filter [11]. More details are given in [8,10].

2.5. Frequency-shift alignment using complex PCA

The PCA method only works well for data sets consisting of single resonant peaks of constant shape but variable amplitude. However, in reality, due to fluctuations in instrumental parameters and other experimental variations, the resonances in an in vivo MRS time series or MRSI experiment are always subject to frequency shifts, phase misadjustment, and lineshape distortions. Assume a set of single resonances $S(\omega)$ subject to frequency shifts $\Delta\omega$ only. This can mathematically be described as

$$S_k(\omega) = A_k f(\omega) = A_k f(\omega_0 + \Delta\omega_k), \tag{14}$$

where $S_k(\omega)$ is the kth spectrum and function of ω, A_k is the complex amplitude of the kth spectrum, $f(\omega)$ is the lineshape function, and $\Delta\omega_k$ represents the frequency shift of the kth spectrum.

We can no longer directly apply PCA to quantify the spectra since now the signal-related variance is not only distributed along the direction of the first PC but also along the direction of the second and the third PC, etc.

Suppose that the frequency shift is so small that we can expand $f(\omega)$ in a Taylor series:

$$S_k(\omega) = A_k \left\{ f(\omega_0) + \frac{\partial f}{\partial \omega}\bigg|_{\omega_0} \Delta\omega_k + \frac{1}{2}\frac{\partial^2 f}{\partial \omega^2}\bigg|_{\omega_0} \Delta\omega_k^2 + \cdots \right\}.$$

If $\Delta\omega_k$ is small enough, we can neglect the second and higher order terms:

$$S_k(\omega) \approx A_k \left\{ f(\omega_0) + \frac{\partial f}{\partial \omega}\bigg|_{\omega_0} \Delta\omega_k \right\}. \tag{15}$$

We can also apply PCA to the same data set to obtain:

$$S_k(\omega) = s_{1k} \cdot \mathbf{v}_1 + s_{2k} \cdot \mathbf{v}_2 + \cdots.$$

The first PC $\mathbf{v}_1$ somehow represents the basic lineshape $f(\omega_0)$ in the data set and $\mathbf{v}_1 \approx a f(\omega_0)$, where a is an arbitrary positive real number. Therefore we can approximate $\frac{\partial f}{\partial \omega}|_{\omega_0}$ by using the derivative of the first PC, $\mathbf{v}_1' \approx a \frac{\partial f}{\partial \omega}|_{\omega_0}$, which is calculated here via numerical differentiation. By applying linear regression (LR), each spectrum $S_k(\omega)$ is approximated by a linear combination of $\mathbf{v}_1$ and $\mathbf{v}_1'$:

$$S_k(\omega) \approx c_1 \cdot \mathbf{v}_1 + c_2 \cdot \mathbf{v}_1'. \tag{16}$$

Equating the coefficient ratios of the expressions [15] and [16] yields the shift:

$$\frac{c_2}{c_1} = \frac{a A_k \Delta\omega_k}{a A_k} = \Delta\omega_k.$$

Since the resonances are only subject to frequency shifts, $\Delta\omega_k$ should be real. However, in practice, other distortions and the approximations made can make the shift complex, in which case the real part of $\Delta\omega_k$ should be taken. We then transform $S_k(\omega)$ to the time domain, correct the frequency by multiplying the nth point by $e^{-j\Delta\omega_k 2\pi t_n}$, and retransform back to the frequency domain.

The above described algorithm, denoted by cPCA-LR(f), where "c" represents complex and "LR" represents linear regression, is applied iteratively to the data until frequency shifts become negligible compared to the linewidth. Now the signal-related variance should mainly lie along the direction of the first PC.

Brown and Stoyanova [9] applied PCA to correct frequency and phase shifts in a data set using only the real part of the spectra. Their basic assumption is that the frequency shift and phase shift are small, i.e., $<45°$ in-phase and less than one linewidth in frequency, as investigated in [12]. This algorithm, denoted here by rPCA(f), was further improved by Witjes et al. [12].

Witjes et al. [12] showed in their paper that their algorithm, denoted rPCA-LR(f,p), was more accurate and robust than rPCA(f). They also demonstrated that the use of high order terms in the Taylor expansion can improve the performance of the algorithm.

cPCA-LR(f) can also be further improved by adding higher order terms in the Taylor expansion, for example, by adding the second-order derivative into the linear regression equation

$$S_k(\omega) \approx c_1 \cdot \mathbf{v}_1 + c_2 \cdot \mathbf{v}_1' + c_3 \cdot \mathbf{v}_1'', \tag{17}$$

where $\mathbf{v}_1'' = \frac{\partial^2 f}{\partial \omega^2}|_{\omega_0}$ is approximated by using numerical differentiation of $\mathbf{v}_1'$. This algorithm is called cPCA2-LR(f).

2.6. Numerical differentiation of PC

In their PCA research, Brown and Stoyanova [9] and Witjes et al. [12] just used first-order finite differences to approximate the derivative of the first PC. Given a set of grid points x_i with $x_{i+1} - x_i = h$ and corresponding first PC values $\mathbf{v}(x_i)$, the derivative can be approximated by:

$$\mathbf{v}'(x_i) \approx \frac{\mathbf{v}(x_{i+1}) - \mathbf{v}(x_i)}{h}. \tag{18}$$

There are more accurate ways to calculate the derivatives [16]. For instance, higher order finite differences such as the fourth-order finite differences yield better approximations:

$$\mathbf{v}'(x_i) \approx \frac{1}{h}\left[-\frac{1}{12}\mathbf{v}(x_{i-2}) + \frac{2}{3}\mathbf{v}(x_{i-1}) - \frac{2}{3}\mathbf{v}(x_{i+1}) + \frac{1}{12}\mathbf{v}(x_{i+2})\right]. \tag{19}$$

Unless stated otherwise first-order finite differences are used in cPCA-LR(f) in order to allow a fair comparison with the PCA methods of [9,12].

2.7. Cross correlation

Cross correlation is a well-known tool for measuring similarity between two different data sets. It can also be used to find the frequency-shift corrections of a given data set. This method, denoted here by cCross(f), is used for aligning spectra in frequency as follows:

1. Calculate the magnitude spectrum of each complex MRS spectrum.
2. Compute the cross correlation between the first spectrum and each remaining spectrum in the data set. The maximal value points out the frequency shift to be applied to each spectrum.
3. Apply the frequency shifts to the original data set (as described in Section 2.3).

Unlike PCA-based methods, which are limited to small frequency shifts because of the first-order Taylor expansion, cross correlation is more robust in the presence of large frequency shifts and should be used prior to PCA-based methods for frequency shift correction, as illustrated in the next section.

3. Simulation results

All the experiments are performed on a Sun Ultra 5 workstation using MATLAB Version 5.3.

