NTNU

Norwegian University of
Science and Technology

MR Spectroscopy

Real-Time Quantification of in-vivo MR Spectroscopic data

Kunal Massé

Master of Science in Electronics

Submission date: April 2009

Supervisor: Lars Magne Lundheim, IET

Norwegian University of Science and Technology
Department of Electronics and Telecommunications

Problem Description

In-vivo MR Spectroscopy uses the resonance properties of certain nuclei when exposed to a magnetic field to determine the patients' metabolism in a region of interest, e.g. brain or prostate. In this thesis the student will work on the SELF-SVD method for quantification of metabolites being feasible for fast, reproducible analysis of data of high spatial and spectral resolution.

Assignment given: 06. October 2008
Supervisor: Lars Magne Lundheim, IET

Abstract

In the last two decades, magnetic resonance spectroscopy (MRS) has had an increasing success in biomedical research. This technique has the faculty of discerning several metabolites in human tissue non-invasively and thus offers a multitude of medical applications.

In clinical routine, quantification plays a key role in the evaluation of the different chemical elements. The quantification of metabolites characterizing specific pathologies helps physicians establish the patient's diagnosis. Estimating quantities of metabolites remains a major challenge in MRS.

This thesis presents the implementation of a promising quantification algorithm called selective-frequency singular value decomposition (SELF-SVD). Numerous tests on simulated MRS data have been carried out to bring an insight on the complex dependencies between the various components of the data. Based on the test results, suggestions have been made on how best to set the SELF-SVD parameters depending on the nature of the data.

The algorithm has also been tested for the first time with in-vivo ^1H MRS data, in which SELF-SVD quantification results allow the localization of a brain tumor.

Acknowledgments

This thesis project would not have been possible without the help of a number of benevolent people.

I would like to express my sincere gratitude to Professor Lars Lundheim for his constant encouragements and his invaluable feedback throughout the project.

I am also much indebted to Fraunhofer MEVIS for making this project possible and to its staff for their outstanding support. I would like to make a special mention of Mr. Bernd Merkel and Mr. Markus Harz for their patient and unfailing guidance as well as for their dedication and commitment to make this project as successful and fruitful as can be.

My thoughts also go to all those who have very patiently helped me with my learning German.

Finally, a special thanks to my family for their love and support.

Abbreviations

AMARES advanced method for accurate, robust and efficient spectral fitting

APES amplitude and phase estimation

ARMA autoregressive moving-average

ATP adenosine triphosphate

Cho choline

Cr creatine

CSI chemical shift imaging

DFT discrete Fourier transform

ESPRIT estimation of signal parameters via rotational invariant techniques

FDM filter diagonalization method

FFT fast Fourier transform

FID free induction decay

FIDO filtering and downsampling

FIR finite impulse response

fMRI functional magnetic resonance imaging

FT Fourier transform

Glx Glutamate-Glutamine

GUI graphical user interface

HIV human immunodeficiency virus

HLSVD Hankel-Lanczos singular value decomposition

HSVD Hankel singular value decomposition

HTLS Hankel total least square

HTLS-PK Hankel total least square with prior knowledge

KNOB-TLS knowledge based total least square

Lac lactate

LAPACK linear algebra package

LCModel Linear Combination of Model

LS least square

LP linear prediction

LPSVD linear prediction singular value decomposition

LPTLS linear prediction total least square

MP matrix pencil

mI myo-Inositol

MMS MeVis Medical Solutions

MODE method of direction estimation

MDL module definition language

MRS magnetic resonance spectroscopy

MRI magnetic resonance imaging

MR magnetic resonance

NAA N-Acetyl-Aspartate

NLLS nonlinear least-squares

NMR nuclear magnetic resonance

ppm parts per million

PRESS point-resolved spectroscopy

ROI region of interest

RF radiofrequency

SELF-MODE selective-frequency method of direction estimation

SELF-SVD selective-frequency singular value decomposition

SNR signal-to-noise ratio

SVD singular value decomposition

SVS single voxel spectroscopy

TE echo time

TLS total least square

VARPRO variable projection method

voxel volumetric pixels

Contents

Chapter 1

Introduction

1.1 Motivation

As life expectancy has increased in developed countries, old-age illnesses have become more and more of an issue. Today, in all countries of the European Union, the life expectancy is above 65 years for men and 75 for women [1]. Figure 1.1 presents the major causes of death in Europe in 2004. The most important cause of death by far is cancer. Other diseases, in particular of the heart, are also to be found as major death causes. Scanning tools, such as *magnetic resonance imaging (MRI)* are widely used and aim to detect many different malignancies. Since the nineties, in-vivo *magnetic resonance spectroscopy (MRS)* is applied medically and is becoming a precious complement to MRI, in particular in the detection and severity grading of cancer.

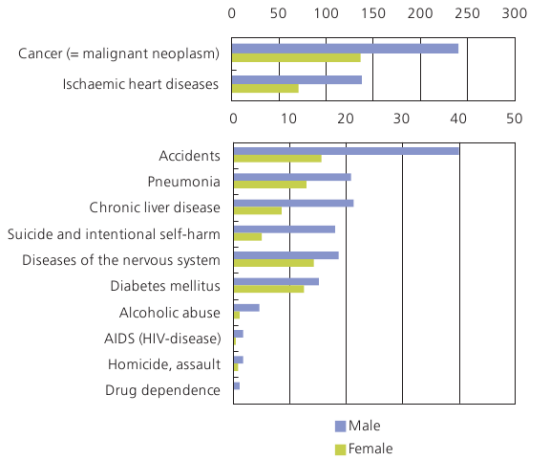

Figure 1.1: Causes of death, EU of 27, 2004, rates per 100 000 persons. Please note the changing of scale. Information taken from [1]

MRS can be applied ex-vivo and in-vivo. In-vivo spectroscopy has more direct applications in hospitals and clinics because diagnosis can be made on-the-fly and has the major advantage of being non-invasive. It is the focus of this document.

MRS and MRI complement each other very well for diagnosis. While MRI gives an anatomic map of the scanned area, MRS gives its chemical content. MRS can be used to back up an assumption made by the reading of an MRI image, or in the case of tumors for example, to evaluate its severity. A typical example of the MRI/MRS data read by clinicians is given in figure 1.2. On the left is an MRI image of the prostate, with a red marker around the *voxel* of interest. On the right is the MRS signal of the voxel marked in the MRI. The MRI image is essential to locate the MRS area to be scanned. Please note that the signal shown in figure 1.2b is the MRS signal after *Fourier transform (FT)* (more on this transformation in 1.2.1). As we shall develop later in this document, the peaks in figure 1.2b provide objective information on specific metabolites encountered in the voxel of interest.

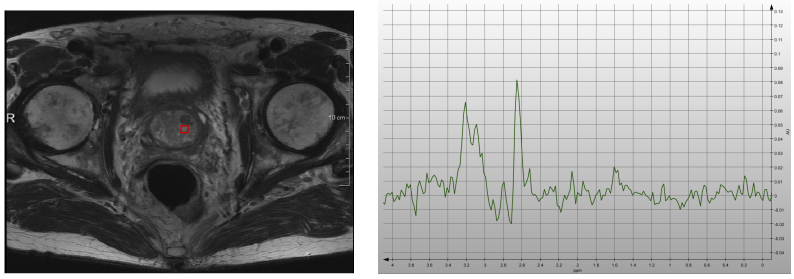

(a) (b)

Figure 1.2: (a) MRI image of the prostate (b) MRS signal (Fourier transformed) corresponding to the red voxel in the MRI

Several techniques based on MRI, such as *contrast-enhancement* MRI and *functional magnetic resonance imaging (fMRI)*, have been developed in order to improve diagnosis. In addition to these MRI techniques, MRS possesses other promissing features.

In contrast-enhancement MRI, a group of contrast agents are used in order to improve the visibility of internal body structures. The contrast agent (generally gadolinium-based) is most often injected to the patient. Compared to this method, MRS has the primary advantage of being a non-invasive technique.

fMRI is another specialized MRI scan. It is used for *functional imaging*, i.e. the detection of changes in the metabolism and their spatial distribution. Functional MRI bases itself on the study of the activation of brain areas whilst the patient undergoes some physical movement activity (e.g. moving a finger). MRS can also be used as a functional imaging tool. Its key additional feature is that it offers the possibility to differentiate metabolites, products of the metabolism.

MRS is also an efficient *differential diagnosis* tool. Differential diagnosis consists in distinguishing a disease from others that present similar symptoms. For example, in [2], MRI and MRS are combined for the differential diagnosis of gliomatosis cerebri.

MRS can also be used to diagnose pathologies where other tools cannot. For instance, the early stages of a *human immunodeficiency virus (HIV)* infection of the brain are not detectable by an MRI image, but can be found through an MRS scan. It has been shown that HIV-positive newborns present abnormal proton MRS spectra already 10 days after birth [3].

MRI scanning is widely applied throughout the world, thus making the magnetic resonance instrument available also for MRS scanning. Nevertheless, the latter is still not systematically used in clinics and hospitals. Unfortunately, the technique is still foreign to many clinicians. This can partially be put on the account of the complexity of MRS spectra analysis. Post treatment of the MRS data is done by software and can still be improved. This document concentrates on the software end of spectroscopic analysis.

The software handling the MRS data should present the spectroscopic results to the medical specialist in the most transparent format possible, thus making medical interpretation easier, speeding up diagnosis and allowing treatment promptly.

One of the major goals of MRS is to measure the concentration of metabolites. The relative alteration of their concentrations helps determine the presence or absence of an illness. It is the base of many diagnoses. Bad or inaccurate quantification means additional risk for the patient. Despite the efforts put into it, qualitative quantification remains today a major difficulty in MR Spectroscopy.

The purpose of this thesis is to implement and extensively test a promising quantification method. It is called *selective-frequency singular value decomposition (SELF-SVD)* [4].

1.2 Magnetic resonance applied in medicine

1.2.1 Magnetic resonance technique

The concept of *nuclear magnetic resonance (NMR)* was first presented in 1938, by Isidor Rabi. It was then brought to application by Felix Bloch and Edward Mills Purcell in 1946 and given the name of nuclear magnetic resonance, also called *magnetic resonance (MR)*. The naming is explained as follows :
nuclear: nuclei of certain atoms react,
magnetic: a magnetic field is required,
resonance: the magnetic resonance frequency of the nuclei is directly proportional to the strength of the applied magnetic field. The concept of resonance is further explained in 1.2.2.

The power of the applied magnetic field differs between in-vivo and ex-vivo MRS. Normally, human scanner magnetic fields range between 1.5 and 3 Tesla,

whilst ex-vivo fields can go up to 21 Tesla. For comparison, the earth's magnetic field at its surface varies between 30 and 60 microteslas and the field created by a ferrite bar magnet does not exceed 0.35 Tesla. The real effect of strong fields, such as those used ex-vivo, on human beings is still unknown and unauthorized. A stronger field brings more accuracy, offering ex-vivo spectroscopy higher resolution characteristics.

In-vivo MRS uses the same scanning apparatus as MRI, with certain constraints. It first needs a strong and homogeneous magnetic field to distinguish resonance peaks. An adjustable radiofrequency transmitter is also required to adapt itself to the resonance frequencies of the studied nuclei. A software is then needed to clean and quantify the MRS data.

The data is acquired through specific methods, namely *single voxel spectroscopy (SVS)* or *chemical shift imaging (CSI)*. In SVS, the signal is measured for a single voxel (volume element). The acquisition can take up to approximately 5 minutes. CSI receives spectra for multiple voxels, forming a 1D projection, 2D slice or 3D volume. The acquisition time is hence very variable, but generally does not exceed 20 to 25 minutes.

MR Spectroscopy is used to discern metabolites. The metabolites are not directly found, but deduced through the determination of nuclei. Proton nuclei (^{1}H) are most often chosen in in-vivo spectroscopy because of their abundance and high nuclear magnetic sensitivity. This document concentrates precisely on proton MRS, also called ^{1}H MRS.

Despite the huge number of different metabolites in human tissue, a relatively low number of them are identifiable in-vivo. Only those present in substantial concentrations give enough signals to be detected. The minimum metabolite concentration detectable by MRS is around 0.5 mmol [5]. Most of the metabolites studies lie in the range 1-10 mmol/L, while water concentrations approaches 90 mol/L [6].

In-vivo MR Spectroscopy is performed on a selected zone of the human body. This area is called *region of interest (ROI)*. The area sampled is divided into voxels, which generally have a minimum size of 1 cm for each dimension. The outcome from a spectroscopic analysis is a series of complex time signals, i.e. one signal per voxel. In order to make analysis easier, a FT of the signal is almost systematically performed. This converts the signal from the initial time domain to the frequency domain. The signal in the frequency domain is commonly called *spectrum*. The spectrum is composed of numerous peaks. Assuming that all noise has been successfully removed, each peak, or group of peaks in the spectra, characterizes a metabolite in the scanned area. The peaks are defined by their resonance frequency, which defines the nature of the metabolite they represent (further explained in 1.2.2). The area under the peak gives a measurement of the concentration of the metabolite.

1.2.2 Physical basics

Quantum mechanics and the study of energy levels will give the complete understanding of magnetic resonance. Several publications are available which thoroughly explain its complex mechanism. In this section, and in order to get familiar to the MRS concept and the underlying physics, the vector model will be presented. It gives a good understanding of pulse-magnetic resonance spectroscopy, the focus of this document.

Some nuclei contain a source of spin angular momentum. Through this spin, the nucleus acts as a little bar magnet, generating a magnetic field, called *nuclear spin magnetic moment*. Placing the nucleus in a magnetic field causes an interaction between its nuclear magnetic moment and the applied field. Their energy of interaction depends on the angle between the two fields. The energy of interaction is minimum when the two fields are parallel, i.e. the individual magnetic moments of the nuclei are globally aligned with the applied field. The magnetic moments have hence a natural tendency to align with the external field. The sample, as a whole, acquires a magnetic moment, i.e. the mean of all the magnetic moments of the nuclei of interest, which is represented by the *bulk magnetization vector*.

First, the sample is placed in a magnetic field. The magnetic moments of the nuclei have a random orientation, as shown in figure 1.3a. After a short period of time (in the range of 1 second) it reaches an equilibrium state, i.e. its bulk magnetization aligns itself with the field. This equilibrium state is represented in figure 1.3b. The process by which the nuclei spins reach this equilibrium is called *relaxation* and is particularly of interest in magnetic resonance techniques.

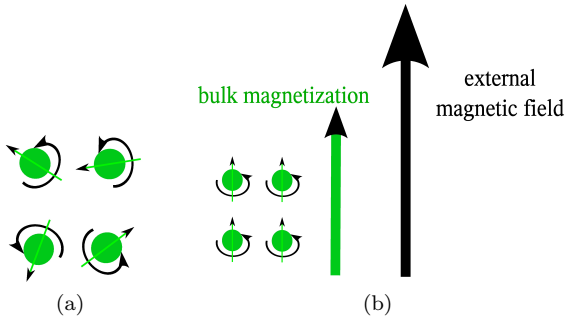

Figure 1.3: (a) Without external field, the magnetic moments have a random orientation. (b) With the external field, they acquire a global magnetization direction.

A much feebler magnetic field is then applied perpendicularly to the first field. This second field is called the *radiofrequency (RF)* field. The description of the RF field and its effects are illustrated by figure 1.4. When applied at a particular frequency, the RF field tips the bulk magnetization vector out of its initial equilibrium state and alignment with the z-axis. In its new misaligned configuration, the magnetization vector will then rotate around the magnetic

field and gradually come back to its equilibrium state. The circular motion it will undertake is called *precession*. The MRS signal is derived from this precession. The frequency at which the bulk magnetization rotates is particularly of interest. It is called the *Larmor frequency*.

The RF field is generated along the x-axis by a coil wound round the sample. Despite the feeble force of the RF field compared to the main magnetic field, it can affect the bulk magnetization vector if applied at a resonance frequency equal or near to the Larmor frequency.

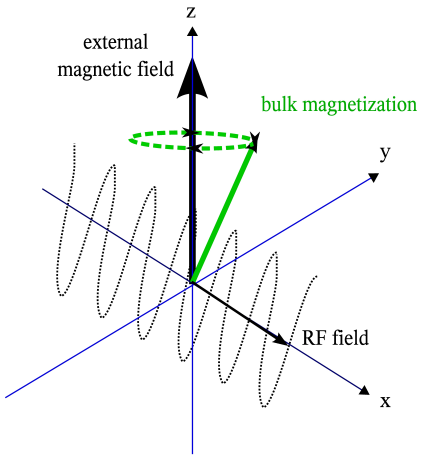

Figure 1.4: The RF field and its effect on the bulk magnetization vector

The Larmor frequency of a nuclei follows the following equation:

$$\nu = \frac{\gamma}{2\pi} B, \tag{1.1}$$

ν: Larmor frequency of the nuclei,
γ: constant *gyromagnetic ratio*, unique for each nuclei type,
B: magnetic field applied to the sample.

The resonance frequency thus depends on the nuclei studied and the magnetic field applied. For instance, at $B = 1.5$ Tesla, phosphorus-31 (^{31}P) has a the resonance frequency of 25.85 MHz, whilst for proton MRS, hydrogen-1 (^{1}H) has a resonance frequency of 63.87 MHz.

A more detailed analysis of the effect of the RF field on the bulk magnetization vector is given below in order to understand how the MRS signal is captured. The equilibrium state is represented in 1.5a, where the magnetization is aligned with the z-axis. When the RF field is applied at the Larmor frequency, it is called an *on-resonance pulse*. This type of pulse is the simplest to understand. The radiofrequency field is applied along the x-axis and makes the magnetization vector tilt in a circular motion in the yz-plane around the RF vector, as shown in figure 1.5b. The rotation angle, or *flip angle*, depends on the duration the RF is applied. A common flip angles are 90 and 180 degrees.

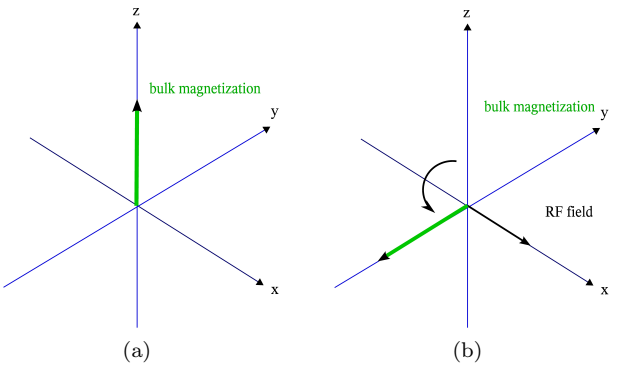

Figure 1.5: (a) Initial equilibrium state (b) 90 degree flip due to the RF field

An MRS signal will always contain more than one resonance frequency, i.e. the area scanned contains more than one type of metabolite. It is not possible to be on-resonance with all frequencies simultaneously. Nevertheless, by using a sufficiently strong RF field, a range of resonances can be covered.

