Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Magnetic Resonance 168 (2004) 259–272
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

Parametric methods for frequency-selective MR spectroscopy—a review☆

Niclas Sandgren,[a,*] Yngve Selén,[a] Petre Stoica,[a] and Jian Li[b]

[a] *Systems and Control Division, Department of Information Technology, Uppsala University, P.O. Box 337, SE-751 05 Uppsala, Sweden*
[b] *Department of Electrical and Computer Engineering, University of Florida, P.O. Box 116130, Gainesville, FL 32611, USA*

Received 17 December 2003; revised 11 March 2004
Available online 12 April 2004

Abstract

Accurate quantitation of the spectral components in a pre-selected frequency band for magnetic resonance spectroscopy (MRS) signals is a frequently addressed problem in the MR community. One obvious application for such a frequency-selective technique is to lower the computational burden in situations when the measured data sequence contains too many samples to be processed using a standard full-spectrum method. Among the frequency-selective methods previously proposed in the literature, only a few possess the two features of primary concern: high robustness against interferences from out-of-band components and low computational complexity. In this survey paper we consider five spectral analysis methods which can be used for MRS signal parameter estimation in a selected frequency band. We re-derive the filter diagonalization method (FDM) in a new way that allows an easy comparison to the other methods presented. Then we introduce a frequency-selective version of the method of direction estimation (MODE) which has not been applied to MR-spectroscopy before. In addition, we present a filtering and decimation technique using a maximum phase bandpass FIR-filter and relate it to a similar ARMA-modeling approach known as SB-HOYWSVD (sub-band high-order Yule–Walker singular value decomposition). Finally, we study the numerical performances of these four methods and compare them to that of the recently introduced SELF-SVD (Singular Value Decomposition-based method usable in a SELected Frequency band) in several examples using simulated MR data, and discuss the benefits and disadvantages of each technique.

Keywords: Frequency-selective spectroscopy; Damped sinusoidal model; Filter diagonalization method; Method of direction estimation; Filtering and decimation; Singular value decomposition; Yule–Walker

1. Introduction and problem formulation

Consider the following model commonly used for MR data:

$$y(t) = \sum_{k=1}^{\bar{n}} \rho_k \lambda_k^t + e(t); \qquad \lambda_k = \mathrm{e}^{-\alpha_k + \mathrm{i}\omega_k} \quad t = 0, \ldots, N-1, \tag{1}$$

where $\bar{n}$ denotes the total number of components, λ_k are the modes of the signal, $\{\rho_k, \alpha_k, \omega_k\}$ are the complex amplitude, damping, and angular frequency of the kth component (note that the sampling period has been absorbed in α_k and ω_k, for notational convenience), $e(t)$ is a noise term and N denotes the number of available samples. The parameters $\bar{n}$ and $\{\rho_k, \alpha_k, \omega_k, k = 1, \ldots, \bar{n}\}$ are unknown. Note that the model in Eq. (1) assumes a Lorentzian lineshape for each of the signal components in $y(t)$. We do not address the case of non-Lorentzian lineshapes in this survey.

In several applications, where the measured data sequence contains many data samples (e.g., between 10,000 and 100,000) and, moreover, when the number of components could be thousands, it would be too computationally intensive to estimate the parameters of all resonances present in Eq. (1) using a standard time-domain method based on, e.g., the singular value decomposition (SVD). For this and other reasons we

☆ This work was partly supported by the Swedish Science Council (VR), the Swedish Foundation for International Cooperation in Research and Higher Education (STINT), and the National Science Foundation Grant OCR-0097114.
* Corresponding author. Fax: +46-18-503611.
E-mail address: niclas.sandgren@it.uu.se (N. Sandgren).

1090-7807/$ - see front matter © 2004 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2004.03.011

might be interested in only a few components of Eq. (1) that lie in a pre-specified frequency band:

$$[\omega_{\mathrm{l}}, \omega_{\mathrm{u}}], \tag{2}$$

where ω_{l} and ω_{u} denote the lower and upper limits of the frequency band of interest. We assume that the number of components of Eq. (1), lying in the frequency band of Eq. (2), which we denote by

$$n \leqslant \bar{n} \tag{3}$$

is given. If n was a priori unknown, it could be estimated from the data by counting the peaks of the FFT (fast Fourier transform) of Eq. (1) that lie in the frequency band of Eq. (2), or as the effective rank of a certain data matrix (see, e.g., [7]).

The problem is to estimate the parameters of the n components of Eq. (1) that lie in the frequency band of Eq. (2). Several methods have been developed for frequency-selective parameter estimation (see, e.g., [1,7–14] and the references therein). The so-called filter diagonalization method (FDM) [1,19–21] is frequently considered in the literature on applications of MR spectroscopy in, e.g., chemistry and biomedicine. As shown in the next section, the theory of FDM can in our opinion be simplified by introducing it from a new signal processing perspective. This makes it easier to compare FDM analytically to the other methods presented in this paper.

Since the performance of FDM has already been compared and found to be superior (see [21]) to that of the LP-ZOOM (see, e.g., [13]), we do not consider the LP-ZOOM in this survey. Instead, we present a new spectroscopic interpretation of the beamspace version of MODE (see, e.g., [4–6]) originally designed for array signal processing applications, which we will refer to as SELF-MODE (a MODE-based technique that is usable in a SELected Frequency band). In addition, we make use of data filtering and decimation to introduce a new approach based on the maximum-phase bandpass FIR-filter method of [9] which is very similar, conceptually, to the ARMA-modeling based filtering and decimation technique called SB-HOYWSVD (see [8]). We call this new approach the FIDO (FIltering and DOwnsampling) method. Finally, we compare these four methods to the recently developed SELF-SVD technique (see [7]) which can be viewed as a frequency-selective implementation of ESPRIT (see, e.g., [15]). Somewhat similarly (see Section 2 and [20]), FDM can be viewed as a frequency-selective implementation of the related state-space and matrix pencil methods of spectral analysis (see, e.g., [2,22]).

In Section 2 we present brief derivations of FDM, SELF-MODE, and FIDO. In addition, the presentation of each of the three methods outlined there is concluded with a step-by-step description, to help the reader write her own implementations. The descriptions of SB-HOYWSVD and SELF-SVD are very brief and located in the Appendix A, as we add very little to what is already known about these two methods. In Section 3 we present several numerical examples to illustrate the gain in using a frequency-selective approach in contrast to a standard full-spectrum method (see, e.g., [17,22]) and to give a fair comparative analysis of some new and some well-established sub-band MR techniques. Finally, in the concluding Section 4, we discuss the similarities and differences between, and the advantages and disadvantages of all five methods (FDM, SELF-MODE, FIDO, SB-HOYWSVD, and SELF-SVD).

2. Overview of three methods

2.1. The FDM revisited

By using a mathematical framework that has roots in the signal processing literature we are able to present the FDM in a practically equivalent but, in our opinion, simpler way compared to how it has been described earlier in e.g. [1] and [20]. Previous derivations of FDM involve no requirement for prior information on n. In this derivation, however, we assume n to be known. This is a relatively minor difference since e.g. [1] computes the parameter estimates of all modes that have an error (see below) lower than some pre-specified threshold, whereas here we estimate the parameters of the n modes with the smallest error (no threshold required). However, this difference indicates that the practical application we are interested in this paper is different than that of e.g. [20] (see more comments in Section 4). We start this derivation by considering the entire spectral range, and later introduce the simple modification required for the sub-band analysis.

Consider the data sequence given in Eq. (1), but disregard the noise term $(e(t))$ for the moment, and rewrite this data sequence as

$$\begin{aligned} y(g+h+j) &= \sum_{k=1}^{\bar{n}} \rho_k \lambda_k^{g+h+j} \\ &= \begin{bmatrix} \lambda_1^g \cdots \lambda_{\bar{n}}^g \end{bmatrix} \begin{bmatrix} \rho_1 \lambda_1^j & & 0 \\ & \ddots & \\ 0 & & \rho_{\bar{n}} \lambda_{\bar{n}}^j \end{bmatrix} \begin{bmatrix} \lambda_1^h \\ \vdots \\ \lambda_{\bar{n}}^h \end{bmatrix}. \end{aligned} \tag{4}$$

Define the $m \times m$ Hankel data matrix $\mathbf{Y}_j$ as

$$\mathbf{Y}_j = [y(g+h+j)]_{g,h=0,\ldots,m-1} = \mathbf{A}\mathbf{D}_j\mathbf{A}^T, \tag{5}$$

where

$$\mathbf{A} = [\mathbf{a}(\lambda_1) \cdots \mathbf{a}(\lambda_{\bar{n}})] = \begin{bmatrix} 1 & \cdots & 1 \\ \lambda_1 & \cdots & \lambda_{\bar{n}} \\ \vdots & & \vdots \\ \lambda_1^{m-1} & \cdots & \lambda_{\bar{n}}^{m-1} \end{bmatrix}, \tag{6}$$

$$\mathbf{D}_j = \begin{bmatrix} \rho_1\lambda_1^j & & 0 \\ & \ddots & \\ 0 & & \rho_{\bar{n}}\lambda_{\bar{n}}^j \end{bmatrix} \tag{7}$$

and where m is an integer user parameter (which we call the "snapshot length"). Note that one could consider several other alternative definitions of the Hankel data matrix $\mathbf{Y}_j$. In the following we will choose m as close to $N/2$ as possible; and hence, for simplicity, we will consider the Hankel data matrix as it is defined in Eq. (5), instead of a perhaps more general non-square matrix.