3.1. Quantitation of MRS data sets with one single Lorentzian line

Simulation procedure. We generate a complex signal of 128 data points composed of one exponentially damped sinusoid (frequency, 260 Hz; damping factor, 100 Hz (or linewidth 31 Hz); amplitude, 100 arbitrary units (a.u.); sampling frequency $f_s = 1/\Delta t = 3003$ Hz; phase, 0°), hereafter called the basic Lorentzian signal. Data sets of P duplicates of the basic Lorentzian are generated and Gaussian distributed white noise (mean 0 and variance σ_v^2) is added to both real and imaginary parts, after which the resulting FIDS are transformed to the frequency domain. The signal-to-noise ratio (SNR) is expressed in decibels and defined as

$$\mathrm{SNR} = 20\log_{10}\frac{A}{\sqrt{2}\,\sigma_v},$$

where A is the amplitude (resonance area). Amplitudes are estimated using cPCA, HTLS, HTLSstack, and HTLSsum. The whole experiment is repeated 500 times in order to compute the sample relative root mean squared error (rrmse) as follows:

$$\mathrm{rrmse}(\tau) = \frac{1}{A_{\mathrm{true}}}\sqrt{\frac{1}{500P}\sum_{i=1}^{500}\sum_{j=1}^{P}(A_{ij}(\tau) - A_{\mathrm{true}})^2},$$

where τ is the considered method, A_{true} is the true amplitude, P is the number of spectra per set, and A_{ij} is the estimated amplitude of the jth spectrum during the ith repetition of the experiment.

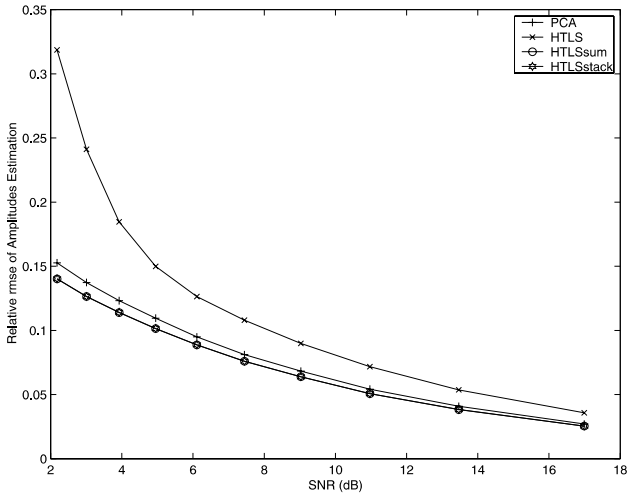

Fig. 1. rrmse versus SNR. Simulated set contains 100 identical single Lorentzians.

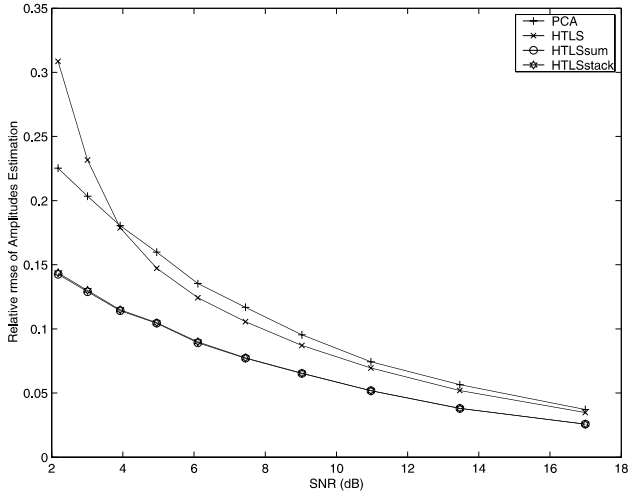

Fig. 2. rrmse versus SNR. Simulated set contains 10 identical single Lorentzians.

In the first two plots, the whole simulation procedure, described above, is repeated for increasing noise variances σ_v^2, ranging from 55^2 to 10^2 and the sample rrmse is plotted versus the SNR. One hundred spectra per data set ($P = 100$) are considered in Fig. 1 and only 10 spectra per data set ($P = 10$) in Fig. 2. Both figures clearly show the advantages of quantifying all spectra in a data set simultaneously. Indeed, HTLS, which quantifies all spectra separately, has the worst performance. HTLSstack and HTLSsum have comparable performances and are better than cPCA. While differences in rrmse are at most 2% at $P = 100$, these differences rise above 8% as the number of spectra per data set decreases to 10 (rrmse(cPCA) is 22.5% compared to a rrmse (HTLSstack) of 14% at the lowest SNR). Even HTLS outperforms cPCA as the SNR > 4 dB, which clearly proves that additional knowledge of the correct lineshape, as used by the HTLS-based methods, is important at low P for guaranteeing a low quantitation error. We also notice the influence of the number of spectra in the data set for cPCA.

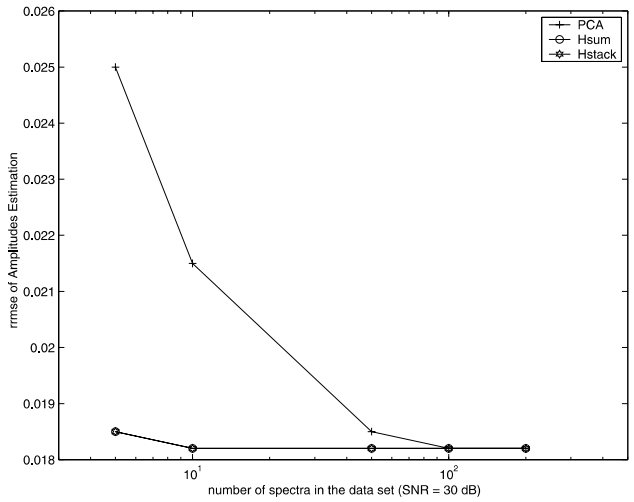

Fig. 3. rrmse versus number of spectra in data set. SNR = 30 dB for each spectrum.

In Fig. 3, the same simulation procedure is repeated for increasing number of spectra ($P = 5$, 10, 50, 100, 200) at SNR 30 dB ($A = 100$ and $\sigma_v = 2$ for each spectrum). It is interesting to notice that the more spectra the data set has, the better cPCA performs. This property makes cPCA a good candidate for quantifying large MRS data sets.

Similar conclusions hold when repeating the simulation procedure on data sets containing spectra with varying amplitudes, as shown in [8,10]. In this case, quantitation of the low SNR spectra benefits from the information provided by the high SNR spectra.

3.2. Quantitation of MRS data sets with one single Gaussian line

Although HTLS-based algorithms are primarily designed to quantify Lorentzians, they can also be applied for quantifying any arbitrary lineshape since each arbitrary lineshape can always be approximated by a number of Lorentzians. Assume J Lorentzians are used, with corresponding signal poles $z_1, z_2, \ldots, z_J$ and amplitudes $c_{1p}, \ldots, c_{Jp}$; then the area a_p under the corresponding lineshape of the pth signal is obtained by linearly combining the amplitudes as [1]

$$a_p = \sum_{i=1}^{J} a_{ip} \cos \phi_{ip}, \tag{20}$$

where $c_{ip} = a_{ip} e^{j\phi_{ip}}$. This computational procedure is illustrated for a Gaussian lineshape. We approximate the Gaussian lineshape by up to 5 Lorentzians lying in its frequency range; i.e., K is set to 5 in Eq. (1) implying that $J \leqslant 5$. At high SNR the five signal poles computed by the HTLS-based methods fall within this range; i.e., $K = J = 5$. However, with decreasing SNR noise peaks are fitted too, resulting in a decrease of J ($J < K$), as shown in our experiments.

