ADVANCING MAGNETIC RESONANCE SPECTROSCOPY AND ENDOSCOPY WITH PRIOR KNOWLEDGE

by
Yi Zhang

A dissertation submitted to Johns Hopkins University in conformity with
the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland
June, 2014

Abstract

Reconstruction is key to the generation of anatomic, functional and biochemical information in the field of Magnetic Resonance (MR) in medicine. Here, prior knowledge based on various conditions is utilized through reconstruction to accelerate current MR techniques and reduce artifacts.

First, prior knowledge from Magnetic Resonance Imaging (MRI) is exploited to accelerate spatial localization in Magnetic Resonance Spectroscopy (MRS). The MRS information is contained in one extra chemical shift dimension, beyond the three spatial dimensions of MRI, and can provide valuable *in vivo* metabolic information for the study of numerous diseases. However, its research and clinical applications are often compromised by long scan times. Here, a new method of localized Spectroscopy with Linear Algebraic Modeling (SLAM) is proposed for accelerating MRS scans. The method assumes pre-conditions that the MRS scan is preceded by a scout MRI scan and that a compartment-averaged MRS measurement will suffice for the assessment of metabolic status. SLAM builds *a priori* MRI-based segmentation information into the standard Fourier-encoded MRS model of chemical shift imaging (CSI), to directly reconstruct compartmental spectra.

Second, SLAM is extended to higher dimensions and to incorporate parallel imaging techniques that deploy pre-acquired sensitivity information based on the use of separate multiple receive-coil elements, to further accelerate scan speed. In addition, eddy current-induced phase effects are incorporated into the SLAM model, and a modified reconstruction algorithm provides improved suppression of signal leakage due to heterogeneity in the MRS signal, especially when employing sensitivity encoding.

Third, prior information from MRI is also used to reduce the problem of lipid artifacts in ^{1}H brain CSI. CSI is routinely used for human brain MRS studies, and low spatial resolution in CSI causes partial volume error and signal 'bleed' that is especially deleterious to voxels near the scalp. A standard solution is to apply spatial apodization, which adversely affects spatial resolution. Here, a novel automated strategy for partial volume correction that employs grid shifting ('PANGS') is presented, which minimizes lipid signal bleed without compromising spatial resolution. PANGS shifts the reconstruction coordinate in a designated region of image space—the scalp, identified by MRI—to match the tissue center of mass instead of the geometric center of each voxel.

Last, prior knowledge of the spatially sparse nature of endoscopic MRI images acquired with tiny internal MRI antennae, and that of the null signal location of the endoscopic probe, are used to accelerate MR endoscopy and reduce motion artifacts. High-resolution endoscopic MRI is susceptible to degradation from physiological motion, which can necessitate time-consuming cardiac gating techniques. Here, we develop acceleration techniques based on the compressed sensing theory, and un-gated motion compensation strategies using projection shifting, to effectively produce faster motion-suppressed MRI endoscopy.

Advisor and first reader: Paul A. Bottomley

Russell H. Morgan Professor, Department of Radiology and Radiological Science, Johns Hopkins University

Second reader: Jinyuan Zhou

Associate Professor, Department of Radiology and Radiological Science, Johns Hopkins University

Acknowledgments

Table of Contents

List of Figures

List of Tables

Abbreviations

1D: one dimension

2D: two dimensions

3D: three dimensions

^{1}H: hydrogen or proton

^{31}P: phosphorus

ADP: adenosine diphosphate

ATP: adenosine triphosphate

Cho: choline

COM: center of mass

Cr: creatine

CS: compressed sensing

CSI: chemical shift imaging

DFT: discrete Fourier transform

FT: Fourier transform

MR: magnetic resonance

MRI: magnetic resonance imaging

IV-MRI: intravascular MRI

PR-MRI: projection reconstruction MRI

MRS: magnetic resonance spectroscopy

NAA: N-acetylaspartate

NMR: nuclear magnetic resonance

PANGS: partial volume correction by grid shifting

PCr: phosphocreatine

PSF: point spread function

PVC: partial volume correction

PVE: partial volume error

SENSE: sensitivity encoding

SLAM: spectroscopy with linear algebraic modeling

fSLAM: fractional SLAM

SLIM: spectral localization by imaging

SLOOP: spectral localization with optimal point spread function

SNR: signal to noise ratio

SRF: spatial response function

dSRF: discrete SRF

T_1: spin-lattice relaxation time

T_2: spin-spin relaxation time

Chapter 1 : Introduction

1.1   Overview

Magnetic Resonance (MR) is one of the most advanced and versatile medical imaging modalities used in the clinic today. MR can provide excellent anatomic, metabolic, and functional information about biological tissues and organs *in vivo*. With regard to medical applications, MR developed historically in two different areas— Magnetic Resonance Imaging (MRI) for providing anatomic information and Magnetic Resonance Spectroscopy (MRS) for providing biochemical and metabolic information. MRI has now evolved to include many specialized applications, such as diffusion MRI, perfusion MRI, functional MRI, angiographic MRI, endoscopic MRI etc. Although the imaging speed and the quality of the MR information obtained have greatly improved since its advent, there is a continuing need to achieve even higher speed and better quality to enable new applications, improve current applications, reduce costs and add value. This dissertation is devoted to utilizing prior knowledge of various sorts to achieve higher speed and reduce artifacts in MR applications via improved reconstruction methodology. It is organized as follows.

Chapter 1 details the organization of this dissertation and the basic background of MR techniques, including both MRI and MRS.

Chapter 2 introduces a new MRS localization method—spectroscopy with linear algebraic modeling (SLAM)—which accelerates the standard chemical shift imaging (CSI) method with prior knowledge from scout MRI. SLAM builds *a priori* MRI-based segmentation information into the CSI model to directly reconstruct spectra for multiple compartments, to achieve significant speed and efficiency advantages. SLAM is

demonstrated with one-dimensional phosphorus MRS in phantoms, the human leg, and the heart, on a 3T clinical scanner. Its SNR performance, accuracy, sensitivity to registration errors, and inhomogeneity are evaluated.

Chapter 3 discusses the extension of SLAM to two and three dimensions, and the incorporation of SLAM with the parallel imaging technique, called sensitivity encoding (SENSE) which has already been implemented for CSI on commercial systems. A modified algorithm, SLAM*, is introduced to improve compatibility with SENSE. Various studies are presented to demonstrate and validate the feasibility and accuracy of SLAM, including phantom studies, as well as normal volunteer and patient studies, both retrospectively and prospectively, resulting in acceleration factors ranging from 4 to 120-fold. SLAM shows the potential to supplant the current MRS standard for clinical and human research applications, by dramatically reducing scan time while maintaining essentially the same quantitative results.

Chapter 4 presents a new technique—partial volume correction by grid shifting (PANGS)—to reduce lipid artifacts in ^{1}H brain MRS by shifting reconstruction coordinates only in the scalp region identified from scout MRI. PANGS shifts the reconstruction coordinates of scalp voxels iteratively to match the 'centers-of-mass' of those voxels rather than the geometric centers. PANGS' performance is evaluated in numerical simulations and in 3T ^{1}H CSI studies of healthy subjects and brain tumor patients. PANGS can significantly reduce lipid artifacts in ^{1}H brain CSI spectra and metabolite maps, and improve metabolite detection in cortical regions, without compromising resolution.

Chapter 5 demonstrates the combination of acceleration and motion-correction techniques for high-resolution endoscopic MRI performed with tiny internal MRI detectors. Acceleration is achieved with the compressed sensing theory, and motion correction is realized by shifting and aligning individual projections in radially encoded images. Acceleration and motion compensation strategies are presented in fruit, human vessel specimens, and a free-breathing ~200 μm *in vivo* rabbit aorta, with up to eight-fold effective acceleration.

Chapter 6 concludes the dissertation, describing some ongoing work and potential future directions.

1.2   Background of MR physics

1.2.1   Nuclear Magnetic Resonance

The Nobel Prize Laureates, Dr. Felix Bloch and Dr. Edward Purcell, first discovered the Nuclear Magnetic Resonance (NMR) effect in liquids and solids in 1946. The NMR phenomenon occurs when certain types of atoms are placed in an external magnetic field, in which they can absorb and re-emit radio waves at well-defined frequencies depending on the field strength. Atoms with an odd number of protons or neutrons can exhibit NMR, such as ^{1}H, ^{13}C, ^{19}F, ^{23}Na, and ^{31}P. The nuclei of these atoms have a special quantum mechanical property, spin angular momentum, which does not follow classical mechanics, although nuclei are often visualized from a classical viewpoint as charged ball particles rotating around their axes. The NMR active nuclei are typically referred to as spins. The frequencies at which NMR occurs follow the Larmor equation,

$$f = \frac{\gamma}{2\pi} B \,,$$
(1.1)

where γ is the nucleus-specific gyromagnetic ratio and B is the external magnetic field strength. The gyromagnetic ratios, $\gamma/(2\pi)$, for common nuclei are: ^{1}H (also called "protons" in the MR field), 42.567 MHz/T; ^{13}C, 10.705 MHz/T; ^{19}F, 40.053 MHz/T; ^{23}Na, 11.262 MHz/T; and ^{31}P, 17.235 MHz/T. The unit of B is in Tesla (T) or Gauss (G) and 1 Tesla = 10^4 Gauss.

Each spin possesses a magnetic dipole moment. In fact, the gyromagnetic ratio is the quotient between the magnetic dipole moment and the spin angular momentum. Although NMR is a quantum mechanical phenomenon, it can be well described from the classical viewpoint macroscopically, which will be used in this dissertation henceforth. When an ensemble of spins in a unit volume are placed in an external magnetic field, **B**, a bulk macroscopic magnetization, **M**, will form due to the alignment of magnetic dipole moments with **B**. The relationship between **M** and **B** can be described by the phenomenological Bloch equation [1],

$$\frac{d\mathbf{M}}{dt} = \mathbf{M} \times \gamma\mathbf{B} - \frac{M_x\mathbf{i} + M_y\mathbf{j}}{T_2} - \frac{(M_z - M_0)\mathbf{k}}{T_1}, \qquad (1.2)$$

where γ is the gyromagnetic ratio; M_x, M_y, M_z are the three components of **M** in **i**, **j**, **k** or x, y, z directions, respectively; M_0 is the equilibrium or initial magnetization at time, t = 0; T_1 is the spin-lattice or longitudinal (z) relaxation time; and T_2 is the spin-spin or transverse (x-y plane) relaxation time.

1.2.2 Magnetic Resonance Imaging

The modern MRI technique is generally considered to have originated with the report of the first NMR-based image [2] by Dr. Paul Lauterbur in 1973, who, along with Dr. Peter Mansfield, won the Nobel Prize in Physiology or Medicine "for their

4

discoveries concerning magnetic resonance imaging" in 2003. MRI has evolved dramatically since that first report [2], but the basic principle remains: magnetic gradient fields are used to encode information, which distinguishes MRI as a unique technique whose resolution (on the order of millimeters) is not limited by the electromagnetic wavelength (on the order of meters) unlike other ultrasound or optical medical imaging techniques.

Three types of magnetic fields

Three types of magnetic fields are required to generate an MRI image: the main field (B_0), the radiofrequency (RF) excitation field (B_1), and linear gradient fields (G). The B_1 excitation field (x-y plane) must be applied perpendicular to the B_0 field (z direction). The G fields, e.g. $G_x = \dfrac{\partial B_0}{\partial x}$, $G_y = \dfrac{\partial B_0}{\partial y}$ and $G_z = \dfrac{\partial B_0}{\partial z}$, are in the B_0 direction, which typically include three subtypes: slice selection (SS); phase encoding (PE); and frequency encoding (FE) gradient fields. The combination of RF and G field pulses applied in some specific order is called an MRI sequence, which has numerous different forms depending on the application, creating a rich variety of contrast and functional information accessible to MRI. A basic Cartesian Fourier encoding MRI sequence is shown in Fig. 1.1.

The B_0 field is employed to enable the NMR phenomenon and to create a macroscopic equilibrium magnetization, M_0, along the z direction. The value of B_0 commonly used in clinical MRI scanners is 1.5 T and 3 T. However, for animal scanners, B_0 can extend to 11.7 T or even 17 T.

RF excitation is needed to tip the equilibrium M_0 away from B_0 designated as the z axis of a Cartesian coordinate system into the x-y plane, where in accordance with the

Bloch equation (Eq. 1.2), a precessing transverse magnetization $M_{xy} = M_x + i\,M_y$, is created. The flip angle or tip angle is defined as the angle between the tipped magnetization vector, **M**, and the z axis in a frame of reference rotating at the NMR frequency about the z-axis. In Eq. 1.2, T_1 characterizes the time taken for M_z to return back to M_0, and T_2 characterizes the rate at which M_{xy} decays in the x-y plane. The strength of a typical B_1 field is around a few micro Tesla. For a constant B_1 of duration—τ, the flip angle (FA) can be computed from the Larmor equation as $FA = \gamma B_1 \tau$ in radians. The precessing M_{xy} magnetization is the signal source by which an MR image is generated.

The G fields play a vital role in spatially localizing of the distribution of the nuclear magnetization signal. According to Eq. 1.1, spins resonate at different frequencies, depending on the external magnetic field strength. An SS gradient field introduces variations of the external B field, for example, along the z direction, $B_z = B_0 + G_z \times z$. If a frequency-selective B_1 field is applied in concordance with the SS gradient field, a thin slice can be specifically selected in the z direction without exciting other locations. To resolve spatial magnetization in the x-y plane, the PE and FE gradients are used sequentially to encode information in, say, the x and y directions, respectively. As shown in Fig. 1.1, the FE gradient is applied coincidently with the readout acquisition (AQ) window. Similar to the SS gradient, the FE gradient assigns a different resonant frequency to each location. The Fourier transform of the received time-domain signal will generate a map of magnetization for each frequency, corresponding to the distribution of the magnetization for each location in the y direction. Similar to SS and FE, PE also makes the resonant frequency location-dependent—this time in the x direction. But, in

contrast to the FE gradient, the PE gradient is applied for a short duration with no RF

pulse but prior to the AQ window. As a result, spins resonating at different frequencies

(i.e., at different locations in the x direction) will appear in-phase and/or out-of-phase

with respect to each other. This phase modulation can be decoded through Fourier

transform to reconstruct the spatial distribution of magnetization in the x direction, as in

the "spin warp" imaging method [3].
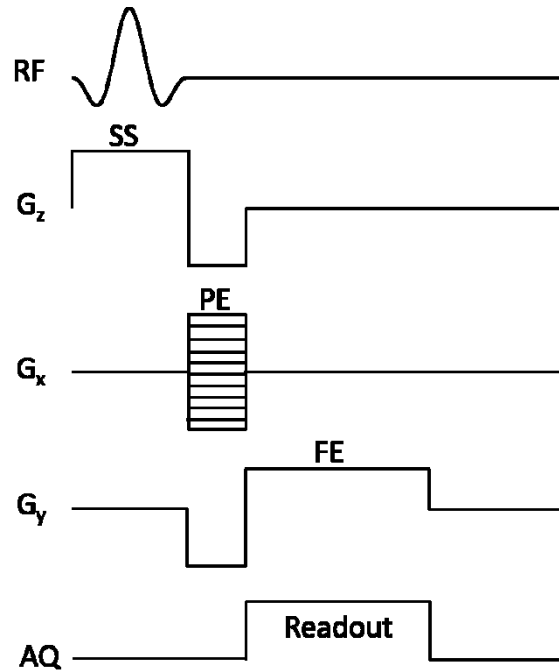
Figure 1.1: A sequence diagram showing the implementation of 2D Cartesian Fourier encoding for MRI.