Construct the *matrix pencil* (see, e.g., [2,3]) $\mathbf{Y}_1 - \lambda\mathbf{Y}_0$. If $m \geqslant \bar{n}$ (which we assume in the following), the solutions to the singular *generalized* eigenvalue problem:

$$(\mathbf{Y}_1 - \lambda\mathbf{Y}_0)\mathbf{v} = \mathbf{A}(\mathbf{D}_1 - \lambda\mathbf{D}_0)\mathbf{A}^T\mathbf{v} = 0 \tag{8}$$

subject to $\mathbf{v} \in \mathcal{R}(\mathbf{Y}_0) = \mathcal{R}(\mathbf{A})$, where $\mathcal{R}(\mathbf{Y}_0)$ denotes the column space of $\mathbf{Y}_0$, are given by

$$\lambda = \lambda_p \text{ (generalized eigenvalue of the matrix pencil } \mathbf{Y}_1 - \lambda\mathbf{Y}_0) \tag{9}$$

$$\mathbf{v} = \mathbf{v}_p \text{ (right generalized eigenvector associated with the eigenvalue } \lambda_p) \tag{10}$$

for $p = 1, \ldots, \bar{n}$. Hence each element in $\{\lambda_p\}_{p=1}^{\bar{n}}$ is a rank reducing number for the matrix pencil $\mathbf{Y}_1 - \lambda\mathbf{Y}_0$ and can be determined as such. We note that $\{\lambda_p\}_{p=1}^{\bar{n}}$ are also the eigenvalues of $\mathbf{Y}_j - \lambda\mathbf{Y}_{j-1}$, a fact which we will use later (it is also used in [1]) to check the validity of each estimated mode.

By writing Eq. (8) as:

$$\left(\tilde{\mathbf{A}}\tilde{\mathbf{D}}\tilde{\mathbf{A}}^T\right)\mathbf{v}_p = 0, \tag{11}$$

where

$$\tilde{\mathbf{A}} = \begin{bmatrix} 1 & \cdots & 1 & 1 & \cdots & 1 \\ \lambda_1 & \cdots & \lambda_{p-1} & \lambda_{p+1} & \cdots & \lambda_{\bar{n}} \\ \vdots & & \vdots & \vdots & & \vdots \\ \lambda_1^{m-1} & \cdots & \lambda_{p-1}^{m-1} & \lambda_{p+1}^{m-1} & \cdots & \lambda_{\bar{n}}^{m-1} \end{bmatrix} \tag{12}$$

and

$$\tilde{\mathbf{D}} = \operatorname{diag}\big(\rho_1(\lambda_1 - \lambda_p), \ldots, \rho_{p-1}(\lambda_{p-1} - \lambda_p), \rho_{p+1}(\lambda_{p+1} - \lambda_p), \ldots, \rho_{\bar{n}}(\lambda_{\bar{n}} - \lambda_p)\big) \tag{13}$$

we easily see that

$$\mathbf{a}^T(\lambda_k)\mathbf{v}_p = 0 \qquad \text{for } k \neq p. \tag{14}$$

Hence, since

$$\hat{\mathbf{y}} = \mathbf{A}\boldsymbol{\psi}, \tag{15}$$

where

$$\hat{\mathbf{y}} = \begin{bmatrix} y(0) \\ \vdots \\ y(m-1) \end{bmatrix} \qquad \boldsymbol{\psi} = \begin{bmatrix} \rho_1 \\ \vdots \\ \rho_{\bar{n}} \end{bmatrix}$$

we get

$$\mathbf{v}_p^T\hat{\mathbf{y}} = \mathbf{v}_p^T\mathbf{a}(\lambda_p)\rho_p, \qquad p = 1, \ldots, \bar{n}. \tag{16}$$

From Eq. (16) we can estimate the amplitudes $\{\rho_k\}_{k=1}^{\bar{n}}$ in Eq. (4) directly by using the previously obtained estimates of $\{\lambda_k\}_{k=1}^{\bar{n}}$. Note that we only use the first "snapshot" ($[y(0)\cdots y(m-1)]^T$) in Eq. (16). One could also compute the amplitude estimates as the average over all available snapshots (i.e., $[y(t)\cdots y(t+m-1)]^T$, $t = 0, \ldots, N-m$). For each $t > 0$ we would then have to compensate for the time delay t, as we would compute estimates of $\{\rho_p\lambda_p^t\}$ rather than $\{\rho_p\}$. However, as we have found empirically, considering all available snapshots in Eq. (16) does not improve the accuracy, but only increases the computational burden (note that for MR signals, the first snapshots have the highest signal-to-noise ratios (SNR) due to the damping of the signal in the future snapshots).

We also outline an alternative approach for the amplitude estimation step. Normalize the generalized eigenvectors $\{\mathbf{v}_p\}_{p=1}^{\bar{n}}$ so that

$$\mathbf{v}_p^T\mathbf{Y}_0\mathbf{v}_k = \begin{cases} 1 & p = k \\ 0 & p \neq k. \end{cases} \tag{17}$$

For $p = k$, we can write Eq. (17) as (see also Eq. (14))

$$\left[\mathbf{v}_p^T\mathbf{a}(\lambda_p)\right]^2\rho_p = 1. \tag{18}$$

Using this result in Eq. (16) gives an expression for the amplitudes which is also used in [1] and [20]:

$$\rho_p = (\mathbf{v}_p^T\hat{\mathbf{y}})^2. \tag{19}$$

In the following we will use Eq. (16) for the amplitude estimation procedure since we have found empirically that it is slightly more accurate than the expression in Eq. (19). This concludes the derivation of the full-spectrum version of FDM.

Now, considering the modifications required for frequency-selectivity, we let $\{\lambda_k\}_{k=1}^{n}$ denote the modes of Eq. (1) whose angular frequencies lie in Eq. (2). We define a frequency grid which approximately spans the interval in Eq. (2) using the following Fourier frequencies:

$$\left\{\frac{2\pi k_1}{m}, \frac{2\pi k_2}{m}, \ldots, \frac{2\pi k_{\bar{K}}}{m}\right\}, \tag{20}$$

where k_j ($j = 1, \ldots, \bar{K}$) are $\bar{K}$ consecutive given integers. We also define the $(m \times \bar{K})$ sub-band transformation matrix $\boldsymbol{\Gamma}$ containing the Fourier vectors corresponding to the frequencies in Eq. (20) as

$$\boldsymbol{\Gamma} = \begin{bmatrix} 1 & \cdots & 1 \\ \gamma_{k_1} & \cdots & \gamma_{k_{\bar{K}}} \\ \vdots & & \vdots \\ \gamma_{k_1}^{m-1} & \cdots & \gamma_{k_{\bar{K}}}^{m-1} \end{bmatrix}, \quad \gamma_{k_j} = \mathrm{e}^{-\mathrm{i}\frac{2\pi}{m}k_j} \tag{21}$$

and solve the generalized eigenvalue problem as outlined above but for the dimension-reduced modified matrix pencil:

$$\mathbf{\Gamma}^T(\mathbf{Y}_1 - \lambda\mathbf{Y}_0)\mathbf{\Gamma}. \tag{22}$$

The pre-multiplication with $\mathbf{\Gamma}^T$ and post-multiplication with $\mathbf{\Gamma}$ constitute the frequency-selective operation that filters out all modes outside the frequency interval in Eq. (20) and selects only the modes of interest. This is due to the fact that $\mathbf{\Gamma}^T\mathbf{A}$ has approximately non-zero columns only for the modes in $\mathbf{A}$ with frequencies that lie within Eq. (2) which, in addition, have low dampings. So, ideally, the pre-multiplication with $\mathbf{\Gamma}^T$ and post-multiplication with $\mathbf{\Gamma}$ reduce the effective rank of $\mathbf{Y}_j$ from $\bar{n}$ to n. The following outline explains all required steps of the procedure in more detail.

2.1.1. Outline of FDM

Step 1. Determine the frequency band of interest $\left[\frac{2\pi k_1}{m}, \frac{2\pi k_{\bar{K}}}{m}\right]$ by inspection of the FFT of $\{y(t)\}_{t=0}^{N-1}$ (possibly the band of interest is pre-specified). Initially we want to estimate K components (where $n \leqslant K \leqslant \bar{K}$) within the selected interval; we will later ignore any estimated modes that lie outside the considered frequency interval and use only the remaining n modes with the lowest *errors* (as explained in Step 3). The choice to estimate $K = \bar{K}$ modes initially, in lieu of n, will result in better parameter estimates for practical MR signals (as we have discovered empirically).

Step 2. Set the user parameter m as large as possible (we use $m = \left\lfloor \frac{N-3}{2} \right\rfloor$ here, where $\lfloor\cdot\rfloor$ denotes the integer part). Construct the three Hankel data matrices $\{\hat{\mathbf{Y}}_j\}_{j=0}^2$ (where $\hat{\mathbf{Y}}_j$ is similarly defined to $\mathbf{Y}_j$, but now based on the noisy data $y(t)$ in Eq. (1) including the noise term $e(t)$) and pre and post multiply them by $\mathbf{\Gamma}^T$ and $\mathbf{\Gamma}$ to compute the dimension-reduced data matrices (see Eq. (22)), which we denote by $\{\mathbf{U}_j\}_{j=0}^2$:

$$\mathbf{U}_j = \mathbf{\Gamma}^T\hat{\mathbf{Y}}_j\mathbf{\Gamma} \qquad \text{for } j = 0, 1, 2. \tag{23}$$

Note that it is possible to implement Eq. (23) in a computationally efficient way (see [1]), which has been taken into account in the results on numerical complexity in Section 3. However, for conciseness reasons, we do not rederive the specific results of [1] for lowering the computational burden, associated with implementing Eq. (23).

Step 3. In order to solve the matrix pencil problem we need to compute the pseudo inverse ($\mathbf{U}_0^+$) of the matrix $\mathbf{U}_0$ [2]. Let

$$\mathbf{U}_0 = \mathbf{V}\mathbf{\Sigma}\mathbf{W}^*, \tag{24}$$

(where $*$ is the complex conjugation operator) denote the SVD of $\mathbf{U}_0$ where $\mathbf{V}$ and $\mathbf{W}$ are unitary and $\mathbf{\Sigma}$ is a diagonal matrix with non-negative main diagonal entries. Compute $\mathbf{U}_0^+$ as:

$$\mathbf{U}_0^+ = \hat{\mathbf{W}}\hat{\mathbf{\Sigma}}^{-1}\hat{\mathbf{V}}^*, \tag{25}$$

where $\hat{\mathbf{\Sigma}}$ is the diagonal matrix containing the K largest singular values of $\mathbf{U}_0$, and $\hat{\mathbf{V}}$ and $\hat{\mathbf{W}}$ contain the K corresponding columns of $\mathbf{V}$ and $\mathbf{W}$, respectively.