Simulation procedure. The same simulation procedure of the previous section is considered here for a basic Gaussian having the same sampling frequency, number of data points, frequency, amplitude, and phase as the basic Lorentzian. Only the damping factor differs and is set to 1000 Hz^2. The HTLS-based methods approximate this resonance by a linear combination of several Lorentzians (e.g., 5, as chosen in Figs. 4 and 5).

These plots show the rrmse versus the SNR for the same noise variances σ_v^2, as chosen in Figs. 1 and 2. Fig. 4 presents the results for 100 spectra per data set ($P = 100$), while Fig. 5 considers only 10 spectra per set ($P = 10$). Both figures show the advantage of using cPCA for a Gaussian lineshape. Without any prior knowledge about the model function. HTLS-based

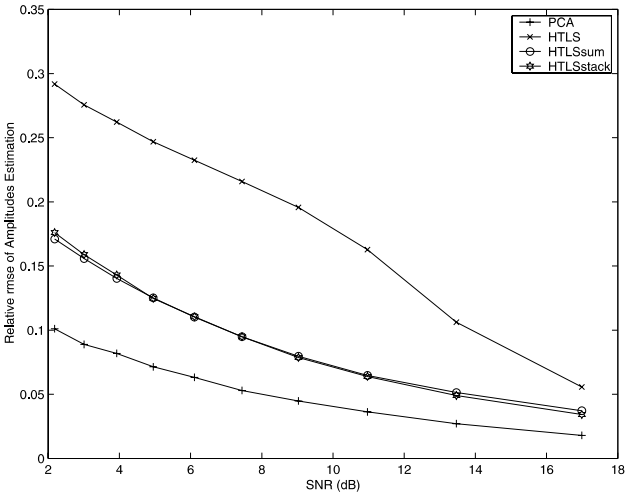

Fig. 4. rrmse versus SNR. Simulated set contains 100 identical single Gaussians.

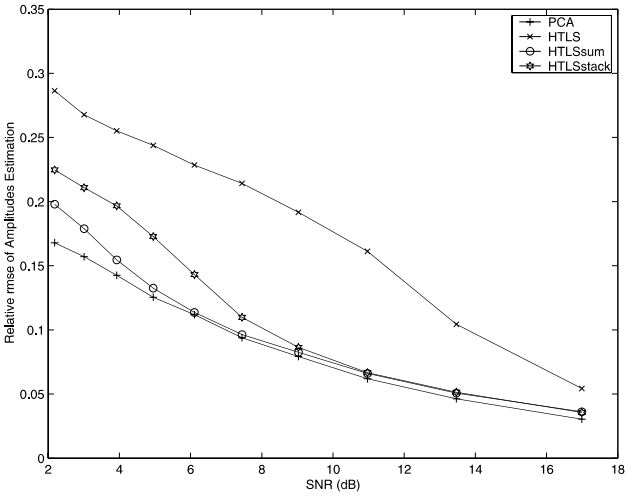

Fig. 5. rrmse versus SNR. Simulated set contains 10 identical single Gaussians.

algorithms do not perform as well as cPCA. This superiority in performance of cPCA clearly decreases for smaller data sets ($P < 20$), as shown in Fig. 5, and even vanishes at high SNR (>18 dB), as shown in [8,10], where HTLSsum and HTLSstack may exhibit similar to slightly better performances.

HTLS exhibits the worst performance while HTLSstack and HTLSsum show comparable results (although at low SNR and small P HTLSstack is slightly worse), indicating again that quantitation improves if all spectra are processed simultaneously. At high SNR each Gaussian is fitted by five Lorentzians ($J = 5$) but, as SNR decreases, the Gaussian is more likely fitted by fewer Lorentzians ($J < 5$) because the HTLS methods start to fit noise peaks as well. It is observed that J decreases to two at the lowest SNR (<4 dB), implying that only two (out of five) computed Lorentzians fall within the frequency range of the Gaussian. Similar conclusions hold when repeating the experiments on data sets with varying amplitudes [8,10].

3.3. Frequency alignment

We compare Witjes et al.'s [12] real-valued PCA method plus linear regression, abbreviated as rPCA-LR(f,p), complex PCA plus linear regression (cPCA-LR(f) and cPCA2-LR(f)) and cross correlation (cCross(f)) on simulated large spectral data sets of single resonances.

Simulation procedure. To enable a fair comparison, the same simulation example considered in [12] is taken here. A set of 100 single Lorentzian resonances ($P = 100$) with equal amplitude A, each containing 512 data points ($N = 512$) and centered at position 256 with linewidth $\tau = 30$ data points, is generated. Uniformly distributed phases (between $-90°$ and $90°$) and frequency shifts ($\pm\Delta\omega$ data points), as well as white Gaussian noise to both real and imaginary parts (mean 0 and variance $\sigma_v^2 = 1$), are added to each resonance of the data set. Amplitudes are estimated after frequency alignment using rPCA-LR(f,p), cPCA-LR(f), and cCross(f). The entire experiment is repeated 500 times in order to compute the sample rrmse of the amplitude A.

3.3.1. Comparison between real PCA-based versus complex PCA-based frequency shift alignment algorithms

In Fig. 6 the frequency shift $\Delta\omega$ is uniformly distributed between plus or minus half the linewidth ($\Delta\omega = \pm 15$ data points). The simulation procedure is repeated for varying amplitudes and unit noise variance. The sample rrmse is plotted versus the resulting SNR, showing that cPCA-LR(f) is better than rPCA-LR(f,p), which is unstable due to the large phase shift.

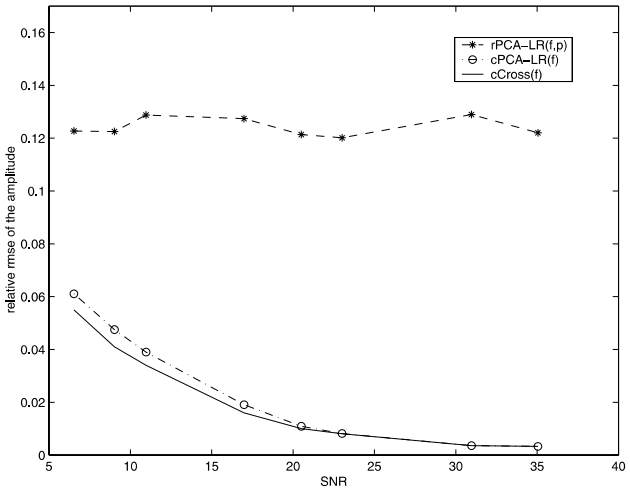

Fig. 6. rrmse of amplitude estimation versus SNR.