Signal equation

The rotating M_{xy} in the transverse plane will cause a time-dependent magnetic

flux, which induces an electromotive force and current in a receiving RF coil oriented

with sensitivity to detect the transverse field, according to Faraday's law of induction. As

a result, the received signal for a 2D Cartesian imaging case (Fig. 1.1) from the

selectively excited slice in the z direction, can be described as [1],

$$s(t) = \omega_0 \int_x \int_y m(x, y) \exp\{-i2\pi[k_x(t)x + k_y(t)y]\} dx dy \ , \qquad (1.3)$$

where

$$\omega_0 = 2\pi f_0 = \gamma \times B_0$$
$$k_x(t) = \frac{\gamma}{2\pi} \int_0^t G_x(\tau) d\tau \ , \qquad\qquad (1.4)$$
$$k_y(t) = \frac{\gamma}{2\pi} \int_0^t G_y(\tau) d\tau$$

and $m(x, y)$ is a function of intrinsic properties, including spin density or proton density

(PD)— $\rho(x, y)$, $T_1(x, y)$, and $T_2(x, y)$ in the slice. As shown in Eq. 1.3, the intended

image, $m(x, y)$, can be reconstructed by Fourier transforming the set of signals, $s(t)$.

The image contrast in $m(x, y)$ can be manipulated by using various kinds of imaging

sequences, resulting in, for example, PD-weighted, T_1-weighted, and T_2-weighted images.

1.2.3 Magnetic Resonance Spectroscopy

MRS measures the same NMR spectroscopy signals widely used in chemistry and

biochemistry to identify and study the properties and structure of molecules, with the

added feature of providing methods of spatially localizing those signals. Like its sister

technology MRI, MRS does not involve use of ionizing radiation and is not considered a

significant risk.

Chemical shift

As the Larmor equation (Eq. 1.1) states, the resonant frequency is linearly

proportional to the external magnetic field strength. However, the external magnetic field

includes not only the main B_0 field, but also local magnetic fields created by the shielding

effect of orbiting electrons. This slight resonant frequency change, due to the local

8

chemical environment, is known as the chemical shift. Chemical shift is usually measured as the frequency change compared to a reference frequency in parts per million (ppm),

$$\delta = \frac{f_{eff} - f_{ref}}{f_{ref}} \times 10^6 , \tag{1.5}$$

where f_{ref} is the reference frequency and $f_{eff} = \gamma / (2\pi) \times B_0 \times (1 - \sigma)$ is the effective frequency after considering the chemical shift effect. σ is the shielding factor, which is dependent only on the local chemical environment, but not on the external main magnetic field. For ^{1}H and ^{13}C, the frequency of tetramethysilane (TMS) is usually used as f_{ref}; for ^{31}P, it is usually phosphoric acid (H_3PO_4) for *in vitro* studies, while phosphocreatine (PCr) is commonly used as a reference for *in vivo* studies.

Chemical shift can provide unique information about *in vivo* metabolism. For ^{1}H studies, some common metabolites of interest are choline (Cho; a cell membrane marker), creatine (Cr; involved in energy metabolism), and N-acetylaspartate (NAA; a neuronal marker). For ^{31}P, some well-studied metabolites are adenosine triphosphate (ATP; the energy fueling many cellular processes), adenosine diphosphate (ADP; which is reversibly converted to ATP), and phosphocreatine (PCr; a reservoir of cellular energy also used to supply ATP).

Chemical shift imaging

The most commonly used spectroscopic localization method for acquiring MRS data from multiple voxels is chemical shift imaging (CSI) [4]. A 2D CSI sequence utilizing Fourier phase encoding is shown in Fig. 1.2. The CSI sequence is quite similar to the MRI sequence in Fig. 1.1. The most distinct difference is that there is no frequency encoding gradient in CSI. Instead, a second phase encoding is used to resolve spatial

information in the y direction. Since the PE gradient must be incremented N times, the

extra PE gradient makes CSI N times slower than the MRI sequence. Compared to MRI,

CSI provides an extra dimension of information, which can be used to specifically

monitor and measure the concentrations of different metabolites *in vivo*. Rather than

generating a single intensity for each voxel, as in an MR image, MRS delivers a spectrum

for every voxel, wherein a metabolite map can be constructed for each spectral frequency.

In Fig. 1.2, the signal acquired in CSI with the FE gradient turned off is called a free

induction decay (FID). The CSI signals can also be described by the signal formula in Eq.

1.3, except that a variable representing the chemical shift dimension must now be added.

This will be discussed further in Chapters 2, 3, and 4.
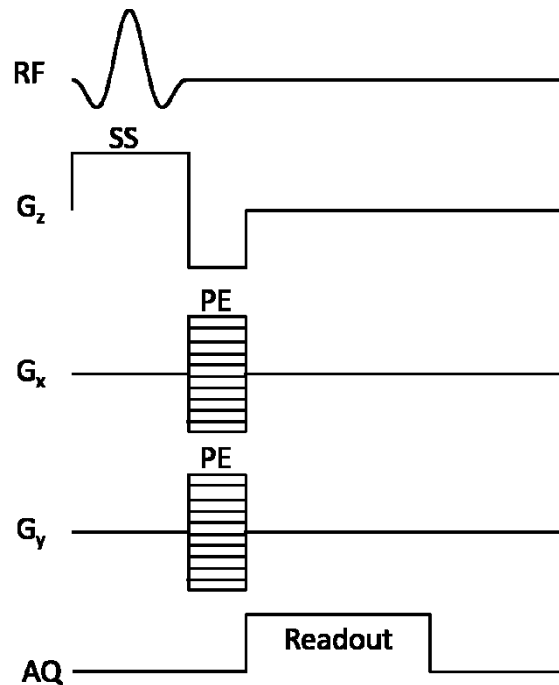
Figure 1.2: A sequence diagram depicting the implementation of Fourier phase encoding for 2D CSI.

1.2.4 Discrete sampling and reconstruction

Equation 1.3 can be reformulated by changing the independent variables as,

$$s(k_x, k_y) = \int_x \int_y m(x, y) \exp\{-i2\pi[k_x x + k_y y]\}dxdy \ . \qquad (1.6)$$

Clearly, Eq. 1.6 is a 2D Fourier transform between the signal, $s(k_x, k_y)$, and the intended

image, $m(x, y)$, where x, k_x and y, k_y are conjugate variable pairs. The signal space, i.e.,

the Fourier space of the image, is called k-space in MR, as in other imaging modalities. In

reality, the signal has to be sampled discretely, as shown in Fig. 1.3. The PE direction (x)

is sampled by increasing the gradient field strength incrementally, ΔG_x, in fixed steps of

duration, τ. The FE direction (y) is sampled as a function of time at increments of Δt, but

in the presence of a fixed gradient field of strength, G_y.
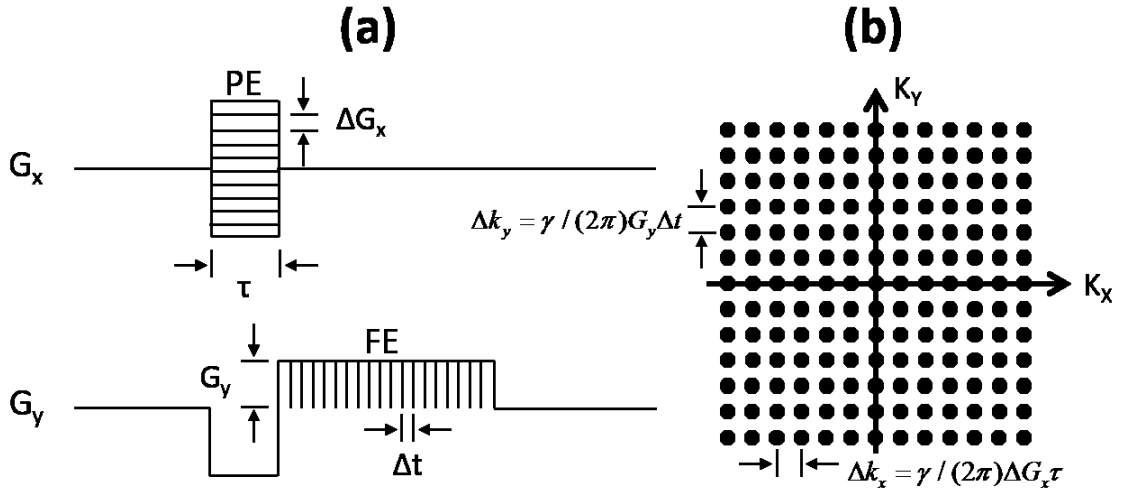
Figure 1.3: Discrete sampling in the phase encoding and frequency encoding directions (a) and corresponding k-space illustration (b).

To reconstruct the final image from the discretely sampled k-space, a discrete Fourier

transform can be readily employed,

$$m(x, y) = \sum_{k_x = -N_{PE}/2}^{N_{PE}/2} \sum_{k_y = -N_{FE}/2}^{N_{FE}/2} s(k_x, k_y) \exp(i2\pi k_x x / N_{PE}) \exp(i2\pi k_y y / N_{FE}), \quad (1.7)$$

where $-N_{PE}/2 \le x \le N_{PE}/2$, $-N_{FE}/2 \le y \le N_{FE}/2$, N_{PE} is the number of phase

encoding steps and N_{FE} is the number of frequency encoding points.

The signal formulae for MRS are quite similar to Eqs. 1.6 and 1.7, and will be

discussed in greater detail in Chapters 2, 3 and 4.

1.2.5 Signal-to-noise ratio

The signal-to-noise ratio (SNR) is of paramount importance for the practical

detection of both MRI and MRS. Without adequate SNR in reasonable scan times, MRI

images and MRS spectra could not be useful. In NMR, SNR is defined as the signal

amplitude divided by the standard deviation of noise. The noise in MRI and MRS is

Johnson noise, which arises from the thermal motions of electrons in conductors in the

detector system and in the sample. The standard deviation of noise, σ_n, can be described

through an effective resistance as [5],

$$\sigma_n = \sqrt{2kTR/\tau},$$ (1.8)

where k is the Boltzmann constant (1.38E-23 J/K), T is the absolute temperature in

Kelvin (K), R is the total effective resistance, and τ is the total acquisition time.

If the noise from the scanner's receiver electronics can be ignored, the effective

resistance primarily includes resistance in the receiver coil, R_c, and effective resistance

of the sample, R_s (e.g., the human body). R_c is related to the skin depth of the receiver

coil [1] and losses in the tuning elements, operating at a relatively high frequency, as

$R_c \propto \sqrt{\omega_0} \propto \sqrt{B_0}$. R_s can be estimated by assuming the human body as shells of

solenoids [5], resulting in $R_s \propto \omega_0^2 \propto B_0^2$. The signal can be derived from Eq. 1.3, after

considering the macroscopic magnetization is proportional to the external magnetic field

strength, as $signal \propto \omega_0 B_0 \propto B_0^2$. Thus, the SNR formula can be derived after adding

weights (*a* and *b*) to R_c and R_s, respectively,

$$SNR \propto \frac{B_0^2}{\sqrt{aB_0^{-1/2} + bB_0^2}} \ . \tag{1.9}$$

In most cases, sample noise is the dominant source, i.e. $a \ll$, leading to $SNR \propto B_0$. If

coil noise happens to be dominant at each field strength ($a \gg$), then SNR will appear to

be more strongly dependent on B₀, as $SNR \propto B_0^{7/4}$ from the first term in the denominator.

In any case, higher SNR will be obtained when a stronger main magnetic field is used.

SNR can be significantly affected by the imaging sequence parameters as well.

The rule of thumb is that SNR is linearly proportional to the voxel size and the square

root of total acquisition time, as

$$SNR \propto \Delta V \times \sqrt{\text{total acquistion time}} \ , \tag{1.10}$$

where ΔV denotes the voxel size.

Moreover, the SNR behavior can be complicated in some other cases, such as

when multiple receiver coils are used in parallel imaging [6], and when imaging artifacts

exist. In practice, a popular way of measuring SNR is to acquire the image or spectrum

twice under identical conditions. The summation of the two measurements provides the

signal, and the subtraction of the measurements provides the noise [7] assuming nothing

changes.

Chapter 2 : Spectroscopy with linear algebraic modeling (SLAM)

2.1   Introduction

Scan-time and signal-to-noise ratio (SNR) are major problems for *in vivo* spatially

localized magnetic resonance spectroscopy (MRS) of low-concentration metabolites.

Because SNR is proportional to voxel size, matching the voxel to the desired anatomical

compartment *a priori* yields the best SNR for a fixed scan time [8]. Consider for example

a first chemical shift imaging (CSI) experiment [9] encoding a voxel *V* with an SNR of

20 per acquisition. Averaging *n*=4 acquisitions yields an SNR of 40 since SNR adds as

$\sqrt{n}$. Now consider a second experiment performed at four times the resolution with *V*/4-

sized voxels. The SNR/voxel is now 5 per acquisition because noise is independent of

voxel size [8]. Phase-encoding is equivalent to averaging, so after 4 gradient steps to

encode the same volume, the SNR per voxel is 10. Adding the 4 signals to make a *V*-

sized voxel now yields an SNR of 20, again because of the $\sqrt{n}$ rule. This compares to 40

from the first experiment. Thus, the SNR for the same scan-time and voxel size is

doubled in the first experiment, just by pre-selecting the correct voxel size to start with

[8].