As mentioned previously, the MRS signal is directly deduced by the relaxation process of the magnetization vector, which takes place just after the RF pulse is applied. The precession motion is acquired in the xy-plane. The projection of the magnetization vector on the xy-plane as well as its rotation can be seen in 1.6a. As the bulk magnetization vector rotates, a complex signal is acquired, with one dimension taken along the x-axis and the other y-axis. This acquisition gives us the MRS data. A simple noiseless and single-mode example of an acquisition taken along the x-axis dimension is shown in figure 1.6b. Note the decay with time of the acquired sinusoid, due to the relaxation of the vector back to the z-axis. This characteristic shape gave the MRS signal the name of *free induction decay (FID)* signal.

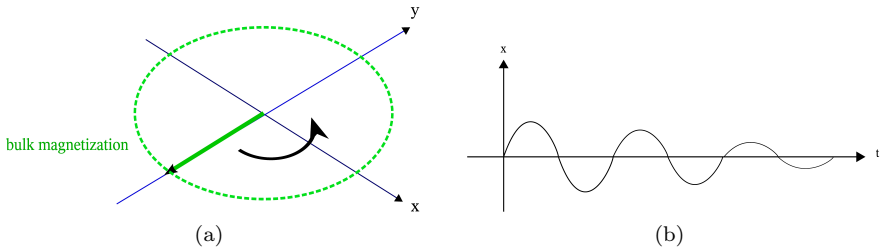

Figure 1.6: (a) Precession motion projected in the xy-plane. (b) The FID is acquired along the x-axis (also possible along the y-axis).

Here is a three-step summary of a pulse MRS data acquisition sequence :

1. Applying the strong magnetic field and reaching equilibrium magnetization

2. Applying a RF pulse

3. Acquiring the free induction decay signal along one of the x or y-axis

Proton MRS technique discerns metabolites through the measurement of protons and thanks to electron *shielding*. As presented previously in this section (see equation (1.1)), the resonance frequency of a proton differs from that of other nuclei. In addition, it has different resonance frequency depending on the field applied to it. The nuclei is surrounded by electrons, which also have a charge and a spin. They create a magnetic field in opposition to the applied magnetic field. Hence, the electrons shield the proton partially from the magnitude of the applied magnetic field. The quantity of electrons surrounding the proton depends on its location, i.e. which metabolite it is in. Hence, the measurement of the proton's frequency indicates in which metabolite it is localized. For more details on the physical foundations of MRS, refer to [7].

The frequency shifts due to electron shielding is very small compared to the resonance frequency differences between different nuclei types. The shifts are generally in the range of one *parts per million (ppm)*. The number of decimal digits needed make the absolute measurement of resonance frequencies impractical. A relative one is preferred, indicating the *chemical shift* or difference in ppm, of a specific metabolite compared to a reference one (placed at 0 ppm).

1.3 MRS applications

MRS can have many different applications. The focus in this section is put on the medical applications of the technique linked to diverse pathologies. There are two main elements to finding new medical applications to MRS. The first is to find the metabolites of interest concerning a specific pathology, i.e. those of which levels vary manifestly between ill and healthy patients. The second is to make sure that the metabolite in question can be discerned by MRS. The applications involve primarily the detection of different cancers, dementia, infections and epilepsies.

MR Spectroscopy is applied to the brain [8], prostate [9], breast [10], muscles [11], as well as to other organs such as the kidney and the liver. The main metabolites observed are N-Acetyl-Aspartate (NAA), choline (Cho), myo-Inositol (mI), lactate (Lac), creatine (Cr), Glutamate-Glutamine (Glx) and lipids.

A study was made in [12] comparing tumor and healthy brain tissue. For this purpose, 28 healthy volunteers and 18 patients with brain gliomas were compared. Table 1.1 gives metabolite concentrations in low-grade and high-grade brain gliomas as well as the concentraion of the same metabolites in healthy brain tissue.

Group	NAA (mM)	Cre (mM)	Cho (mM)
Healthy brain	23.59 ± 2.62	13.06 ± 1.8	4.28 ± 0.8
Low-grade glioma	6.28 ± 1.38	8.41 ± 2.67	3.55 ± 0.74
High-grade glioma	1.23 ± 2.54	4.74 ± 2.18	3.84 ± 2.36

Table 1.1: Metabolite concentrations in the brain

Concentrations given as: mean value ± standard deviation.
The absolute concentrations were calculated by relative comparison to the water level, itself given by external reference (method detail in chapter 2).

This study shows that the NAA concentration levels between healthy brain tissue (23.59 ± 2.62 mM), low-grade gliomas (6.28 ± 1.38 mM) and high-grade gliomas (1.23 ± 2.54 mM) are very different. Hence NAA quantification can be used to distinguish these three cases. Creatine values, although less far apart, can also be used to distinguish the three types of brain tissue. Choline values, on the other hand, are too close to each other to offer differentiation for the current cases evaluated.

Another example of MRS usage is the Alzheimer disease, most common form of dementia. This illness often develops for years before being diagnosed. Alzheimer is characterized by a reduction of NAA level, an increase in mI and unchanged Cr and Ch levels [13].

1.4 Outline of the thesis

Chapter 2 presents the particularities of quantification applied to in-vivo MRS data. A review of the existing quantification methods is given, displaying their main advantages and drawbacks. Then, the focus is put on the method chosen for implementation, namely SELF-SVD. After giving the principles of the method, follows a detailed explanation of the algorithm. The implementation of the algorithm as well as that of its complete testing environment in the MeVisLab software is then put forth.

Chapter 3 explains the experiments that were carried out and their results. First, the tests concerning simulated data are detailed. The rigorous testing protocol is presented, followed by the different simulation results. Propositions are made on how to set the different SELF-SVD parameters in order to increase quantification performances, depending on the characteristics of the given data. The second part of the chapter presents an application of the algorithm for the first time on in-vivo data. SELF-SVD is used to measure NAA and choline concentrations in order to attempt localization of a tumor in the brain region.

Chapter 4 sums up the findings emanating from the different tests made and puts forth ideas for future improvements of the current algorithm implementation.

Chapter 2

Methods

2.1 Quantification of MRS data

This section describes the singular characteristics of ^{1}H MRS data and presents general notions related to MRS quantification.

2.1.1 In-vivo ^{1}H MRS signal characteristics

MRS data acquisition produces a raw time signal (FID) in the form of damped oscillations in the presence of noise. The analysis of the signal under its spectral form (after FT) is often preferred.

The individual peaks in the spectra are characterized by their resonance frequency, lineshape, width and amplitude. These characteristics are a result of the content of the scanned tissue, but also of the acquisition conditions.

In-vivo ^{1}H MRS signals present characteristics making quantification particularly difficult. Overlapping peaks in the spectra are a main issue. For this reason, for instance in brain MRS, it is difficult to distinguish a low-grade neoplasm[1] from non-neoplastic lesions [14].

The *signal-to-noise ratio (SNR)* is generally relatively low. It can be improved by increasing scanning time and with extra signal averaging. The latter is performed through several acquisitions of the signal and the averaging of the results. A balance has to be found between growing scanning time and signal quality.

The voxel size is normally quite large (1cm per dimension) in order to be able to localize metabolites of low concentration. Furthermore, choosing smaller voxels also increases scanning time.

Issues to be taken into account in order to reduce noise in the signal, include minimization of patient motion, the obtaining of field homogeneity through shimming and age-related changes in brain metabolites [15].

[1] A neoplasm is an abnormal proliferation of cells. If it forms a lump, it is called a tumor.

2.1.2 General MRS quantification notions

Pre-processing is done before quantification in order to standardize the MRS signals. Standard pre-processing steps are water peak removal, zero-filling, baseline correction, phase correction and apodization. These steps have been implemented in the MRS MeVisLab software assistant and are detailed in section 2.4.1. Some quantification algorithms can also be used to process typical pre-processing steps (e.g. SVD-based methods can be used for water suppression [16]).

Quantification can be relative or absolute. Absolute quantification gives the absolute values of the metabolite concentrations. A wide-spread technique for this is the *external reference* method. In this method a phantom containing a known reference solution is placed near the area of the patient to be scanned, as shown in Figure 2.1.

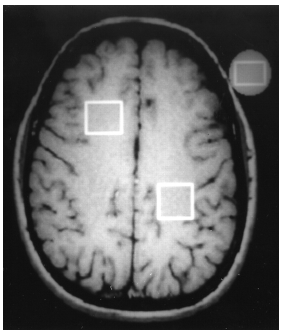

Figure 2.1: Scanning of brain and reference solution

The scanned ROIs are represented in the image by rectangles. There are two in the brain and one in the phantom reference solution (at the top right). The absolute concentrations of the ROIs of the brain are calculated by comparison with the known reference concentration value. The major drawbacks of the method are augmentation of patient scanning time and possible introduction of distortions of the constant magnetic field, due to the container used for the external phantom solution [8]. The latter makes pre-processing of the signal more complex. Other external methods also exist, such as the *replace-and-match* method and the *water signal reference* method, see [8] information on these.

Relative quantification on the other hand calculates peak ratios. One of the peaks is taken as a reference for the measurement of the others. Commonly, reference metabolites are NAA, creatine and choline. In CSI, the quantification measurements of a single metabolite can be compared between different voxels. This method is faster than absolute quantification but also has its drawbacks. In particular, measurements show that the concentration of the reference metabolite can fluctuate in an unexpected manner, hence frauding results. This could lead to ambiguous result interpretations. This is the case for example for choline levels in temporal lobe epilepsy [17] and frontal lobe epilepsy [18].

The inclusion of *prior knowledge* about the scanned metabolites can significantly improve quantification robustness and speed, in particular in cases where data is difficult to model (e.g. low SNR). Assumptions on lineshape model is a basic type of prior knowledge. Other knowledge includes amplitude ratios, relative resonance frequencies and scalar coupling between different metabolite peaks.

Physiological conditions, such as pH and temperature for instance, influence spectral patterns. Interactions between metabolites are another factor. Quality prior data is generally difficult to obtain, but nonetheless essential. If the knowledge is not a good representation of the data, it will systematically impoverish the quantification results instead of improve them. More prior knowledge is often reflected by more user-interaction, which can make the application more time consuming and complex for the physicians to use. Another point is that this knowledge is sometimes unfortunately not available.

2.2 Review of quantification methods

Quantification can be performed either on the raw MRS signal or its spectrum after FT. There exists some advantages and disadvantages to both approaches which will be presented in this section. No method stands out from all the others as the ideal quantification solution. The method to use depends on the application.

2.2.1 Time-domain methods

The time-domain methods can be divided into two categories, namely the interactive and the black-box techniques.

Interactive methods

These methods require a lot of user interaction thus giving them the name of *interactive methods*. The objective is to obtain the minimization of the difference between the model function and the MRS data. It consists typically in the resolution of a *nonlinear least-squares (NLLS)* problem. The vector notation of the function to be minimized is as follows:

$$\|y - \overline{y}\|^2, \tag{2.1}$$

$\|\|$: the Euclidean vector norm,
y: the vector of the original FID data points,
$\overline{y}$: the vector of the estimated FID data points.

The NLLS problem can be solved using either local or global optimization procedures. Local methods often lead to local minima, instead of the global one searched. Global optimizers take care of this problem. The drawback of the latter is poor computational efficiency, such as for the methods of *simulated*

annealing [19] and *genetic algorithms* [20]. Most methods in use are thus based on local optimizers. Local methods can be improved by obtaining good starting values for the algorithm. This speeds up computation and reduces the risk of resulting in a local minimum. Some local optimization methods are presented below.

The *variable projection method (VARPRO)* [21] was the first widely used method in MRS quantification. It is based on an algorithm called Levenberg-Marquardt. It was then replaced by the *advanced method for accurate, robust and efficient spectral fitting (AMARES)*, which improves the former method in terms of robustness and flexibility [22]. AMARES integrates several types of prior knowledge. It also offers, for each peak in the signal, the choice between Lorentzian and Gaussian typed lineshapes.

A selective-frequency version of AMARES [23] was created in order to model lineshapes contained in selected frequency zones. This version is preferred, for instance, if the water signal has not previously been eliminated in pre-processing stages. The method is not very effective in treating *nuisance peaks*, in particular to remove unwanted peaks with frequencies close to the ones of peaks to be retained [23, 24].

These model-fitting methods are very flexible and it is easy to include prior knowledge information. The major drawback of these methods resides in the requirement of much user-interaction. In addition, the time required for a local minimization algorithm to converge to a global minimum can be very variable.

Black-box methods

Black-box methods are called such because they are more automated than interactive ones. They usually incorporate little prior knowledge. These methods are generally either based on the *linear prediction (LP)* principle or *state-space* formalism.

LP methods proceed by estimating the prediction coefficients of the following equation:

$$y_n \approx -\sum_{m=1}^{M} p_m y_{n-m} + e'_n, \qquad with\ n = 0, ..., N-1 \tag{2.2}$$

y_n: MRS data point,
p_m: prediction coefficient,
e'_n: circular, white Gaussian noise,
M: prediction order,
N: number of data points.

The popular *linear prediction singular value decomposition (LPSVD)* method [25] is based on SVD decomposition (more details on SVD in section 2.3.1). By truncating the singular value matrix of the SVD decomposition, significant noise contribution is removed. The linear prediction coefficients are then estimated in a least square sense. Different variants of the method were developed. The *linear*

prediction total least square (LPTLS) method for instance, improves parameter estimation by calculating them in a *total least square (TLS)* sense [26].

All LP methods are unfortunately statistically suboptimal because they are based on a non statistically optimal model. Additionally, they carry out polynomial rooting, which is increasingly time-consuming with growing prediction orders. The root selection accomplished by LP methods can be particularly difficult in low SNR cases.

State-space methods on the other hand avoid polynomial rooting and root selection, by estimating the signal poles as eigenvalues of a matrix. The *Hankel-Lanczos singular value decomposition (HLSVD)* method [27] is the most popular of these. It improves the standard *Hankel singular value decomposition (HSVD)* method [28] by using the Lanczos algorithm and computing only part of the SVD algorithm. Another version of HSVD is the *Hankel total least square (HTLS)* method [29]. It improves quantification with spectra containing close peaks.

The *matrix pencil (MP)* [30] principle is very much alike the state-space algorithms. In [31] a comparison between the two is made. With noiseless data, results are the same, while in the presence of noise they are similar. MP is implemented in the MRS software discussed in section 2.4.1.

Inclusion of prior knowledge has been introduced through the *Hankel total least square with prior knowledge (HTLS-PK)* method [32, 33] derived from HTLS. Another more robust method was then presented, called *knowledge based total least square (KNOB-TLS)* [34]. It provides parameter estimations close to those found with AMARES. It has been used to give good starting values to the latter [34].

Black-box methods are also used to quantify and remove the water signal, if this has not been done during pre-processing. The water peak is usually modeled as a sum of Lorentzian peaks.

The main advantage of black-box methods resides in the minimal amount of user interaction required. The parameters are estimated without the need of a starting value sometimes difficult to determine. Their drawback is the limited amount of prior knowledge that can be incorporated, making them sometimes lack physical meaning. Another drawback for all SVD-based methods is that they need the data to be modeled as containing only Lorentzian-shaped signals.

2.2.2 Frequency-domain methods

The frequency-domain techniques can be divided into the peak integration, the model-fitting and the SVD-based methods.

Peak integration method

The simplest and most wide-spread quantification technique is peak integration. It consists in integrating the area under the peaks of interest in the frequency domain [35].

This technique is very sensitive to peak overlapping, to the presence of a baseline and to low SNR. In [36], it is shown that curve fitting outperforms by far

the peak integration technique, when applied with in-vivo data. An important factor to take into account is the choice of the boundaries for integration, which can greatly influence results [35]. Furthermore, when operating on the real part of MRS spectra, it is necessary to perform the delicate phasing process.

LCModel

Linear Combination of Model (LCModel) is an application developed for automatic quantification of in-vivo MR spectra. It is mainly applied to brain data. The method differs from other methods by its use of a complete model of the spectra rather than only individual resonances. LCModel analyses spectra as a linear combination of a set of model spectra of metabolite solutions in-vitro. A constrained regularization technique deals with complications in the data, which can be for instance due to the presence of a baseline, phase or to different metabolite lineshapes. LCModel includes a high quantity of prior knowledge, thus rendering the processing fully automatic, i.e. no parameters needs to be adjusted by the user. Refer to [37] for further information on its application on in-vivo ^1H spectra.

Model-fitting methods

Alongside the time-domain fitting methods, similar frequency-domain methods have been developed.

The FT and its inverse allow us to easily switch the result from the time to the frequency domain and vice-versa. The numerical version of the FT is called the *discrete Fourier transform (DFT)*. Its input and output are discrete functions of finite duration. In the case of the fitting methods, with a model different from the Lorentzian one, the frequency domain model (after DFT) does not have any simple and exact analytical expression. These fitting methods are hence faster and more precise when applied in the time domain.

It can be interesting to model other lineshapes than the typical Lorentzian. Ideally, the raw MRS signal is a sum of damped sinusoids, the FT of which is a series of Lorentzian lineshapes. Imperfections in the signal, due to the MR instrument or an inhomogeneous sample, can also lead to the presence of different lineshapes in the signal [38]. In [39] it is even suggested to use Gaussian lineshapes instead of Lorentzian ones.

The model-fitting techniques seek to solve the NLLS problem, presented in equation (2.1). Here again, most of the methods proceed by local optimization techniques.

An advanced method is presented here to illustrate what can be done in the frequency domain. It is called the *total-line-shape* method[2] [40]. This technique is based on the Gauss-Newton algorithm. It uses a Lorentzian lineshape for modeling, with the possibility of distorting the model in order to follow more accurately the MRS spectra. This has been used for example to deal with

[2]The abbreviation TLS exists in literature. It has not been adopted in this document to avoid confusion with the *total least-square* and its TLS abbreviation already in use.

magnetic field inhomogeneities [41]. A combination of Lorentzian components can also be used to model other lineshapes. It has been applied successfully with high resolution MRS [41, 42, 43].

SVD-based methods

The formal approach to these methods is the same as the SVD-based methods presented for the time domain in the previous section. A primary advantage is the possibility to implement frequency selection more easily than in the time domain. Frequency selection reduces the data-set to the points in the selected frequency-range, resulting in faster algorithms.

A comparison of five different frequency-selective SVD methods is proposed in [44]. The first method is *filter diagonalization method (FDM)* [45], in a selective-frequency version derivation. Then comes the *selective-frequency method of direction estimation (SELF-MODE)*, based on the *method of direction estimation (MODE)* [46], and *SB-HOYWSVD* [47], which follows the concept of *autoregressive moving-average (ARMA)* [48]. The *filtering and downsampling (FIDO)* [44] method consists of passband filtering and a downsampling approach. Finally, SELF-SVD is the frequency selective version of ESPRIT [49]. The performances of these methods are also compared to the standard time-domain HSVD method presented in section 2.2.1. The latter is the only non-frequency selective method. It hence covers the whole frequency range.