Consider the singular generalized eigenvalue problem of interest (see Eq. (8), Eq. (22), and Eq. (23)):

$$(\mathbf{U}_1 - \lambda_k\mathbf{U}_0)\mathbf{v}_k = 0, \tag{26}$$

where $\mathbf{v}_k$ is the right generalized eigenvector associated with the eigenvalue λ_k, $k = 1, \ldots, K$. Left multiplication with the truncated pseudo inverse $\mathbf{U}_0^+$ of $\mathbf{U}_0$ and using the fact that $\mathbf{U}_0^+\mathbf{U}_0 \approx \mathbf{I}$ (where $\mathbf{I}$ is the identity matrix) yield

$$\mathbf{U}_0^+\mathbf{U}_1\mathbf{v}_k \approx \lambda_k\mathbf{v}_k, \quad k = 1, \ldots, K. \tag{27}$$

Let $\mathbf{l}_1 = \{\lambda_k^{(1)}\}_{k=1}^K$ denote the K eigenvalues of the matrix $\mathbf{U}_0^+\mathbf{U}_1$. In a similar way, let $\mathbf{l}_2 = \{\lambda_k^{(2)}\}_{k=1}^K$ denote the K eigenvalues of $\mathbf{U}_1^+\mathbf{U}_2$ ($\mathbf{U}_1^+$ is computed similarly to $\mathbf{U}_0^+$). Below we will use the fact that $\lambda_k^{(1)} \approx \lambda_k^{(2)}$ for the $(k = 1, \ldots, n)$ modes which we are interested in. Eliminate any modes in $\mathbf{l}_1$ which lie outside the interval $\left[\frac{2\pi k_1}{m}, \frac{2\pi k_K}{m}\right]$ and define the error $(\epsilon_{\bar{k}})$ of each remaining estimated mode in $\mathbf{l}_1$ as:

$$\epsilon_k = \min_j \left|\lambda_k^{(1)} - \lambda_j^{(2)}\right|. \tag{28}$$

Finally, obtain $\{\lambda_k\}_{k=1}^n$ by selecting the n modes in $\mathbf{l}_1$ which have the smallest error and lie within the interval of interest. Note that it is possible (but not likely), that we end up with less than n modes inside the interval of interest. This, however, has never occurred in our simulation study.

Step 4. Conclude by estimating the amplitudes $\{\rho_k\}_{k=1}^n$ as (of Eq. (16)):

$$\rho_k = \frac{\mathbf{v}_k^T\mathbf{\Gamma}^T\hat{\mathbf{y}}}{\mathbf{v}_k^T\mathbf{\Gamma}^T\mathbf{a}(\lambda_k)}. \tag{29}$$

2.2. The SELF-MODE

We consider the MODE method (see [6]) for estimating the parameters of damped sinusoids from noisy data [5]. The following discussion is inspired by [4], which deals with beamspace array signal processing, a problem related to that stated in Section 1. In order to apply the MODE estimation technique to this problem, we form a snapshot vector $\breve{\mathbf{y}}(t)$ of the measured signal in Eq. (1)

$$\breve{\mathbf{y}}(t) = \begin{bmatrix} y(t) \\ \vdots \\ y(t+m-1) \end{bmatrix} \tag{30}$$

for some $m > \bar{n}$. We also define:

$$\mathbf{x}(t) = \begin{bmatrix} \rho_1\lambda_1^t \\ \vdots \\ \rho_{\bar{n}}\lambda_{\bar{n}}^t \end{bmatrix} \qquad \breve{\mathbf{e}}(t) = \begin{bmatrix} e(t) \\ \vdots \\ e(t+m-1) \end{bmatrix}.$$

Then we can write

$$\breve{\mathbf{y}}(t) = \mathbf{A}\mathbf{x}(t) + \breve{\mathbf{e}}(t), \tag{31}$$

where $\mathbf{A}$ is the Vandermonde matrix defined in Eq. (6). We form the following $(m \times m)$ sample covariance matrix to which we will apply our frequency-selective operations:

$$\mathbf{R} = \sum_{k=0}^{M-1} \breve{\mathbf{y}}(k)\breve{\mathbf{y}}^*(k) \approx \mathbf{A}\left[\sum_{k=0}^{M-1} \mathbf{x}(k)\mathbf{x}^*(k)\right]\mathbf{A}^*, \tag{32}$$

where the approximation holds for high SNRs (i.e., when $\breve{\mathbf{e}}(t)$ is close to zero), and where M is the total number of snapshots defined as

$$M = N - m + 1. \tag{33}$$

Our goal is to find a sub-band transformation matrix $\mathbf{C}$ which approximately eliminates the data components which lie outside the interval of interest and which in addition reduces the dimensionality of Eq. (32) to allow for a faster computation (see [4]).

Define, similarly to Eq. (20), the L Fourier frequencies which approximately span the interval in Eq. (2) as:

$$\left\{\frac{2\pi k_1}{m}, \frac{2\pi k_2}{m}, \ldots, \frac{2\pi k_L}{m}\right\}. \tag{34}$$

Also, let $\{\varphi_k\}_{k=1}^{m-L}$ denote the remaining $(m-L)$ Fourier frequencies which are *not* within the interval of interest. Note that the number of frequencies in Eq. (34) is lower than for FDM (see Eq. (20)) for the same frequency interval Eq. (2). This is due to the filtering operation below on the vectors $\{\breve{\mathbf{y}}(t)\}$ of length m.

Now consider the matrix $\mathbf{C}$, which is a $(m \times L)$ banded Toeplitz matrix of the form

$$\mathbf{C} = \begin{bmatrix} c_0 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ c_{m-L} & & c_0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & c_{m-L} \end{bmatrix}, \tag{35}$$

where $m > L$. The coefficient vector

$$\mathbf{c} = [c_0 \cdots c_{m-L}]^T \tag{36}$$

is computed as follows. Define the polynomial

$$c(z) = c_{m-L}z^{m-L} + \cdots + c_1 z + c_0 \tag{37}$$

with $(m-L)$ roots at $z_k = e^{i\varphi_k}, k = 1, \ldots, m-L$. The coefficients of $c(z)$ can be found by

$$\mathbf{c} = \mathrm{poly}\left[\left\{e^{i\varphi_k}\right\}_{k=1}^{m-L}\right], \tag{38}$$

where the operator poly[·] converts the roots to the polynomial coefficients. Note that Eq. (37) can be seen as an FIR filter which has nulls placed at the $m-L$ Fourier frequencies $\{\varphi_k\}$ which lie outside the interval of interest. We can thus expect that the filtering operation $\mathbf{C}^*\breve{\mathbf{y}}(t)$ should attenuate any components in $\breve{\mathbf{y}}(t)$ which are not within Eq. (2).

We will now show how to determine the parameter estimates of interest using the sub-band sample covariance matrix, obtained by pre and post multiplication of $\mathbf{R}$ by $\mathbf{C}^*$ and $\mathbf{C}$, which we will denote by $\mathbf{R}_s$:

$$\begin{aligned}\mathbf{R}_s = \mathbf{C}^*\mathbf{R}\mathbf{C} &= \sum_{k=0}^{M-1} \mathbf{C}^*\breve{\mathbf{y}}(k)\left[\mathbf{C}^*\breve{\mathbf{y}}(k)\right]^* \\ &\approx \mathbf{C}^*\mathbf{A}\left[\sum_{k=0}^{M-1} \mathbf{x}(k)\mathbf{x}^*(k)\right]\mathbf{A}^*\mathbf{C}.\end{aligned} \tag{39}$$

This is the key equation in the SELF-MODE approach. We observe that $\mathbf{R}_s$ can be seen as the $L \times L$ sample covariance matrix of the data sequence filtered by $c(z)$. We also show below the important fact that $\mathbf{C}^*\mathbf{A}$ is also a Vandermonde matrix (where each column is the corresponding column in $\mathbf{A}$ but truncated to length L) multiplied with a diagonal matrix where the diagonal entries will be "small" (close to zero) for all frequencies outside the interval of interest.

$$\begin{aligned}\mathbf{C}^*\mathbf{A} &= \begin{bmatrix} c_0^* & \cdots & c_{m-L}^* & \cdots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \cdots & c_0^* & \cdots & c_{m-L}^* \end{bmatrix}\begin{bmatrix} 1 & \cdots & 1 \\ \lambda_1 & \cdots & \lambda_{\bar{n}} \\ \vdots & & \vdots \\ \lambda_1^{m-1} & \cdots & \lambda_{\bar{n}}^{m-1} \end{bmatrix} \\ &= \begin{bmatrix} 1 & \cdots & 1 \\ \lambda_1 & \cdots & \lambda_{\bar{n}} \\ \vdots & & \vdots \\ \lambda_1^{L-1} & \cdots & \lambda_{\bar{n}}^{L-1} \end{bmatrix}\begin{bmatrix} \sum_{k=0}^{m-L} c_k^*\lambda_1^k & & 0 \\ & \ddots & \\ 0 & & \sum_{k=0}^{m-L} c_k^*\lambda_{\bar{n}}^k \end{bmatrix}.\end{aligned} \tag{40}$$

Now, let $\{b_k\}_{k=0}^{n}$ be the coefficients of the following polynomial:

$$b_0 z^n + \cdots + b_{n-1}z + b_n = \prod_{k=1}^{n}(z - \lambda_k), \tag{41}$$

where $\{\lambda_k\}_{k=1}^{n}$ are the modes of interest, and let $\mathbf{B}^*$ be the $(L-n) \times L$ banded Toeplitz matrix defined as

$$\mathbf{B}^* = \begin{bmatrix} b_n & \cdots & b_0 & & 0 \\ & \ddots & & \ddots & \\ 0 & & b_n & \cdots & b_0 \end{bmatrix}. \tag{42}$$

Also define the eigenvalue decomposition of $\mathbf{R}_s$ as

$$\mathbf{R}_s = [\mathbf{S}_s\ \mathbf{S}_r]\begin{bmatrix} \Sigma_s & \mathbf{0} \\ \mathbf{0} & \Sigma_r \end{bmatrix}\begin{bmatrix} \mathbf{S}_s^* \\ \mathbf{S}_r^* \end{bmatrix} \approx \mathbf{S}_s\Sigma_s\mathbf{S}_s^*, \tag{43}$$

where $\mathbf{S}_s$ is the matrix whose columns are the n principal eigenvectors of $\mathbf{R}_s$, and Σ_s is a diagonal matrix with the corresponding eigenvalues on the diagonal. Σ_r is a diagonal matrix with "small" diagonal elements due to the noise in the data and the presence of components outside the interval of interest. In the approximation in Eq. (43) it is assumed that the noise level is relatively low and that the nuisance components have been well attenuated by the filtering operation in Eq. (39), and thus that the elements in Σ_r are close to zero.