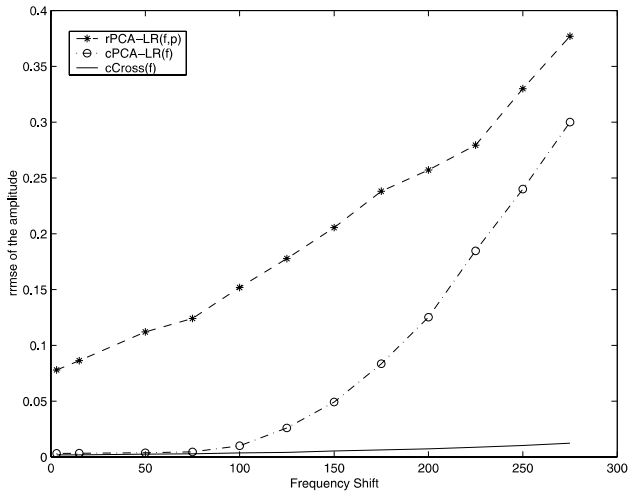

Fig. 7. Frequency Shift of the horizontal axis by Frequency Shift (in data points).

In Fig. 7, the amplitude A is set to 50 and $\sigma_v = 1$ but the frequency shifts $\Delta\omega$ are varied uniformly from 0 up to 10 times the linewidth (300 data points). The sample rrmse is plotted versus these frequency shifts. At large frequency shifts, it is obvious that rPCA-LR(f,p) as well as cPCA-LR(f) fail but rPCA-LR(f,p) performs clearly worse than cPCA-LR(f) at all frequency shifts.

3.3.2. Comparison between PCA-based algorithms and cross correlation

The advantages of PCA-related frequency alignment methods are based on processing the whole data set simultaneously, which can be described as using the common (global) information in the data set. On the other hand, cross correlation can only find frequency shifts one by one, where the accuracy mainly depends on the SNR of each individual spectrum (local information).

In Fig. 6, the cross-correlation algorithm demonstrates the lowest relative rrmse because the frequency shifts are fairly large, although the differences with cPCA-LR(f) are (almost) negligible. In Fig. 7, the performance of cross correlation almost remains the same no matter how large the frequency shift is, demonstrating the robustness of the method. Also observe that cPCA-LR(f) performs worse than cCross(f) for large frequency shifts (>2 times the linewidth).

We now consider again the same simulated data set, as described in the beginning of this section ($\Delta\omega = \pm 15$ data points) but only include 50 resonances ($P = 50$) and increase the Gaussian noise variance σ_v^2 to 25. The amplitudes of the first 25 resonances are set to 100 and the remaining ones are set to 1000. cCross(f) takes the

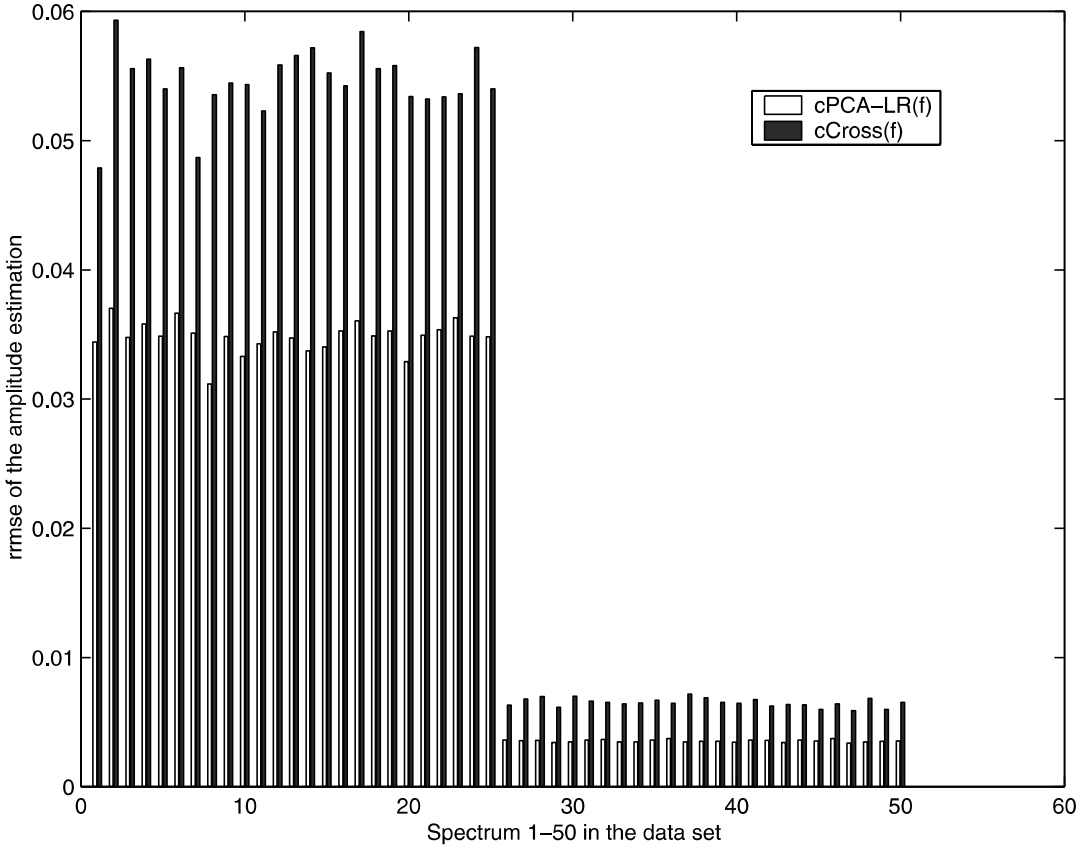

Fig. 8. Comparison in rrmse of 50 amplitudes, as estimated by cPCA-LR(f) and cCross(f). Amplitudes of the first 25 resonances are set to 100 and the remaining ones are set to 1000. $\sigma_v^2 = 25$.

Table 1
Comparison in rrmse of 50 amplitudes, as estimated by cPCA-LR(f), $cCross_1(f)$ and $cCross_{50}(f)$, where subscript refers to the number of the resonance taken as reference

	cPCA-LR(f)	$cCross_1(f)$	$cCross_{50}(f)$
rrmse of amplitudes of resonances 1 to 25	0.0344	0.0528	0.0356
rrmse of amplitudes of resonances 26 to 50	0.0035	0.0062	0.0036

first resonance, which has an amplitude of 100, as its reference.

Fig. 8 and Table 1 compare the performance of cCross(f) and cPCA-LR(f) in rrmse and clearly show that cPCA-LR(f) yields better amplitude estimates than cCross(f).

The reason is twofold. On the one hand, it is due to the advantage of using global information. On the other hand, it is due to the choice of the reference for cCross(f). We have selected the worst possible choice for the reference spectrum to demonstrate the importance of choosing the reference for cCross(f). In applications, one should always choose the resonance of highest SNR as a reference. We select now another reference for the cross-correlation algorithm with amplitude equal to 1000 and repeat the same experiment. The result is given in the last column of Table 1.