Tests were made with simulated data. The most interesting example concerns an attempt to simulate in-vivo ^{31}P MRS data of the brain. Eleven metabolite peaks are represented in the presence of Gaussian noise. The quality of the estimation of a metabolite peak (the first of the two *adenosine triphosphate (ATP)* peaks) was measured in order to compare the performance of the methods. The parameters of each algorithm were tuned to their experimentally found optimal values.

The parameter estimation of FIDO and SB-HOYWSVD are generally superior to the other methods. Nonetheless, FDM and SELF-SVD present a more stable performance behavior, thus offering more reliable estimations.

In the presence of noise, HSVD offers good results but then radically breaks down when the noise standard deviation is strong. With elevated noise, SELF-SVD is the most precise for amplitude and frequency estimation, while SB-HOYWSVD and FIDO have the better damping estimation.

With small separation between peaks in the spectra, amplitude estimation becomes difficult. HSVD performance is the best and remains constant even with short separation values. FIDO presents the second best estimation overall and SELF-SVD enjoys one of the most stable performances as peak separation is reduced.

SB-HOYWSVD offers the largest number of parameters, rendering it less user-friendly but offering the possibility to incorporate prior knowledge.

The computational speed of the techniques vary, depending on the width of the scanned frequency zone. The reduced scanning zone makes the frequency selective methods faster than HSVD. The fastest method is SELF-SVD.

2.3 Quantification with SELF-SVD

The quantification method of ***sel****ective-****f****requency quantification based on* ***s****ingular* ***v****alue* ***d****ecomposition* (SELF-SVD) will be detailed in this section. First the principle of SELF-SVD will be given, followed by the description of the algorithm and finally a discussion on the possibility of automating the selection of the number of modes to estimate.

The following notations will be used:

- A^T: transpose of matrix A,
- A^*: conjugate transpose of matrix A,
- $\lceil x \rceil$: integer part of x,
- $effrank(x)$: effective rank of x.

2.3.1 Principle

SELF-SVD is used to estimate the modes of the signal lying in a predefined frequency range. The modes, also called *formants*, are the basic components forming the signal. SELF-SVD is based on *singular value decomposition (SVD)*.

Frequency-selection reduces the computational burden compared to other formant estimation methods [44], such as *HSVD* [50] or *estimation of signal parameters via rotational invariant techniques (ESPRIT)* [49].

The FID signal is modeled as a sum of damped sinusoids in the presence of white noise, where each damped sinusoid is a mode of the signal. The model is given by the following equation:

$$x(t) = \sum_{j=1}^{m} \rho_j \lambda_j^t + \epsilon(t), \qquad \lambda_j = e^{-\alpha_j + i\omega_j} \; and \; t = 0, ..., N-1 \tag{2.3}$$

N: number of data points in the signal,
m: number of components,
λ_j: j^{th} mode,
ω_j: frequency of component j,
α_j: damping of component j,
ρ_j: complex amplitude of component j,
$\epsilon(t)$: noise term.

The Fourier-transformed spectrum corresponding to a sum of damped sinusoids gives a series of *Lorentzian* peaks, modeling the characteristic shape of the peaks in a typical MRS spectrum. The transformation of a single damped sinusoid from the time to the frequency domain is shown in fig 2.2.

In equation (2.3) the values of m and $(\omega_j, \alpha_j, \rho_j)$ are unknown. Nevertheless, it is generally unnecessary to evaluate all of the signal parameters. Owing to

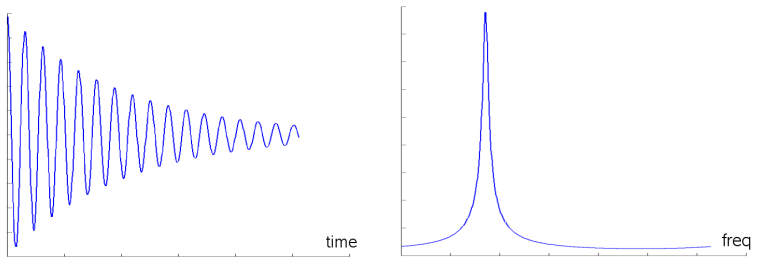

(a) (b)

Figure 2.2: (a) Time domain: a single damped sinusoid (b) Frequency domain: the corresponding FT is a Lorentzian peak.

its frequency-selective nature, SELF-SVD models only the n modes ($n \leq m$), comprised in this restricted band. The value of n can be defined by the user or estimated automatically by a method described in section 2.3.3. The objective of SELF-SVD is the modeling of the n components lying in the predefined frequency band.

SVD plays an important role in the formant-modeling process of the algorithm, where it is coupled with the *least square (LS)* method to carry out order reduction. SVD offers the possibility to single-out the n modes requested. The ones retained correspond to the highest energy peaks in the spectrum, i.e. presenting the largest areas under the curve (see figure 2.2b above), or highest quantification results.

The SVD of a matrix consists in performing a particular factorization in three sub-matrices. For example, when applied to a complex rectangular $i \times j$ matrix A, it results in a decomposition of the form:

$$A = USV^*,$$

where U is an $i \times j$ orthogonal matrix, V a $j \times j$ orthogonal matrix and S an $i \times j$ diagonal matrix. The columns of U and of V are respectively called the *left* and *right singular vectors*. The diagonal elements of S are real and non-negative. They are called the *singular values* of A.

In the SELF-SVD algorithm, a SVD decomposition is performed on a particular matrix, named Z and introduced in section 2.3.2 and representing the FID data points on which a frequency-selection process has been performed.

From Z and equation (2.3), the SVD decomposition is put in a canonical representation, called the *modal decomposition form*. In this configuration, the σ_i singular values of matrix S represent the energy levels of the spectral peaks corresponding to the different modes in the signal. They are organized in S in a decreasing order, as follows:

$$\sigma_1 \geq \sigma_2 \geq ...\sigma_m \geq 0.$$

S is finally truncated in order to keep only the n highest energy modes. Please refer to [51] for more information on the use of SVD for formant estimation.

2.3.2 The SELF-SVD algorithm

The formant model used by SELF-SVD for the MRS signal has been presented in the previous section, in equation (2.3). SELF-SVD has for objective the estimation of the parameters of the formants in the equation. The number of formants chosen for estimation, n, is inferior or equal to the m present in the signal, as mentioned in 2.3.1. For each of the n formants to be modeled, the following are estimated:

ρ: complex amplitude
α: damping
ω: frequency

SELF-SVD offers the following user parameters, which can be fixed before computation:

- k_1 and k_M, defining the frequency band where SELD-SVD should be performed. The algorithm will operate on the M frequency values:

$$\{\frac{2\pi k_1}{N}, \frac{2\pi k_2}{N}, ..., \frac{2\pi k_M}{N}\}, \text{ with } k_{i+1} = k_i + 1$$

- n: the number of modes to estimate in the chosen frequency band. It can be set manually or automatically by the algorithm. For more detail refer to 2.3.3.

- P: the accuracy variable. P is an integer, which will play a role in equation (2.18), presented further. It represents the length of the FIR filter and is used during the complex amplitude estimation of the modes. Depending on its value, either the estimation is more precise or computational speed is increased. The constraint on it is to have $P > 1$.

- S: should be chosen in the following interval:

$$S \in [\lceil \tfrac{M}{3} \rceil, \lceil \tfrac{M}{2} \rceil]$$

  S thus respects the conditions for equations (2.4) and (2.13), detailed later on in this section. S is taken as $\lceil \frac{M}{3} \rceil$, value found to reduce the computational burden of the method and proposed in [4].

The choice of the user parameters n, P and the choice of the frequency band on which to operate (variables k_1 and k_M) are further discussed in chapter 3.

Here follows a step-by-step description of how to apply SELF-SVD, according to [4]:

1. Compute {X(k)}, with k=0,...,N-1
   $\{X(k)\}_{k=0}^{N-1}$ is the *fast Fourier transform (FFT)*[3] of the MRS signal $\{x(t)\}_{t=0}^{N-1}$.

[3] A fast Fourier transform is an efficient algorithm to compute the discrete Fourier transform and its inverse.

2. Compute $\Pi_U^{\perp}$
   $\Pi_U^{\perp}$ is the orthogonal projection matrix onto the null space of U. Under the condition of $(M \geq S + n)$, it is defined as:

$$\Pi_U^{\perp} = I - U^*(UU^*)^{-1}U \tag{2.4}$$

   Matrix U is formed as follows:

$$U = [u_{k_1} ... u_{k_M}] \; and \tag{2.5}$$

$$u_k = [w_k^1 ... w_k^S]^T, \; with \; w_k = w^{\frac{i2\pi k}{N}} \tag{2.6}$$

   $\Pi_U^{\perp}$ is calculated through the QR decomposition of U^*:

$$U^* = \begin{bmatrix} \bar{Q} & Q \end{bmatrix} \begin{bmatrix} R \\ 0 \end{bmatrix} \; and \; Q^*Q = I_{M-S} \tag{2.7}$$

   $\Pi_U^{\perp}$ is then directly derived:

$$\Pi_U^{\perp} = QQ^* \tag{2.8}$$

3. Estimate the modes of the signal, $\{\lambda_k\}_{k=1}^n$

   (a) Compute Z as follows:

$$Z = X\Pi_U^{\perp}, \; with \; X = [u_{k_1} X_{k_1} ... u_{k_M} X_{k_M}] \tag{2.9}$$

   (b) Compute SVD(Z) and adjust n if needed
   Check the following constraint:

$$effrank(Z) = n \tag{2.10}$$

   If the condition is not met, adjust n accordingly.

   (c) Build W
   W is the $S \times n$ matrix whose columns are the left singular values of Z associated with the n largest singular values.

   (d) Estimate $\{\lambda_k\}_{k=1}^n$ from W by a SVD and LS method, as follows:

   - Derive W_l and W_u from W:

$$W_l = \begin{bmatrix} 0 & I_{S-1} \end{bmatrix} W \tag{2.11}$$

$$W_u = \begin{bmatrix} I_{S-1} & 0 \end{bmatrix} W \tag{2.12}$$

   - Introduce the singular values:
     A is the matrix defined by:

$$A = \begin{bmatrix} \lambda_1 & ... & \lambda_n \\ \vdots & & \vdots \\ \lambda_1^S & ... & \lambda_n^S \end{bmatrix}$$

The following approximation is shown in [4]:

$$A \approx WK, \tag{2.13}$$

with K a nonsingular transformation matrix

It is based on several assumptions, only briefly described here:

i. the noise is relatively small compared to the signal.

ii. the out-of-bound components do not influence too strongly the components in the selected frequency band.

iii. $S \geq n$.

With the assumption that these conditions are met, we derive the following:

$$W_l P \approx A_l = A_u \begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix} = W_u K \begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix} \tag{2.14}$$

By setting:

$$\phi = K \begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix} K^{-1}, \tag{2.15}$$

equation 2.14 can be rewritten as:

$$W_l \approx W_u \phi \tag{2.16}$$

- Estimate ϕ from equation (2.16) by the LS method:

$$\hat{\phi} = (W_u^* W_u)^{-1} W_u W_l \tag{2.17}$$

- Insert $\hat{\phi}$ in equation (2.15) and derive the estimation of the modes as follows:

$$\begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix} = K^{-1} \phi K$$

4. Estimate the complex amplitude $\{\rho_k\}_{k=1}^n$ by the *amplitude and phase estimation (APES)* method. APES will use the the signal modes $\{\lambda_k\}_{k=1}^n$ estimated in the previous point. The method was chosen over others because of its accuracy in amplitude estimation and reasonable computational cost [52]. Here below is a brief description of the method.

   $\{z(t)\}_{t=0}^{L-1}$, contains the MRS signal and is constructed such:

$$z(t) = [x(t)...x(t+P-1)]^T, \; and \; t = 0, ..., L-1 \tag{2.18}$$

   with $L = N - P + 1$,
   and $P > 1$ a user parameter discussed in chapter 3

A *finite impulse response (FIR)* filter is applied to the signal. Its output is $h^*(\alpha,\omega)z(t)$, where $h(\alpha,\omega)$ is the coefficient vector of the filter. The filter has to fulfill the following conditions:

(a) the damped sinusoid $\{\rho e^{(-\alpha+i\omega)t}\}$ should be unchanged by the filter,

(b) the output of the filter should be as close as possible in the LS sense to a damped sinusoid with given frequency and damping (ω,α), and an amplitude (ρ) that minimizes the LS fitting error.

Condition (a) is satisfied if and only if:

$$h^*(\alpha,\omega)s(\alpha,\omega) = 1,\ with\ s(\alpha,\omega) = [e^{-\alpha+i\omega}...e^{(-\alpha+i\omega)(P-1)}]^T \tag{2.19}$$

Condition (b) is satisfied by minimizing the following LS criterion, with respect to $h^*(\alpha,\omega)$ and ρ:

$$\sum_{t=1}^{L} |h^*(\alpha,\omega)z(t) - \rho e^{(-\alpha+i\omega)t}|^2,\ with\ given\ \alpha,\omega \tag{2.20}$$

The estimation of the complex amplitude $\hat{\rho}$ is given by resolving 2.20. It is resolved in [53]. The coefficient vector of the filter can be expressed as:

$$h(\alpha,\omega) = \frac{Q^{-1}(\alpha,\omega)s(\alpha,\omega)}{s^*(\alpha,\omega)Q^{-1}(\alpha,\omega)s(\alpha,\omega)},\ with \tag{2.21}$$

$$Q\ = \sum_{t=1}^{L} z(t)z^*(t) - \sum_{t=1}^{L} e^{-2\alpha t}Z(\alpha,\omega)Z^*(\alpha,\omega) \tag{2.22}$$

Z is expressed as follows:

$$Z(\alpha,\omega) = \frac{\sum_{t=1}^{L}[z(t)e^{-\alpha t}]e^{-i\omega t}}{\sum_{t=1}^{L} e^{-2\alpha t}} \tag{2.23}$$

Thus, the amplitude is estimated by:

$$\hat{\rho}(\alpha,\omega) = h^*(\alpha,\omega)Z(\alpha,\omega) \tag{2.24}$$

The core of the SELF-SVD algorithm resides in equation (2.13). Its origin is explained in detail in [4]. The equation is used in order to estimate the $\{\lambda_k\}_{k=1}^n$ modes of interest via a SVD-based method. Their associated amplitudes $\{\rho_k\}_{k=1}^n$ can then be derived, for example by the APES method.

2.3.3 Selection of the number of modes to estimate

SELF-SVD offers the possibility to calculate n, the number of modes present in the frequency range of interest. This is shown through equation (2.10) of the previous section, where the *effective rank* of a matrix (referred to as Z) gives the number of modes in the frequency band. The assumptions made to get to this result will not be detailed here, but can be found in [4].

In theory, the calculation of the rank of matrix Z should be sufficient to get the number of modes in the frequency band of interest. In practice, however, due to finite precision of numerical operations, restriction of the frequency band and additional noise, all the singular values of the matrix have non-null values, hence making the rank estimation method unusable.

The effective rank must thus be used. The calculation of the effective rank of a matrix representing noisy data has for objective to give its underlying or *effective* dimension. This is performed by excluding the extra dimensions added by errors due to calculation imprecisions and additional noise. Note that the effective rank is noted $effrank$, which should be distinguished from the $erank$ notation, also called effective rank but referring to a different mathematical notion.

Effective rank calculation of a matrix representing noisy data is non-trivial and constitutes a research field on its own. It goes beyond the scope of this thesis. An approach to it can be found in [54]. Tests were hence brought out with manual selection of the number of modes and not through $effrank$ calculation. A simple automated version of n selection was nevertheless implemented for the in-vivo tests and is described in section 3.2.3.

2.4 Implementation of SELF-SVD in MeVisLab

This section presents the implementation of the SELF-SVD algorithm. After a description of the MeVisLab software and development platform, as well as its MRS software assistant, the implementation of SELF-SVD as well as its complete testing environment is given.

2.4.1 Overview of MeVisLab and its MRS software assistant

MeVisLab has been developed by *Fraunhofer MEVIS* (formerly MeVis Research GmbH) and *MeVis Medical Solutions (MMS)*. It is a software for computer assistance in medicine, in particular in radiology. It includes computer aided diagnosis, therapy planning and monitoring. The MeVisLab Basic version, containing the image processing tools and visualization functionality, can be obtained free of charge. For more information refer to *http://www.mevislab.de*.

MeVisLab is also a development platform where various data/image processing modules can be created. The numerous modules range from the implementation of mathematical functions, to graph visualization and image processing

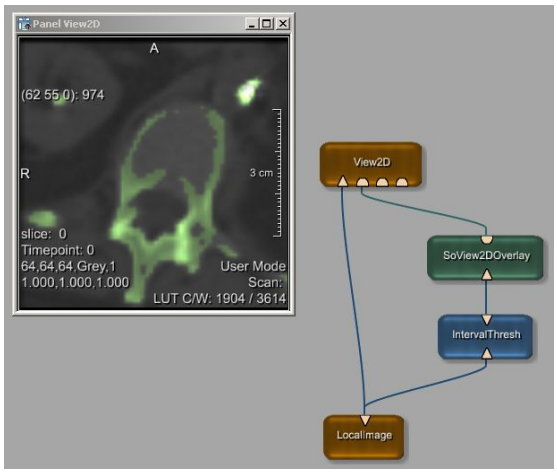

Figure 2.3: MeVisLab network example

modules. These modules are put together in a flow chart in order to perform a more complex task. A simple flow chart network is shown in figure 2.3.

The data flow is directed from bottom to top. Often in image processing networks such as this one, the data flow starts from an image data opening module (*LocalImage*) and finishes with a visualization module (*View2D*). In this example, the *IntervalThresh* module allows the user to choose a grey-level threshold value. Thus only the part of the image respecting this restriction is given to the output of the module. The following module in the network, *SoView2DOverlay*, allows the superposition of the thresholded image on the initial image in the View2D viewer. The result has been displayed. The original image is in grey-scale, whilst the threshold image is in green.

New modules can either be created as macro modules, i.e. composed of other modules in a network (e.g. figure 2.3), or directly programmed in C++ language. The *graphical user interface (GUI)* layouting can either be generated automatically or defined with the internal *module definition language (MDL)*. Batch processing and application control can be performed through Python or Java scripting.

A software assistant for the analysis of MRS data has been developed under MeVisLab [55]. Its goal is to provide a fast and robust interpretation of MRS signals. More specifically, it aims to be a reliable cancer diagnosis tool for physicians.