The important conclusion from the decomposition in Eq. (43) is that $\mathbf{S}_s$ approximately spans the signal subspace of interest and hence that $\mathcal{R}(\mathbf{S}_s) \approx \mathcal{R}(\mathbf{C}^*\mathbf{A})$ (see Eqs. (39) and (43)). Consequently, as $\mathbf{B}^*[\mathbf{C}^*\mathbf{A}] = 0$, which is easily verified considering Eqs. (41) and (42) and the "truncated proportionality" between $\mathbf{A}$ and $\mathbf{C}^*\mathbf{A}$ (see above) we also have

$$\mathbf{B}^*\mathbf{S}_s = \mathbf{0}. \tag{44}$$

An estimate of the vector $\mathbf{b}$

$$\mathbf{b} = [b_n \ \cdots \ b_0]^T \tag{45}$$

can be obtained by minimizing the least-squares (LS) cost function $\mathrm{tr}[\mathbf{B}^*\mathbf{S}_s(\mathbf{B}^*\mathbf{S}_s)^*] = \mathrm{tr}\left[\mathbf{B}\mathbf{B}^*\mathbf{S}_s\mathbf{S}_s^*\right]$ with respect to $\mathbf{b}$.

In [5] a minimization of a more refined weighted LS cost function was suggested:

$$f(\mathbf{b}) = \mathrm{tr}\left[\mathbf{B}\mathbf{F}_1\mathbf{B}^*\mathbf{S}_s\mathbf{F}_2\mathbf{S}_s^*\right], \tag{46}$$

where $\mathbf{F}_1$ and $\mathbf{F}_2$ are weighting matrices. It can be shown (see [5]) that good choices of the weighting matrices are $\mathbf{F}_2 = \Sigma_s$ and $\mathbf{F}_1 = (\mathbf{B}^*\mathbf{B})^{-1}$. This indicates that a two-step procedure is required for the minimization of Eq. (46), where an estimate of $\mathbf{B}$ is obtained in the first step by minimizing Eq. (46) with $\mathbf{F}_1 = \mathbf{I}$. The roots that are the solutions to the polynomial equation in Eq. (41) correspond to the n modes that we are interested in. The steps required for the suggested algorithm are outlined in the following.

2.2.1. Outline of SELF-MODE

Step 1. Construct the full sample covariance matrix $\mathbf{R}$ (Eq. (32)) of the measured data sequence.

Step 2. Compute the coefficients of the polynomial $c(z)$ using the $(m-L)$ Fourier frequencies $\{\varphi_k\}_{k=1}^{m-L}$ outside the interval of interest and construct the banded Toeplitz matrix $\mathbf{C}$. Use Eq. (39) to obtain the dimension-reduced covariance matrix $\mathbf{R}_s$.

Step 3. Compute the eigenvalues and the corresponding eigenvectors of $\mathbf{R}_s$. Let Σ_s denote the diagonal matrix containing the n principal eigenvalues and let $\mathbf{S}_s$ be the $(m \times n)$ matrix of the corresponding eigenvectors.

Step 4. Obtain an initial estimate of $\mathbf{b}$ by minimizing Eq. (46) with $\mathbf{F}_1 = \mathbf{I}$ and $\mathbf{F}_2 = \Sigma_s$, where $\mathbf{I}$ is the $(L-n) \times (L-n)$ identity matrix. The minimization step can be implemented as outlined below (see [5]).

Rewrite Eq. (46) as

$$f(\mathbf{b}) = \mathrm{vec}(\mathbf{B})^*\left(\mathbf{F}_1^T \otimes \mathbf{S}_s\mathbf{F}_2\mathbf{S}_s^*\right)\mathrm{vec}(\mathbf{B}), \tag{47}$$

where $\otimes$ denotes the *Kronecker* product and

$$\mathrm{vec}(\mathbf{B})^* = [\,\mathbf{b}^T\ \mathbf{0}^T\ \mathbf{b}^T\ \mathbf{0}^T\ \cdots\,]. \tag{48}$$

Let $\bar{\mathbf{\Omega}}$ denote the matrix $\mathbf{F}_1^T \otimes \mathbf{S}_s\mathbf{F}_2\mathbf{S}_s^*$ in Eq. (47), where the rows and columns corresponding to the zeros in Eq. (48) have been eliminated, and let $\mathbf{\Omega}$ be defined by

$$\mathbf{\Omega}^T = [\,\mathbf{I}\ \cdots\ \mathbf{I}\,]\bar{\mathbf{\Omega}}\begin{bmatrix}\mathbf{I}\\ \vdots \\ \mathbf{I}\end{bmatrix}. \tag{49}$$

Then Eq. (47) can be rewritten as

$$f(\mathbf{b}) = \mathbf{b}^*\mathbf{\Omega}\mathbf{b}. \tag{50}$$

Minimizing Eq. (50) with respect to $\mathbf{b}$, under a unit-norm constraint on $\mathbf{b}$, gives the total least square (TLS) solution, which is easily obtained as the eigenvector of $\mathbf{\Omega}$ corresponding to the smallest eigenvalue:

$$\hat{\mathbf{b}} = \text{the smallest eigenvector of} \mathbf{\Omega}. \tag{51}$$

Step 5. Derive the enhanced estimate of $\mathbf{b}$ by minimizing Eq. (46) as outlined in Step 4, but with $\mathbf{F}_1 = (\hat{\mathbf{B}}^*\hat{\mathbf{B}})^{-1}$, where $\hat{\mathbf{B}}$ is constructed using the previously obtained estimate $\hat{\mathbf{b}}$ of $\mathbf{b}$. Compute estimates of $\{\lambda_k\}_{k=1}^n$ from this enhanced estimate of $\mathbf{b}$ by using Eq. (41).

Step 6. Finally, use the estimates of $\{\lambda_k\}$ obtained in the previous step and the APES method [16] to estimate the amplitudes $\{\rho_k\}_{k=1}^n$ similarly to [7] (the acronym APES stands for Amplitude and Phase EStimation). Note that there are several other methods that could be used for amplitude estimation, once $\{\lambda_k\}$ have been estimated. Our choice of APES is motivated by the fact that this method provides the most accurate amplitude estimates of all methods of which we are aware, at a reasonable computational cost (see, e.g., [18]).

2.3. The FIDO method

A straightforward implementation of frequency-selective spectroscopy would be to simply make a passband filtering of the signal, passing only the frequency band of interest, and then applying any standard method to the filter output to estimate the parameters within that band. This approach was considered in [9]. The downside, however, is that the signal length is generally only slightly decreased by this filtering, and thus the computational complexity will still be very high. Here we refine the method of [9] by using a *downsampling* approach.

The idea of downsampling, or *data decimation*, is to divide the original data sequence into several shorter sequences with lower sampling frequency than the original data sequence. Specifically, by letting the positive integer d denote the *downsampling factor* we can form d different vectors of downsampled data sequences:

$$\begin{aligned}\mathbf{z}_k &= [y(k)\ y(d+k)\cdots y(d(N_k-1)+k)]^T \\ &= [z_k(0)\ z_k(1)\cdots z_k(N_k-1)]^T,\end{aligned} \tag{52}$$

where $k = 0, \ldots, d-1$ and $N_k \approx \frac{N}{d}$ is the length of vector $\mathbf{z}_k$ which is chosen to include as many samples of $\{y(t)\}$ as possible. Using these d sequences, a weighted average of the covariance matrix estimates can be computed as

$$\hat{\mathbf{R}}_D = \frac{1}{d}\sum_{k=0}^{d-1}\sum_{t=0}^{N_k-m}\frac{\tilde{\mathbf{z}}_k(t)\tilde{\mathbf{z}}_k^*(t)}{N_k-m+1}, \tag{53}$$

where $\tilde{\mathbf{z}}_k(t) = [z_k(t)\ z_k(t+1)\cdots z_k(t+m-1)]^T$ and m is usually chosen close to $\max_k\lfloor N_k/2\rfloor$ (in our numerical examples we set $m = \max_k\lfloor N_k/2\rfloor$). Using the estimated covariance matrix $\hat{\mathbf{R}}_D$ we can choose any method from the wealth of available sub-space methods operating on the data covariance matrix, to obtain the modes of the signal. In this paper we use the HSVD method [25].

As demonstrated in [23] and [24], downsampling may *increase* the resolution of sub-space based methods. Even more interesting, for our purposes, is that it can *decrease* the computational burden significantly, as the size of the covariance matrix estimate is reduced (m is chosen close to $\max_k\lfloor N_k/2\rfloor \approx \lfloor N/2d\rfloor$ instead of $\lfloor N/2\rfloor$ which would normally be the case if we did not apply decimation).