The same principle applies in general wherever the CSI voxel size is smaller than

the object of interest. The SNR gain factor for a fixed scan-time obtained by correctly

encoding a compartment at the outset, as compared to adding signals from individual CSI

voxels to form the equivalent-sized compartment post-acquisition, is:

$$g = \sqrt{\frac{compartment\ size}{CSI\ voxel\ size}} , \qquad (2.1)$$

notwithstanding the effects of nonuniform sensitivity and concentration distributions, or differences in the integrated spatial response function (SRF). This differential g-fold SNR gain vs. CSI can be seen as arising from the time lost by CSI in acquiring the low SNR, high gradient-strength, high k-space signals that are unnecessary for resolving the larger compartment.

Prior phase-encoding gradient based MRS localization methods such as SLIM [10], GSLIM [11] and SLOOP [12], could realize the g-fold SNR gain if the desired compartments were prescribed from scout MRI prior to acquisition, and if an appropriately SNR-optimized gradient set were then applied. In SLIM, the compartment's signal is modeled as the integral of phase-encoded signal contributions in each compartment, assumed homogeneous. The approach is prone to inter-compartmental [13] and intra-compartmental errors when metabolite distributions are non-uniform between and within each compartment, and as the number of phase-encoding gradient steps are reduced. GSLIM [11] and SLOOP [12] were introduced to minimize the inter-compartmental errors. GSLIM does this by applying non-Fourier generalized series modeling to the SLIM result [11, 13]. SLOOP minimizes the inter-compartmental error by optimizing the SRF for the desired compartment, ideally by specifically tailoring the phase-encoding gradient set for the acquisition [12]. Several other proposed improvements add constraints to deal with inhomogeneity in the main (B_0) field [14-16], registration errors [16], and multi-element receivers [17].

Even though all of these techniques can generate spectra from multiple compartments from the same data set, they are seldom used pro-actively for human MRS. Thus, SLIM was applied retroactively to [1]H CSI data sets acquired from the human calf

[10, 18] and brain [16], and both GSLIM and SLIM were used in ^{1}H MRS CSI

acquisitions from a gerbil brain [19]. Although SLOOP ^{1}H MRS *was* initially performed

with proactively optimized gradients on an excised rabbit kidney [12], all subsequent

applications to human heart applied SLOOP retroactively to phosphorus (^{31}P) MRS data

acquired with regular CSI gradients [20-23]. Because all of these human applications

employed conventional CSI gradient sets and uniform *k*-space sampling, a *g*-fold SNR

advantage versus CSI, beyond that obtained by simply summing the signals from the

constituent CSI voxels or accounting for differences in the integrated SRF, was not

realized, or at least reported. The lack of pro-active implementation and absence of a

demonstrated SNR advantage have likely contributed to the failure of these methods to

supplant routine CSI. In any case, the prescribing of compartments and tailoring of

gradient encoding steps to match the desired compartment and achieve the full SNR gain

predicted by Eq. (2.1) has, to the best of our knowledge, never been realized *in vivo* or in

humans.

Here we apply a sharply-reduced SNR-optimized gradient set to perform localized

*S*pectroscopy with *L*inear *A*lgebraic *M*odeling (SLAM) to acquire and reconstruct average

spectra from *C* signal-generating anatomical compartments that are identified by the

scout MRI that is routinely acquired during a spatially-localized MRS exam. Spectral

reconstruction for this new SLAM method differs from SLIM, GSLIM and SLOOP in

that it solves, by matrix analysis, a set of linear simultaneous equations essentially equal

to *C* (provided that all signal-generating tissues are included) by eliminating un-needed

phase-encoding steps from the standard CSI algorithm. The SLAM pulse sequence differs

in that the number of phase-encoding steps is also ~*C,* and they are always selected from

Table 2.1: Comparison between SLIM, GSLIM, SLOOP, SLAM and fSLAM

Comparison	SLIM	GSLIM	SLOOP	SLAM*	fSLAM*
Method	CSI phase-encodes and MRI-based segmentation	CSI phase-encodes and MRI-based segmentation	CSI phase-encodes may be optimized to minimize bleed when compartments are non-uniform.	Only the M central *k*-space CSI phase-encodes are used to yield optimum SNR.	Central *k*-space phase-encodes are selected to optimize SNR, errors and contamination.
Assumption	Every compartment is uniform.	SLIM's non-uniform compartments are addressed with non-Fourier generalized series (GS) basis sets.	Uniform compartments, but gradients adjustable to minimize inter-compartment bleed for non-uniform compartments.	The compartment spectrum is the average of the CSI spectra from each compartment.	The compartment spectrum is the average of the CSI spectra from each compartment, with SNR optimized and errors minimized vs. CSI.
Gradient Optimization	Not optimized	Not optimized	May be optimized for SNR and inter-compartment bleed.	SNR optimized using minimum number of central *k*-space phase-encodes	Optimized for SNR, inter-compartment leakage and intra-compartment error
Gradient Step	Stepped as $\sim n/FOV$, n $=\ldots-2,-1,0,1,2...$ (integer)	Stepped as $\sim n/FOV$, n $=\ldots-2,-1,0,1,2...$ (integer)	Stepped as $\sim p/FOV$, p integer or non-integer (if optimized)	Stepped as $\sim n/FOV$, with $-M/2 \leq n \leq M/2$, n integer, $C \leq M$	Stepped per optimization, maximum gradient limited by compartment size
Reconstruction	Compartment signal computed from the integrated CSI signal contributions	Same as SLIM, followed by GS constrained modeling to minimize errors	Compartment signal computed from the integrated CSI signal contributions	Starting from CSI, phase-encodes are eliminated down to C compartments.	Signals acquired with pro-actively optimized phase-encodes are reconstructed same as SLAM.

Errors	Exact for infinite n. Inter- and intra-compartment errors arise for small n.	Reduces leakage for finite n by applying GS modeling to SLIM result.	Inter-compartment leakage reduced vs. SLIM. Errors from intra-compartment non-uniformity remain.	Errors increase as the number of phase encodes decrease.	Both inter-compartment bleed and intra-compartment error are suppressed.
In vivo human use	Applied to retroactive [1]H CSI data in human calf	None	Applied to retroactive [31]P CSI data from human heart	Applied both retroactively to CSI data, and pro-actively to human cardiac [31]P and brain [1]H MRS. Quantitation validated	fSLAM applied pro-actively in human cardiac [31]P MRS.

* Shaded squares indicate significant differences vs. other methods

the center of the integer-stepped *k*-space of CSI where SNR is highest. Other than

determining the number, *C*, the need for image-guided gradient optimization, prescription

and implementation at the scanner-side prior to acquisition, is avoided. Using SLAM, *g*-

fold SNR gains of 30-200% are demonstrated in 3T [31]P studies of the human leg and

heart *in vivo*, compared to conventional [24-29] one-dimensional (1D) CSI spectra from

the same net volume and scan-time. Moreover, we show that application of SLAM to raw

[31]P 1D CSI data acquired from heart patients and scout MRI-based segmentation yields,

after discarding 75% of the data, essentially the same *quantitative* measures of adenosine

triphosphate (ATP) and phosphocreatine (PCr), four-times faster.

Finally, we extend the SLAM approach to allow for fractional gradient increments

instead of conventional, integer-stepped, CSI gradients. In this "fSLAM" method, the

phase-encoding gradients are pro-actively optimized at the scanner-side to maximize

SNR and/or minimize *both* the inter-compartmental leakage as well as the intra-compartmental errors produced by nonuniform signal distributions. Intra-compartmental errors have not been addressed in prior methods [10-12]. fSLAM is demonstrated in pro-active human cardiac ^{31}P studies.

2.2   Theory

Consider the basic equation for 1D CSI:

$$s(k,t) = \int \int \rho(x,f)e^{-i2\pi(kx+ft)}dfdx, \qquad (2.2)$$

where k is the spatial frequency, $s(k,t)$ is the acquired time-domain signal and $\rho(x,f)$ is the spectrum to reconstruct. Since localization is in the spatial domain which is independent of the frequency domain, we denote the spectrum at a spatial position x after $s(k,t)$ is Fourier transformed (FT), as $\rho(x)$ in the spectral frequency domain. Assuming there are M phase encoding steps, $k_1...k_M$, Eq. (2.2) is discretized as:

$$
\begin{bmatrix} s(k_1) \\ s(k_2) \\ \vdots \\ s(k_M) \end{bmatrix}_{M*N}
\begin{bmatrix} e^{-i2\pi k_1 x_1/M} & e^{-i2\pi k_1 x_2/M} & \cdots & k_1 x_M/M \\ e^{-i2\pi k_2 x_1/M} & e^{-i2\pi k_2 x_2/M} & \cdots & k_2 x_M/M \\ \vdots & \vdots & \ddots & \vdots \\ e^{-i2\pi k_M x_1/M} & e^{-i2\pi k_M x_2/M} & \cdots & k_M x_M/M \end{bmatrix}_{M*M}
\begin{bmatrix} \rho(x_1) \\ \rho(x_2) \\ \vdots \\ \rho(x_M) \end{bmatrix}_{M*N}
$$
(2.3)

Each row of the known signal matrix, $\mathbf{S}_{M*N}$, on the left side of the equation is an N-point array, where N is the number of time-domain data points. The first matrix on the right side is the phase-encoding FT operator (**PE**), and each term of the unknown spectral matrix, $\boldsymbol{\rho}$, is also an N-point array. For simplicity, we write Eq. (2.3) as:

$$\mathbf{S}_{M*N} = \mathbf{PE}_{M*M} \times \boldsymbol{\rho}_{M*N}.$$

2.2.1   Localized spectroscopy using a linear algebraic model (SLAM)

The goal of the CSI experiment is to reconstruct the M unknown spectra in matrix ρ of Eq. (2.3), from the M known signals ($\mathbf{S}$) acquired with M different phase-encodes. However, from scout MRI we learn that ρ has just $C\ (< M)$ MRS compartments of interest, as well as the spatial position of each compartment. Theoretically, only C measurements with C phase-encoding steps are needed to unambiguously solve ρ and reconstruct the C spectra.

To illustrate, consider a 4-voxel 1D CSI experiment. Denoting the exponential terms by $e_{i,j}$, Eq. (2.3) becomes:

$$\begin{bmatrix} e_{11} & e_{12} & e_{13} & e_{14} \\ e_{21} & e_{22} & e_{23} & e_{24} \\ e_{31} & e_{32} & e_{33} & e_{34} \\ e_{41} & e_{42} & e_{43} & e_{44} \end{bmatrix} \times \begin{bmatrix} \rho_1 \\ \rho_2 \\ \rho_3 \\ \rho_4 \end{bmatrix} = \begin{bmatrix} S_1 \\ S_2 \\ S_3 \\ S_4 \end{bmatrix}. \tag{2.4}$$

Now suppose that from prior information, the second and third rows of ρ are the same ($\rho_2 = \rho_3$). Then we need only solve:

$$\begin{bmatrix} e_{11} & e_{12} + e_{13} & e_{14} \\ e_{21} & e_{22} + e_{23} & e_{24} \\ e_{31} & e_{32} + e_{33} & e_{34} \\ e_{41} & e_{42} + e_{43} & e_{44} \end{bmatrix} \times \begin{bmatrix} \rho_1 \\ \rho_2 \\ \rho_4 \end{bmatrix} = \begin{bmatrix} S_1 \\ S_2 \\ S_3 \\ S_4 \end{bmatrix}. \tag{2.5}$$

Eq. (2.5) is now over-determined and the minimum number of phase-encoding steps required can be reduced from 4 to 3.

The same theory shows that we can reconstruct C spectra from C homogeneous compartments, with only C phase-encoding steps instead of M steps, regardless of k-space truncation. In general, prior information is incorporated via a $\mathbf{b}$-matrix which zeros out identical rows in the ρ-matrix to retain only one spectrum for each compartment:

20

$$\mathbf{S}_{M*N} = \mathbf{PE}_{M*M} \times \mathbf{b}^{-1}_{M*M} \times \mathbf{b}_{M*M} \times \boldsymbol{\rho}_{M*N}, \tag{2.6}$$

where **PE** is the phase-encoding operator from Eq. (2.3). For SLAM based on the 1D CSI experiment, the **b**-matrix is an identity matrix with "-1" elements inserted to zero out identical rows in **ρ.** For example, for an 8-voxel CSI experiment performed on a two-compartment sample in which the first compartment extends from voxels 1-3 and the second extends from voxels 4-8,

$$\mathbf{b} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 0 & 1 \end{bmatrix}. \tag{2.7}$$

Here, only the spectra in the compartments in voxels 1 and 4 are kept after dimensional reduction.

If we choose $M' \geq C$ pre-defined phase-encoding steps, and eliminate identical rows to reduce the dimension of $\mathbf{b}_{M*M} \times \boldsymbol{\rho}_{M*N}$ from M to C, Eq. (2.6) shrinks to,

$$\mathbf{S}_{M'*N} = \mathbf{PE}^r_{M'*C} \times \boldsymbol{\rho}^r_{C*N}, \tag{2.8}$$

where $\boldsymbol{\rho}^r_{C*N}$ is a submatrix of $\mathbf{b}_{M*M} \times \boldsymbol{\rho}_{M*N}$ retaining the C non-eliminated rows; $\mathbf{PE}^r_{M'*C}$ is a submatrix of $\mathbf{PE}_{M'*M} \times \mathbf{b}^{-1}_{M*M}$ that retains the C columns corresponding to the C non-eliminated rows; and $\mathbf{S}_{M'*N}$ is a submatrix of $\mathbf{S}_{M*N}$ acquired from the sample using a subset of $M' \ll M$ phase-encoding steps. Solution of Eq. (2.8) results in a set of spectra, each of which closely approximates the average spectrum of each 1D CSI compartment.

2.2.2 The SLAM recipe

In summary, the SLAM experiment is performed with Steps 1-5 as follows:

1. Acquire an MRI to extract the prior knowledge of the number of compartments ($C<<M$), and the spatial position of each compartment for SLAM reconstruction.

2. Choose $M' \geq C$ phase-encoding steps. Theoretically, these can be chosen arbitrarily, but different choices will lead to different SNR and different condition numbers for the matrix $\mathbf{PE}^r_{M'*C}$ which affect computational accuracy [30]. Of the M original CSI phase-encoding steps, selecting the M' steps that are closest to the center of k-space generally yields the best SNR. Because the set of CSI steps are discrete, fixed and finite, choosing only those from central k-space results in a SLAM phase-encoding gradient set that is determined only by the number M' or C. Moreover, because C is typically the same for a given study protocol (eg, $C = 2$ to 4 in the present cardiac studies with chest muscle, heart, and/or adipose and ventricular blood compartments), the same SLAM gradient set can be used for all the studies, eliminating the need for scanner-side gradient optimization or image-based gradient prescription.