The software assistant proposes a novel approach by combining the display of MRS and MRI data. Figure 2.4 is an image taken from the software. On the top is the MRI image of the prostate, where the MRS scanning grid is superposed. Below is the quantification viewer. The colors represent the strong (red) or weak (green) ratio between the two metabolites studied (1^{st} and 2^{nd}). With this viewer, the physician can have an overall picture of metabolite distribution and localize malignancies faster than by having to look at the spectra of each

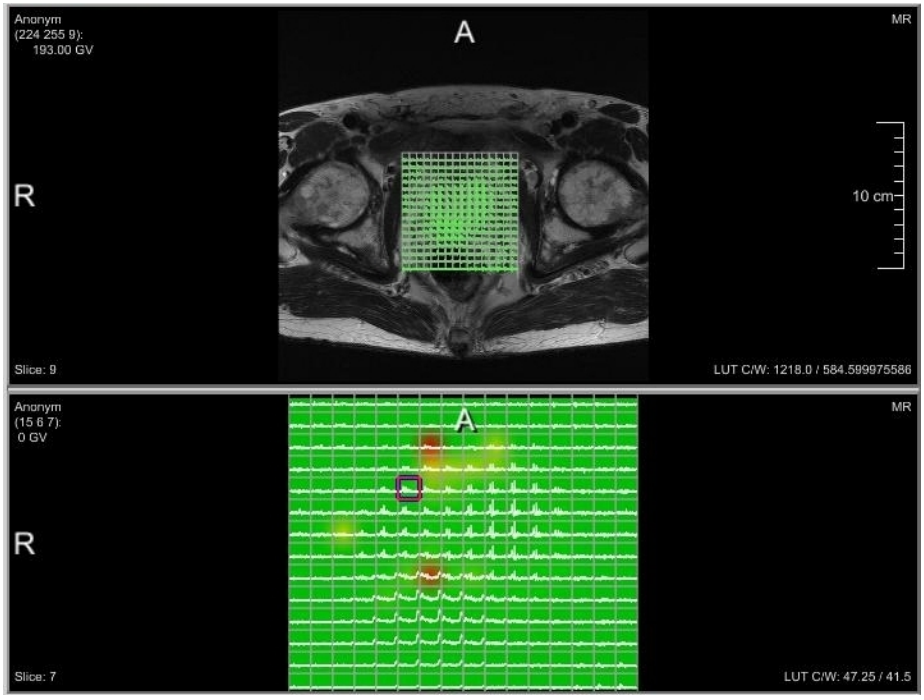

Figure 2.4: MeVisLab MRS software viewers

voxel individually. Looking closer, one can notice that superposed to each voxel in white is its corresponding MRS spectra. This is also the case in the voxels of the MRI image, where the spectra are printed in green.

By clicking on a voxel in the quantification viewer, it is marked in red and its spectrum is magnified and displayed in the viewer shown in figure 2.5. Baseline and phase correction have been applied to the data to make it easier to analyse. The two colored areas selected in the spectrum are the areas of interest, representing the two metabolites of interest.

The software assistant provides several preprocessing tools, listed here below:

- Apodization
  Restrictions of measurement time force the acquired FID to be truncated before it has decayed to zero. This truncation causes the presence of disturbance oscillations at the foot of the peaks in the spectra. This effect is suppressed by multiplying the FID by an apodization function, which smoothes the truncated edges of the signal. Several apodization functions are available (Hamming, exponential, etc.).

- Extrapolation
  Extrapolation is the process of constructing new data points outside a discrete set of known data points. It is an important process for ^{1}H MRS

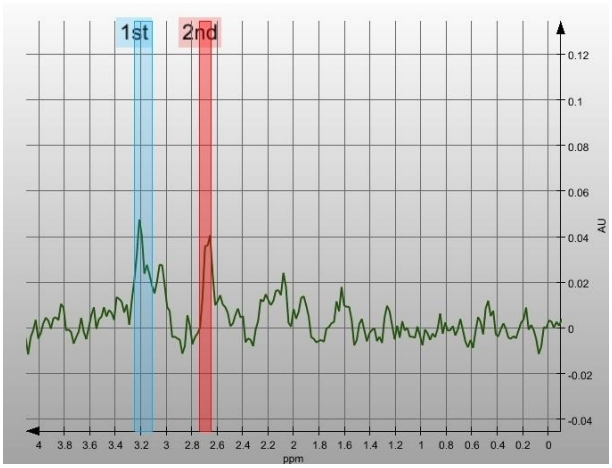

Figure 2.5: MeVisLab MRS software spectrum

because of the presence of close narrow peaks in its spectra. Two solutions are implemented. Zero-filling is the first. It consists in extending the sampled time domain points with points of zero amplitude. It results in higher resolution spectra. The second solution provided is linear prediction (theory discussed briefly in 2.2.1).

- Phase correction
  The spectrometer produces time-domain data with arbitrary phase. Both manual and automated phase correction are proposed in the software.

- Baseline correction
  Depending on the quality of the acquisition, baseline correction is required. Different filter types are proposed for this purpose. For instance, when used in its *dominant filter mode*, it can perform water peak removal.

The MRS software also has two quantification tools implemented. For these, relative quantification has been preferred to insure scale-invariance. The available tools are the following:

- Matrix-Pencil method
  The MP method estimates the parameters of the components that compose the signal. In this method, the FID signal is approximated as a sum of decaying sinusoids. It was discussed in section 2.2.1.

- Area under the curve
  It consists in integrating the area under the FT curve to quantify metabolites. The method was discussed in section 2.2.2.

2.4.2 The implementation

A module was created in the MeVisLab platform for SELF-SVD calculation. It has been written in C++ language and named `SelfSVD`. Its GUI can be seen in figure 2.6 alongside the module.

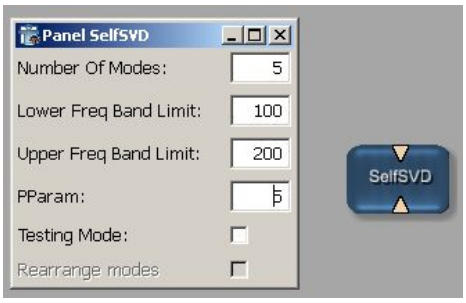

Figure 2.6: SELF-SVD module

The module takes as entry the complex time points of the FID data. As output, it gives the parameter estimations for the components of the damped sinusoid sum function. For each component in the spectra, the returned parameters are the damping, frequency and complex amplitude (presented in section 2.3.1). The following user parameters are to be defined:

- *Number of modes*: the number of modes to be modeled,
- *Lower / Upper frequency band limit*: these define the interval of the spectrum on which to perform the SELF-SVD algorithm.
- *P parameter*: it is the amplitude estimation precision parameter, corresponding to the length of the FIR filter.

The parameters have been detailed in section 2.3.

Two additional options were implemented for facilitating testing. They make comparison of the results obtained more straightforward. The options are the following:

- *Testing mode*: when checked, the output of the modes is given in a format matching the second output of the `MRSdataSimulator` module (detailed below in this section).
- *Rearrange the modes*: when checked, the output of the modes is given in increasing frequency order. This mode is only available if the *Testing mode* box is checked.

In order to calculate certain mathematical functions, such as the SVD decomposition, an existing mathematical library was to be chosen. CLAPACK[4]

[4]Refer to http://www.netlib.org/clapack for more information

is the C version of the mathematical library *linear algebra package (LAPACK)*. CLAPACK was chosen because it is the only free library found to offer SVD and QR calculations on complex matrices.

The `SelfSVD` module follows the input/output data types set in MeVisLab's MRS software. The module `SinusoidsFromParameters` in the MRS application takes parameters from the `MatrixPencil` module and created curves, i.e. a series of complex data points. The curves are generated as a sum of damped sinusoids. This module has been adapted to generate curves taking input either by `MatrixPencil`, as previously, or from `SelfSVD`.

A seperate testing module, called `TestSelfSVD`, was created in order to test `SelfSVD`. The test module tries to be as complete as possible in order to identify the influence of the different parameters on its performance and to reveal the potential and limits of the SELF-SVD algorithm. `TestSelfSVD` is a macro module composed of a number of sub-modules, two of which were implemented in C++ and are described here below.

The first is an FID generator, called `MRSdataSimulator`. It takes as parameters the number of modes to generate and their parameters (identical to those evaluated by `SelfSVD` in order to compare results, see section 2.3.2 for more detail). The number of data points to be created can also be selected. The generator offers a second output where it gives out the initial parameters mentioned above. This is used to compare the `SelfSVD` parameter estimations to the initial parameters known to be in the signal.

Additionally, `MRSdataSimulator` offers the possibility to simulate the error created by the sampling of the data, frequently called *aperture effect*. This additional error is not reflected on the parameters given in second output, as this output gives the underlying perfect signal characteristics. The aperture effect is an error due to the fact that the sampling is obtained as a time average within a sampling region, as opposed to at the sampling instant.

The other C++ testing module implemented is `MatchMRSmodes`. This module has for objective to match the SELF-SVD estimated modes to those given to `MRSdataSimulator`. Depending on the quality of the estimation, the matching process can be more or less complex. Matching modes manually by inspection would be the most precise technique. Nevertheless, when performing a series of tests with many different parameter values and many signal modes, manual matching becomes too time consuming.

The method used for automating the matching is based on the estimations of the frequency and damping. The amplitude was left out as its estimation sometimes lacks accuracy and can lead to bad matching. The method of minimization of the least square error was implemented in order to choose the modes to match. The least-square error for each match made is printed in order to enable the user to detect possible matching abnormalities.

`TestSelfSVD` is composed of a number of modules. A screenshot of the network is shown in figure 2.7.

Modules have been put in functional groups, that can be seen in color in the network. Each and every module will not be detailed, but the general pipeline will be explained.

At the bottom of the network is the generator of simulated MRS data. On the left, in yellow, are the modules that add Gaussian noise to the signal, if wanted by the user, in order to simulate real data. Above this block, is the `SelfSVD` module. Once it has estimated the modes of interest, the red block compares them to the original modes in the signal. This block contains the `MatchMRSmodes` module, mentioned above, and an arithmetic module, `Arithmetic2`, that calculates the estimation error for each mode in the signal.

At the top left, the purple group enables visualization of the estimation error for each parameter and mode estimated. This group has been simplified for clarity reasons to the modules necessary for the visualization of a single parameter estimation, i.e. frequency estimation. The visualization pipelines for the other parameters (damping, real part of amplitude and imaginary part of amplitude) are very similar.

On the right is the group for calculation and visualization of the signals. Here again the module has been simplified. The pipeline presented is linked to the visualization of the original MRS spectrum. Other similar branches perform the visualization of the signals issued from the `SelfSVD`- and `MatrixPencil`-estimations. A FT module, `FFT1D1`, is used in order to visualize the signals in the spectral domain. The three spectra are then superposed in a single display in order to compare results.

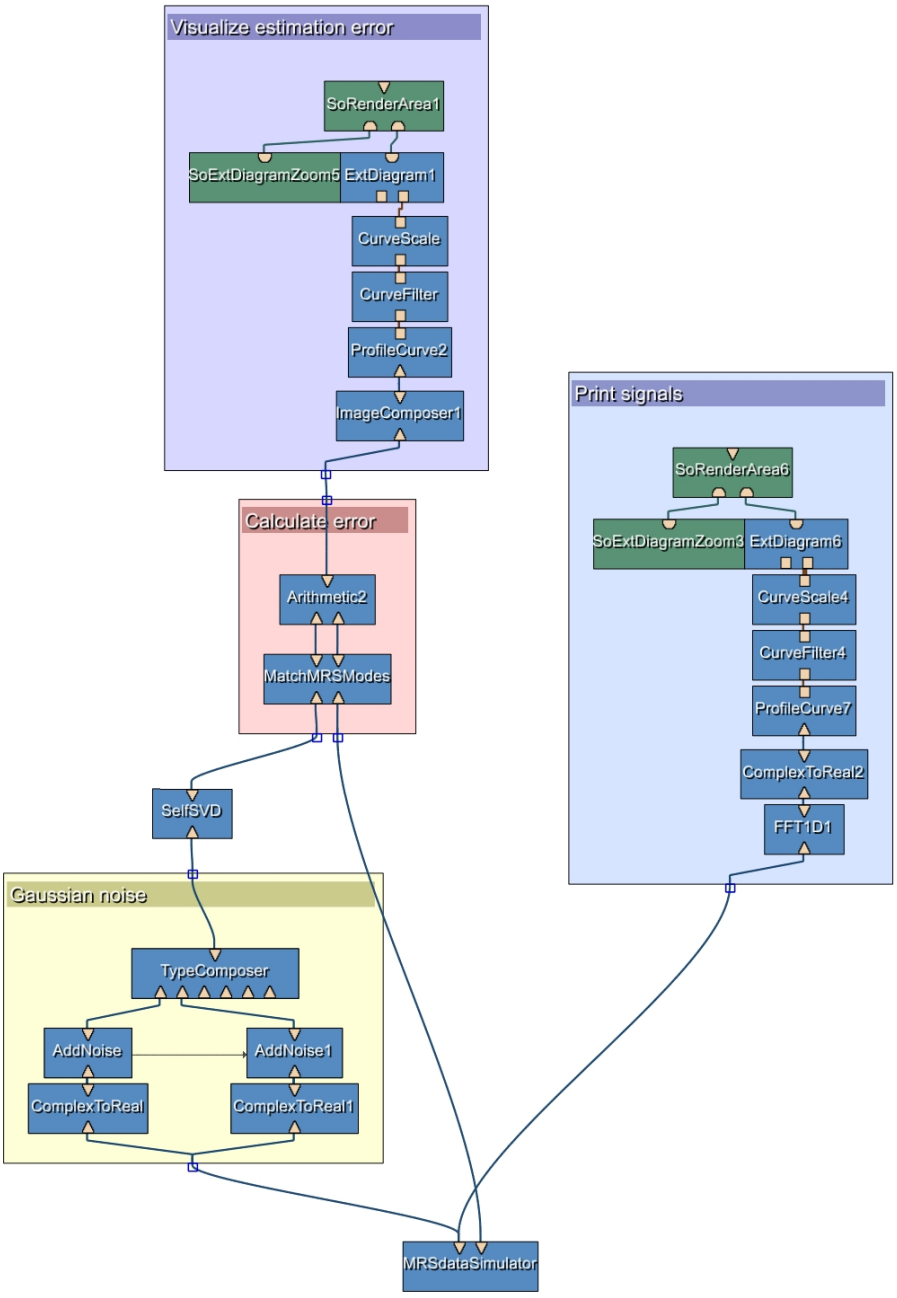

Figure 2.7: TestSelfSVD network

The user interface related to `TestSelfSVD` will be detailed below, in order to present the possibilities offered by the testing module. The GUI has been divided on several tabs. The first is shown in figure 2.8. The number of modes of the testing signal, as well as their parameters, have to be detailed. Furthermore, the user can choose to simulate the sampling aperture effect or add Gaussian noise to the signal.

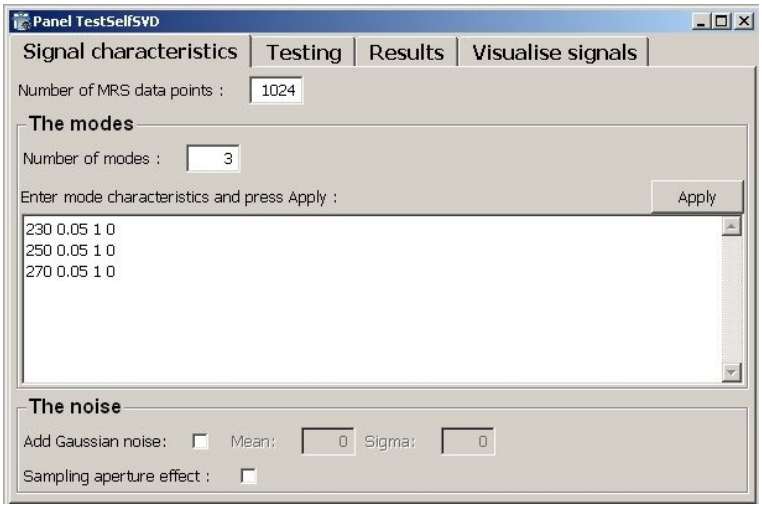

Figure 2.8: TestSelfSVD: Signal characteristics tab

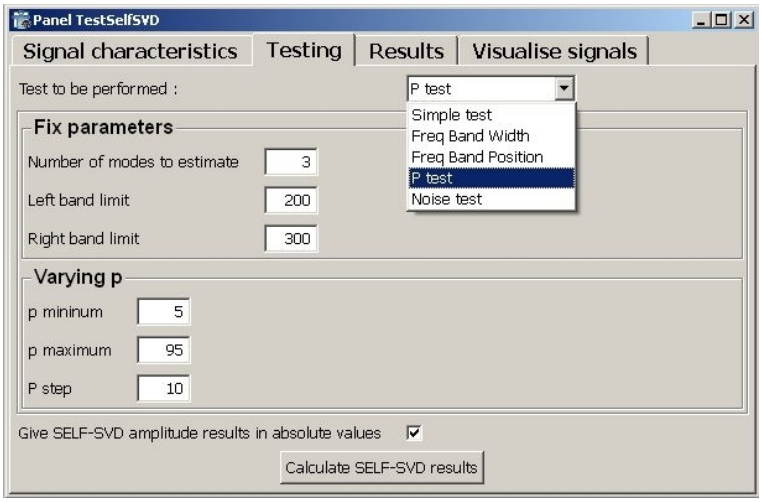

Figure 2.9: TestSelfSVD: Testing tab

In the second tab, different testing modes are proposed. It is shown in figure 2.9. The user can choose between testing the influence of the band width, that of the band position or that of the FIR length (parameter P). A standard *simple test* mode is also available. Depending on the test performed, different parameters are requested. Once these have been entered, the *Calculate SELF-SVD results* button should be pressed for calculation. Just above, the user can choose to have the amplitudes given in absolute values, as opposed to as complex numbers. In the *simple test* mode, it is also possible to print the `MatrixPencil` results, for direct comparison to those of `SelfSVD`.

The third tab presents the results of the test. An example is shown in figure 2.10. The estimation results are presented in two formats, as text and in graphs. The graphical representation was created in order to promote the finding of correlations between the parameter values and the estimation precision. At the top of the tab, the mode to be visualized can be chosen. Each graph then presents the errors in the estimation of a parameter (frequency, damping, real part of amplitude or imaginary part of amplitude). If in the second tab, the user chose to estimate the amplitude in absolute value, as in this example, there are three parameters to visualize instead of four. The error is printed on the y-axis in function of the value of the tested parameter, given on the x-axis (e.g. P parameter).

The estimation results are also presented with their exact values in the text window below the graphs. Here, the matching with the initial signal modes is presented. The *least-square matching error* is given (detailed above in this section) in order to enable the user to discern potential matching errors.

The fourth tab can be used in the case of a *Simple test*. A screenshot is shown in figure 2.11. Here the original signal is superposed to the SELF-SVD-estimated signal. Several viewing options are proposed. Amongst these, the signals can be seen in their real or absolute values, in the time domain or the Fourier domain.