The problem with downsampling is that it can introduce aliasing. When downsampling the signal we increase the angular frequency of the signal components (as the sampling frequency is artificially reduced) by the downsampling factor d. This means that any peaks not lying within the angular frequency band $[-\pi/d, \pi/d]$ will cause aliasing, so for our decimation to be as effective as possible (i.e., to enable choosing a high d) we would like to have all peaks of interest within a small frequency band centered around the origin (i.e., the zero frequency) and we would also like to suppress the out-of-band components to reduce the effect of the aliasing *and* achieve frequency-selectivity. This can be achieved by first modulating the data so that the center of the band of interest is mapped to the origin, and then applying a passband filter to attenuate the out-of-band components (assuming these components are well separated from the band of interest). In this paper we use the passband FIR filtering algorithm of [9], as it has been shown to give good results for MR spectroscopy data. This filtering algorithm basically consists of using FIR filters of increasing orders until the out-of-band components are "sufficiently" suppressed. Below we give an outline of the FIDO method.

2.3.1. Outline of FIDO

Step 1. Modulate the data $\{y(t)\}$ by using the frequency $-\bar{\omega}$, where $\bar{\omega} = (\omega_u+\omega_l)/2$ denotes the *center frequency* of the band of interest $[\omega_l, \omega_u]$, so that a new data series $\{\bar{y}(t)\}$ is obtained

$$\bar{y}(t) = y(t)\mathrm{e}^{-\mathrm{i}\bar{\omega}t}, \quad t = 0,\ldots,N-1 \tag{54}$$

for which the frequency band of interest is centered around the zero frequency.

Step 2. Apply the FIR filtering algorithm of [9] to the data sequence $\{\bar{y}(t)\}$ to obtain the filtered data sequence $\{\bar{y}_F(t)\}$ which has most of its power in the frequency interval of interest $[-(\omega_u-\omega_l)/2, (\omega_u-\omega_l)/2]$ (the chosen passband for the FIR filter).

Step 3. Downsample the data sequence $\{\bar{y}_F(t)\}$ using the downsampling factor $d = \lfloor 2\pi/(\omega_u-\omega_l)\rfloor$ to obtain the d vectors of downsampled data sequences $\{\mathbf{z}_k\}$ in Eq. (52). We recommend choosing a lower value of d if there are interferences very close to the band of interest.

Step 4. Compute the weighted average covariance matrix estimate $\hat{\mathbf{R}}_D$ using Eq. (53).

Step 5. Employ a standard sub-space method operating on the covariance matrix estimate $\hat{\mathbf{R}}_D$ to obtain the modes $\bar{\lambda}_{1,D},\ldots,\bar{\lambda}_{n,D}$ of the downsampled data. We use the HSVD algorithm [25] in our numerical examples.

Step 6. Correct the modes obtained in the previous step so that the corrected modes correspond to the modes of the filtered signal $\{\bar{y}_F(t)\}$ *before* downsampling:

$$\bar{\lambda}_{k,F} = \bar{\lambda}_{k,D}^{1/d}, \quad k = 1,\ldots,n. \tag{55}$$

Step 7. Use the estimated modes from the previous step and the filtered signal $\{\bar{y}_F(t)\}$ from Step 2 to get the amplitudes $\bar{\rho}_{k,F}$ of the components in $\{\bar{y}_F(t)\}$ using LS:

$$\begin{bmatrix}\bar{\rho}_{1,F}\\ \vdots\\ \bar{\rho}_{n,F}\end{bmatrix} = (\mathbf{A}_{\bar{N}_F}^*\mathbf{A}_{\bar{N}_F})^{-1}\mathbf{A}_{\bar{N}_F}^*\bar{\mathbf{y}}_F, \tag{56}$$

where $\mathbf{A}_{\bar{N}_F}$ is defined similarly to Eq. (6) but with $\bar{N}_F$ rows, where $\bar{N}_F$ is the number of elements in the vector $\bar{\mathbf{y}}_F$ consisting of the filtered data sequence $\{\bar{y}_F(t)\}$, and with n columns where the kth column corresponds to $\bar{\lambda}_{k,F}$ from the previous step ($k = 1,\ldots,n$).

Step 8. Finally, obtain the modes of the original *unmodulated* signal $\{y(t)\}$ via

$$\lambda_k = \bar{\lambda}_{k,F}\mathrm{e}^{\mathrm{i}\bar{\omega}}, \quad k = 1,\ldots,n. \tag{57}$$

Also correct the estimated complex amplitudes $\{\bar{\rho}_{k,F}\}$ obtained in the previous step by using the correction factor given in [9], to remove the distortion from the FIR filter, and obtain the amplitudes $\{\rho_k\}$ corresponding to the modes $\{\lambda_k\}$ of the original signal $\{y(t)\}$.

3. Numerical examples

Five frequency-selective methods, namely FDM, SELF-MODE, FIDO, SB-HOYWSVD, and SELF-SVD, will be compared statistically by showing the parameter estimation performance for different component configurations and levels of noise. To show the advantages and disadvantages of the frequency-selective approach we also present the performance obtained by using a standard time-domain SVD method (HSVD) (see, e.g., [15] and [17]) covering the entire spectral range. A simple two-component example is considered in detail to compare the estimation accuracy of the five

methods, their computational complexity and the interference problems induced by a nuisance component. We will also explore a simulated 11 component example where the signal is derived from a typical in vivo ^{31}P spectrum measured in the brain. The added noise is zero mean, white and circular Gaussian distributed with standard deviation σ. The quality of the different parameter estimates is measured as the relative root mean square error (RRMSE) for each component $k = 1, \ldots, n$ [in percent]:

$$\mathrm{RRMSE}_k \triangleq 100\sqrt{\frac{1}{\Upsilon}\sum_{\upsilon=1}^{\Upsilon}\frac{(\xi_k - \hat{\xi}_k^{\upsilon})^2}{\xi_k^2}}, \tag{58}$$

where Υ is the number of Monte-Carlo runs (we use 1000), ξ_k denotes the relevant parameter and $\hat{\xi}_k^{\upsilon}$ is its estimate obtained in the υth run.

Consider a data sequence consisting of two exponentially damped sinusoids of which only one (component 1) is of interest. The two components have the following true parameter values (where au stands for amplitude units):

$$\omega_1 = 20\,\mathrm{Hz} \quad \omega_2 = \text{varied},$$

$$\alpha_1 = 10\,\mathrm{Hz} \quad \alpha_2 = 10\,\mathrm{Hz},$$

$$\rho_1 = 20\,\mathrm{au} \quad \rho_2 = 320\,\mathrm{au}.$$

The sampling frequency (f_s) is 1 kHz and the number of data points N is 512. Unless otherwise specified, the frequency of component 2 (ω_2) is 100 Hz. Note that the interference component (component 2) is significantly larger than component 1. The FFT spectrum of a typical example of this data sequence is presented in Fig. 1 where the noise standard deviation (σ) is equal to 5.

Throughout all simulations we consider a relatively small frequency interval compared to the full spectral range. The lower and upper frequency limits are in this example set to about 10 and 30 Hz, respectively. This interval corresponds to $\frac{1}{50}$ of the full spectral range. For SELF-MODE we set the snapshot length (m) to 64 samples. This choice of m results in a sparser Fourier grid compared to FDM and SELF-SVD, which forces us to choose the limits of the frequency interval of interest as 0 and 32 Hz, since the number of Fourier frequencies (L) is equal to 3 (with 16 Hz separation) for this small interval. Empirical evidence showed that a larger m would not provide any significant improvement in the parameter estimation accuracy, but would increase the computational complexity. We set the FIR-filter length (see, e.g., [7]) for the methods that use APES in the amplitude estimation step (SELF-MODE and SELF-SVD) to $\lfloor \frac{N}{100} \rfloor$. The user parameters of SB-HOYWSVD were, after some testing, set to $p = 2$, $q = 4$, and $c = 16$ (see the Appendix A). We also set $K = n$ (instead of $K = \bar{K}$) for FDM in this example, a choice which resulted in better parameter estimates, in the low SNR case, for this particular (rather artificial) data signal. For high SNR, however, this choice of K resulted in a small bias in the parameter estimates, which in our opinion was small enough to be neglected. We consider this simple data example useful for showing the numerical performances of the frequency-selective techniques for different data conditions.

The frequency interval should generally be chosen so that as much as possible of the spectral tails of the

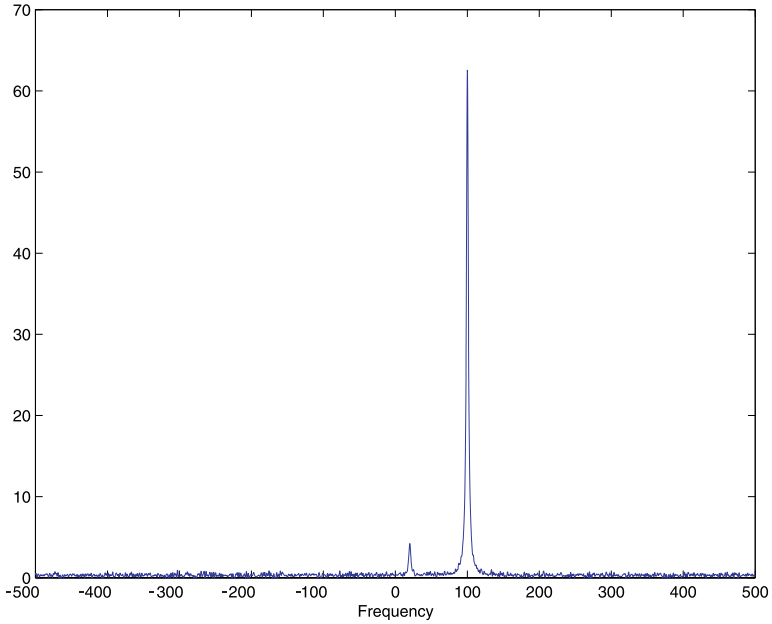

Fig. 1. FFT spectrum of a typical data sequence in the two-component example where the noise standard deviation (σ) is equal to 5.

component(s) of interest are included but without including nuisance components. If there is a "spectral gap" (meaning an interval comprising mostly noise with low spectral energy) between the region of interest and nuisance components, then including this "gap" might increase the estimation accuracy, but at the expense of increased computational complexity. Our experience is that none of the methods presented herein is very sensitive to small fluctuations in the choice of frequency interval, which in our opinion is a sound feature of any frequency-selective approach.