Comparing the last two columns in Table 1, we observe that the performance of cCross(f) is improved by selecting a high SNR resonance as a reference. But cPCA-LR(f) still exhibits the best accuracy, in particular for quantitation of low SNR spectra.

3.3.3. Frequency shifts smaller than the spectral resolution

We consider here frequency shifts which are smaller than the spectral resolution Δf $(=\frac{1}{N}\Delta t)$ of the given spectra. Apart from quantitation purposes, removal of small frequency shifts could play an important preprocessing role for any further processing tasks, such as classification. Under these conditions, the differences in alignment precision between cCross(f) and cPCA-LR(f) become very well pronounced. Fig. 9 shows that cCross(f) cannot recognize frequency shifts smaller than Δf. It can only move the peaks to positions which are integer multiples of Δf. But cPCA-LR(f) does not have this limitation.

Using the same simulation procedure described in the beginning of this section (with $\Delta\omega = \pm 15$ data points) and repeated for varying amplitude A, we can mathematically demonstrate the advantage of cPCA-LR(f) by calculating the sample variance of the final central positions of the resonances after frequency-shift alignment versus the SNR in Fig. 10. The better the resonances are aligned, the lower the variance should be.

The advantages of cPCA-LR(f) over cCross(f) are clearly shown. At all SNR, the variance of the final central positions of the resonances after frequency correction by cPCA-LR(f) is lower than that of cCross(f).

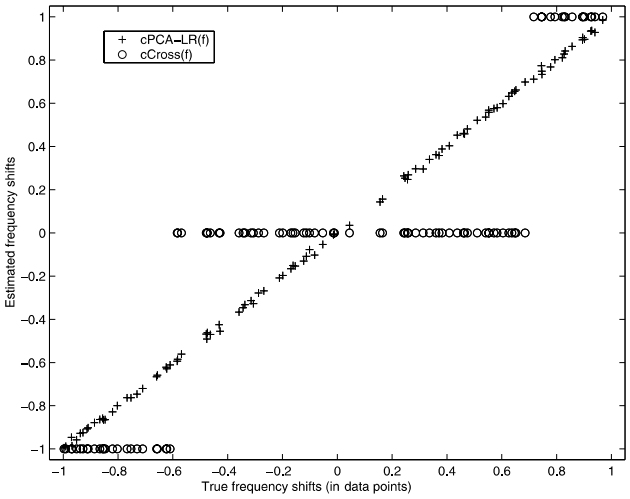

Fig. 9. Estimated versus true frequency shift variations below spectral resolution.

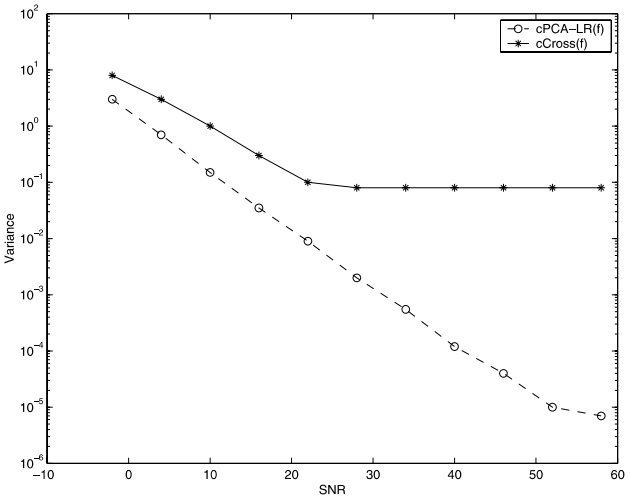

Fig. 10. Variances of the final central positions of the resonances after frequency correction by cPCA-LR(f) and cCross(f) versus SNR.

Furthermore, as SNR increases, cPCA-LR(f) continuously lowers the variance in contrast to cCross(f) of which the variance reaches a lower bound at a certain SNR (20 dB). This is because the limit of frequency resolution is reached, which implies that the variance of the final central positions of the resonances after correction by cCross(f) can no longer be lowered by increasing the SNR.

Zero-filling is a technique often used to improve digital resolution. By simply adding zeros at the tail of the time-domain signal, we can get a smaller Δf since N becomes bigger ($\Delta f = \frac{1}{N}\Delta t$). The following experiment tries to show how zero-filling can be adopted to increase the performance of cCross(f). A set of 100 noise-free ($\sigma_v^2 = 0$) single Lorentzian resonances with equal amplitude $A = 100$, each containing 16 data points and centered at position 8 with linewidth $\tau = 2$ data points, is generated. Uniformly distributed frequency shifts ($\Delta\omega = \pm 1$ data point) are added. We apply cCross(f) to the data set and obtain a result, similarly to the performance of cCross(f) in Fig. 9. We then zero-fill each signal in the data set to 32, 64, 128, 256, and 512 data points and then observe that the digital resolution of cCross(f) improves. If the number of data points reaches 512, cCross(f) visually yields a performance similar to cPCA-LR(f) in Fig. 9. However, after calculating the sample variance over 500 repetitions of the experiment, as we did in Fig. 10, we find out that zero-filling failed to decrease the variance of the final central positions of the resonances after frequency-shift alignment, as shown in Fig. 11. We can conclude from Fig. 11 that zero-filling cannot improve the accuracy of frequency-shift alignment algorithms. For small frequency shifts cPCA-LR(f) always demonstrates better performance than cCross(f).

3.3.4. Two ways to improve the performance of PCA-related algorithms

In this section, the simulated data sets generated for Figs. 6 (respectively, 7) are exactly the same as those used in Figs. 12 and 14 (respectively, 13 and 15).

One way to improve the performance of PCA-related algorithms is to add the second-order derivative of the

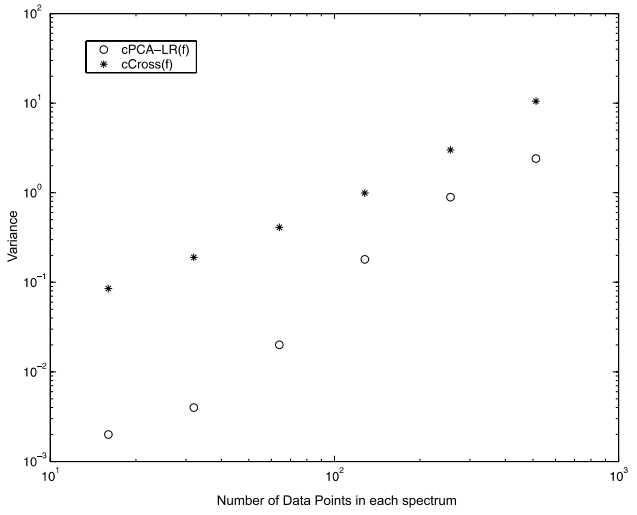

Fig. 11. Sample variance of the final central positions (expressed in data points) of the resonances after frequency correction by cPCA-LR(f) and cCross(f) on zero-filled data sets.