To conclude the testing implementation section, here is a summary of the different possibilities offered by the `TestSelfSVD` module:

- generation of an MRS signal with the number of modes and characteristics entered,
- simulation of the sampling aperture effect,
- addition of Gaussian noise to the signal,
- manual choosing of the number of modes to estimate,
- automatic matching of the simulated modes with those in the original signal,
- visualization of the estimation error for each parameter estimated,
- automated testing of the influence of frequency band position, width and FIR filter length parameter.

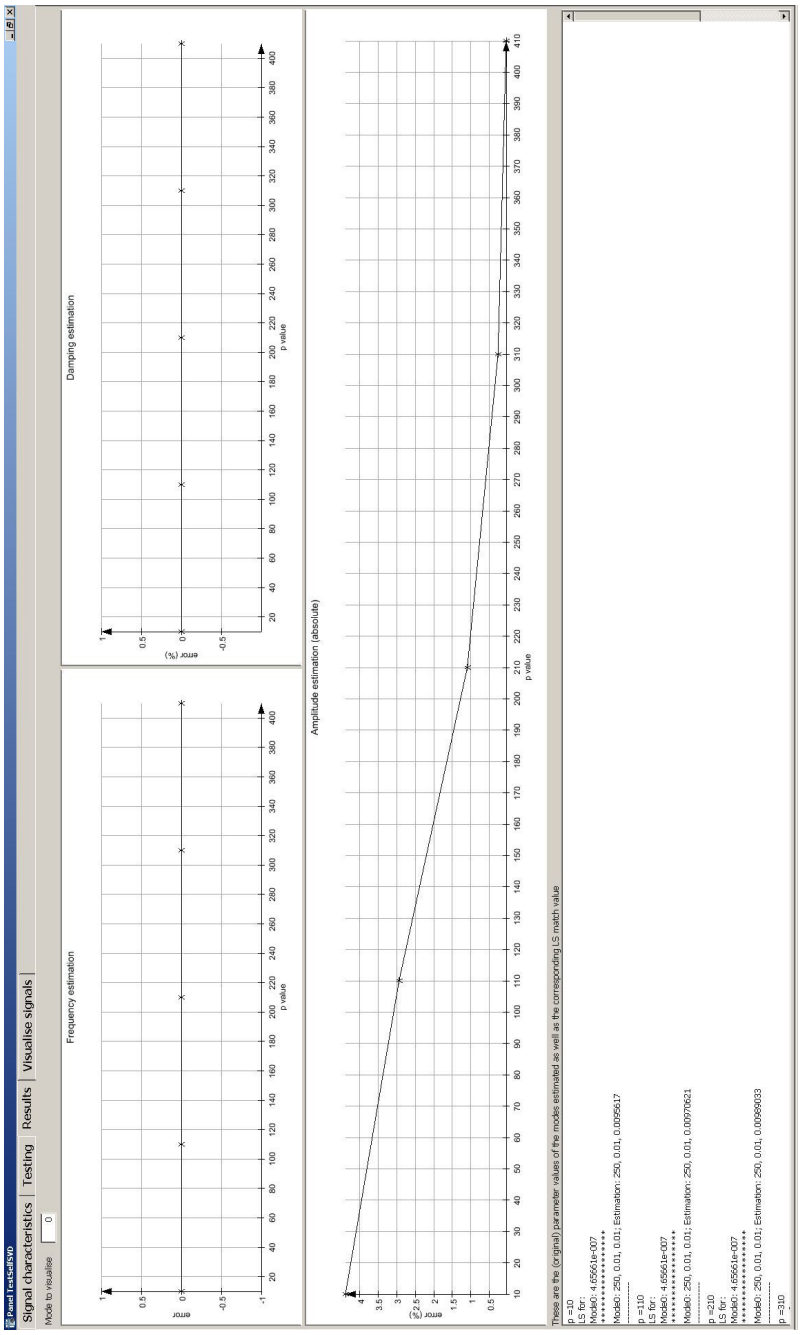

Figure 2.10: TestSelfSVD: Results tab

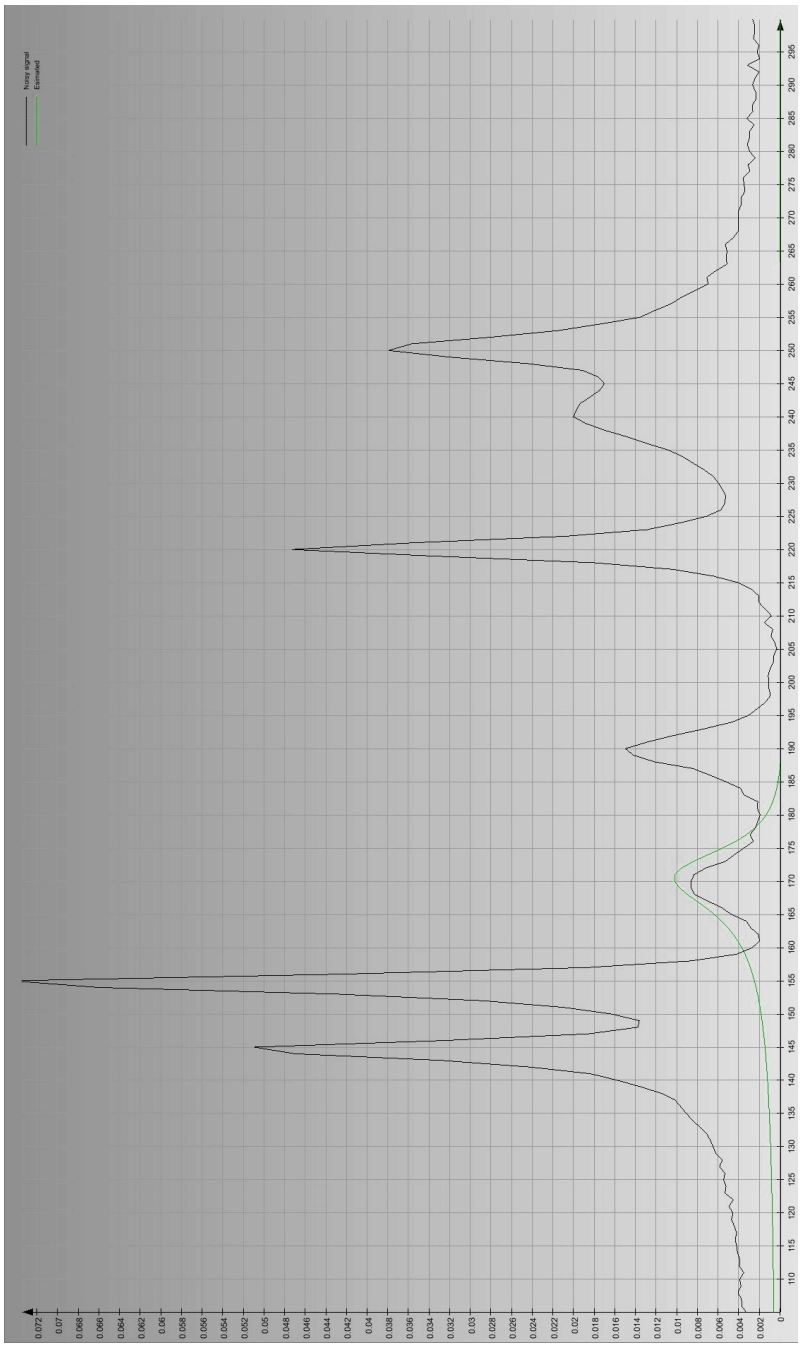

Figure 2.11: TestSelfSVD: Visualize signals tab
Line in black: Simulated signal
Line in green: SELF-SVD estimation

Chapter 3

Experiments and results

3.1 Tests on simulated data

Simulated data was used in order to test the absolute performance of the algorithm. The results throw light on the response of SELF-SVD to different types of input signals and different parameter settings. The characteristics found can then give hints for optimal use of the algorithm with real data.

Result accuracy is measured by comparing estimation results with the parameters initially used to create the signal (refer to section 2.4 for more details).

This section presents the protocol used for testing, followed by the results, given in five subsections. The first concerns the modeling of a single mode, where it is not influenced by others. The second and third sections concern respectively the total modeling of two- and three-mode signals and the modeling of a single mode out of a multiple-mode signal. The advantages and disadvantages of estimating one peak in a certain frequency band or several peaks in a larger band is also discussed. The fourth section presents the testing with additional Gaussian noise, and the fifth focuses on the particular case of overestimating the number of modes to be modeled. Finally, a complete simulation example is detailed.

3.1.1 Protocol

The modes generated for testing are based on values taken from a MRS data from the prostate. The data used comes from a Siemens scanner and a $16 \times 16 \times 16$ 3D CSI measurement.

In the voxels of the data that were analyzed, the characteristics of the peaks of interest were used to derive a range of possible values for the three parameters estimated by SELF-SVD (presented in equation (2.3) of section 2.3.1). Standard values have thus been determined, as well as a range which encompasses the minimum and maximum values measured.

The parameter values were determined as follows:

- frequency (w): the proximity of the peaks is here of interest. The closest peaks analyzed are creatine (at 3.0270 ppm) and choline (at 3.1850 ppm). This type of proximity will be tested. On a scale of 1024 data points, this precision corresponds to a peak separation of 16 points.
- damping (α): the range taken is $0.005 <= damping <= 0.05$ and 0.01 *du* as the standard value. *du* stands for *damping units*.
- amplitude (ρ): the range taken is $0.005 <= amplitude <= 0.05$ and 0.01 au as the standard value. *au* stands for *amplitude units*.

The simulated signal is generated with 1024 data points. Hence, the position of the peaks as well as the separation between peaks is given in number of data points.

In the systematic testing procedure adopted, estimations of 1, 2 and 3 modes were carried out. First, a *standard test* is performed. In this test the mode parameters are set to the standard values mentioned above. One parameter is then modified to an extreme value, e.g. damping set to 0.05 au. When modeling several modes, the peak separation values taken are 20, 15 and 10 data points, in order to evaluate its effect. The results of these tests are then compared to the standard test and thus the effect of the parameter changed is evaluated.

3.1.2 Single mode estimation in unconstrained data

These tests determine the best estimations one can expect in *single mode estimation*. The selected mode is taken in a signal with a single mode or well-spaced peaks, i.e. peaks spaced in order not to influence each other. A minimum spacing of 50 data points is required with the parameters used.

By changing the parameters of frequency (ω), damping (α) and amplitude (ρ) in the generated data, the variation in estimation quality is observed. In addition, the effects of the parameters of SELF-SVD are examined, namely the frequency band of interest and parameter P.

With standard settings for damping and amplitude, the frequency band on which to calculate SELF-SVD has little influence on the estimated parameters. Estimation of frequency and damping are precise, while amplitude estimation presents an inaccuracy.

The graph 3.1 presents the error on the amplitude estimation in function of the FIR filter length parameter, P. The graph will be named *amplitude error graph*. It will be very much of interest and discussed in the following sections too.

The graph shows that the greater P is taken, the more the amplitude estimation is stable and its value accurate. For example, in order to have less than a 1% error on the amplitude, P should be set to 220 or above. The shape of the curve flattens progressively at the end. Hence, varying P from a value of 100

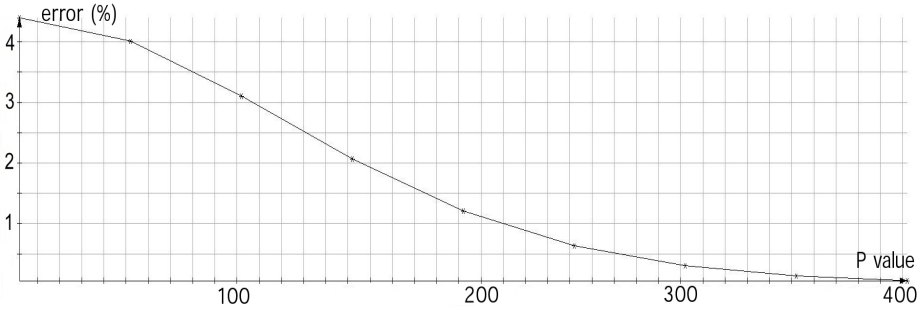

Figure 3.1: Amplitude estimation of a single mode signal of parameters: damping=0.01 du and amplitude=0.01 au.

to 150 improves the estimation of 1%, while a variation from 300 to 350, only results in a 0.1% improvement.

Different values of amplitude or frequency of the single mode does not affect the quality of its estimation. The damping factor on the other hand has an influence on the quality of the amplitude estimation. A bigger damping value will improve it, while a smaller value makes the estimation poorer. For instance, with P set to 10, for a damping of 0.01 du, the error on amplitude estimation is 4.4%, while with a damping of 0.005 du, it is 13%. Hence, in this basic scenario with one peak, with higher damping values, P can be set to a smaller value in order to get equivalent accuracy on amplitude estimation.

3.1.3 Total mode estimation

In *total mode estimation*, all the peaks close enough to influence the peak(s) to be estimated, are equally modeled. The tests which follow are made on a first signal containing two modes and a second one containing three modes. To start with, all modes have a damping of 0.01 du and an amplitude of 0.01 au.

The parameters mentioned in 3.1.1 are varied and their repercussions both on the estimation of the same peak, which value has been varied, and on neighboring peaks are analyzed.

From a peak separation under 20 data points, estimation of amplitude starts to degrade noticeably. 10 data points was considered the limit for estimation. In this configuration and with the three-mode signals tested, the estimation error reaches 10%, even with optimal settings of all parameters. If noise is added to this signal, the estimation will be further affected. Hence a 10 data-point separation between peaks was considered the limit for reasonable amplitude estimation.

The most difficult parameter to estimate is the amplitude. In total-mode estimation, the errors on damping and frequency estimations are close to null and will hence not be considered here.

In the previous section it was shown that with one mode estimation without noise, increasing parameter P resulted systematically in a better modeling of the amplitude parameter. With multiple mode estimation the effect is more

complex. The amplitude error graph shape can be put in two general categories: the *delayed improvement* and *progressive improvement* graphs. An example for each shape is shown in figure 3.2. Both have oscillations and are characterized by a maximum error point.

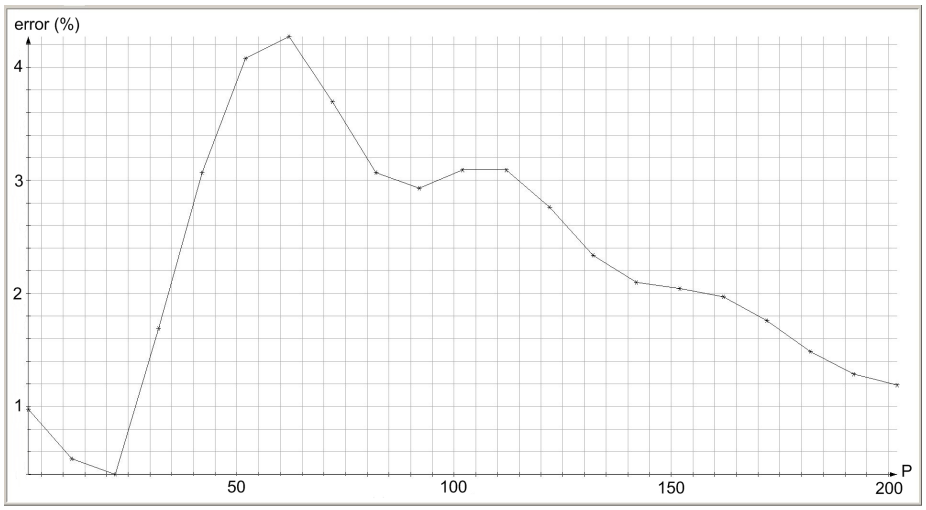

(a)

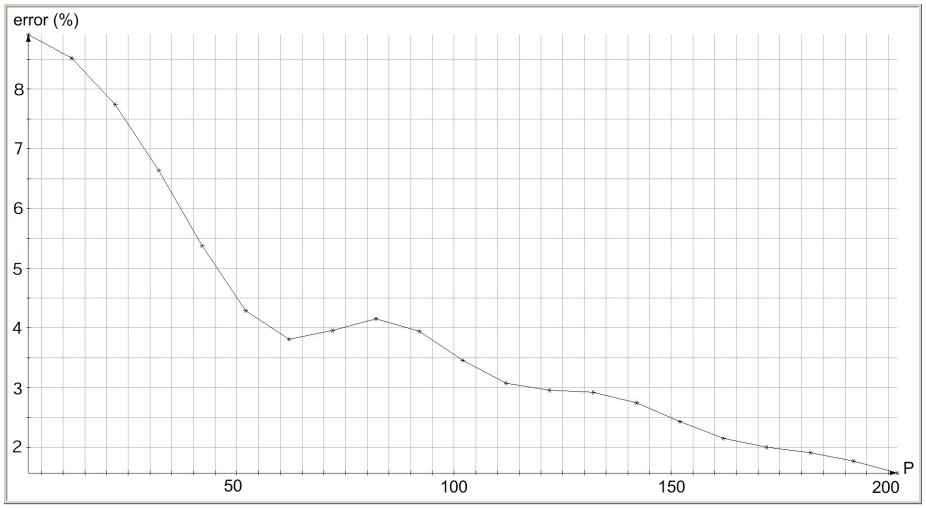

(b)

Figure 3.2: General amplitude error graph shapes. P is varied up to a value of 200. These are the modes of a two-mode signal with 20 data points peak separation and standard characteristics. They are named:
(a) the delayed improvement graph shape,
(b) the progressive improvement graph shape.

The delayed improvement type of graph has a maximum error in a P difficult to estimate. In the present graph its value is 62. The second graph presents maximum error in the minimum P value, 2. After the global maximum point, the error decreases progressively and finally stabilizes to a minimum value. Hence, it is important to avoid the maximum error area. The second constraint could be to set P if possible such as to satisfying a maximum error value.

The progressive improvement graph is similar to the shape seen with single mode estimation in the previous section, with added oscillations. It is much easier in this case to choose a P value, as the error will overall tend to decrease with P. Unfortunately, the scenario type is a priori unknown.

Larger amplitude values increase the oscillations in the amplitude error graph and the highest peak value. Hence, it takes a bigger P value to bring down the estimation error. Overall, it results in decreased amplitude precision.

High values of damping broadens the peak in the spectrum. In the amplitude error graph, it results in reduced oscillations. In the case of a delayed improvement graph type, it also shifts the global maximum point to an earlier P value. Hence, the graph gradually resembles the progressive improvement graph. Hence, the estimation is stabler even with shorter P values. The drawback is that the final error level, even at high P values, is worse than with lower damping values.

The effect of different data characteristics on neighboring peaks has also been analyzed. While the amplitude does not have a clear impact on the neighboring peak estimations, the damping does. A large damping systematically reduces the amplitude estimation quality of its adjacent peaks.