3. Apply the chosen M' encoding gradients and acquire the M' signals.

4. Determine the **b** matrix from the spatial position of each compartment identified by MRI.

5. Reduce the dimensions from M to C and compute the C spectra in the $\mathbf{\rho}^r$ matrix using:

$$\mathbf{\rho}^r_{C*N} = \mathbf{PE}^+_{C*M'} \times \mathbf{S}_{M'*N},$$
(2.9)

where $\mathbf{PE}^+_{C*M'}$ is the inverse ($M'=C$) or pseudo-inverse ($M'>C$) of $\mathbf{PE}^r_{M'*C}$.

A flow diagram of the reconstruction algorithm appears in Fig. 2.1.
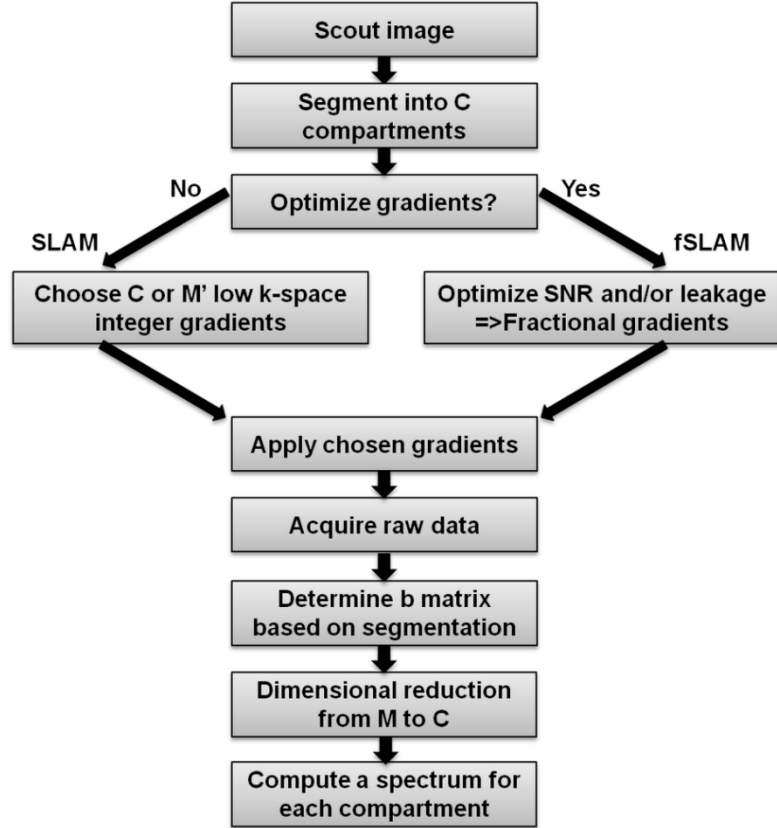
Figure 2.1: Flow chart depicting implementation of SLAM (left pathway) or fSLAM (right pathway).

2.2.3 SLAM with fractional gradients (fSLAM)

Choice of the *M'* phase-encoding steps need not be limited to the original basis set of *M* CSI steps corresponding to integer *k*'s in Eq. (2.3). The *M'* phase-encoding gradients can be chosen to optimize desired properties of the reconstruction. For example as we now show, the gradients can be optimized to maximize the SNR, and/or minimize the inter-compartmental signal contamination, and/or minimize the intra-compartmental error due to nonuniform signal sources. This effectively involves allowing for fractional *k*'s in the CSI Eq. (2.3), with all other experimental parameters left unchanged. Unlike SLAM, this fractional SLAM method, denoted fSLAM, does require scanner-side gradient optimization and prescription.

23

2.2.4 fSLAM with maximum SNR

To maximize the SNR, Eq. (2.9) is modified to include noise terms $\boldsymbol{\varepsilon}_{M'*N}$ in the

time-domain signal:

$$\boldsymbol{\rho}^r_{C*N} + \boldsymbol{\xi}_{C*N} = \mathbf{PE}^+_{C*M'} \times (\mathbf{S}_{M'*N} + \boldsymbol{\varepsilon}_{M'*N}), \tag{2.10}$$

where $\boldsymbol{\xi}_{C*N}$ is the noise in the reconstructed spectra. The noise in the time-domain signal

and the noise in the spectra are related via the linear transformation,

$\boldsymbol{\xi}_{C*N} = \mathbf{PE}^+_{C*M'} \times \boldsymbol{\varepsilon}_{M'*N}$. Assuming the standard deviation (SD), σ, of $\boldsymbol{\varepsilon}_{M'*N}$ is constant,

the SNR of the spectrum reconstructed from the ith compartment is:

$$SNR_i = \frac{\boldsymbol{\rho}^r_{C*N}(i)}{\sqrt{\sum_{m=1}^{M'}[|\mathbf{PE}^+_{C*M'}(i,m)|^2 \times \sigma^2]}}, \tag{2.11}$$

where $\mathbf{PE}^+_{C*M'}(i,m)$ is the element corresponding to the mth signal. To maximize the SNR

of the ith spectrum in Eq. (2.11), we numerically minimize the cost-function

$$\Gamma_i = \sum_{m=1}^{M'}[|\mathbf{PE}^+_{C*M'}(i,m)|^2] / I_{cond}, \tag{2.12}$$

where I_{cond} is '1' when the condition number [30] of $\mathbf{PE}^+_{C*M'}$ is less than a user-

predefined threshold, u, and '0' otherwise. This logic function ensures the equation

system is well-conditioned. Minimization of Γ_i, yields the best SNR of the ith spectrum

for the fSLAM experiment, or indeed the SLAM experiment when the gradients in

$\mathbf{PE}^+_{C*M'}$ are limited to integer steps.

For comparison, the SNR of the CSI experiment is given by:

$$SNR_i^{CSI} = (L_i / M)^{1/2} \times \boldsymbol{\rho}^{CSI}_{C\times N}(i) / \sigma, \tag{2.13}$$

where L_i is the size of the ith compartment with average spectrum $\boldsymbol{\rho}_{C\times N}^{CSI}(i)$. Note that the quotient of Eqs. (2.11) and (2.13) approximates Eq. (2.1) for SLAM and fSLAM when multiplied by $\sqrt{M/M'}$ to account for scan-time differences.

2.2.5 fSLAM with minimum inter-compartmental leakage

So far we have assumed that every compartment is homogeneous. However, spectra in the CSI basis set that deviate from the compartmental averages can generate signals that propagate between and within each compartment following reconstruction. To optimize the fSLAM experiment with M' phase-encoding steps to suppress leakage, Eq. (2.6) is reformulated to separate the original $\boldsymbol{\rho}$ matrix into an average and an inhomogeneous part:

$$\begin{aligned} \mathbf{S}_{M'*N} &= \mathbf{PE}_{M'*M} \times \mathbf{b}_{M*M}^{-1} \times \mathbf{b}_{M*M} \times (\boldsymbol{\rho}_{M*N}^{avg} + \boldsymbol{\rho}_{M*N}^{inhom}) \\ &= \mathbf{PE}_{M'*M} \times \mathbf{b}_{M*M}^{-1} \times \mathbf{b}_{M*M} \times \boldsymbol{\rho}_{M*N}^{avg} + \mathbf{PE}_{M'*M} \times \boldsymbol{\rho}_{M*N}^{inhom} \end{aligned} \tag{2.14}$$

where each row in $\boldsymbol{\rho}_{M*N}^{avg}$ is an average spectrum of its compartment and each row in $\boldsymbol{\rho}_{M*N}^{inhom}$ is the deviation of the true spectrum from its compartmental average. For example, assume we have a 3-voxel compartment with single-point spectra with magnitudes [1.1, 1.0, 0.9]. The average spectrum in this compartment will be '1' and the inhomogeneity will be [0.1, 0, -0.1]. Note that by definition the inhomogeneity terms for the same compartment sum to zero.

On the right side of Eq. (2.14), the first part $(\mathbf{PE}_{M'*M} \times \mathbf{b}_{M*M}^{-1} \times \mathbf{b}_{M*M} \times \boldsymbol{\rho}_{M*N}^{avg})$ satisfies the ideal homogeneity assumption of SLAM, and the second part $(\mathbf{PE}_{M'*M} \times \boldsymbol{\rho}_{M*N}^{inhom})$ is the source of signal leakage and errors. The solution to Eq. (2.14) after dimensional reduction is:

$$\mathbf{PE}^+_{C*M'} \times \mathbf{S}_{M'*N} = \mathbf{\rho}^{avg}_{C*N} + \mathbf{PE}^+_{C*M'} \times \mathbf{PE}_{M'*M} \times \mathbf{\rho}^{inhom}_{M*N}. \qquad (2.15)$$

Clearly, we need to minimize $(\mathbf{PE}^+_{C*M'} \times \mathbf{PE}_{M'*M} \times \mathbf{\rho}^{inhom}_{M*N})$ to suppress leakage. In the

absence of control over $\mathbf{\rho}^{inhom}_{M*N}$, a reasonable strategy is to minimize the coefficients in

$\mathbf{PE}^l_{C*M} = \mathbf{PE}^+_{C*M'} \times \mathbf{PE}_{M'*M}$. Because the inhomogeneity terms in the same compartment

sum to zero, their mean can be subtracted. In the example above, if the three coefficients

corresponding to inhomogeneity [0.1, 0, -0.1] are [1/2, 1/3, 1/6], they will generate the

same errors as coefficients [1/6, 0, -1/6] after subtracting the mean value of 1/3. This

coefficient set has a smaller sum-of-the-squares and is not affected by differences in the

mean coefficient of each compartment.

Let $\mathbf{PE}^{ll}_{C*M}(i)$ denote the new matrix of coefficients that results from subtracting

the mean from $\mathbf{PE}^l_{C*M}(i)$, for each compartment. Then, to minimize the inter-

compartmental leakage into the i[th] compartment, we minimize the sum-of-the-squares of

the coefficients in $\mathbf{PE}^{ll}_{C*M}(i)$ that derive from outside of the i[th] compartment, analogous to

SLOOP [12]:

$$\phi_i = \sum_{\substack{j \neq i}}^{C} \sum_{\substack{m \in \text{compartment } j}}^{M} w_{ij} \times \left| \mathbf{PE}^{ll}_{C*M}(i,m) \right|^2. \qquad (2.16)$$

Here, w_{ij} is the weight of inter-compartment leakage from the j[th] compartment into the i[th]

compartment. The w_{ij} can reflect, for example, intrinsic differences in metabolite

concentrations between compartments.

2.2.6 Minimizing intra-compartmental errors in fSLAM

To minimize the errors due to inhomogeneity within the i^{th} compartment in the fSLAM experiment, we minimize the sum-of-the-squares of the coefficients that originate from inside of the i^{th} compartment itself:

$$\varphi_i = \sum_{m \in \text{compartment } i}^{M} w_{ii} \times \left| \mathbf{PE}_{C*M}^{ll}(i,m) \right|^2 , \tag{2.17}$$

where w_{ii} is the weight of intra-compartment error in the i^{th} compartment.

To perform a numerical optimization that minimizes both the inter- and intra-compartmental errors, in practice we minimize the cost-function:

$$\Lambda_i = (\phi_i + \varphi_i) / I_{cond} , \tag{2.18}$$

for the i^{th} compartment.

2.2.7 Summary of the fSLAM experiment

In summary, the fSLAM experiment is performed using the same Steps 1-5 as the SLAM protocol (Fig. 2.1) except that the phase-encoding gradients in Step 2 are obtained by minimizing either the SNR cost-function in Eq. (2.12) or the error cost-function in Eq. (2.18). In general, the different optimizations will result in different sets of phase-encoding gradients. If a gradient set optimized for both SNR and minimum error is being sought, minimization of the sum of the cost-functions in Eqs. (2.12) and (2.18) cannot be used because their scales differ. Instead, minimization of a weighted sum of the ratio of cost functions for fSLAM to those for SLAM can suffice. The choice of the weighting will depend on the application and error tolerance. The phase-encoding gradients in Step 2 are typically fractional.

27

Because $\mathbf{PE}_{C*M}^{ll}(i)$ is derived from **b** and therefore requires knowledge of compartment location and size, and the choice of gradients is not constrained to the CSI integer gradient steps, optimization and selection of the fSLAM gradient set must be performed scanner-side as part of the MRS set-up in order to achieve any SNR advantage compared to the summed CSI spectra from the same compartment volume.

2.2.8 Spatial response function

In accordance with Eqs. (2.9) and (2.12) of references [20] and [31] respectively, we define a spatial response function for the heart compartment corresponding to the row $\mathbf{PE}_{C*M'}^{+}(h)$ as:

$$SRF_h(x) = \sum_k \mathbf{PE}_{C*M'}^{+}(h) \cdot \exp(-i2\pi kx) .$$ (2.19)

The heart compartment spectrum is

$$\rho_h = \int_{FOV} SRF_h(x) \times f(x)dx ,$$ (2.20)

where $f(x)$ is the true continuous signal. $f(x)$ can be decomposed into signals from heart, $f_h(x)$, chest, $f_c(x)$, and "other", $f_r(x)$:

$$\rho_h = \int_{heart} SRF_h(x) \times f_h(x)dx + \int_{chest} SRF_h(x) \times f_c(x)dx + \int_{other} SRF_h(x) \times f_r(x)dx .$$ (2.21)

The second integral in Eq. (2.21) is the chest-to-heart leakage, $L_{c \to h}$.

We express $f_c(x)$ as a mean $\overline{f_c}$ plus an inhomogeneity $\Delta f_c(x)$. Then:

$$L_{c \to h} = \int\limits_{chest} SRF_h(x) \times [\overline{f_c} + \Delta f_c(x)]dx$$

$$= \overline{f_c} \times \int\limits_{chest} SRF_h(x)dx + \int\limits_{chest} SRF_h(x) \times \Delta f_c(x)dx$$

$$\leq \overline{f_c} \times \int\limits_{chest} SRF_h(x)dx + \int\limits_{chest} |SRF_h(x)| \times |\Delta f_c(x)|dx \qquad (2.22)$$

$$\leq \overline{f_c} \times \int\limits_{chest} SRF_h(x)dx + \max(|\Delta f_c(x)|) \int\limits_{chest} |SRF_h(x)|dx$$

The right hand side of the last line of Eq. (2.22) is the upper limit of the contamination of the heart spectrum from chest signal. The upper limit of leakage from the other compartments, $L_{r \to h}$, can be computed similarly.