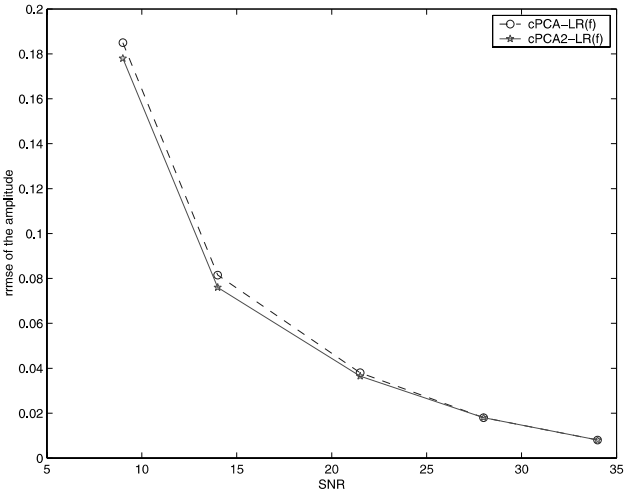

Fig. 12. rrmse of amplitude estimation versus SNR for cPCA-LR(f) and cPCA2-LR(f).

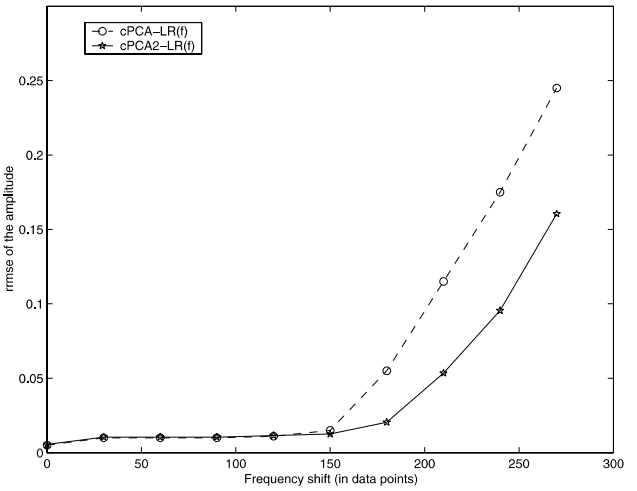

Fig. 13. rrmse of amplitude estimation versus frequency shift for cPCA-LR(f) and cPCA2-LR(f).

Taylor series to the linear regression equation. In Figs. 12 and 13, where both algorithms are tested for different SNR and frequency shifts, the advantage of cPCA2-LR(f) is obvious. cPCA2-LR(f) clearly yields better results when the frequency shift is large and the SNR is low.

The advantages of using fourth-order finite differences to calculate the derivatives of PCs are shown in Figs. 14 and 15. At low SNR and large frequency shifts, we notice the improvement of cPCA-LR(f) by using fourth-order finite differences instead of first order.

4. In vivo MRSI data of human brain

We have applied the frequency-alignment algorithms discussed above to an in vivo MRSI data set of the

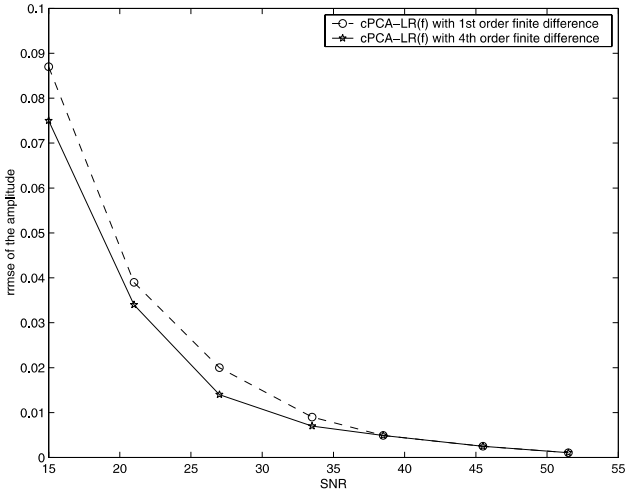

Fig. 14. rrmse of amplitude estimation versus SNR for cPCA-LR(f) using 1st and 4th order differences.

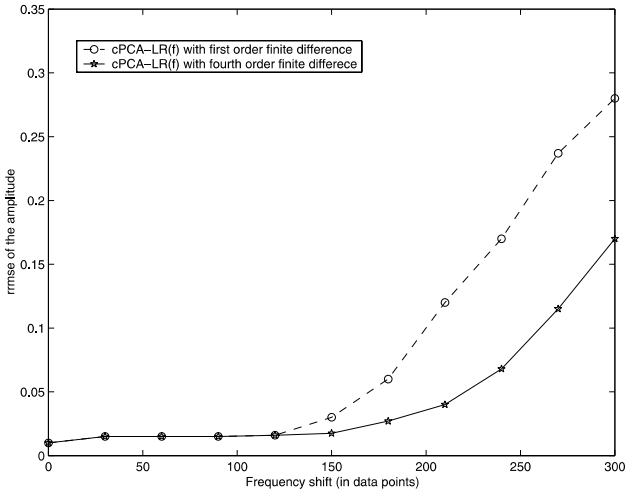

Fig. 15. rrmse of amplitude estimation versus frequency shift for cPCA-LR(f) using 1st and 4th order differences.

human brain, containing 16×16 in vivo signals measured on a 1.5 T Siemens Vision system, using a 2D CSI sequence with TR/TE 2000/135 ms, slice thickness 10 mm, and localization PRESS box, as shown in Fig. 16. There are 64 voxels inside the white box. Some voxels contain mainly water, since they are positioned in the ventricles of the human brain (>99% water). Ventricles can be seen in the middle of the MR image as the areas with the highest intensity. The dominant residual water resonances are located at 4.7 ppm. It is observed that some spectra mainly contain water and some neutron activation analysis (NAA) resonances are greatly distorted. After removal of all these greatly distorted spectra, the residual water resonance is first removed in the remaining spectra by means of HSVD [18]. By cutting the NAA resonance region from these filtered spectra, we obtain a data set of 21 NAA resonances. These peaks are located around 2 ppm (see Fig. 17).

We have applied cCross(f) and cPCA-LR(f) to the data set. The results are shown in Fig. 18. We observe that frequency shifts of the original data set are successfully removed by cCross(f) and cPCA-LR(f) (using four iterations). We then apply cPCA to check the variance distribution of the corrected spectra. After correction by cCross(f), the first principal component PC_1 explains 87.37% of the variance in the complete data set. After correction by cPCA-LR(f), PC_1 explains 88.02% of the variance. This small difference in variance is due to the low SNR and cannot be improved here by first applying cCross(f) followed by cPCA-LR(f). The first six PCs, together with their corresponding normalized eigenvalues (squared singular values), are displayed in Fig. 19. Due to remaining lineshape variations, more than one PC still contains resonance information. According to Eq. (13) in [14], amplitude estimates of the NAA resonances can be computed in the presence of lineshape variations, as the magnitude value of the sum of weighted scores of the K most informative PCs, i.e.,

$$A = \left| \sum_{j=1}^{K} S_j \left(\sum_{i=1}^{N} v_{ij} \right) \right|, \tag{21}$$

where $|.|$ denotes the magnitude value and $S_j = u_j \sigma_i$ is the jth score.