The RRMSEs (Eq. (58)) for the three parameters ($|\rho_1|$, ω_1, α_1) are shown in Fig. 2 versus the noise standard deviation. As references we also show the RRMSE obtained using HSVD and the (full-spectrum) Cramér-Rao bounds (CRB). We see that there is a severe increase in the RRMSEs of HSVD when the signal is buried in strong noise (i.e., a noise standard deviation above 20). The reason for this appears to be the occurrence of a "noise component" somewhere in the spectrum, with a slightly higher amplitude than component 1, which results in a false estimate. A similar performance is obtained for FDM when the noise level is high. A possible explanation for this is that FDM does not include any additional step to reduce the interferences from noise or out-of-band components after the windowing in the frequency domain; an operation which is carefully carried out using sub-band projection operations in SELF-SVD. We also note that, for a noise standard deviation above 10, we find a similar phenomenon for FIDO. A possible explanation for this is that the FIR-filtering step (Step 2 in Section 2.2) colors the noise; and colored noise of strong power may badly affect the accuracy.

Another important aspect to be studied concerns the interference from a nuisance component. We consider the case where the frequency separation between component 1 and component 2 is varied from 20 to 200 Hz. The resulting RRMSEs for the amplitude estimates of component 1 are shown in Fig. 3 for a fixed noise standard deviation of $\sigma = 5$. We note that for a frequency separation as low as 20 Hz, the interference from

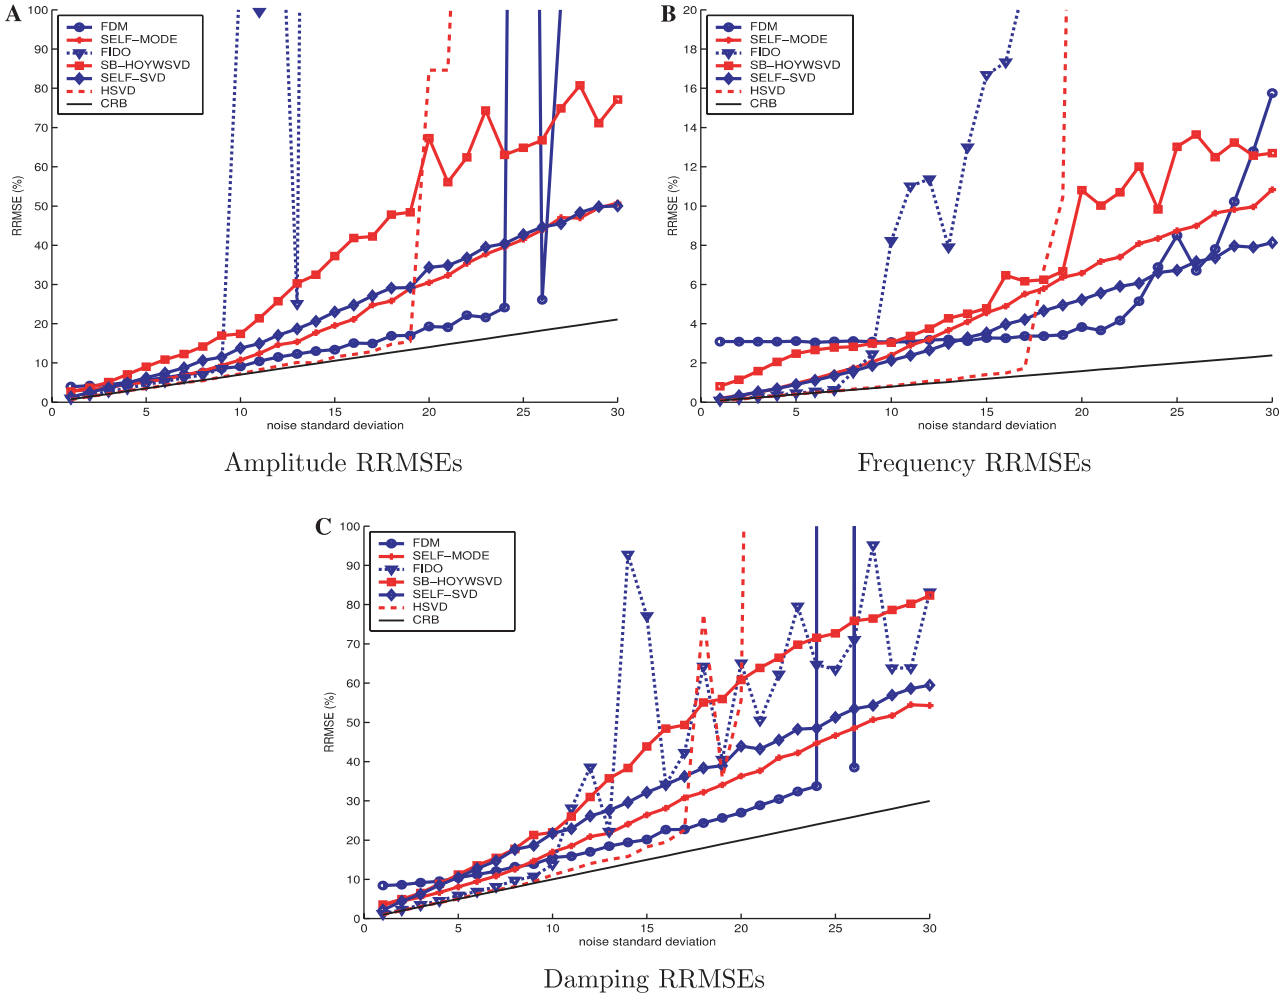

Fig. 2. RRMSEs for component 1 in the two-component example using 1000 Monte Carlo runs.

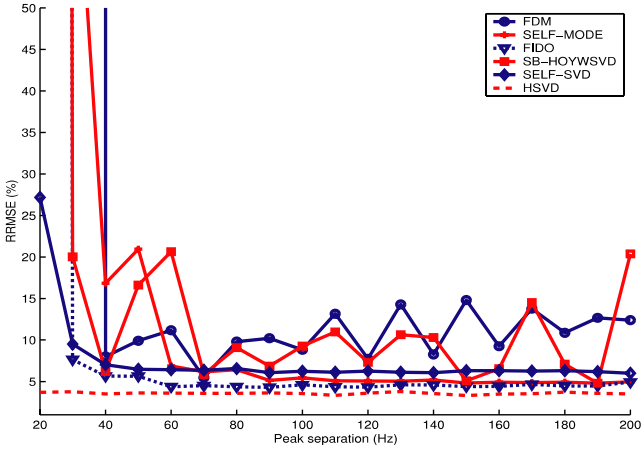

Fig. 3. Amplitude RRMSEs for component 1 when the frequency of component 2 is varied and the noise standard deviation (σ) is equal to 5.

Table 1
Parameters of the 11 damped sinusoids in the ^{31}P magnetic resonance spectroscopy data; α_k are the dampings, ω_k are the frequencies and $|\rho_k|$ denote the amplitudes

k	ω_k [Hz]	α_k	ρ_k ($e^{i0.75\pi}$)
1	−86	0.0167	75
2	−70	0.0167	150
3	−54	0.0167	75
4	152	0.0167	150
5	168	0.0167	150
6	292	0.0167	150
7	308	0.0167	150
8	360	0.0083	150
9	440	0.0951	1400
10	490	0.0083	60
11	530	0.0666	500

the larger component is rather strong, which usually results in increased RRMSEs. However, HSVD is not affected, as expected, unless the separation is very small, since it estimates the parameters of both components using the full spectral range.

Finally, we consider a more practical example that mimics a Magnetic Resonance data analysis (see, e.g., [10]), for which there are 11 components whose parameters are presented in Table 1. The ^{31}P components corresponding to the brain tissue, phosphomonoesters, inorganic phosphate, phosphodiesters, phosphocreatine, and α-, β-, and γ-ATP (adenosine triphosphate) are present in this signal. The number of samples N is 512 and the sampling frequency (f_s) is 3 kHz. The FFT spectrum of a typical example of this data sequence is presented in Fig. 4 where the noise standard deviation (σ) is equal to 5. We focused on the first of the two ATP doublets (components 4 and 5 in Table 1). The lower and upper frequency bounds were set to about 120 and 200 Hz, respectively for FDM, FIDO, SB-HOYWSVD, and SELF-SVD. For SELF-MODE the selected frequency interval is slightly larger for the same reasons as discussed in the previous example; here too, m is equal to 64. The RRMSEs for the three parameters of interest ($|\rho_k|$, ω_k, α_k) are shown in Fig. 5 for component 4 for increasing noise standard deviation. The results for component 5 were very similar to those of component 4 and have therefore been omitted. The length of the FIR filter for the amplitude estimation step (APES) in SELF-MODE and SELF-SVD was set to $\lfloor \frac{N}{15} \rfloor$ in this example. The user parameters of SB-HOYWSVD were, after some testing, set to $p = 8$, $q = 12$, and $c = 16$ (see the Appendix A). For FDM we set $K = \bar{K}$ as suggested for practical MR signals in Section 2.