2.3   Methods

2.3.1   Computer simulations

Computer simulations were performed to investigate the accuracy of SLAM applied to human cardiac [31]P MRS, where 1D CSI has served as a work-horse in our laboratory [24-28]. Three compartments were assumed: the heart, chest skeletal muscle, and 'other'. In practice, the 'other' compartment is needed because any signal generated outside of the designated compartments that is not assigned a compartment, will end up in the chest and heart, introducing errors depending on its magnitude. The chest and heart spectra are shown in Fig. 2.2(a, c). Signals are generated from these spectra with predefined compartment distributions using a 16-voxel 1-cm resolution 1D CSI model.

Monte Carlo simulations were done to quantify errors in SLAM arising from imperfections in the homogeneity assumption for this model. Based on experience[24-27], we assumed a metabolite-bearing chest muscle thickness of 2-3 voxels, a heart muscle thickness of 2-6 voxels, and zero or a single voxel separation between the chest and heart compartments, and zero signal in the 'other' compartment. This yielded 20 possible

anatomical combinations. To accommodate the combined effect of differences in concentration and surface coil sensitivity, two scenarios were investigated. In the first, we assumed a constant chest to heart signal ratio of 4. In the second, we assumed a chest PCr concentration 2.5 times higher than heart [28], and scaled the result by the experimental surface coil spatial sensitivity profile as shown in Fig. 2.3(a). A random inhomogeneity of ±15% (30% total) in the resultant signal was then simulated for both scenarios. The mean signal was determined for each compartment by adding signals from the corresponding voxels of the full CSI set to serve as a reference. Then, white noise was added such that the SNR in the heart compartment was 20. The FT of the data set was used to generate a set of time-domain CSI acquisitions from which $M' = 4$ central k-space acquisitions were selected. SLAM reconstruction from these 4 phase-encoding steps was implemented, and the percentage error relative to the reference CSI value was calculated. The mean error and the SD of the error were determined after 1000 Monte Carlo simulation runs.

Monte Carlo simulations were also performed to compare the sensitivity of SLAM with SLIM [10], with respect to registration errors. A 1D cardiac [31]P model with chest from -60 mm to -30 mm, heart from -30 mm to 10 mm, and a chest-to-heart signal ratio of 4 was assumed as in scenario-1, above (Fig. 2.3a). A random segmentation error between -2 mm and +2 mm was introduced at the edges of either compartment: (i) with the chest and heart stationary (no partial volume error); and (ii) with the chest and heart also moved by ±2 mm (partial volume error).The chest was constrained never to overlap the heart. Both SLAM and SLIM were simulated with four CSI phase-encodes from central k-space. SLIM reconstruction was performed as prescribed [10], by integrating

the phase-encoding coefficients over the 3-compartment model of heart, chest and 'other' and generating a 4x3 'G'-matrix [10]. The mean (±SD) % error between the reconstructed signal and the true or the CSI result was calculated for 1000 runs.

The SNR and the root-of-the-sum-of-the-squares of the inter- and intra-compartment errors, $\sqrt{\phi_i + \varphi_i}$, were computed for the model heart, assuming 3- and 4-voxel cardiac compartments and a 2-voxel chest compartment for both SLAM and fSLAM, and that both techniques yield the same compartmental average. The SNR was measured relative to the compartment average SNR of the 16-voxel 1D CSI (Eq. (2.13)), using the $M' = 3$ to 16 central k-space acquisitions for SLAM, and fractional (low k-space) phase-encodes for fSLAM. Optimization was performed using the simplex method implemented via the Matlab *"fminsearch"* routine (The MathWorks, Natick, MA) on a lap-top computer with $u = 50$ in Eqs. (2.12) and (2.18), and with all the leakage weighting factors, w_{ij} set to '1' in Eqs. (2.16) and (2.17).

SRF$_h$ was calculated from Eq. (2.19) for 4-step SLAM, 4-step fSLAM, 16-step CSI and 4-step CSI (zero-filled to 16 steps) for the 3-voxel chest/ 4-voxel heart model. The upper bound of chest contamination of the heart spectrum for the four cases was calculated from Eq. (2.22) assuming an effective chest to heart ratio of 4 and an intra-compartmental inhomogeneity of ±15% (30% total) for the chest.

2.3.2 Experiments

[31]P 1D CSI, SLAM, and fSLAM were implemented in a 3T Philips Achieva MRI/MRS system on phantoms, the human leg, and the human heart. The phantom studies were done with a 14-cm diameter single loop transmit/receive coil, and the human studies used a 17-cm/11-cm diameter dual loop transmit and a 8-cm diameter single loop

receive ^{31}P coil set described previously [29]. All human studies were approved by the Johns Hopkins Medicine Institutional Review Boards and all participants provided informed consent. The individual CSI spectra from all of the volume elements constituting each compartment were co-added post-acquisition for all comparisons of spectra from the equivalent volumes reconstructed using SLAM and fSLAM.

Phantom studies were performed on two standard Philips ^{31}P test disks 15-cm in diameter and 2.5-cm thick. One contained 300 mM H_3PO_2, the other had 300 mM H_3PO_4. A standard 1D CSI protocol using frequency-sweep-cycled (FSC) adiabatic half passage (AHP) pulses was applied (field-of-view, FOV = 160 mm; voxel/slice thickness, SL = 10 mm; repetition time, TR = 6 s; CSI phase-encoding steps, k = -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7; acquisition delay, 1.4ms) [29]. The SLAM protocol (Fig. 2.1) was then implemented with the same CSI parameters except for the phase-encoding gradients, which were reduced to a subset of 4 of the same steps (-2, -1, 0, 1). A 3-compartment model comprised of the two disks plus an 'other' compartment was assumed.

The leg was studied with the 300 mM H_3PO_4 disk phantom on top to create an additional compartment. 1D CSI was first performed with FSC AHP excitation (FOV = 160 mm; SL = 10 mm; TR = 8 s; phase-encoding steps, k = -8, -7, …, 7). This was followed by 3-compartment SLAM with the same total scan time and gradient-step increments but using only the 4 central k-space steps (-2, -1, 0, 1 repeated 4 times).

Human cardiac ^{31}P MRS studies comparing SLAM and CSI were performed on 8 normal volunteers and 16 patients with non-ischemic cardiomyopathy using the same protocol (FOV = 160 mm; SL = 10 mm; TR = 15.7 s, cardiac triggered). For each subject, CSI data reconstructed from all 16 phase-encoding steps, was compared with SLAM

reconstruction employing only the middle 4 phase-encoding steps of the same CSI data sets. This effectively reduced the scan time by 4-fold. The effect of using just 2 phase-encoding steps from central k-space corresponding to chest and heart compartments only, was also investigated. The resulting spectra were fit by the circle-fit method [32] to provide a quantitative comparison of PCr and γ-ATP peak areas measured by SLAM with those from conventional CSI (the localization and spectral analysis method are independent). Spectra were exponential-filtered (15-Hz line-broadening) and zero-filled 4 times to 2048 points.

The performance of fSLAM with respect to SNR and compartmental leakage was compared with that of CSI and SLAM in proactive cardiac ^{31}P MRS studies of 6 additional healthy volunteers. Sequentially, a first CSI, a SLAM, an fSLAM, and a repeat last CSI scan were acquired from each subject. CSI utilized the standard 16 phase-encoding steps from -8 to 7 (FOV = 160 mm; SL = 10 mm; TR = 15.7 s, cardiac triggered). SLAM used the same 4 middle k-space phase-encoding steps for each exam, repeated four times for the same total scan-time as CSI. fSLAM phase-encoding employed 4, typically-fractional gradient steps, specifically optimized for minimum compartmental leakage in the heart compartment for each volunteer, after manual segmentation of the scout MRI using the scanner's cursor function. As in the simulations, optimization was performed using Matlab on a lap-top computer at the scanner-side, with weighting factors set to unity. The four gradient values were manually entered as experimental parameters in the fSLAM pulse sequence on the scanner. The four steps were repeated four times for the same total scan-time as the CSI.
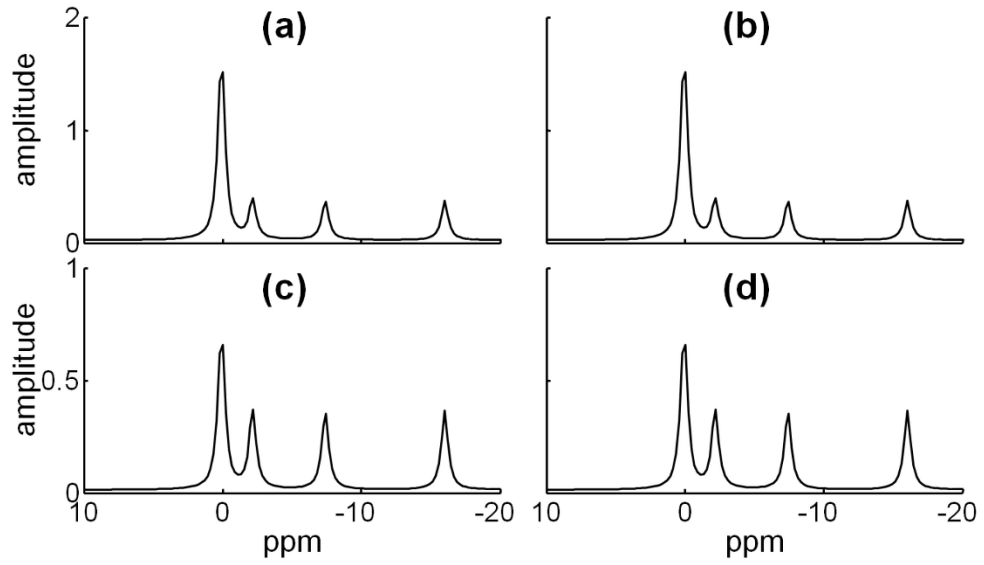
Figure 2.2: Simulated 16-step ^{31}P 1D CSI spectra of a model chest with 3 skeletal muscle voxels (a) and 4 heart voxels (c). The reconstructed SLAM chest (b) and heart (d) spectra are indistinguishable from the originals.
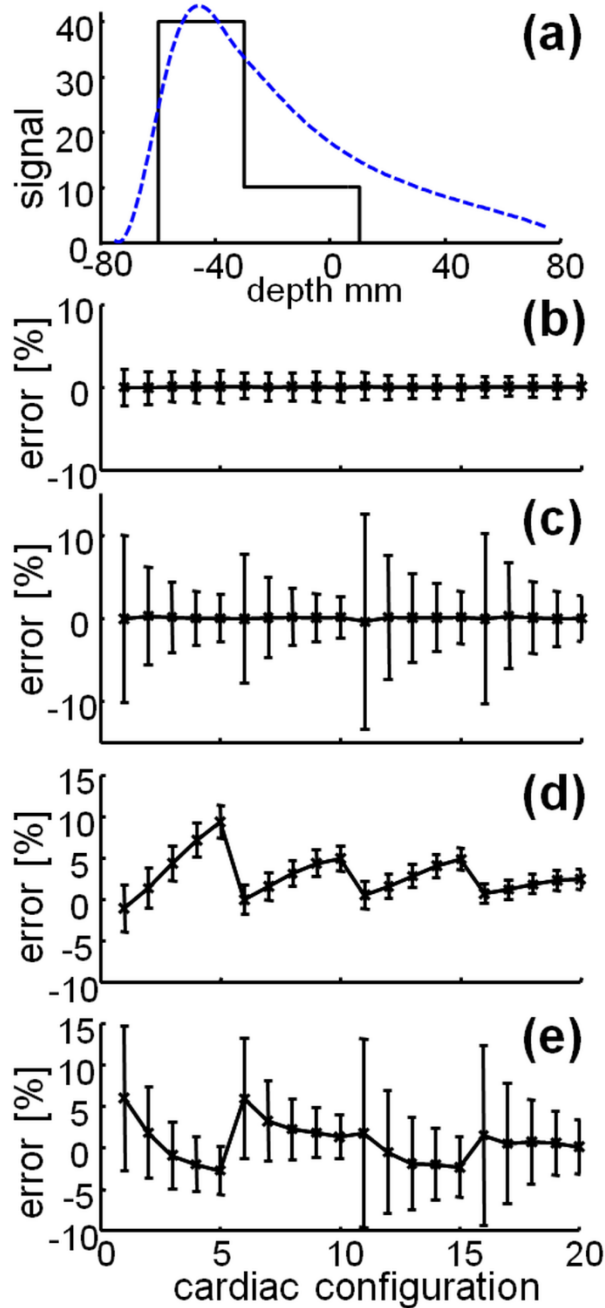
Figure 2.3: Cardiac model (a) and Monte Carlo simulation of the effect of noise and 30% (±15%) inhomogeneity on the accuracy of SLAM signal reconstruction vs. CSI (b-e). The chest-to-heart signal ratio is held constant at 4 in (b,c) depicted by the dark continuous curve in (a). In (d,e) the ratio is 2.5 scaled by the experimental surface coil sensitivity profile depicted by the blue dashed curve in (a). Errors are mean ±SD in the chest (b,d) and heart (c,e), calculated with cardiac SNR=20. The largest errors in the heart correspond to configurations #1 (2cm chest, 2cm heart, no separation between chest and heart), #6 (2cm chest, 2cm heart, 1cm gap), #11 (3cm, 2cm, 0cm), and #16 (3cm, 2cm, 1cm).

2.4   Results

2.4.1   Computer simulations

Fig. 2.2, shows that SLAM spectra of the chest and heart, reconstructed using only the three middle (k-space) phase-encoding steps of the original 16, are indistinguishable from the original simulated spectra in the absence of inhomogeneity or noise. The effect of adding noise and inhomogeneity on SLAM spectra reconstructed for a range of different chest-muscle and heart compartment distributions, is illustrated by the Monte Carlo simulations for both models of concentration and sensitivity variations in Fig. 2.3. These show that the accuracy of the reconstruction, as indexed by the mean of the error <10% for all chest/heart anatomical combinations. As might be expected, the higher the concentration or larger the compartment size, the smaller the error SD. For the heart, the simulations predict highest errors when the effective extent of the cardiac compartment is smallest.

The effect of small errors in the registration of compartments for CSI, SLAM and SLIM, as compared to the true value and to CSI, are summarized in Table 2.1. The Monte Carlo simulations show that small segmentation errors of just ±2 mm can introduce random errors approaching 10% for SLIM when the object is stationary, while SLAM is virtually unaffected and is less sensitive to partial volume errors. SLAM's relative insensitivity to small segmentation errors is critical for real applications since perfect segmentation is rarely possible in practice–especially in cardiac ^{31}P MRS.