If the goal of the analysis is resonance quantitation, no lineshape corrections are needed [14]. Fig. 20 shows that the NAA amplitude estimates, computed by cPCA-LR(f) using two PCs ($K = 2$ in Eq. (21)), are comparable ($R^2 = 0.7$) to those computed by a commonly used nonlinear curve fitting method AMARES [5] (assuming Gaussian lineshapes), thereby confirming the applicability of cPCA-LR(f) in MRSI. Further increasing the number of PCs in Eq. (21) no longer improves R^2 since the higher PCs mainly describe noise. Note that an additional lineshape correction procedure can further increase the variance contribution in the first PC. This is currently under investigation.

5. Conclusions

In this paper we have presented new complex-valued PCA-based methods (cPCA, cPCA-LR(f)) and analyzed their advantages and limitations for quantifying and aligning complex-valued single resonances in large MR spectral sets originating from time series or MRSI experiments. The advantage of PCA-based algorithms lies in their ability to process the data set as a whole, thereby improving accuracy of quantitation compared to methods who process the spectra separately such as HTLS. Moreover, cPCA-based methods are insensitive to any

Fig. 16. MR image of the human brain and a grid indicating the positioning of the voxels in the MRSI data set. Only the voxels within the box indicated with the thick white line (8×8) contain useful information.

phase shift and, in addition, the use of the entire complex-valued data set improves their parameter accuracy compared to that of their real-valued counterparts which only use the real part of the spectra (rPCA). For quantitation purposes, cPCA demonstrates good performance in terms of both accuracy and computational efficiency for MRS data quantitation of spectral data sets consisting of one single resonance. However, as soon as the data set contains more than one resonance, cPCA can no longer be directly applied for quantitation. In these cases it is still possible to apply cPCA after filtering out one single resonance at the expense of an increase in the bias of the amplitude estimate. HTLS-stack and HTLSsum simultaneously quantify multiple resonances and perform best if the lineshapes are Lorentzian. The same results hold for unphased sets. In addition, PCA-based frequency-shift correction is shown to improve amplitude estimation across large data sets of frequency-shifted single resonances by using complex data. Furthermore, it is insensitive to phase shifts, which makes correction of phase shifts unnecessary. The use of higher order derivatives of Taylor series and higher order finite differences can further improve the performance of PCA-related frequency-shift algorithms. The cross-correlation algorithm is robust and

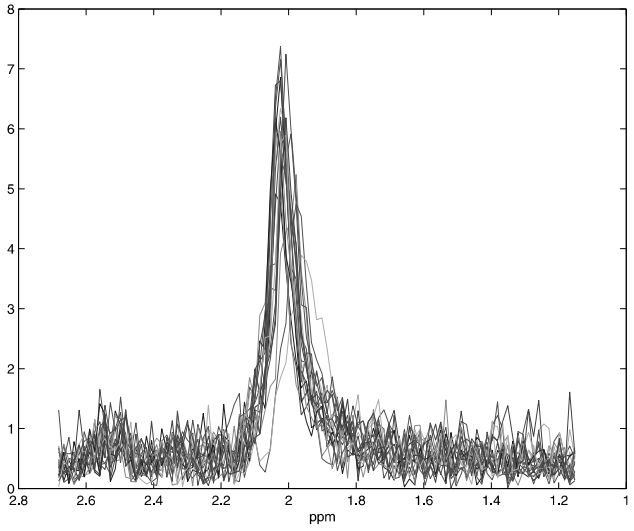

Fig. 17. NAA resonances from the original data set.

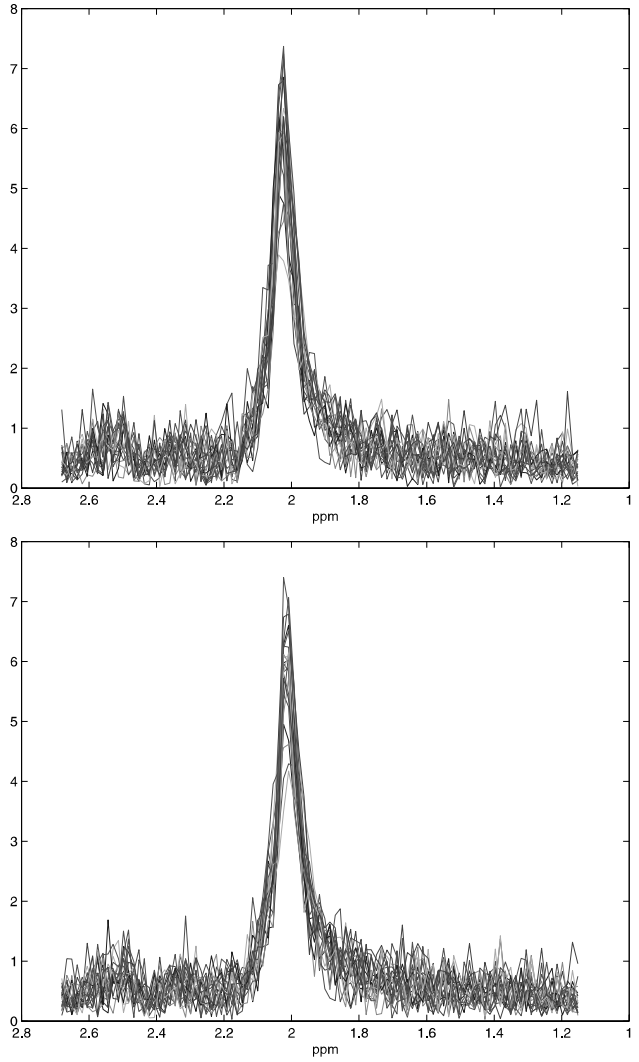

Fig. 18. NAA resonances after frequency correction by cCross(f) (top) and cPCA- LR(f) (bottom).

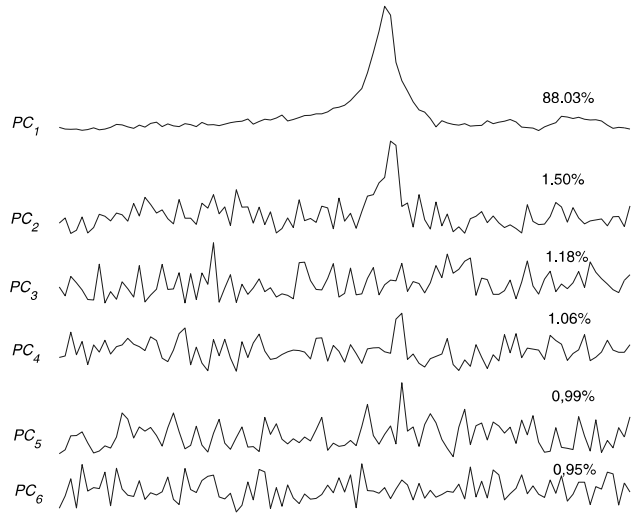

Fig. 19. Magnitude spectrum of the first six principal components, together with their normalized eigenvalues.