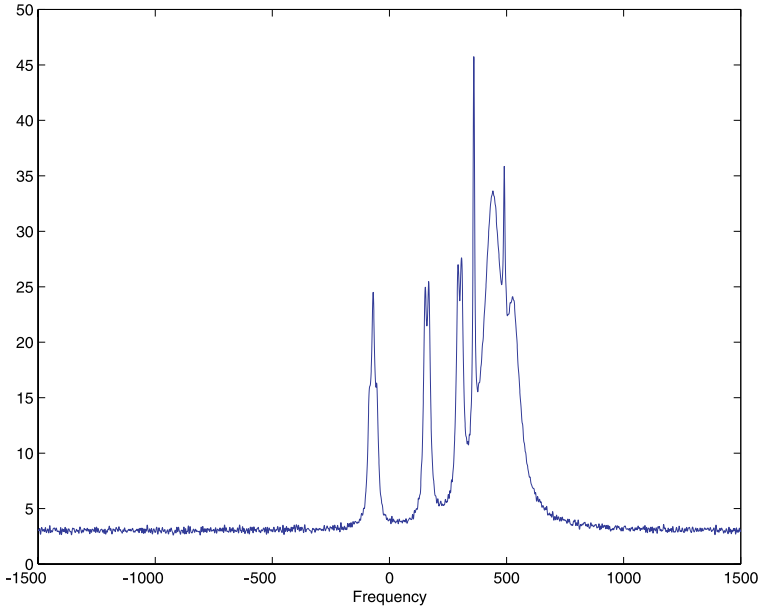

Fig. 4. FFT spectrum of a typical data sequence in the 11-component ^{31}P example where the noise standard deviation (σ) is equal to 5.

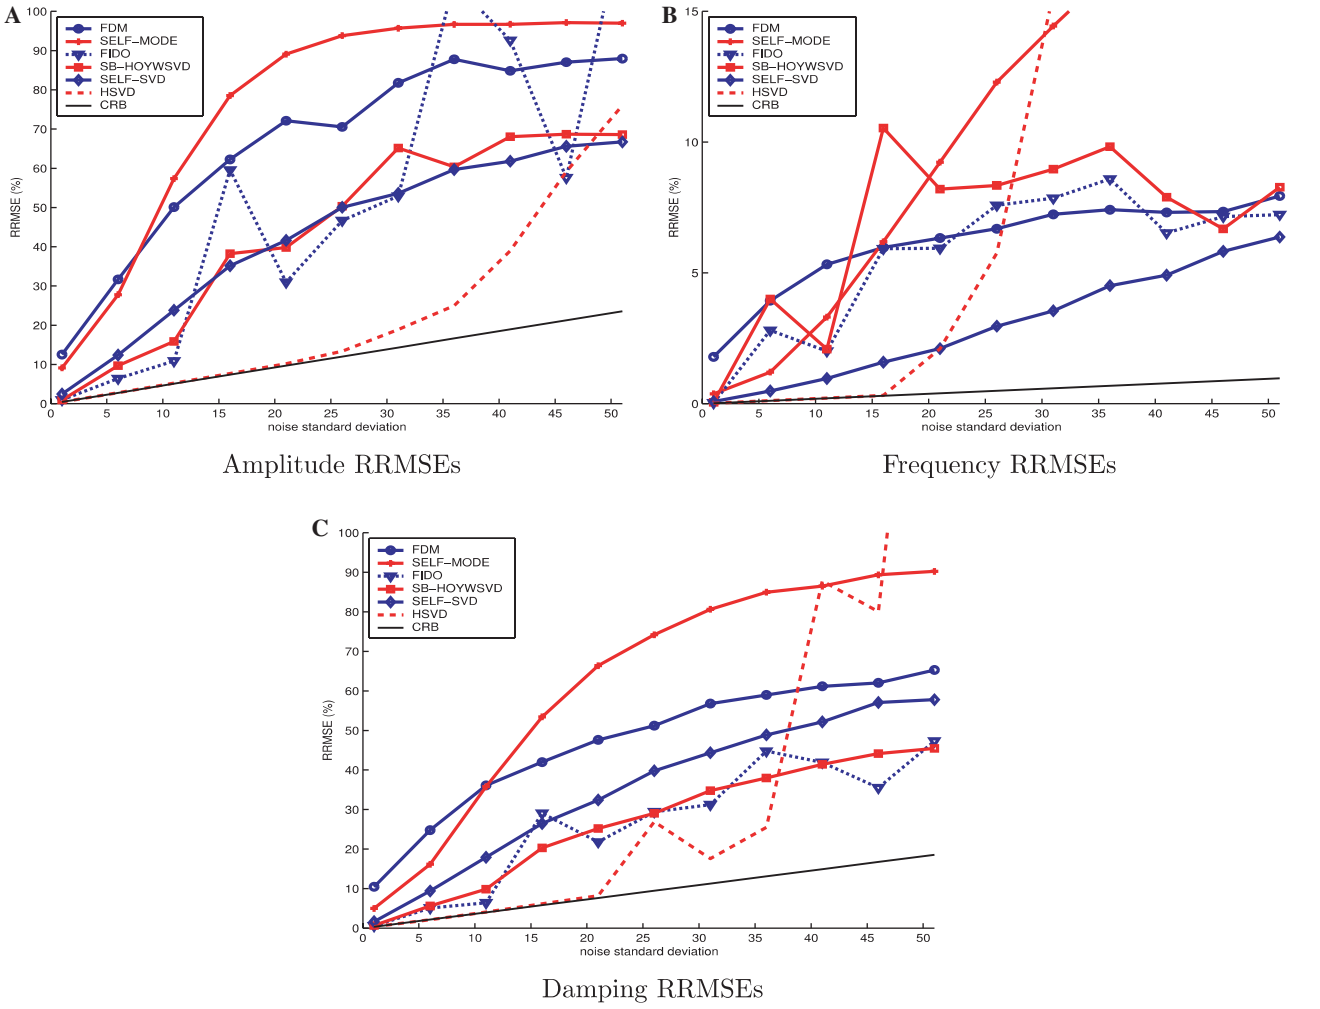

Fig. 5. RRMSEs for component 4 in the 11-component ^{31}P MRS example using 1000 Monte Carlo runs.

The computational speed of the five frequency-selective methods is generally superior to that of HSVD. Table 2 shows the required number of flops to perform one simulation run for each method in the above two given examples. The magnitudes of these values of course depend on the size of the frequency interval of interest, the number of components and the total number of data samples. For SB-HOYWSVD, this number also depends heavily on the user parameters. We note that the computational complexity of any method involving a large SVD can be reduced by using the fast Lanczos algorithm (see [26]). This technique has however not been considered in this paper.

Table 2
Required number of flops for the different methods for the two-component example (Ex. 1) and the 11-component example (Ex. 2)

Method	# Flops (Ex. 1)	# Flops (Ex. 2)
FDM	1.005×10^7	1.367×10^7
SELF-MODE	1.872×10^7	2.479×10^7
FIDO	1.048×10^7	2.908×10^7
SB-HOYWSVD	0.988×10^7	2.795×10^7
SELF-SVD	7.202×10^5	6.897×10^8
HSVD	7.491×10^8	7.271×10^8

4. Conclusions

It has been suggested previously (see, e.g., [20]) that there are primarily two different cases to consider in a parametric spectroscopic analysis. One alternative is to estimate only the parameters of a few genuine components in an MR-signal corrupted by noise, and obtain the quantities of interest directly from the parameters. Another alternative is to estimate a very large set of spectral lines (which is usually used to construct the spectrum) from the measured data, where the SNR is high, without paying any explicit attention to which parameters that correspond to the true signal components. The main goal is often to determine the concentration of each chemical substance in a sample, and this can be carried out in both cases using directly the

estimated amplitudes in the first case or, e.g., by area integration for each spectral component (or groups of components) in the second case. In our opinion, the first case should be preferred if we consider the frequency-selective parameter estimation problem stated in Section 1, since the integration necessarily introduces subjectivity in the choice of integration interval. One can also argue that the first approach should be a better choice if, e.g., the spectrum contains overlapping components. Should the goal of the spectroscopic analysis be to estimate the frequencies of the spectral components, these are also directly obtained from the parameter estimates, and there is no reason to construct the spectrum.

Of the methods considered we note that FDM was previously (in, e.g., [1,20]) designed for the second of the two cases mentioned above, and the remaining four can be considered as approaches for estimating directly the parameters of the resonances. This also suggests that the major difference between the methods presented here should lie between FDM and the other techniques. We have presented an interpretation of FDM that allows parameter estimation for a desired number of components lying in a pre-specified frequency interval, in the same manner as for the other methods. This interpretation can be seen as an addition to the previous derivations of FDM, and as a suggestion on how to extend the usefulness of this method.

Considering the two-component example in Section 3, we are able to draw some conclusions regarding the parameter estimation accuracy for these five methods. For moderately high SNRs, FDM seems to give better parameter estimates than any of the other four frequency-selective techniques. The FIDO and SB-HOYWSVD filtering and decimation techniques (especially FIDO) show good performance when the SNR is high but for low SNR they seem to be rather erratic (again, especially FIDO). SELF-MODE and SELF-SVD have a stable parameter estimation accuracy with RRMSEs lying somewhere in between FDM and the two filtering and decimation methods in this example. In addition the two latter methods give satisfactory results even for low SNRs. From Fig. 3 we see that none of the methods are severely affected by closely spaced nuisance components, but SELF-SVD seems to be more stable than the other frequency-selective techniques when the frequency separation between the component of interest and the strong nuisance component is small. The computational complexity is rather similar for FDM, SELF-MODE, FIDO, and SB-HOYWSVD in this example, while SELF-SVD is about 10 times faster than the other four. In addition, all five frequency-selective methods outperforms HSVD in speed, as expected.