36

Table 2.2: Monte Carlo analysis of the effect of ±2 mm misregistration on accuracy of cardiac PCr measurements for a 30 mm chest/40 mm heart model.

Simulation (1000 runs)	Error (mean ±SD), %		
	CSI	SLAM	SLIM
With model fixed on CSI grid, error vs. true	2.7 ± 0.0	0.7 ± 0.0	1.1 ± 8.4
With ±2 mm partial volume shift, error vs. CSI	0.0 ± 0.0	-2.2 ± 5.1	-1.1 ± 9.8
With ±2 mm partial volume shift, error vs. true	2.4 ± 5.2	0.3 ± 8.2	1.1 ± 8.7
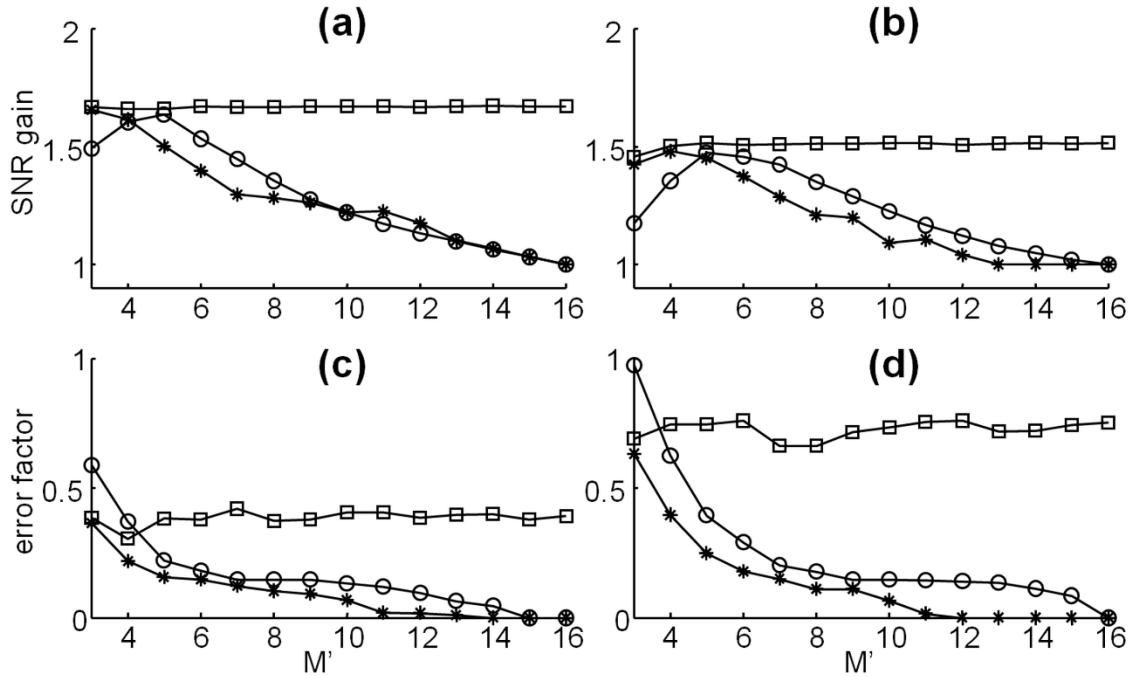
Figure 2.4: The SNR gain for the same volume (a, b), and the total inter- and intra-compartment error factor , $\sqrt{\phi_i + \varphi_i}$, (c, d) for SLAM and fSLAM in the heart as a function of the number of phase encodes, M', of the original M =16 that are allowed. For comparison, CSI has an SNR =1 with zero error assumed. Points depict results for three sets of gradients (square points, fSLAM with maximum SNR; stars, fSLAM with minimized inter/intra-compartmental errors; circles, SLAM). Here, (a) and (c) are for a 4-voxel thick heart; (b) and (d) are for a 3-voxel-thick heart compartment, all with a 2-voxel thick chest compartment.

37

The results of the analysis of SNR gain and the combined inter- and intra-compartment error factor, $\sqrt{\phi_i + \varphi_i}$, for SLAM and fSLAM, as compared with 16-voxel 1D CSI of the heart, are shown in Fig. 2.4 with 3- and 4-voxel thick cardiac compartments. The maximum SNR results from choosing the phase-encoding steps closest to central k-space. Because SLAM is confined to the CSI's integer phase-encoding set, its SNR advantage fades as more-and-more of the high k-space phase-encodes are used. Optimum SNR performance for SLAM occurs when the number of phase-encodes approximates the number of compartments, wherein its performance approximates that of fSLAM. Thus for SLAM, the best strategy is to choose the $M' \approx C$ non-equal CSI phase-encoding steps at or closest to the center of k-space, and repeat or average the acquisitions up to the allotted scan time, rather than add any higher k-space phase-encodes. On the other hand, fSLAM always achieves 1.5-1.8 times the SNR of standard CSI independent of the number of phase-encoding steps that are allowed. This reflects the fact that fSLAM is free to choose an array of fractional phase-encodes that all fall close to central k-space. For fSLAM, the additional phase-encodes offer the added benefit of reduced signal bleed (Fig. 2.4c, d). The errors, $\sqrt{\phi_i + \varphi_i}$, for fSLAM decay faster than SLAM as phase-encodes are added, indicating better error suppression, with larger compartments generating less error than smaller ones.

Fig. 2.5 plots SRF$_h$ for 16- and 4-step CSI, 4-step SLAM and 4-step fSLAM. It is important to recognize that the signal derives from the integral of the curve over each compartment, resulting in cancellation of signal outside the heart. When the chest signal is uniform, the cancellation is essentially perfect in the case of SLAM and fSLAM but not CSI (Table 2.2, first row). When the signal in the chest compartment varies by up to

30% peak-to-peak, the upper bound for contamination of the heart compartment from chest muscle rises to 12-14% for SLAM and fSLAM. This compares to 9% for 16-step CSI, while the 4-step CSI is basically unusable (Table 2.2, second row).
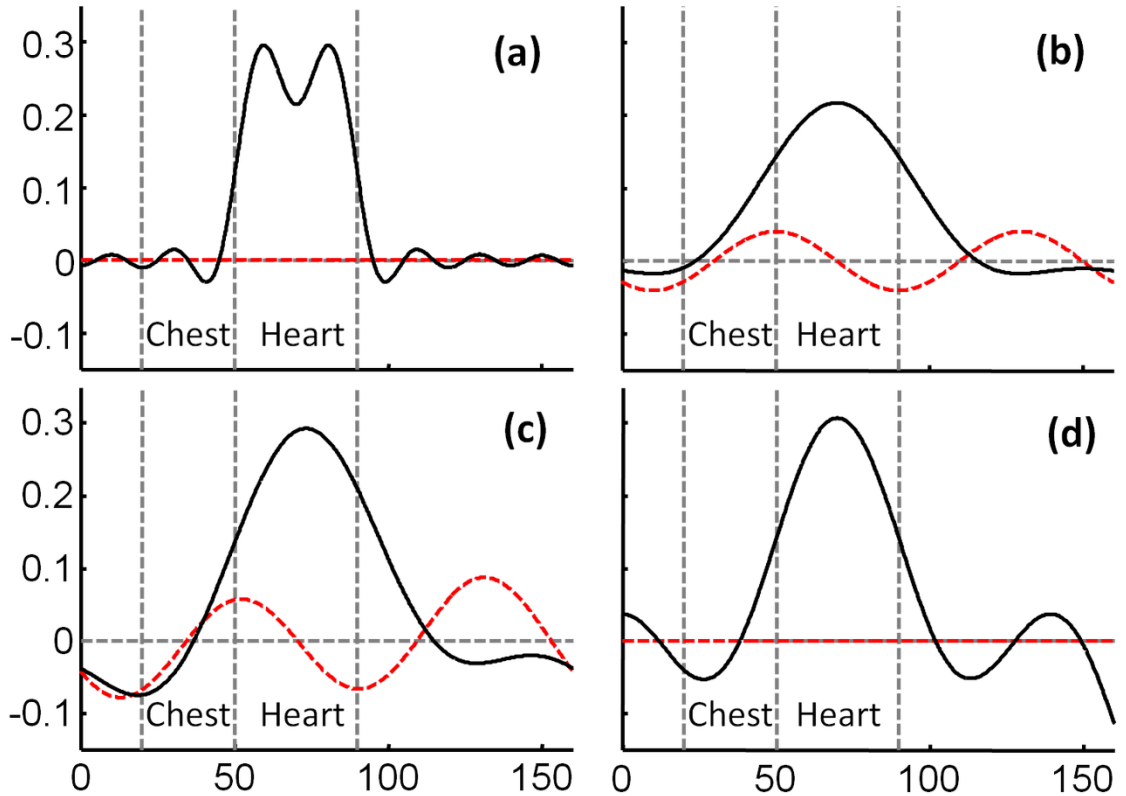
Figure 2.5: SRF_h for (a) 16- and (b) 4- step CSI (zero-filled to 16 steps), (c) 4-step SLAM and (d) 4-step fSLAM, computed for a model comprised of 3 chest voxels adjoining 4 heart voxels (black lines, real part; dashed red, imaginary component). Vertical dashed lines delineate the chest and heart compartments, as labeled. The signal contribution from each compartment derives from the integral of the curve over that compartment.

Table 2.3: Integral of SRF_h and upper bound of chest contamination for CSI, SLAM and fSLAM[a]

	16-step CSI	4-step CSI	SLAM	fSLAM
Integral of SRF_h over chest	0.0138	0.1654	0.0045	0.0073
Upper bound of $L_{c \to h}$ [b]	9.0%	77.3%	13.9%	12.0%

[a]Computed for 3-slice chest/ 4-slice heart model.
[b]Computed with chest/heart signal ratio of 4 and ±15% (total 30%) chest inhomogeneity.

39

2.4.2 Experiments

Spectra from the two-disk inorganic phosphate phantom reconstructed using CSI and SLAM are shown in Fig. 2.6. H_3PO_4 has a single ^{31}P peak at 2.9 ppm, while the H_3PO_2 resonance is a triplet centered about 13.5 ppm (coupling constant, 545 Hz), due to heteronuclear coupling with hydrogen. Despite the 4-fold reduction in scan-time, the SLAM spectra from the two disks are very similar to the summed CSI spectrum from the same compartment volumes, with negligible leakage consistent with the simulations (Fig. 2.3).

Spectra from the same-sized leg compartment obtained by CSI (averaging n=6 voxels) and SLAM (same volume) are presented in Fig. 2.7(a) normalized to constant noise. The SLAM spectrum has 2.1 times better SNR than CSI, and shows negligible signal contamination or bleed from the H_3PO_4 phantom positioned above the leg. Fig. 2.7(b) shows ^{31}P heart spectra from a 16-step 1D CSI (averaging 4 voxels), 4-step SLAM and 2-step SLAM from the same volume. The baseline roll is due to the acquisition delay for the phase-encoding gradient. Again, negligible bleed is evident in the 4-step SLAM spectrum, either from adjacent chest skeletal muscle or from an embedded coil marker (at ~23 ppm). Importantly, while SLAM faithfully reproduces the CSI heart spectrum and the SNR of CSI and SLAM are comparable, the SLAM spectrum was acquired 4-times faster. Even with only two steps, the SLAM reconstruction remains surprisingly good as shown in Fig. 2.7(c). With 2 phase encodes, just two signal-generating compartments, chest muscle and heart, are allowed, resulting in some signal bleed from the external coil marker in a spectrum acquired 8-times faster than the CSI standard.

Table 2.4: The cardiac ^{31}P MRS SNR of PCr for the same cardiac voxel volumes and scan time using CSI, SLAM, and fSLAM in n = 6 healthy volunteers (left-to-right).

	Vol 1	Vol 2	Vol 3	Vol 4	Vol 5	Vol 6	Average
CSI	61	30	33	18	18	26	31 ± 16
SLAM	86	53	41	30	20	39	45 ± 23^{a}
fSLAM	71	45	44	32	21	31	41 ± 17^{b}