Reduced peak separation increases oscillations in the amplitude error graph and reduces the quality of estimation even with large P values. For instance, with three-peak estimation and standard parameter values, except for the first peak, which has an amplitude of 0.05 au, the error on amplitude estimation of the first mode is tripled.

3.1.4 Partial mode estimation

Instead of modeling all the peaks in a signal, as done in the last section, partial mode selection consists in modeling only a part of the formants, although neighboring peaks may be close by. To do this, the same two- and three-mode signals as in the previous section were used, but only one of the modes is modeled by SELF-SVD. In the three-mode signal, the most difficult peak to estimate is chosen, i.e. the one in the middle, as it is affected by two neighbors.

The choice of the frequency band on which to operate is of particular importance in partial-mode selection. Tests to find ideal band width and band positioning have been carried out. For these tests, P has been set to 200, in order to insure a relatively good and stable estimation of the amplitude.

A frequency band centered in the frequency to estimate proved to be the safest choice, always resulting in the best results.

A band width test is shown in figure 3.3. The test is made with a three-mode signal. The peak-spacing is 10 data points. All the modes are set to the standard

damping and amplitude settings. The band is centered in the estimated mode. The errors on the estimation of the three parameters of the central peak are given in function of the band width.

It results that the estimation of all the parameters follow similar patterns. The tests have shown that too short band widths never bring good estimation results. Even with peak-spacing of 10 data points, as in this example, the band width should not be taken below 8 data points. It is clear by the graphs that above a certain width, estimation errors degrade. The threshold value in this case is 14 data points for all three parameters. This upper limit is due to the influence of neighboring peaks. With more spacing between the peaks, the threshold is higher.

Once the frequency band is selected, SELF-SVD is applied. The results of selecting one peak out of the two or three in the signal have been compared to those in total mode estimation.

With a two-mode signal, and even with the smallest peak-spacing of 10 data points, the estimation results are very similar to those with total estimation. None of the two result as systematically better or worse.

With a three-mode signal, the estimation of the central peak is poorer in respect to total estimation when peak spacing is reduced to 15 data points and below. For example, peak separation is set to 10 data points and modes have their standard parameter values. The amplitude estimation at P=200 in single-mode estimation of the central peak amplitude has an error of 6%, while it is of half the value with total mode estimation.

If the mode on the left, the right, or at a greater extent both modes neighboring the estimated peak, are set to a big amplitude value, 0.05 au for instance, the amplitude estimation is degraded to a stronger degree in partial estimation than in total estimation. On the other hand, the setting of specific damping values of an adjacent peak to that estimated does not show particular bias to the election of one strategy over the other.

Damping and frequency estimations can also become an issue in the case of partial mode estimation. Frequency estimation with the different parameters tested presents an error from a peak separation of 10 data points. For instance, it is the case with three modes in the data, peak separation of 10 and all peaks with standard parameter values, except the first, which has an amplitude of 0.05 au. In this example the single estimation of the frequency of the central peak gives an error of 1 data point. On the ppm scale, it corresponds to around 0.01 ppm, which appears quite reduced for such an extreme case.

Damping estimation degrades as peaks are closer in the signal. With a standard three-mode signal for instance, the damping error on the estimation of the middle peak in single mode estimation goes from 0.18% with peak separation of 20 data points, to 4.4% with peak separation of 10 data points.

Hence, while partial peak estimation is robust with peak separation of 20 data points and above, it can become an issue when below this threshold. In this case, large amplitude values of neighboring peaks are the main factor for the degradation of amplitude estimation. On the other hand, partial estimation has the advantage of being less computationally expensive. Partial estimation

is translated by the limitation of the frequency band width of interest and the selection of less peaks for modeling. These two characteristics will make SELF-SVD estimation much faster.

3.1.5 Tests with additional Gaussian noise

Tests have been performed with additional *Gaussian noise* (noted ϵ in the FID equation (2.3) in section 2.3.1). 1-, 2- and 3-mode data has been used for this purpose. Parameters have been tested with the values given in 3.1.1.

Gaussian noise is statistical noise with amplitude values following a *Gaussian distribution*, also known as normal distribution. The mean value has been set to null in order to distribute the noise around a null noise value. Summarily, σ, or the square root of the variance, represents the variation of the noise around the mean value. σ is set to a percentage of the amplitude of the peak to be modeled. Its value is then varied proportionally to the quantity of noise wanted. The ratio between the value of σ and the amplitude of the peak to be modeled will be referred to simply as the noise value.

Additive noise affects the estimations of frequency, damping and amplitude, and this even from a noise value of 1%.

First, the effect on amplitude estimation will be detailed. The presence of Gaussian noise modifies the amplitude error graph in a singular manner. The noise reflects itself as an additional peak in the graph. This is shown in figure 3.4, with from top to bottom 1%, 2% and 5% noise. The shape of the peak seen in the graph is characteristic. It is very regular and presents a slow climb and a steep descend. After the steep descent, the amplitude gradually tends to a stable estimation value, as it did in the case seen in 3.1.2 without the additional noise. In order to get a stable estimation one must thus estimate the amplitude with a minimum P value taken above the maximum error point.

Increasing the noise shifts the position of the maximum error peak to a higher P value, makes it base broader and increases its maximum value. The error on the estimation is hence globally higher, even with high P values. The first graph, with 1% noise, presents a maximum error of 28%, a noise peak width of 20 P and the position of its maximum is in P=14. In the third graph, with 5% noise, the maximum error is above 50%, the noise peak width is approximately 100 P and the position of its maximum is in P=82. The scale is selected short in order to show with more detail the variations around the noise peaks, but at higher P values, the error in the 1% noise graph goes below that of the 5% one.

Tests have also been made with two- and three-mode data. For instance, figure 3.5 represents the amplitude error graph of the first of the two modes in a two-mode signal modeled with total estimation. The noise level is 2% and the maximum P value taken is 100. In this graph, the delayed improvement shape can be recognized (shape presented in figure 3.2a), with its characteristical oscillations, and in P=26, the peak due to the additional noise can be found.

Hence, by knowing the approximate noise level in real data, it is possible to localize its peak in the amplitude error graph and choose an appropriate P in order to optimize amplitude estimation.

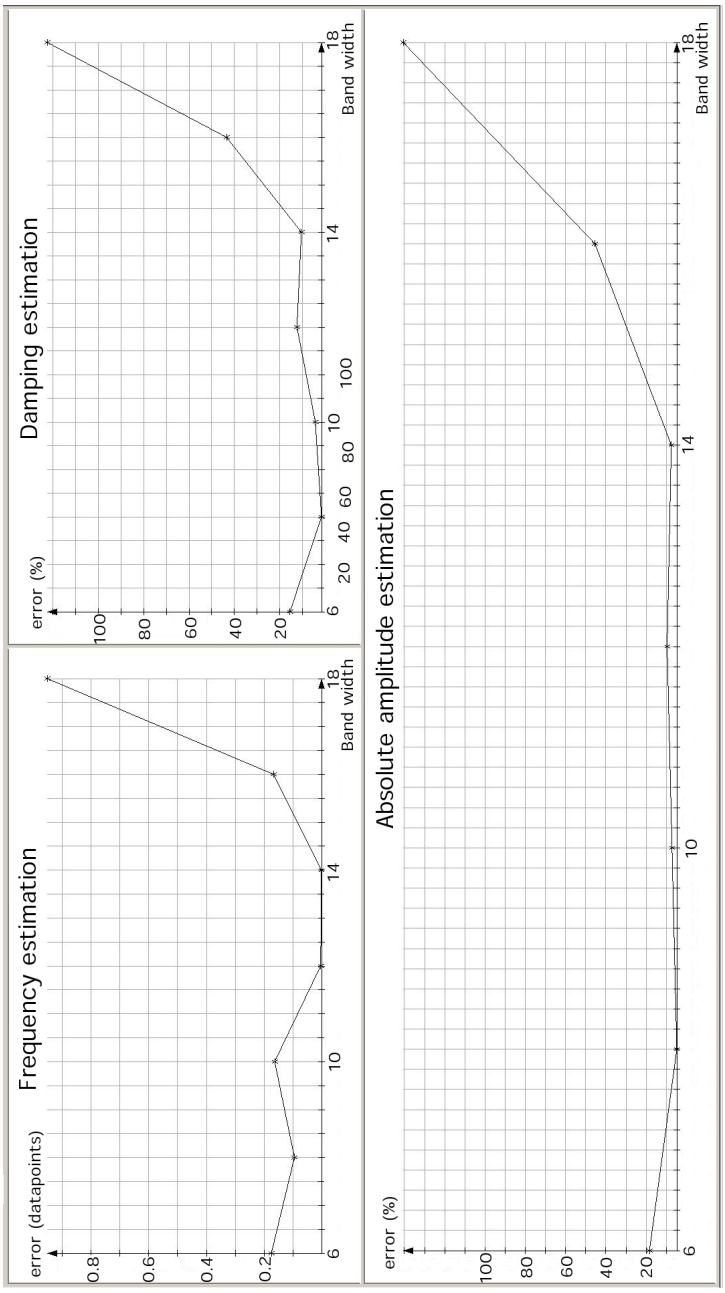

Figure 3.3: Estimation of the middle peak in a three-mode signal with peaks spaced of 10 data points and having identical standard characteristics for damping and amplitude.

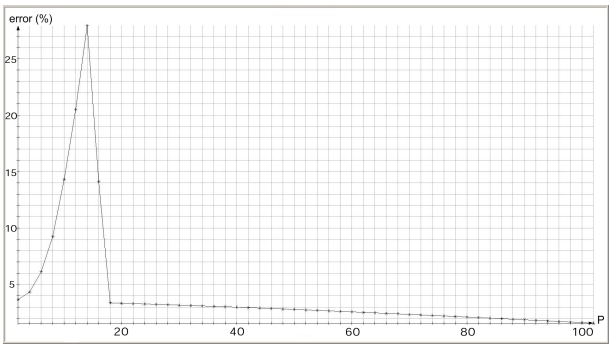

(a)

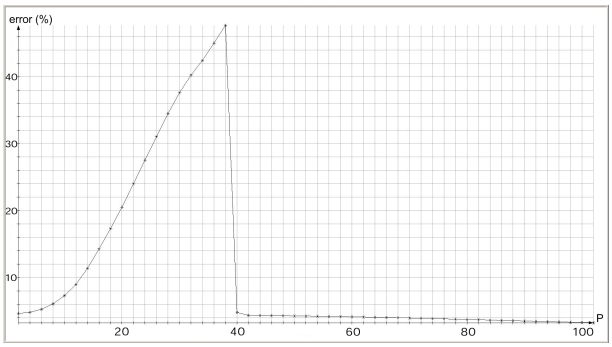

(b)

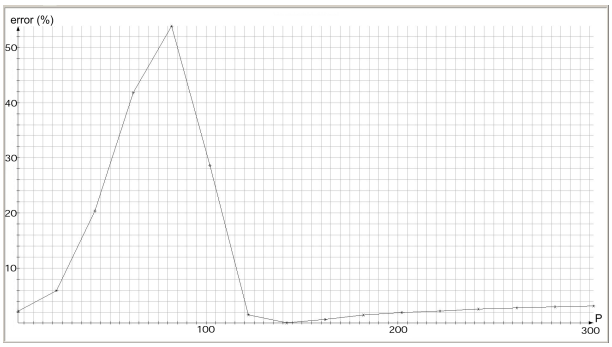

(c)

Figure 3.4: Estimation of a single mode with damping=0.01du and amplitude=0.01au:
(a) with Gaussian noise of 1%, scaled on P=100
(b) with Gaussian noise of 2%, scaled on P=100
(a) with Gaussian noise of 5%, scaled on P=300

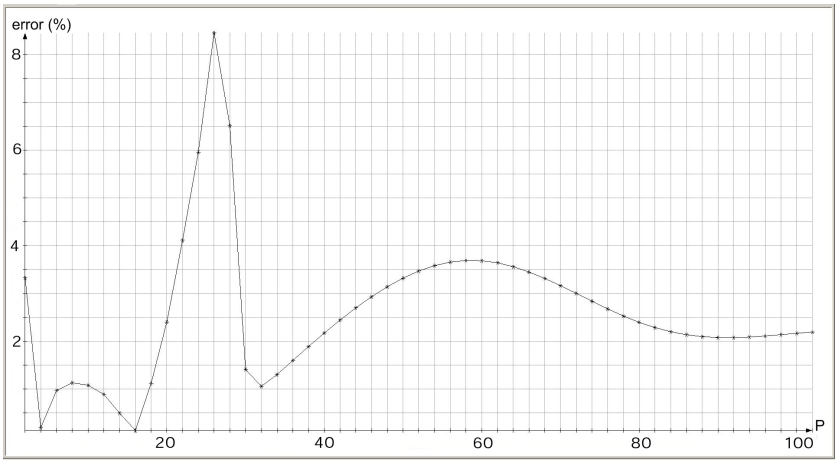

Figure 3.5: Amplitude error graph of the first of the two modes modeled in a two-mode signal with total estimation. The noise level is 2% and P is varied between 2 and 100.

The other two parameters estimations, namely frequency and damping, are also degraded as the noise level is increased. An example of the modeling of a single peak, with P set to 100, and the error graphs on its three parameters in function of the noise level is given in figure 3.6. Here, the noise level is varied from 1% to 10%.

In this example it can be seen that all the parameters have a tendency to have an increased modeling error as noise rises. The irregularities in the increase is due to the random nature of the noise. Due to this characteristic, one would have to take a mean value on a high number of runs to be precise on its effect on the signal.

In this example, at 10% noise level, the frequency error reaches 0.25 data points, the damping error 10% and the amplitude 57%. It is interesting to note the sudden increase in the amplitude estimation error at noise between noise level of 4% and 5%, where the estimation error goes from 11% to 54%. This is due to the setting of P to a value of 100. With such a setting, low noise values result in good estimations because the peak in the amplitude noise graph is avoided (see figures 3.4a and 3.4b). From a noise level of 5%, the peak of the error noise graph is situated around the value of P=100 (see 3.4c), which explains the drop in estimation quality. With a P set to a higher value, this change is situated for a higher noise level.

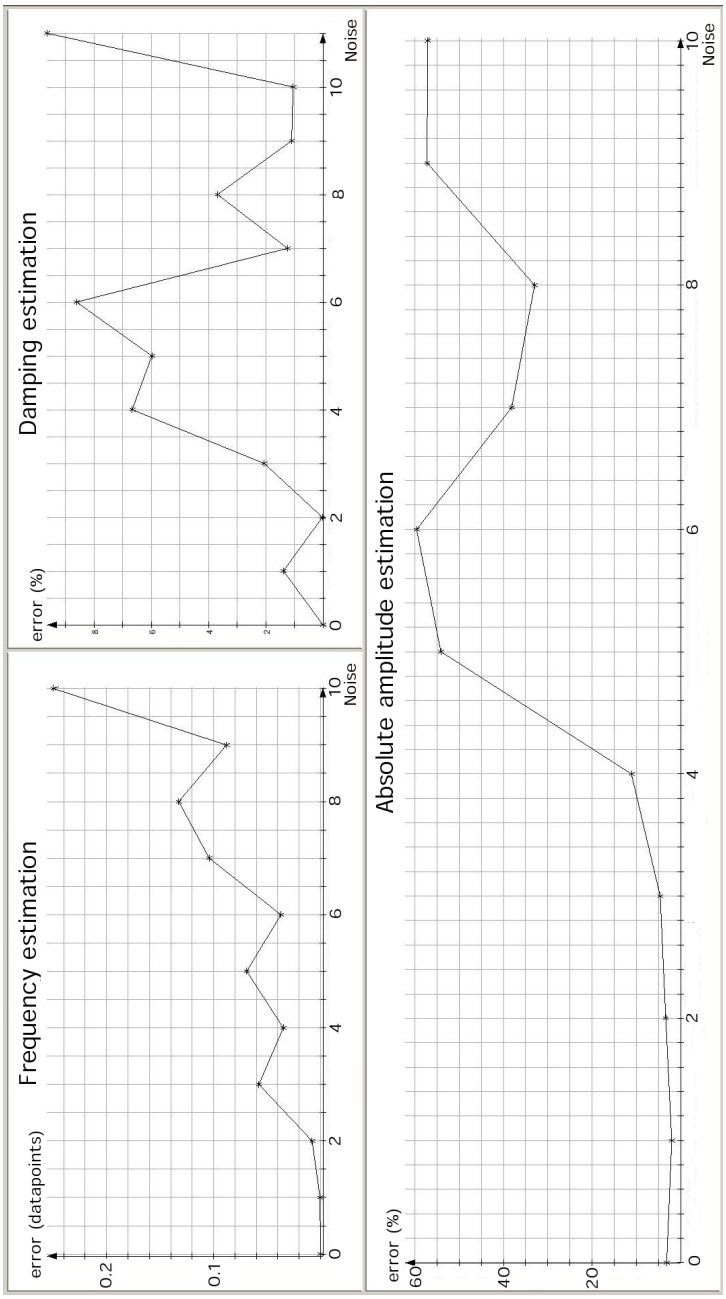

Figure 3.6: Estimation of a single mode with standard characteristics in increasing noise level.

3.1.6 Overestimation of the number of modes

This section deals with the use of SELF-SVD to model more modes than present in the signal. The tests were made with signals containing one to three modes. The number of estimated modes ranges from one to 5 extra. Tests were also made in the presence of additional Gaussian noise.

The tests without additional noise systematically result in the correct estimation of the peaks really present in the data, as well as extra estimations. The extra estimated modes have varying values for frequency, but most often these are out of the predefined frequency range selected. Furthermore, they are characterized by abnormally high amplitude and damping values. These appear to be random and always above 1000 au for the amplitude and 1000 du for the damping. If modeled in the frequency domain, they are represented by heavy baseline components.

These false estimations can be easily detected due to their very large values of damping and amplitude and thus also be discarded. The remaining modes are estimated identically as in the case without overestimation and thus the signal can be reconstructed. It appears hence that the additional modes requested are hence not actually taken into account by the SELF-SVD for the signal estimation.

When adding Gaussian noise, overestimation gives completely different results. All the demanded modes now contribute to the estimation of the signal as a sum of damped sinusoids. For instance, let us take an initial signal composed of m damped sinusoids and where k, with $k > m$, modes are requested for modeling. The adding of noise alters the signal and hence it can no more be perfectly estimated as m damped sinusoids. Thus, overestimation leads to an attempt from the algorithm to estimate the noisy signal as the k modes demanded, and not to a degenerated case as previously.

Overestimation with Gaussian noise leads to the modeling of modes that are not in the original underlying signal. The original modes are thus most of the time not found any more amongst the modes that have been estimated. Real data always presents noise and hence overestimation should be avoided.