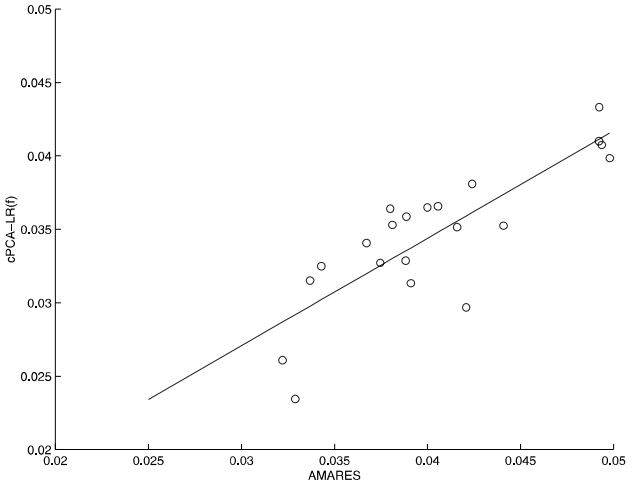

Fig. 20. cPCA-LR(f) amplitude estimates (computed from the first 2 PCs) versus AMARES amplitude estimates of 21 NAA resonances.

computationally efficient. Furthermore, cross correlation can align arbitrary large frequency shifts. These properties make it a good candidate as a starting algorithm for automatic frequency alignment. After that, PCA-related algorithms could be applied to further refine the results (provided the SNR is high enough) since they are able to correct extremely small shifts, which are even smaller than the discrete spectral resolution. These results are confirmed on simulated as well as on in vivo MRSI data sets. Future research will be focused on lineshape correction procedures using complex PCA. Although lineshape corrections are not needed for resonance amplitude quantitation [14], these corrections are recommended when the goal is to end up with a single signal-related PC, as may be the case prior to pattern recognition, e.g., for tumor classification.

Acknowledgments

The authors thank R. Stoyanova, T. Brown, and H. Witjes for discussions on PCA-based methods and Drs. F. Schubert and A. Link for providing information on their research of cross correlation [17]. Leentje Vanhamme is a postdoctoral researcher of the F.W.O. (Fund for Scientific Research—Flanders). This paper presents research results of the Belgian Programme on Interuniversity Poles of Attraction (IUAP Phase V-10-29), initiated by the Belgian State, Prime Minister's Office—Federal Office for Scientific, Technical and Cultural Affairs, of the Concerted Research Action (GOA) projects of the Flemish Government MEFISTO-666 (Mathematical Engineering for Information and Communication Systems Technology), of the IDO/99/03 project (K.U. Leuven) and the FWO Projects G.0269.02, G.078.01 and G.0360.98.

References

[1] C. Decanniere, P. Van Hecke, F. Vanstapel, H. Chen, S. Van Huffel, C. Van Der Voort, B. Van Tongeren, D. Van Ormondt, Evaluation of signal processing methods for the quantitation of strongly overlapping peaks in 31P NMR spectra, J. Magn. Reson. B 105 (1994) 31–37.

[2] T.R. McKnight, S.M. Noworolski, D.B. Vigneron, S.J. Nelson, D.R. Nat, An automated technique for the quantitative assessment of 3D-MRSI data from patients with glioma, J. Magn. Reson. Imaging 13 (2001) 167–177.

[3] M.C. Preul, Z. Caramanos, R. Leblanc, J.G. Villemure, D.L. Arnold, Using pattern analysis of in vivo proton MRSI data to improve the diagnosis and surgical management of patients with brain tumors, NMR Biomed. 11 (1998) 192–200.

[4] S. Van Huffel, H. Chen, C. Decanniere, P. Van Hecke, Algorithm for time-Domain NMR data fitting based on total least squares, J. Magn. Reson. A 110 (1994) 228–237.

[5] L. Vanhamme, S. Van Huffel, Multichannel quantitation of biomedical magnetic resonance spectroscopy signals, in: F.T. Luk (Ed.), Advanced Signal Processing Algorithms. Architectures, and Implementations VIII, vol. 3461, SPIE Press, 1998, pp. 237–248.

[6] R. Stoyanova, A.C. Kuesel, T.R. Brown, Application of principal-component analysis for NMR spectral quantitation, J. Magn. Reson. A 115 (1995) 265–269.

[7] M.A. Elliott, G.A. Walter, A. Swift, K. Vandenborne, J.D. Schotland, J.S. Leigh, Spectral quantitation by principal component analysis using complex singular value decomposition, Magn. Reson. Med. 41 (1999) 450–455.

[8] Y. Wang, S. Van Huffel, L. Vanhamme, P. Van Hecke, "Magnetic Resonance Spectroscopic Quantitation via Complex Principal Component Analysis and Hankel Total Least Squares-Based Methods," Internal Report 00-81, ESAT-SISTA, Katholieke Universiteit Leuven, Leuven, Belgium, 2000.

[9] T.R. Brown, R. Stoyanova, NMR Spectral quantitation by principal-component analysis II, determination of frequency and phase shifts, J. Magn. Reson. B 112 (1996) 32–43.

[10] Y. Wang, S. Van Huffel, E. Heyvaert, L. Vanhamme, N. Mastronardi, P. Van Hecke, Magnetic resonance spectroscopic quantitation via complex principal component analysis, in: Y. Baozong, T. Xiaofang (Eds.), Proceedings of the 2000 5th International Conference on Signal Processing, World Computer Congress, Beijing, China, 2000, vol. III, IEEE Press, New York, 2000, pp. 2074–2077.

[11] L. Vanhamme, T. Sundin, P. Van Hecke, S. Van Huffel, R. Pintelon, Frequency selective quantitation of biomedical magnetic resonance spectroscopy data, J. Magn. Reson. 143 (2000) 1–16.

[12] H. Witjes, W.J. Melssen, H.J.A. in 't Zandt, M. van der Graaf, A. Heerschap, L.M.C. Buydens, Automatic correction for phase shifts, frequency shifts and lineshape distortions across a series of single resonance lines in large spectral datasets, J. Magn. Reson. 144 (2000) 35–44.

[13] R. Stoyanova, T.R. Brown, NMR Spectral quantitiation by principal-component analysis, NMR Biomed. 14 (2001) 271–277.

[14] R. Stoyanova, T.R. Brown, NMR Spectral quantitiation by principal-component analysis III. A generalized procedure for determination of lineshape variations, J. Magn. Reson. 154 (2002) 163–175.

[15] S. Van Huffel, J. Vandewalle, in: The total Least Squares Problem: Computational Aspects and Analysis, Frontiers in Applied Mathematics Series, vol. 9, SIAM, Philadelphia, PA, 1991.

[16] L.N. Trefethen, Spectral Methods in Matlab, SIAM, Philadelphia, PA, 2000.

[17] F. Schubert, F. Seifert, C. Elster, A. Link, M. Walzel, H. Rinneberg, Improvement of the analytical quality of MR spectroscopy data by frequency corrected averaging, MAGMA 11 (Suppl. 1) (2000) 189.

[18] 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 (1992) 122–134.