[a]p<0.01 vs. CSI, paired t-test. [b]p<0.002 vs. CSI, paired t-test.
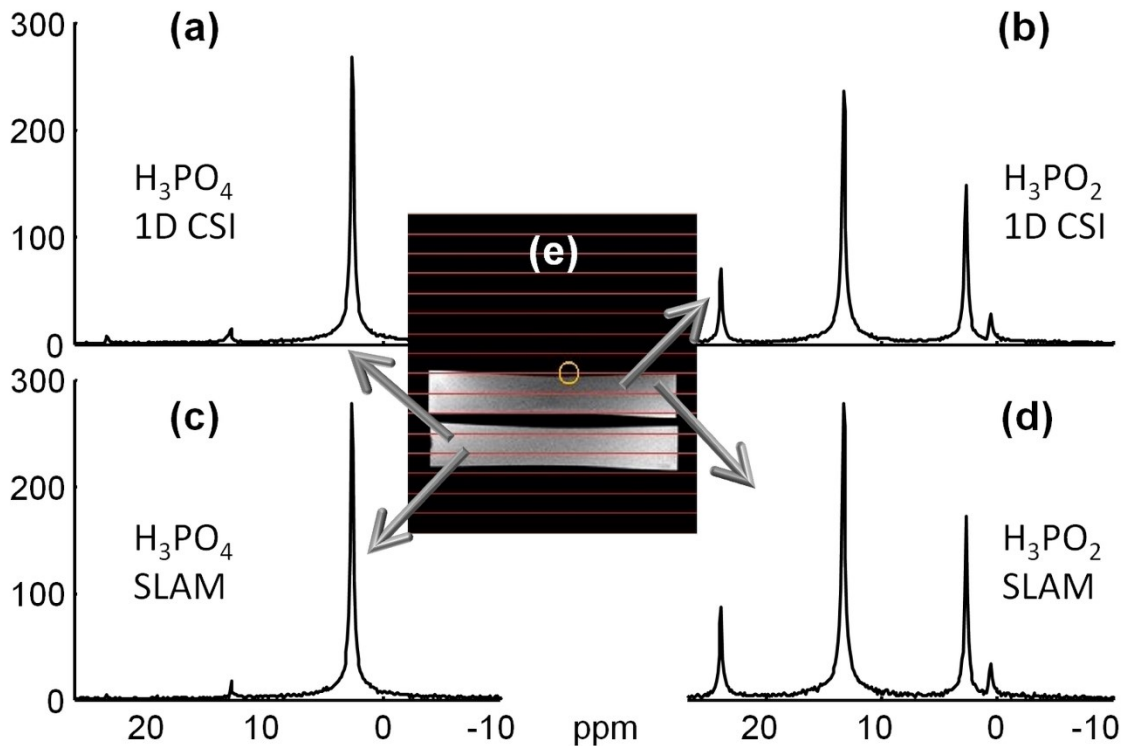
Figure 2.6: CSI and SLAM spectra reconstructed from the standard Philips ^{31}P test phantom comprised of a H_3PO_4 disk on the bottom (a, c), and a H_3PO_2 disk on top (b, d), as shown in the image (e). The CSI spectra (a, b) are the sum of the spectra from the voxels (red horizontal lines) containing the disks and were acquired with 16 phase-encoding gradients (-8…+7). The SLAM spectra (c, d) were acquired 4-times faster with just 4 phase-encodes (-1,-2,0,1). The SNR for the CSI spectra are 660 (a) and 638 (b), compared to 528 (c) and 482 (d) for SLAM. The signal at ~0 ppm is a contaminant present only in the H_3PO_2 disk.
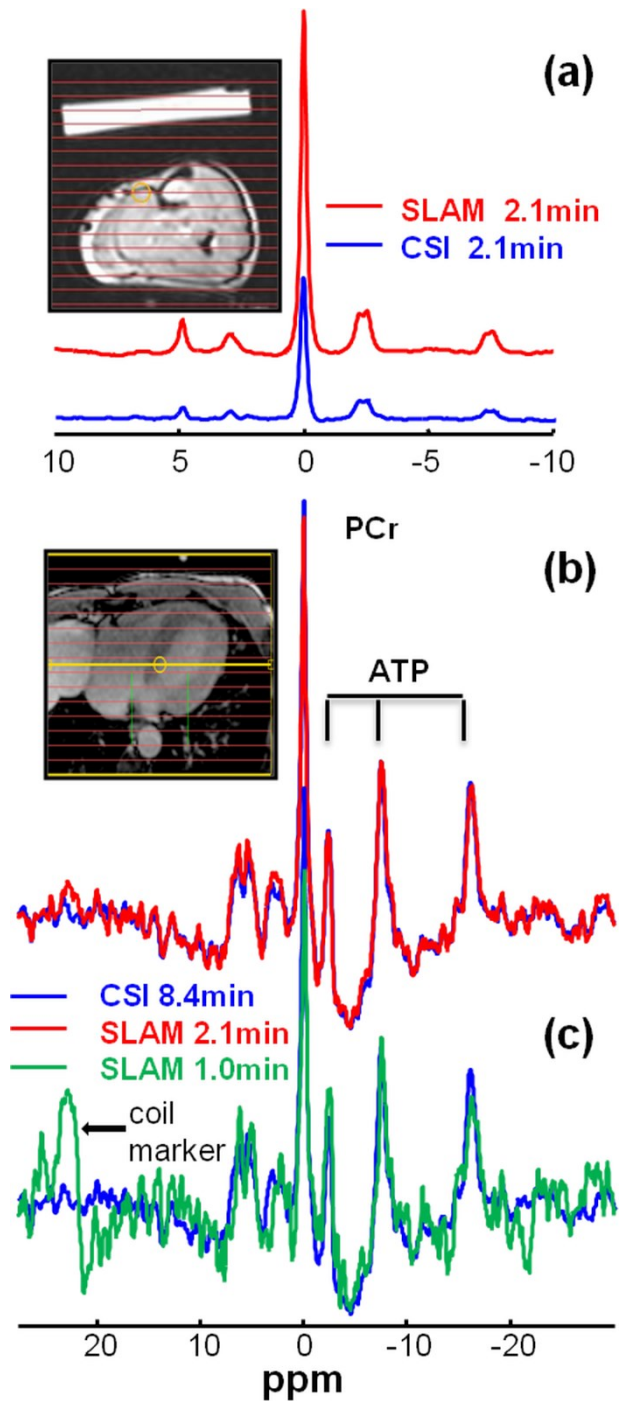
Figure 2.7: (a) Human leg ^{31}P spectrum acquired by SLAM (red, top) and CSI (blue, lower) from the same 6-voxel volume in the same scan time (2.1 min). (b) ^{31}P spectra acquired from a normal human heart from the same 4-voxel volume, using 1D CSI (blue) in 8.4 min, and SLAM spectra reconstructed with a subset of 4 central k-space phase-encodes and a 3-compartment model, (red spectrum). (c) Spectra acquired with just two phase-encodes (green) and a 2-compartment model (chest and heart). The effective SLAM acquisition times were 1/4th and 1/8th of CSI.

The comparison of fitting results from the cardiac 4-step SLAM and 16-step CSI spectra from 8 healthy subjects and 16 patients are presented in Fig. 2.8. In these data sets, the ratio of PCr signal in chest to that in heart compartments was at or below ~5. The PCr and γ-ATP peak areas from the SLAM reconstruction agree with those from CSI reconstruction (Fig. 2.8a, b). The myocardial PCr/ATP ratio for the pooled patients and healthy subjects was the same (1.94 ± 0.60 in CSI vs. 1.90 ± 0.67 in SLAM; p=0.42, paired t-test), consistent with negligible contamination from chest muscle with its much higher PCr/ATP ratio of ~4 [26]. In addition, Fig. 2.8(c) shows that the total PCr in the chest-plus-heart compartments measured by SLAM, is the same as that measured by CSI, indicating that the total signal is conserved. Because the fraction of cardiac PCr to the total chest-plus-heart PCr measured by SLAM is also equal to that measured by CSI (Fig. 8d), the contamination of heart spectra from chest muscle in SLAM is indistinguishable from that in CSI. This result is consistent with Table 2.2. Importantly, all these SLAM results correspond to acquisitions effectively taking 1/4[th] of the scan-time of CSI.

Fig. 2.9 compares CSI, SLAM and fSLAM [31]P cardiac spectra proactively acquired in the same total scan-time from the same volume size in the same healthy volunteer. The time taken to implement fSLAM scanner-side was 1-2 min to manually segment the scout MRI, plus several seconds to optimize the gradient set on the lap-top computer. The SLAM and fSLAM spectra both have higher SNR than CSI from the same volume, while a possible bleed signal from the coil marker in the SLAM spectrum is eliminated by either designating a separate compartment for the coil phantom (Fig. 2.9b), or using fSLAM (Fig. 2.9c). There is no obvious chest muscle contamination of either the SLAM or fSLAM spectra. The PCr/ATP ratio for SLAM and fSLAM was not

significantly different from that measured in either the first or the repeated last CSI scans, and the absolute metabolite signal levels do not change (Fig. 2.9d).

Table 2.3 lists the SNR of human cardiac PCr in same-sized voxels for CSI, SLAM and fSLAM in 6 volunteers acquired in the same scan-time. The mean SNR improvement for SLAM vs. CSI for the six studies is 1.42 ± 0.23. The mean SNR improvement for fSLAM vs. CSI for the six studies is 1.34 ± 0.19. According to Eq. (2.1), this SNR gain would be consistent with a cardiac compartment equivalent to two of the 1-cm CSI voxels even though the reconstruction assumed a 4-voxel cardiac compartment. This likely reflects the combined effect of the decline in surface coil sensitivity with depth, and the 1-2 cm thickness of the anterior myocardial wall.
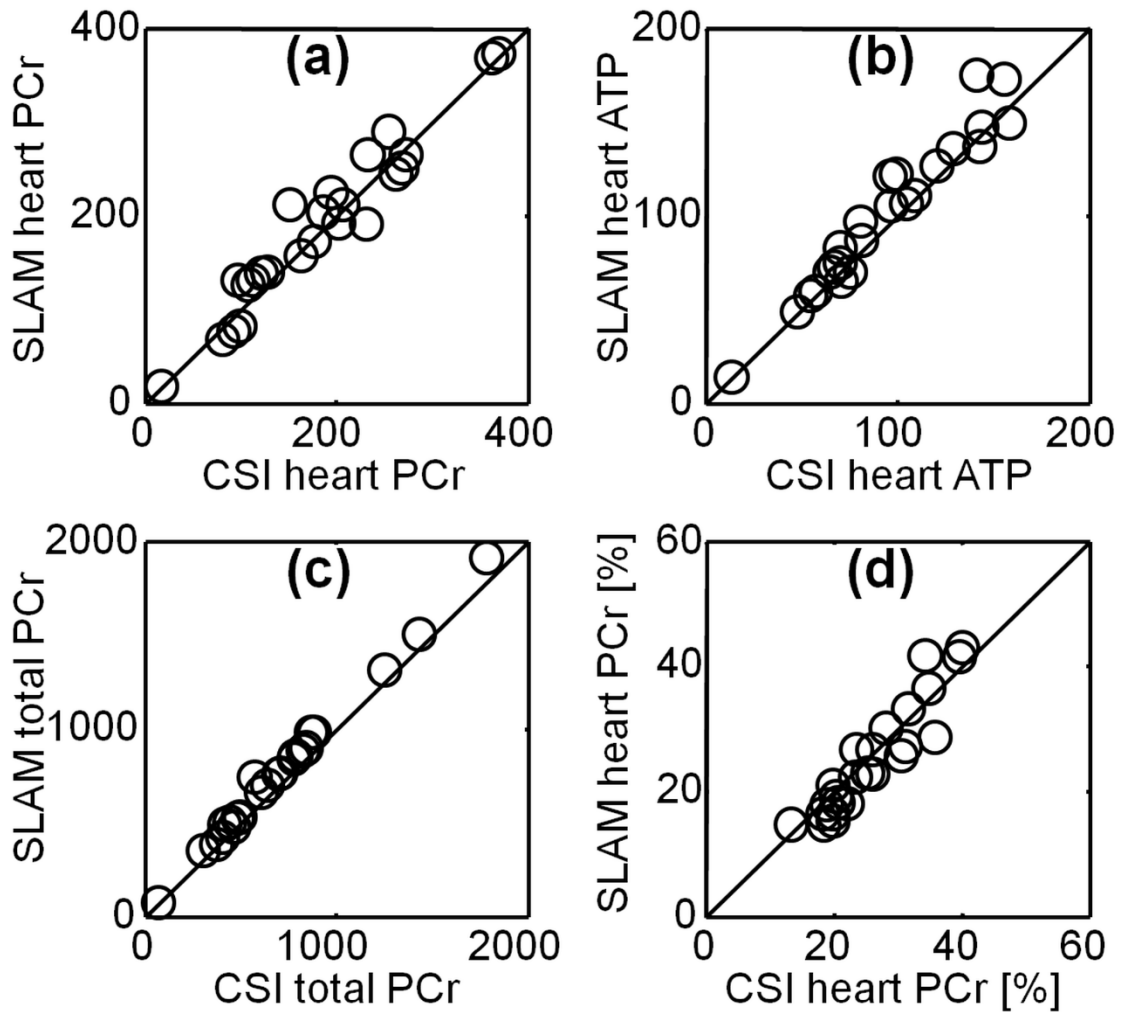
Figure 2.8: Fitting results reconstructed by SLAM from a subset of 4 of the 16 CSI phase encoding steps acquired from the 24 heart patients and control subjects, as compared to the CSI results. (a) PCr and (b) γ-ATP peak areas quantified in the cardiac compartment. (c) The total PCr from both heart and chest compartments. (d) The ratio of heart PCr to the total PCr from both chest and heart compartments. The correlation coefficients are r >0.97 in all cases, and the solid line is the identity line.
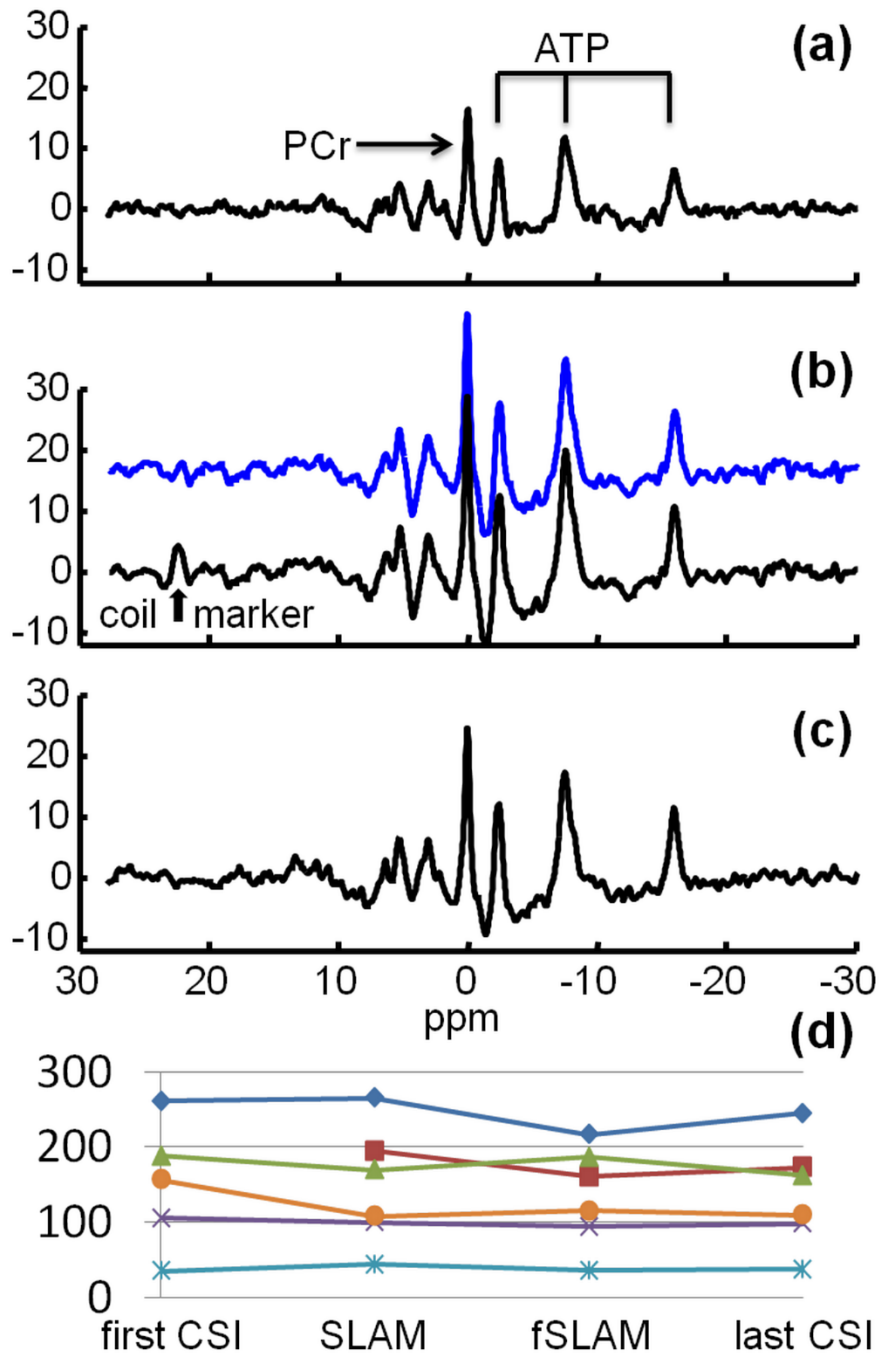
Figure 2.9: (a) CSI, (b) 4(blue)- and 3(black)-compartment SLAM and (c) error-minimized 3-compartment fSLAM spectra, all normalized to constant noise on the same volunteer with the same total scan time and total voxel volume. Gradient encoding steps of -8 to +7 (integer) were used for standard CSI; integer steps -2, -1, 0, 1 repeated 4 times were used for SLAM; and fSLAM used non-integer steps -2.13, -0.73, +0.73, +2.13 repeated 4 times. (d) Cardiac PCr peak area from proactive [31]P MRS studies of all 6 subjects in the first CSI, the SLAM, the fSLAM and the repeated CSI scan (no significant difference between exams at paired t-testing; lines connect measurements from the same subjects).