3.1.7 Example

Here is an example to illustrate the estimation of a simulated signal and its estimation by SELF-SVD. The following characteristics are used to build a signal with 7 modes:

mode	frequency	damping	amplitude
1	145	0.01	0.025
2	155	0.01	0.025
3	170	0.03	0.01
4	190	0.02	0.01
5	220	0.01	0.015
6	240	0.04	0.025
7	250	0.015	0.015

Noise is set to 4%. The mode to be estimated is mode 3, at a frequency of 170 data points. We will use partial estimation in order to model only this peak.

4 peaks are in the range of 50 data points from our mode of interest and hence will influence the amplitude estimation. The two modes that should mainly have influence on it are the one in 155 and 190 data points, respectively 15 and 20 data points away. Frequency and damping estimations will present an uncertainty due to the presence of noise. The damping will further be perturbed by the presence of the 2 peaks at less than 20 data points distance.

The frequency band should be chosen relatively short to avoid too much influence from neighboring modes. The interval taken is $[160, 180]$ data points.

Since the mode to estimate has a broad damping value, its amplitude error graph will stabilize quite rapidly. The 4% noise will create a peak in this graph, which should have decayed by P=100. With these characteristics in mind, a value of 100 for P should be a relatively good value, avoiding unstable errors on amplitude estimation and keeping P as short as possible to calculate as fast as possible.

The signal and its SELF-SVD estimation can be seen in figure 3.7. The error on the frequency estimation is 0.4 data points, the damping 11.9%, and the amplitude 16.6%.

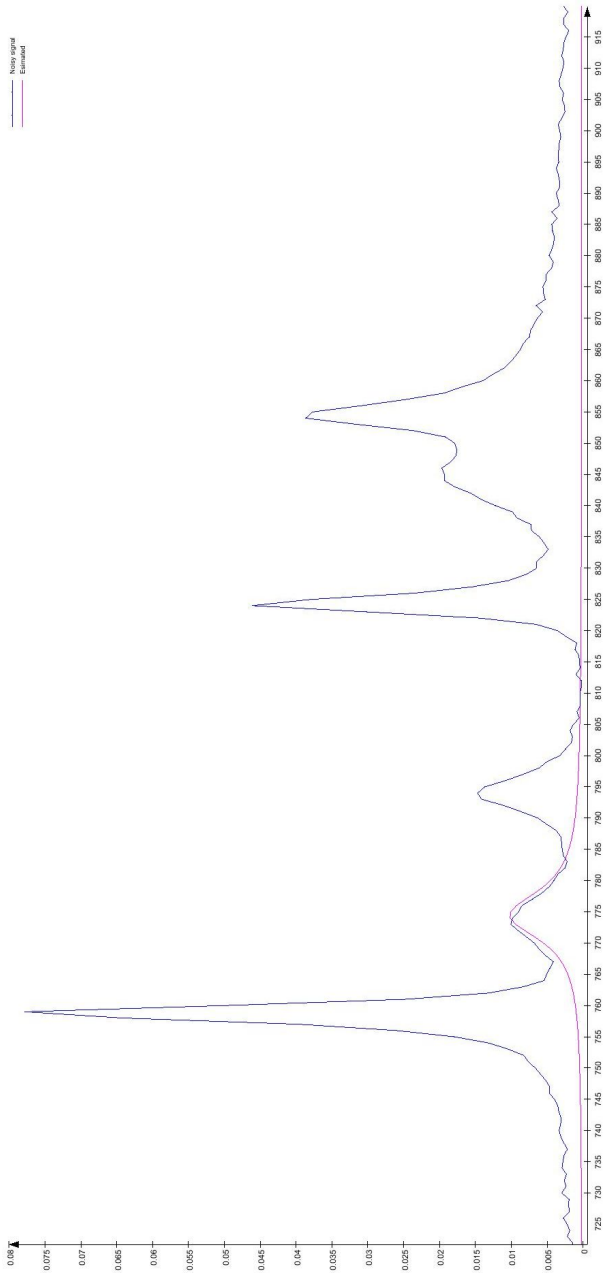

Figure 3.7: Example signal with noise level of 4% and SELF-SVD estimation of the peak positioned at data point 170.

3.2 Tests on in-vivo data

This section presents the tests made on in-vivo data. The data has been taken from a CSI 2D-scan of the brain, acquired with a Siemens 3T scanner[1]. The data concerns a patient diagnosed with a brain tumor.

After describing the testing protocol used, the results concerning quantification of NAA and choline are put forth. Finally, a simple method which automates the choice of the number of peaks for modeling is given and tested on the in-vivo data.

3.2.1 Protocol

In this concrete application, the capacity of SELF-SVD to locate the tumor region, by estimating the concentrations of choline and NAA, is evaluated. The algorithm results will be compared to those of a vendor software from Siemens, called *Syngo MR Spectroscopy Evaluation*. For more information on the Siemens application, please refer to their homepage: *http://www.medical.siemens.com*.

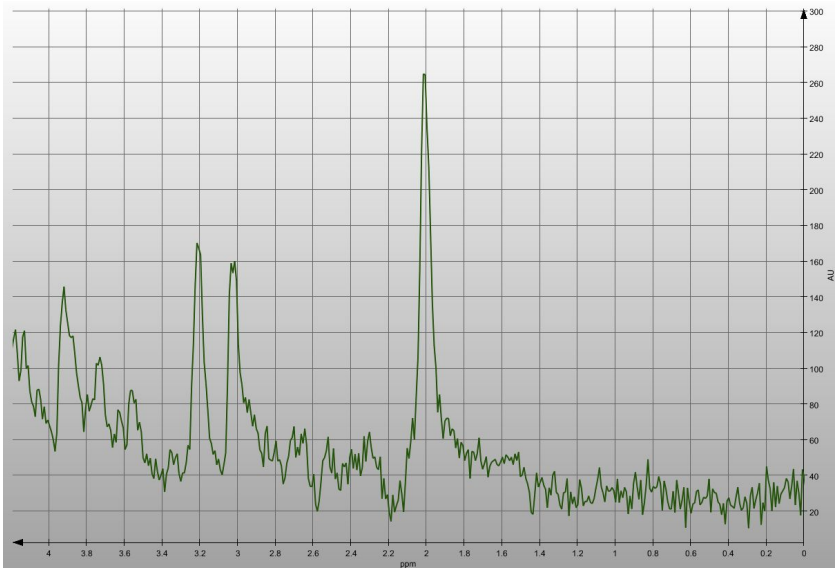

Figure 3.8: Spectrum of a voxel in the middle zone from the brain data.

The 16×16 voxels where acquired in a *point-resolved spectroscopy (PRESS)* sequence with an *echo time (TE)* of $TE = 135ms$. 1024 data points were recorded with a bandwidth of 1.2 MHz. The region of interest is composed of a region containing 8×8 of these voxels.

[1]Courtesy of MR- und PET/CT-Zentrum Bremen Mitte, Germany.

The quantification results given by the Syngo workstation follow the standard clinical protocol used, which encompasses water removal, baseline correction, zero filling, apodisation, Fourier transformation for visualization and relative time-domain based quantification.

The spectrum of a voxel in the middle area of the grid is shown in figure 3.8. This spectrum illustrates the presence of noise and a baseline in the data. SELF-SVD has been performed without any pre-processing, in order to avoid other processes to influence results. On the other hand, this will make the results more sensitive to noise than the vendor software.

The three highest peaks in the spectra from left to right, are choline (3.18 ppm) creatine (3.03 ppm) and NAA (2.01 ppm), the biggest of the three. The focus will be put on the choline and NAA levels in order to localize the tumor. In both cases, several settings for SELF-SVD have been tested in order to find the most appropriate ones.

SELF-SVD hence results in the modeling of NAA and choline peaks. The concentrations of the metabolites were then derived by the simple measuring of the peak heights.

3.2.2 NAA quantification

The NAA metabolic map of the brain data can be seen in figure 3.9. In order to get a spatial view of the MRS results, they are superposed to an MRI image of the scanned region. The color scale shown on the left represents the concentration levels of NAA. These range from low (blue) to high (red) values compared to the rest of the data.

Low NAA values are characteristic of tumor regions. The tumor can thus be localized on the left of the image, where the non-colored regions indicate a particularly poor concentration of NAA.

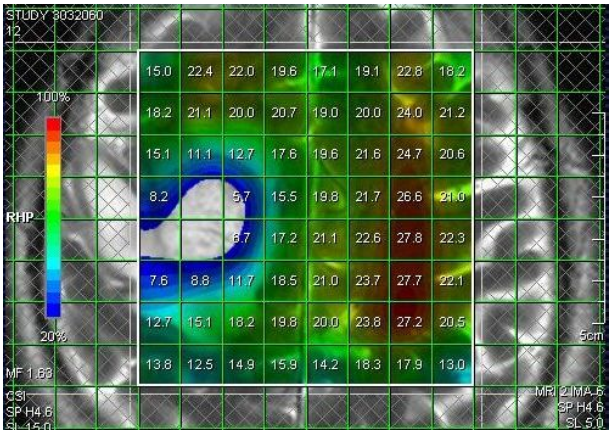

Figure 3.9: NAA metabolic map.

SELF-SVD was set in order to estimate one peak in the $[1.9, 2.1]$ ppm range. According to the simulated results, the presence of moderate noise in the data should lead to the presence of a peak in the *amplitude error graph*, which should decrease and stabilize at a certain P level (refer to section 3.1.5 for more detail). P is thus fixed to the moderate value of 100.

In order to compare results between SELF-SVD and Syngo, a linear regression was made between the two datasets. This can be seen in figure 3.10. The voxels (1,3), (0,4) and (1,4) were removed to improve results. These would cause the linear regression to miss too strongly the origin of the graph for reasons explained later on. SELF-SVD estimations were rescaled to match those given by the Siemens software by using the slope of the linear regression equation (the equation is shown in the figure). The correlation coefficient (R^2) is 0.46, indicating a moderate correlation between the datasets.

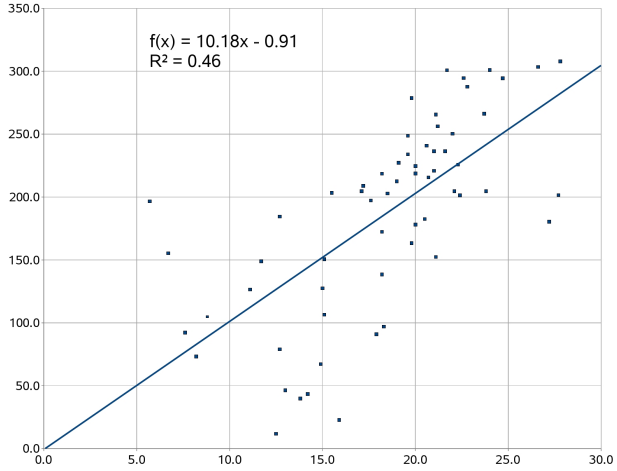

Figure 3.10: Linear regression between the SELF-SVD and Syngo results.

The results with the Siemens software and those with SELF-SVD are given in figure 3.11. The main difference between the datasets is found in the last row. In the tumor area, SELF-SVD presents particularly low concentrations of NAA, while the vendor software displays null concentration values for certain voxels in the area. These concern the voxels removed to improve linear regression analysis. The null values in the Syngo application are most probably linked to the setting of a threshold concentration value, under which results are set to null.

The main difference between the results resides in the last row. Here, the Syngo estimations are up to 2 times greater than those of SELF-SVD. By looking at the spectra, one notes that the last row voxels present a significant decrease in signal-to-noise ratio. This is shown in the spectrum in figure 3.12, where the difference between the spectrum taken in a voxel in the last row is compared to its neighboring voxel in the one-but-last.

y/x	0	1	2	3	4	5	6	7
0	15.0	22.4	22.0	19.6	17.1	19.1	22.8	18.2
1	18.2	21.1	20.0	20.7	19.0	20.0	24.0	21.2
2	15.1	11.1	12.7	17.6	19.6	21.6	24.7	20.6
3	8.2	0.0	5.7	15.5	19.8	21.7	26.6	21.0
4	0.0	0.0	6.7	17.2	21.1	22.6	27.8	22.3
5	7.6	8.8	11.7	18.5	21.0	23.7	27.7	22.1
6	12.7	15.1	18.2	19.8	20.0	23.8	27.2	20.5
7	13.8	12.5	14.9	15.9	14.2	18.3	17.9	13.0

(a)

y/x	0	1	2	3	4	5	6	7
0	12.5	19.8	24.6	24.4	20.1	22.3	28.3	21.5
1	13.6	15.0	17.5	21.2	20.9	22.1	29.6	25.2
2	10.5	12.4	18.1	19.4	23.0	23.2	28.9	23.7
3	7.2	9.6	19.3	20.0	27.4	29.5	29.8	21.7
4	6.3	5.8	15.3	20.5	26.1	28.9	30.2	22.2
5	9.1	10.3	14.6	19.9	23.2	26.1	19.8	20.1
6	7.8	14.8	16.9	16.1	21.5	20.1	17.7	17.9
7	3.9	1.2	6.6	2.2	4.3	9.5	8.9	4.6

(b)

Figure 3.11:
(a) Concentration values with the Siemens Syngo software.
(b) Concentration values with SELF-SVD and the settings: 1 mode, frequency range = $[1.9, 2.1]$ ppm, P=100.
The blue voxels are in the tumor region and the faded blue in its immediate surrounding.

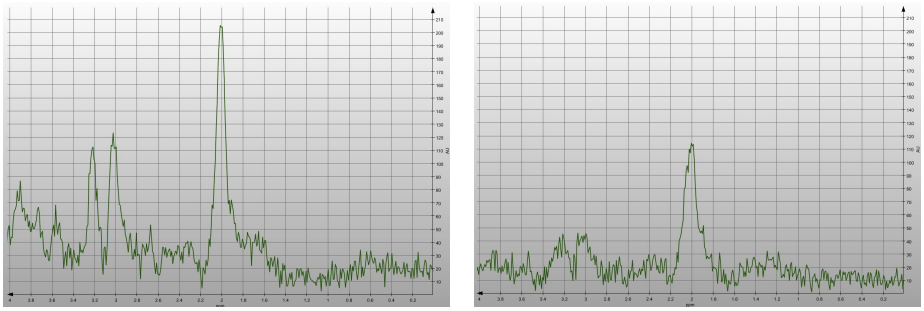

(a) (b)

Figure 3.12: Spectrum of a voxel taken in :
(a) voxel (5,7) in the one-but-last row,
(b) voxel (5,8) in the last row.

Big P values do not lead to better results, suggesting that P was well chosen. Estimating more than one peak in the frequency range generally degraded quantification results.

On the other hand, broadening the frequency range to $[1.9, 2.2]$ ppm leads to better estimations for the last row. Yet, even with these settings, the Syngo estimation remain 30% to 50% higher than those of SELF-SVD. By looking directly at the spectra in the last row, one notes that the NAA peak (the tallest

in the spectrum) is typically halved compared to the on-but-last row. This is not reflected in the Syngo data, suggesting that the concentrations may have been overestimated.

The broader frequency range of $[1.9, 2.2]$ ppm was not retained because it leads to less good quantification estimations for the other rows.

3.2.3 Choline quantification

Next, choline concentrations are evaluated. The corresponding metabolic map can be seen in figure 3.13. The same coloring scheme is used as that for NAA, detailed in the previous above. Choline is involved in the development of new cells and hence a high level can point to tumor regions. The tumor is thus localized on the left of the image in the red-colored zone.

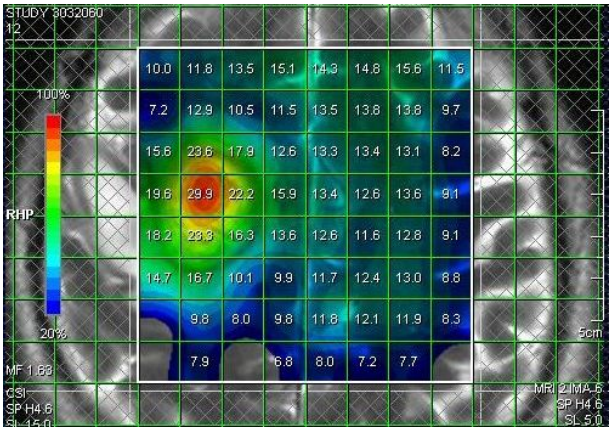

Figure 3.13: Choline metabolic map.

The SELF-SVD settings are similar to those set for NAA quantification, as the MRS dataset is identical. P is fixed to 100 and a single peak is estimated in the range of $[3.1, 3.3]$ ppm.

Using the same technique as with NAA, the SELF-SVD results are rescaled by using the slope of the linear regression equation shown in figure 3.10. The correlation coefficient (R^2) between the data is much better than with NAA.

The results with the Siemens software and the rescaled SELF-SVD ones are given in figure 3.15. The tumor zone on the left can be localized by SELF-SVD through the unusually high choline concentration values.

Here again, the main difference between the results resides in the last row, where SELF-SVD estimations are close to null. Nevertheless, even the Syngo estimations are particularly low in this area. By looking at the spectra shown in 3.12 one notes that the SNR level is so low for choline that it becomes almost difficult to distinguish it from the noise. In this case, Syngo seems to approximate the signal better than SELF-SVD.

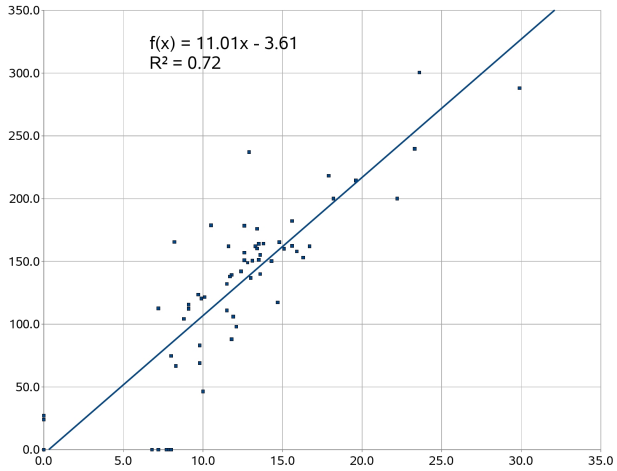

Figure 3.14: Linear regression between the SELF-SVD and Syngo results.

y/x	0	1	2	3	4	5	6	7
0	10.0	11.8	13.5	15.1	14.3	14.8	15.6	11.5
1	7.2	12.9	10.5	11.5	13.5	13.8	13.8	9.7
2	15.6	23.6	17.9	12.6	13.3	13.4	13.1	8.2
3	19.6	29.9	22.2	15.9	13.4	12.6	13.6	9.1
4	18.2	23.3	16.3	13.6	12.6	11.6	12.8	9.1
5	14.7	16.7	10.1	9.9	11.7	12.4	13.0	8.8
6	0.0	9.8	8.0	9.8	11.8	12.1	11.9	8.3
7	0.0	7.9	0.0	6.8	8.0	7.2	7.7	0.0

(a)

y/x	0	1	2	3	4	5	6	7
0	4.2	12.7	14.9	14.5	13.6	15.0	14.8	10.1
1	10.2	21.5	16.2	12.0	13.7	14.9	14.9	11.2
2	16.5	27.3	19.8	13.7	14.7	14.6	13.7	15.0
3	19.5	26.2	18.2	14.4	16.0	16.2	14.1	10.5
4	18.2	21.8	13.9	12.7	14.3	14.7	13.5	10.2
5	10.7	14.7	11.0	10.9	12.5	12.9	12.4	9.5
6	2.2	6.3	6.8	7.6	8.0	8.9	9.6	6.1
7	0.0	0.0	2.5	0.0	0.0	0.0	0.0	0.0

(b)

Figure 3.15:
(a) Concentration values with the Siemens Syngo software.
(b) Concentration values with SELF-SVD and the settings: 1 mode, frequency range = $[3.1, 3.3]$ ppm, P=100.
The blue voxels are in the tumor region and the faded blue in its immediate surrounding.