The results obtained for the 11-component MRS data example indicate that for this more complex and realistic signal, the difference in accuracy between the five frequency-selective methods is generally smaller than in the first example (especially FIDO shows a more stable performance). We note that the performance of SELF-MODE and FDM is slightly worse than in the previous example compared to that of the other methods. From Table 2 we see that FDM is faster than SELF-MODE, FIDO, and SB-HOYWSVD, while SELF-SVD is not as superior in speed for this example as for the previous two-component signal. SELF-SVD is, however, still the fastest. The parameter estimation accuracy for FIDO and SB-HOYWSVD is sometimes better than that of the other methods in this example. However, their rather erratic behavior indicates that it might be better to use a technique with a more stable performance, such as, e.g., SELF-SVD, which in summary seems to have the most reliable parameter estimation accuracy if we consider all examples presented in Section 3.

The general issue of user-friendliness of a method is certainly also quite important, as the choice of a specific method from the wealth of available techniques often depends on the individual preferences of the user. For example, it can be important to design a method with as few user parameters as possible, so that it can be used as a ''black-box'' alternative to the FFT. The methods considered here have one or several user parameters (in addition to the specification of the frequency interval and the number of components, n). However, the choice of these user parameters does not seem to severely affect the results as long as they are set within reasonable intervals. We note that SB-HOYWSVD has the largest number of user parameters of the approaches considered here. In the numerical examples presented herein we have not tried hard to find the optimal parameter settings for this method or any of the others. There might be parameter settings that result in lower RRMSEs for some of these techniques, but if such an optimization is required for a method to work properly we can certainly consider it to be less useful in practice since we generally do not have the prior information necessary to perform such an optimization.

Acknowledgments

The authors thank Dr. El-Hadi Djermoune for providing the code for the SB-HOYWSVD method and Dr. Leentje Vanhamme for the maximum phase FIR filtering code. We are also grateful to Dr. Vladimir A. Mandelshtam for useful discussions about some practical details of MR spectroscopy.

Appendix A

For the reader's convenience we outline below the main parts of SB-HOYWSVD and SELF-SVD without going into details about the derivations of these meth-

ods. The SB-HOYWSVD approach is conceptually rather similar to FIDO. The major differences include an ARMA signal model instead of the model in Eq. (1) for the measured signal and, since the AR part contains all useful spectral information for that ARMA model (see [8]), estimation of the AR parameters using an SVD-based procedure for solving the high-order Yule–Walker equations. The approach of SELF-SVD, however, is based on the signal model in Eq. (1).

A.1. Summary of SB-HOYWSVD

It can be shown (see [8,15]) that the signal model in Eq. (1) can be rewritten as

$$y(t) = -\sum_{k=1}^{\bar{n}} g(k)y(t-k) + \sum_{k=0}^{\bar{n}} g(k)e(t-k), \qquad (A.1)$$

which, under the assumption of a white noise sequence $e(t)$, is an ARMA process having the same coefficients $\{g(k)\}$ for the AR and MA parts. For the general case of colored noise, Eq. (A.1) can be written as

$$y(t) = -\sum_{k=1}^{p} g(k)y(t-k) + \sum_{k=0}^{q} h(k)e(t-k), \qquad (A.2)$$

where $h(k)$ incorporates the whitening filter of the noise, which is assumed to be of MA type. It is well known that for Eq. (A.2)

$$r(k) = -\sum_{j=1}^{p} g(j)r(k-j), \quad k \geqslant q, \qquad (A.3)$$

where $r(k) = E[y^*(t)y(t+k)]$ is the autocorrelation sequence of $y(t)$. Writing Eq. (A.3) for $q \leqslant k \leqslant q+p$ yields the so-called modified Yule–Walker equations. In the case where $q \leqslant k \leqslant q+c$, where $c > p$, the over-determined Yule–Walker equation system is obtained, which can be used to estimate the coefficients $\{g(k)\}$. After this very brief introduction to ARMA-modeling of exponentially damped sinusoids (see, e.g., [15] for more details) we outline the structure of SB-HOYWSVD. Note that in [8], SB-HOYWSVD is used for the entire spectral range by dividing it into several sub-bands and analyzing them one at a time, whereas here we only use one sub-band, corresponding to the frequency interval in Eq. (2).

Step 1. Choose a value for the decimation factor d (similarly as for FIDO). Fix the user parameter p and set q and c to be at least equal to p (see above). The operating parameter (or model order) p is preferably chosen to be larger than the number of estimated components in the frequency interval of interest (i.e., $p > n$).

Step 2. Generate a sub-band signal by modulating the original signal (see Section 2.3) followed by a bandpass filtering, focusing on the interval of interest using an FIR filter (in our numerical examples we have used the FIR filter of [9]), and decimation. In SB-HOYWSVD, only the first decimated sequence is considered (i.e., $k = 0$ in Eq. (52)). Note that the FIR filtering procedure colors the noise in an MA manner, a fact that is dealt with in the signal model of Eq. (A.2).

Step 3. Solve the (over-determined) Yule–Walker equation system for the sub-band signal using the HOYWSVD technique (see [8]) and obtain estimates of the signal poles. Reflect the poles lying outside the unit circle, compute the sub-band complex amplitudes and discard all poles outside the filter's bandpass.

Step 4. Correct the parameters obtained in Step 3 to compensate for the filtering and decimation, similarly to Step 8 in FIDO.

A.2. Summary of SELF-SVD

For conciseness reasons, the discussion about the ideas behind SELF-SVD is kept as short as possible (see [7] for details).

Step 1. Multiply the measured signal $y(t)$ with the Fourier vectors corresponding to the Fourier frequencies comprising the interval in Eq. (2) (similarly as for FDM in Eq. (20)) to focus on the spectral region of interest.

Step 2. Use sub-band projection operations (see [7]) to suppress the interferences from noise and out-of-band components in the filtered data sequence from Step 1.

Step 3. Use an SVD-based technique to obtain estimates of the n signal modes located in the interval of interest.

Step 4. Finally, estimate the (complex) amplitudes of the n peaks using APES [16] and the signal modes estimated in Step 3.

References

[1] V.A. Mandelshtam, H.S. Taylor, Harmonic inversion of time signals and its applications, J. Chem. Phys. 107 (1997) 6756–6769.
[2] Y. Hua, T.K. Sarkar, Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise, IEEE Trans. Acoust. Speech Signal Process. 38 (1990) 814–824.
[3] Y. Hua, T.K. Sarkar, On SVD for estimating generalized eigenvalues of singular matrix pencil in noise, IEEE Trans. Signal Process. 39 (1991) 892–900.
[4] Z. Tian, Beamspace iterative quadratic WSF for DOA estimation, IEEE Signal Process. Lett. 10 (2003) 176–179.
[5] M. Cedervall, P. Stoica, R. Moses, Mode-type algorithm for estimating damped, undamped or explosive modes, Circ. Syst. Signal Process. 16 (1997) 349–362.
[6] P. Stoica, K.C. Sharman, Maximum likelihood methods for direction of arrival estimation, IEEE Trans. Acoust. Speech Signal Process. 38 (1990) 1132–1143.
[7] P. Stoica, N. Sandgren, Y. Selén, 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.
[8] M. Tomczak, E.-H. Djermoune, A subband ARMA modeling approach to high-resolution NMR spectroscopy, J. Magn. Reson. 158 (2002) 86–98.

[9] 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.
[10] I. Dologlou, S. Van Huffel, D. van Ormondt, Frequency-selective MRS data quantification with frequency prior knowledge, J. Magn. Reson. 130 (2) (1998) 238–243.
[11] 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.
[12] S. Cavassila, B. Fenet, A. van den Boogaart, C. Rémy, C. Briguet, D. Graveron-Demilly, ER-filter: a preprocessing technique to improve the performance of SVD-based quantitation methods, J. Magn. Reson. Anal. 3 (1997) 87–92.
[13] J. Tang, J.R. Norris, LP-ZOOM, a linear prediction method for local spectral analysis of NMR signals, J. Magn. Reson. 79 (1988) 190–196.
[14] J. Slotboom, C. Boesch, R. Kreis, Versatile frequency domain fitting using time domain models and prior knowledge, Magn. Reson. Med. 39 (1998) 899–911.
[15] P. Stoica, R. Moses, Introduction to spectral analysis, Prentice Hall, Upper Saddle River, NJ, 1997.
[16] P. Stoica, T. Sundin, Nonparametric NMR spectroscopy, J. Magn. Reson. 152 (2001) 57–69.
[17] H. Barkhuysen, R. de Beer, D. van Ormondt, Improved algorithm for noniterative time-domain model fitting to exponentially damped magnetic resonance signals, J. Magn. Reson. 73 (1987) 553–557.
[18] P. Stoica, H. Li, J. Li, Amplitude estimation of sinusoidal signals: survey, new results and an application, IEEE Trans. Signal Process. 48 (2000) 338–352.
[19] H. Hu, Q.N. Van, V.A. Mandelshtam, A.J. Shaka, Reference deconvolution, phase correction, and line listing of NMR spectra by the 1D filter diagonalization method, J. Magn. Reson. 134 (1998) 76–87.
[20] V.A. Mandelshtam, FDM: the filter diagonalization method for data processing in NMR experiments, Progr. NMR Spectrosc. 38 (2001) 159–196.
[21] V.A. Mandelshtam, H.S Taylor, Multidimensional harmonic inversion by filter-diagonalization, J. Chem. Phys. 108 (1998) 9970–9977.
[22] S.Y. Kung, K.S. Arun, B. Rao, State space and singular value decomposition based on approximation methods for harmonic retrieval, J. Opt. Soc. Am. (1983) 1799–1811.
[23] B. Halder, T. Kailath, Efficient estimation of closely spaced sinusoidal frequencies using subspace-based methods, IEEE Signal Process. Lett. 4 (1997) 49–51.
[24] P. Stoica, A.E. Nordsjö, Subspace-based frequency estimation in the presence of moving-average noise using decimation, Signal Process. 63 (1997) 211–220.
[25] 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.
[26] T. Laudadio, N. Mastronardi, L. Vanhamme, P. Van Hecke, S. Van Huffel, Improved lanczos algorithms for blackbox MRS data quantitation, J. Magn. Reson. 157 (2002) 292–297.