When setting too short frequency bands, SELF-SVD modeling can sometimes result in the evaluation of a peak at the wrong frequency. This could be due to the presence of noise, but also to a slight shifting of the peaks in the data.

The tests made with SELF-SVD in the studied in-vivo case has proven to give results leading to the localization of the tumor region both with choline and NAA concentrations. The results are also comparable to those in a vendor software, indicating that they are relatively reliable.

Time-wise, SELF-SVD appears to be an interesting solution. The modeling of choline or NAA in the conditions described in this section takes less than 1 second per metabolite and per voxel. The total 8×8 grid is hence estimated for the two metabolites of interest in less than 2 minutes.

For comparison, Syngo estimates the two metabolite concentrations in around 5 seconds per voxel. Hence, in this particular application, quantification of the whole grid would take around 5 minutes, or 2.5 times more than SELF-SVD.

3.2.4 Automated selection of the number of modes to model

The automated choice of the number of modes to estimate has been implemented on the in-vivo dataset. It has been shown in section 2.3.3 that this is possible through the calculation of the *effective rank* of a certain matrix (Z) representing the noisy data. This calculation is complex and has been replaced by a rough estimation performed with a method described below.

The matrix on which to calculate the effective rank is decomposed by SVD. The singular values of its singular value matrix represent the energy levels of the spectral peaks corresponding to the different modes in the signal (see section 2.3.1 for more details). By fixing an appropriate threshold value, the singular values representing the signal modes were distinguished from those corresponding to noise. The thresholding method is conceivable in this dataset because the modes representing the noise correspond to lower energy levels than those of the signal.

On the brain data tested, it was found that by discarding the singular values of energy level below 1 resulted in good estimation results. In this configuration the peaks of interest (choline, creatine and NAA) were kept and most of the noise was not modeled. This configuration leads to the choice of an optimal number of modes in the frequency bandwidths used in the in-vivo tests presented above. Narrow frequency bandwidths generally gave better results. The method was successful also with bandwidths encompassing two modes. For instance, in the $[2.9, 3.3]$ ppm range, the method made the best choice, which was to estimate two peaks, i.e. choline and creatine.

Thresholding is too rigid to work in all datasets and does not take into account the signal-to-noise ratio, but these tests put forth the potential of automated selection of the number of modes to model.

Chapter 4

Conclusion and outlook

Within the scope of this thesis, the MRS quantification algorithm of SELF-SVD has been successfully implemented and tested both on simulated and in-vivo data.

The implementation was made in the MeVisLab environment, where a module was written in C++ and made compatible with the MeVisLab MRS assistant pipeline. The next stage would be to integrate the module in the MRS assistant and create a graphical user interface for it.

The algorithm has been tested through a seperate testing module. This module enables multiple automated tests, including for instance the influence of the P parameter and the noise level in the signal. Many observations were derived from these tests, in particular the complex influences between modes in the data. Based on the test results, suggestions have been made on how best to set the SELF-SVD parameters depending on the nature of the data.

Further tests could be made with the same test module. The simulated data could be built to match several typical in-vivo datasets and hence possibly lead to new hints on how to optimize the SELF-SVD parameters. Also, with a complete simulated spectrum reflecting real in-vivo data, it would be possible with this testing module to assess the quantification error made by SELF-SVD in in-vivo cases.

Tests on in-vivo data have been made here for the first time with the SELF-SVD algorithm. They have shown promising results on the brain data tested. The inclusion of pre-processing stages should further improve results by reducing the noise level in the signal.

Processing time is an important issue when dealing with 2D or 3D CSI data. SELF-SVD becomes particularly interesting in these situations because of its frequency-selective characteristic. The APES method uses an FIR filter for amplitude estimation. It is worth noting that the FIR filter length, represented by parameter P, can notoriously influence calculation time. The ideal dimension of the filter depends on the precision required for the modeling of the amplitude of the modes. The simulation tests have also shown that the amplitude estimation depends on the proximity of peaks in the spectrum, their parameters, but also

and mainly of the noise level.

The in-vivo data tests have shown good results with a P of 100, which represented a calculation time of less than 2 minutes for the estimation of choline and NAA peaks for the entire 8×8 grid.

Knowing the quantity of noise in the data a priori, or an estimation of it, would lead to a better choice of P. The process of estimating it could nevertheless be too time consuming. Another possibility would be to estimate noise levels and P values for each voxel separately. It appears nevertheless to have several potential drawbacks, such as the additional time needed for noise estimation and the arising question of result consistency between the different voxels.

User interaction could be diminished in several ways. First of all, the frequency bands on which to estimate the peaks of interest could be fixed, since these are known a priori. The problems linked to peak-shifts in the data will have to be dealt with precaution.

The second and major step towards reduced user interaction would be for SELF-SVD to automatically give the number of peaks to be modeled in the frequency band of interest. It has been shown that this is possible through the calculation of the effective rank of a matrix which represents the noise-affected data. A simple thresholding system has been implemented for this purpose. It successfully estimates the number of modes in typical frequency ranges used for the modeling of NAA, choline and creatine, but is only applicable to the given in-vivo dataset. The next stage would be to implement a robust calculation of the effective rank. This would avoid the problems linked to underestimation or overestimation of the number of peaks.

The tests made on simulated data have shown that one of the most sensitive parameters to estimate is the amplitude, which has direct repercussions on quantification results. Testing has further shown that amplitude estimation not only depends on P, but on the separation between peaks in the spectrum and the damping of the peak estimated.

It is conceivable to create a function that approximates the amplitude error and takes as parameters damping and peak separation. In consequence P could be set to an optimal value to balance speed and precision requirements. Moreover, this approach should not greatly increase computational time due to the fact that the APES algorithm used for amplitude estimation is launched only after frequency and damping have been estimated. Hence, P could be regulated on the fly in the algorithm. Furthermore, an estimation of the error made on the amplitude measurement could avoid misinterpretations linked to the reading of mediocre estimation results.

Bibliography

[1] Eurostat. Key figures on europe. 2007.

[2] A. Yu, K. L., and H. Li. Value of diagnosis and differential diagnosis of MRI and MR spectroscopy in gliomatosis cerebri. *European Journal of Radiology*, 59(2):216 – 221, 2006.

[3] A. Cortey, J. G. Jarvik, R. E. Lenkinski, R. I. Grossman, I. Frank, and M. Delivoria-Papadopoulos. Proton MR spectroscopy of brain abnormalities in neonates born to HIV- positive mothers. *AJNR Am J Neuroradiol*, 15(10):1853–1859, 1994.

[4] P. Stoica, N. Sandgren, Y. Selén, L. Vanhamme, and S. Van Huffel. Frequency-domain method based on the singular value decomposition for frequency-selective NMR spectroscopy. *Journal of Magnetic Resonance*, 165(1):80 – 88, 2003.

[5] B. D. Ross and T. Michaelis. MR spectroscopy of the brain: neurospectroscopy. *Clinical Magnetic Resonance Imaging*, pages 928–981, 1996.

[6] Y. Rosen and R. E. Lenkinski. Recent Advances in Magnetic Resonance Neurospectroscopy. *Neurotherapeutics*, 4(3):330 – 345, 2007.

[7] J. Keeler. Understanding NMR Spectroscopy. 2005.

[8] J. F. A. Jansen, W. H. Backes, K. Nicolay, and M. E. Kooi. 1H MR Spectroscopy of the Brain: Absolute Quantification of Metabolites. *Radiology*, 240(2):318–332, 2006.

[9] J. Kurhanewicz and D. B. Vigneron. Advances in mr spectroscopy of the prostate. *Magnetic Resonance Imaging Clinics of North America*, 16(4):697–710, 2008.

[10] D. K. W. Yeung, G. M. K. Tse, and W.-T. Yang. Breast Cancer: In Vivo Proton MR Spectroscopy in the Characterization of Histopathologic Subtypes and Preliminary Observations in Axillary Node Metastases. *Radiology*, 225(1):190–197, 2002.

[11] T.-J. Hsieh, C.-K. Wang, H.-Y. Chuang, Y.-J. Jong, C.-W. Li, and G.-C. Liu. In Vivo Proton Magnetic Resonance Spectroscopy Assessment for Muscle Metabolism in Neuromuscular Diseases. *The Journal of Pediatrics*, 151(3):319 – 321, 2007.

[12] Z. Tong, K. Harada, and K. Houkin. In vivo quantification of the metabolites in normal brain and brain tumors by proton MR spectroscopy using water as an internal standard. *Magnetic resonance imaging*, 22(5):735–742, June 2004.

[13] M. Castillo, L. Kwock, and S. K. Mukherji. Clinical applications of proton MR spectroscopy. *AJNR Am J Neuroradiol*, 17(1):1–15, 1996.

[14] K. M. Cecil and B. V. Jones. Magnetic resonance spectroscopy of the pediatric brain. *Top Magnetic Resonance Imaging*, (12):435–452, 2001.

[15] R. Kreis, T. Ernst, and B. D. Ross. Development of the human brain: In vivo quantification of metabolite and water content with proton magnetic resonance spectroscopy. *Magnetic Resonance in Medicine*, 30(4):424–437, 1993.

[16] M. Sansone, M. Cesarelli, P. Bifulco, M. Bracale, G. Tedeschi, S. Bonavita, and F. di Salle. Use of different time-domain techniques for metabolite maps reconstruction in 1H-MRSI. volume 3, pages 2002–2003 vol.3, 2000.

[17] A. Connelly, G. D. Jackson, J. S. Duncan, M. D. King, and D. G. Gadian. Magnetic resonance spectroscopy in temporal lobe epilepsy. *Neurology*, 44(8):1411–, 1994.

[18] N. Lundbom, E. Gaily, K. Vuori, R. Paetau, E. Liukkonen, J. C. Rajapakse, L. Valanne, A.-M. Häkkinen, and M.-L. Granström. Proton Spectroscopic Imaging Shows Abnormalities in Glial and Neuronal Cell Pools in Frontal Lobe Epilepsy. *Epilepsia*, 42(12):1507–1514, 2001.

[19] G. J. Metzger, M. Patel, and X. Hu. Application of Genetic Algorithms to Spectral Quantification. *Journal of Magnetic Resonance, Series B*, 110(3):316 – 320, 1996.

[20] F. DiGennaro and D. Cowburn. Parametric estimation of time-domain NMR signals using simulated annealing. *Journal of Magnetic Resonance 96*, pages 582–588, 1992.

[21] A. J. W. C. van der Veen, R. de Beer, P. R. Luyten, and D. van Ormondt. Accurate quantification of in vivo ^{31}P NMR signals using the variable projection method and prior knowledge. *Magnetic Resonance in Medicine*, 6(1):92–98, 1988.

[22] L. Vanhamme, A. van den Boogaart, and S. Van Huffel. Improved Method for Accurate and Efficient Quantification of MRS Data with Use of Prior Knowledge. *Journal of Magnetic Resonance*, 129(1):35–43, 1997.

[23] L. Vanhamme, T. Sundin, P. van Hecke, S. van Huffel, and R. Pintelon. Frequency-selective quantification of biomedical magnetic resonance spectroscopy data. *J Magn Reson*, 143(1):1–16, March 2000.

[24] R. Romano, A. Motta, S. Camassa, C. Pagano, M. T. Santini, and P. Luigi Indovina. A New Time-Domain Frequency-Selective Quantification Algorithm. *Journal of Magnetic Resonance*, 155(2):226–235, 2002.

[25] R. Kumaresan and D. Tufts. Estimating the parameters of exponentially damped sinusoids and pole-zero modeling in noise. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 30(6):833–840, Dec 1982.

[26] P. Koehl. Linear prediction spectral analysis of NMR data. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 34(3-4):257–299, 1999.

[27] W. W. F. Pijnappel, A. van den Boogaart, R. de Beer, and D. van Ormondt. SVD-based quantification of magnetic resonance signals. *Journal of Magnetic Resonance*, 97(1):122–134, March 1992.

[28] S. Y. Kung, K. S. Arun, and D. V. B. Rao. State-Space and Singular-Value Decomposition-Based Approximation Methods for the Harmonic Retrieval Problem. 2002.

[29] S. Van Huffel, H. Chen, C. Decanniere, and P. Van Hecke. Algorithm for time-domain NMR data fitting based on total least squares. *Journal of Magnetic Resonance*, 110(1):228–237, 1994.

[30] Y. Hua and T. K. Sarkar. Matrix pencil method for estimating parameters of exponentiallydamped/undamped sinusoids in noise. *Acoustics, Speech and Signal Processing*, 38(5):814–824, 1990.

[31] B. D. Rao. Relationship between matrix pencil and state space based harmonicretrieval methods. *Acoustics, Speech and Signal Processing*, 38(1):177–179, 1990.

[32] H. Chen, S. van Huffel, D. van Ormondt, and R. de Beer. Parameter estimation with prior knowledge of known signal poles for the quantification of NMR spectroscopy data in the time domain. *J. Magn. Reson.*, 119:177–179, 1996.

[33] H. Chen, S. van Huffel, and J. Vandewalle. Improved methods for exponential parameter estimation in the presence of known poles and noise. *Signal Processing, IEEE Transactions on*, 45(5):1390–1393, May 1997.

[34] T. Laudadio, Y. Selén, L. Vanhamme, P. Stoica, P. van Hecke, and S. van Huffel. Subspace-based MRS data quantitation of multiplets using prior knowledge. *Journal of Magnetic Resonance*, 168(1):53–65, 2004.

[35] R. A. Meyer, M. J. Fisher, S. J. Nelson, and T. R. Brown. Evaluation of manual methods for integration of in vivo phosphorus NMR spectra. *NMR in Biomedicine*, 1(3):131–135, 1988.

[36] E. B. Cady. A reappraisal of the absolute concentrations of phosphorylated metabolites in the human neonatal cerebral cortex obtained by fitting lorentzian curves to the ^{31}P NMR spectrum. *Journal of magnetic resonance*, 91(3):637–643, 1991.

[37] S. W. Provencher. Automatic quantitation of localized in vivo ^{1}H spectra with LCModel. *NMR in Biomedicine*, 14(4):260–264, 2001.

[38] S. Mierisová and M. Ala-Korpela. MR spectroscopy quantitation: a review of frequency domain methods. *NMR in Biomedicine*, 14(4):247–259, 2001.

[39] I. Marshall, S. D. Bruce, J. Higinbotham, A. MacLullich, J. M. Wardlaw, K. J. Ferguson, and J. Seckl. Choice of spectroscopic lineshape model affects metabolite peak areas and area ratios. *Magnetic Resonance in Medicine*, 44(4):646–649, 2000.

[40] P. Soininen, J. Haarala, and J. Veps˙Strategies for organic impurity quantification by 1H NMR spectroscopy: Constrained total-line-shape fitting. *Analytica Chimica Acta*, 542(2):178–185, 2005.

[41] R. Laatikainen, M. Niemitz, W. J. Malaisse, M. Biesemans, and R. Willem. A computational strategy for the deconvolution of NMR spectra with multiplet structures and constraints: Analysis of overlapping ^{13}C-^{2}H multiplets of ^{13}C enriched metabolites from cell suspensions incubated in deuterated media. *Magnetic Resonance in Medicine*, 36(3):359–365, 1996.

[42] R. Laatikainen. Automated analysis of NMR spectra. *J. Magnetic Resonance*, 92:1–9, 1991.

[43] R. Laatikainen, M. Niemitz, U. Weber, J. Sundelin, T. Hassinen, and J. Veps˙General strategies for total-lineshape-type spectral analysis of NMR spectra using integral-transform iterator. *J. Magnetic Resonance*, 120:1–10, 1996.

[44] N. Sandgren, Y. Selén, P. Stoica, and J. Li. Parametric methods for frequency-selective MR spectroscopy, a review. *Journal of Magnetic Resonance*, 168(2):259–272, 2004.

[45] V. A. Mandelshtam. FDM: the filter diagonalization method for data processing in NMR experiments. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 38(2):159 – 196, 2001.

[46] P. Stoica and K. C. Sharman. Maximum likelihood methods for direction-of-arrival estimation. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 38(7):1132–1143, Jul 1990.

[47] E. H. Djermoune, M. Tomczak, and P. Mutzenhardt. A new adaptive subband decomposition approach for automatic analysis of NMR data. *J. Magnetic Resonance*, 169:73–84, 2004.

[48] M. R. Smith, S. T. Nichols, R. M. Henkelman, and M. L. Wood. Application of Autoregressive Moving Average Parametric Modeling in Magnetic Resonance Image Reconstruction. *Medical Imaging, IEEE Transactions on*, 5(3):132–139, Sept. 1986.

[49] R. Roy and T. Kailath. Estimation of Signal Parameters via Rotational Invariance Techniques. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP(37):984–995, 1989.

[50] J. M. A. Scherpen and W. S. Gray. On singular value functions and Hankel operators for nonlinear systems. In *Proc. ACC'99*, pages 2360–2364, 1999.

[51] R. E. Slyh and R. L. Moses. Formant estimation using singular value decomposition. *Circuits and Systems, 1989., Proceedings of the 32nd Midwest Symposium on*, pages 109–112 vol.1, Aug. 1989.

[52] P. Stoica, H. Li, and J. Li. Amplitude estimation of sinusoidal signals: survey, new results,and an application. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 48:338–352, Feb. 2000.

[53] P. Stoica and T. Sundin. Nonparametric NMR Spectroscopy. *Journal of Magnetic Resonance*, 152:57–69(13), September 2001.

[54] K. Konstantinides and K. Yao. Statistical analysis of effective singular values in matrix rank determination. *IEEE transactions on acoustics, speech, and signal processing*, 36(5):757–763, 1988.

[55] B. Merkel, M. T. Harz, O. Konrad, H. K. Hahn, and H.-O. Peitgen. A novel software assistant for the clinical analysis of MR spectroscopy with MeVisLab. *Medical Imaging 2008: Computer-Aided Diagnosis*, 6915(1):69152R, 2008.