2.5   Discussion

Single voxel methods such as PRESS [33], STEAM [34] or ISIS [35] are good localization choices for performing MRS of a single compartment, but do not offer optimum SNR for a fixed scan time for MRS of multiple compartments. In addition, their sensitivity to relaxation effects (both T_1 and T_2) and motion, presents real problems for quantification, especially in ^{31}P MRS [8, 33]. CSI, being a simple pulse-and-acquire experiment that collects all-of-the-signal from all-of-the-sample, all-of-the-time, currently offers the cleanest approach to quantitative MRS, with potentially the highest SNR efficiency. It is however, limited by the minimum scan-time required to encode the entire sensitive volume or FOV of the detector coil. This can limit the direct translation of SNR gains, such as those afforded by higher B_0 magnetic field strengths, to reductions in scan-time. In addition, the highest SNR efficiency of CSI is only realized when the spatial resolution imposed at the time of acquisition, matches the desired compartment size [8]. Unfortunately, CSI's spatial resolution is usually set not by the size of the desired compartment, but by the geometry of the tissue that it must be distinguished from (e.g., the chest in heart or liver studies, the scalp in brain studies).

Alternative approaches that localize spectra to pre-selected compartments based on anatomical MRI information, are not new. The SLIM, GSLIM and SLOOP methods were originally proposed some 20 years ago [10-12], but see little use today compared to CSI or even PRESS, STEAM or ISIS. When SLIM, GSLIM and SLOOP are applied to regular CSI acquisitions, without pro-active implementation or gradient selection criteria that place a premium on SNR–as is most often the case [10, 11, 13-23], they cannot deliver the highest SNR achieved by matching the resolution to the compartment, *a priori*.

Although not previously documented, matching the SLIM, GSLIM and SLOOP compartments *a priori* would yield significantly higher SNR than that gained from summing signals from the constituent CSI voxels, by a factor of ~g (Eq. 2.1). Similarly, a many-fold speed-up in the minimum CSI scan-time could result if the phase-encoding gradient set were cut.

Here for the first time we have exploited differences in volume sizes between desired MRS compartments and CSI resolution, to realize and document a g-fold SNR gain consistent with Eq. (2.1), using a new MRS localization method, SLAM. SLAM differs from SLIM, GSLIM and SLOOP in both the pulse sequence that is applied, and in MRS reconstruction. Simply put and unlike other methods, the SLAM pulse sequence is based on a CSI sequence from which essentially all of the high-order gradient phase-encoding steps are eliminated except for the C phase-encoding steps closest to central k-space. Because the CSI gradient set is discretized, this means that the only *a priori* information needed to run the sequence is the number C, which is generally fixed for a given study protocol. Compared to pro-active implementation of SLOOP [12], this has the advantage of avoiding image-guided gradient optimization, prescription, and implementation at the scanner-side prior to acquisition. On the other hand, SLIM and GSLIM utilize standard CSI sequences [10, 11, 13, 18, 19].

Like prior methods, reconstruction of SLAM spectra *does* require a scout MRI to identify and segment the compartments which are assumed uniform. However, SLAM reconstruction *differs* from SLIM, GSLIM and SLOOP in that it solves a set of C linear simultaneous equations by eliminating un-needed phase-encoding steps from the standard CSI algorithm. SLAM aims to generate spectra that are at best equal to the compart-

mental average CSI spectra, whereas SLIM, GSLIM and SLOOP use MRI-based constrained reconstruction or SRF optimization to obtain optimally-localized compartment spectra. Because of the relatively coarse resolution of CSI, this renders SLAM relatively insensitive to registration errors in segmenting the compartments compared to SLIM for example (Table 2.1), where problems were noted previously [18, 36].

With SLAM, we demonstrate many-fold reductions in the minimum scan-time compared to CSI in theory (Fig. 2.4) and in practice (Fig. 2.7, 2.8), and substantial SNR gains in human *in vivo* studies on a standard clinical MRI/MRS scanner at 3T (Fig. 2.9, Table 2.3). Importantly, in 1D [31]P human cardiac applications, SLAM delivers qualitative and quantitative results (Figs. 2.8 and 2.9) that are practically indistinguishable from results obtained from conventional CSI, other than being 4-times faster or higher in SNR (Figs 2.2, 2.6 – 2.9). Even so, significant inter-compartmental contamination may arise when signals from adjacent compartments differ greatly or are not segmented. This can occur in [31]P MRS heart studies, for example, when chest skeletal muscle compartment signals are many-fold higher (e.g. >5-fold) than cardiac signals due to the higher muscle metabolite concentrations, and/or its thickness, and/or proximity to surface coil detectors with nonuniform sensitivity. Conventional CSI, used here as a standard, is not immune from this problem [37] (Table 2.2). Despite the uniform compartment assumption, both the numerical results (Fig. 2.3, Tables 2.1, 2.2) and the experiments (Figs. 2.6 – 2.9) suggest that SLAM is relatively robust to the variations in signal that arise in practical applications such as cardiac surface-coil MRS.

The SLAM acquisition pulse sequence with integer k-space phase-encodes was surprisingly simple to implement, at least for the 1D case. For our cardiac ^{31}P MRS studies, we chose the same 4 central gradient steps to provide a fixed SLAM acquisition sequence suitable for up to 4 compartments, extracted from a standard 16-step CSI sequence with the other 12 steps discarded. Further reductions in C and the number of phase-encodes, to 2 for example (Fig. 2.7c), risk leakage from unaccounted-for signal sources that lie outside of any designated compartment. This may be tolerable if the leakage does not interfere with the spectral region of interest. For validation or test purposes, SLAM can be performed retroactively on raw data sets that are accompanied by a scout MRI, simply by applying the algorithm to a subset of frames in each CSI data set. The result can be compared with the summed CSI from the same-sized compartments analogous to Fig. 2.8.

fSLAM extends SLAM by removing the limitation that the phase-encodes be selected from the set of integer-stepped CSI gradients. Instead, they are adjusted to minimize leakage or errors due to inhomogeneity and/or maximize SNR. We observed that maximizing SNR alone can produce unacceptable error if the clustering of phase-encodes at the center of k-space is unchecked (Fig. 2.4). Minimization of inter- and intra-compartmental errors alone yields acceptable results, albeit at the expense of a small reduction in SNR (Fig. 2.9c; Table 2.3). Thus, inter- and intra-compartmental error was substantially eliminated with fSLAM, also using only four phase-encoding gradient steps. When SNR is low also, inter-compartmental leakage could become problematic relative to the compartment signal. Adjustment of the weighting factors in Eq. (2.18) from the values of unity used herein may help attenuate bleed from specific adjacent

compartments, depending on the particular application. The gradient optimization in fSLAM derives from tracking errors through the reconstruction process and includes those due to both inter- and intra-compartmental signal inhomogeneity. This differs from SLOOP's use of the SRF to minimize only the inter-compartmental leakage [12], while SLIM and GSLIM do not use optimized gradients.

Note also that the SRF is not global but is specific to the cardiac model. Inter- and intra-compartmental leakage occurs only when the integral over the entire compartment is non-zero or in the presence of significant heterogeneity. Compartmental segmentation in SLAM ensures that the integral of the SRF vanishes over other compartments, while fSLAM minimizes the effect of heterogeneity within the compartment of interest as well. Ultimately however, the spatial responses for SLAM and fSLAM and their compartmental contamination are fully characterized by determining the accuracy of the solutions and leakage errors, for which CSI is used as the standard in the current work (Figs 2.4, 2.6 – 2.9, Tables 2.1, 2.2).

In conclusion, the SLAM and fSLAM methods yield spectra comparable to the average of same-sized CSI compartments but with large scan-time reductions, SNR gains, and manageable, if not insignificant, bleed artifacts. The SNR gains predicted by Eq. (2.1) will be moderated in practice by the depth-dependence of the surface coil sensitivity, as well as the actual metabolite distribution (in our case, the myocardial wall thickness). Independent of the SNR gain, SLAM and fSLAM reduce the minimum scan-time required for localization from M acquisitions in CSI, to C or $M' <<M$. We believe that this efficiency advantage alone can dramatically reduce MRS scan-times for patient MRS studies employing CSI in global disease such as cardiomyopathies [22-27], large lesions,

or where single voxel methods are limited by relaxation, motion or other considerations

[28]. In addition, the significant reductions in minimum scan time provided by the SLAM

and fSLAM methods compared to CSI, provides a practical pathway for translating the

higher SNR afforded by increases in magnetic field strength, into faster MRS exams.

Chapter 3 : Sensitivity-encoded SLAM in higher dimensions

3.1   Introduction

Spatially localized magnetic resonance spectroscopy (MRS) has provided valuable insight into many normal and diseased human conditions [27, 38-47]. Compared to single voxel MRS techniques, such as PRESS [48], STEAM [49] and ISIS [35], the standard multi-voxel chemical shift imaging (CSI) method [4] has the advantages of accessing multiple regions simultaneously with higher signal-to-noise ratio (SNR) per unit scan-time. However, the clinical application of CSI is limited by the long scan time required to apply a complete set of phase-encoded acquisitions before the individual voxel spectra can be reconstructed. In addition, due to the way that SNR adds in magnetic resonance imaging (MRI) and MRS, SNR lost by setting the voxel size smaller than needed during acquisition, cannot be entirely restored post-acquisition [8, 50].

Many methods have been proposed to speed-up CSI acquisitions. These include methods that combine spectral and spatial encoding [51, 52]; and those that actually reduce the amount of image k-space sampled by gradient-encoding [6, 50, 53-55]. In general, the latter fall into three categories: (a) methods employing parallel imaging [6, 53-55]; (b) those incorporating prior knowledge based on scout MRI [10, 14, 15, 50, 56-58]; and (c) those using compressed sensing and sparse reconstruction [59-62]. Of these, only Category (a) methods are currently in widespread use in clinical MRI/MRS scanners. Of Category (b), SLIM [10], GSLIM [56] and SLOOP [57] are strictly speaking, non-CSI methods that in theory offer superior resolution to CSI by eliminating Fourier bleeding. Indeed, this has been the focus of their *in vivo* applications to date, which have been limited to retroactively acquired CSI data [18, 20, 21, 23, 63, 64]. Thus, although SLIM

and SLOOP could provide faster scanning than CSI [10, 57, 58], to our knowledge no speed advantage has been demonstrated proactively *in vivo* or in humans. Category (c) methods are relatively new, and have seen even less use, perhaps reflecting the low SNR of most non-hyperpolarized MRS applications, which is not ideal for compressed sensing.

In Chapter 2, we introduced another Category (b) MRS localization method, "spectroscopy with linear algebraic modeling" or SLAM, wherein compartmental-average spectra are acquired using a greatly reduced CSI phase-encoding gradient set selected from central *k*-space where SNR is highest. The spectra are reconstructed from C compartments segmented from scout MRI, which are incorporated into the standard CSI model using an auxiliary "b" matrix. SLAM was demonstrated on both retroactively and proactively acquired one-dimensional (1D) phosphorus (^{31}P) human cardiac CSI data, yielding either a 4 to 8-fold acceleration in scan-time with the same quantitative results, or a ~40% SNR improvement for the same scan-time, as compared to our standard protocol [27, 39, 41].

In this chapter, SLAM is extended to two- (2D) and three dimensions (3D), and in addition, combined with parallel imaging techniques, specifically SENSE [6], to achieve dramatic speedup factors of 5-120 compared to CSI and SENSE CSI [53]. A modified SLAM reconstruction algorithm is introduced that improves accuracy by reducing the method's sensitivity to signal inhomogeneity within compartments. Additional improvements are provided to incorporate spatial and temporal main (B_0) and RF (B_1) field inhomogeneity terms, including eddy-current correction. These advances are implemented on 2D and multi-slice proton (^{1}H) MRS studies of the brains of healthy subjects and patients with tumors, both retroactively and proactively. Brain

compartmental average metabolite levels and ratios from CSI and SENSE CSI are

determined and quantitatively compared with those from corresponding high-speed

SLAM spectra. Finally, 3D SLAM is applied to ^{31}P MRS in a phantom and in human

heart, with speedup factors of 100 and 7, respectively.

3.2   Theory

The conventional CSI [4] reconstruction can be cast as a linear equation:

$$\mathbf{s}_{M*N} = \mathbf{PE}_{M*M} \times \boldsymbol{\rho}_{M*N} , \tag{3.1}$$

where $\mathbf{s}$ is the known vectorized signal matrix, $\mathbf{PE}$ is the phase-encoding operator, and

$\boldsymbol{\rho}$ is the unknown vectorized spectral matrix. For 1D CSI, $\mathbf{PE}$ is simply a discrete

Fourier transform (DFT) operator. For 2D or 3D CSI, $\mathbf{PE}$ is the Kronecker product [65]

of double or triple serial DFT operators, respectively. M is the total number of phase-

encoding steps or spatial voxels, and N is the number of chemical shift domain data

points.

When sensitivity encoding [6, 53, 66] is used, Eq. (3.1) is rewritten as:

$$\mathbf{s}_{M'*N} = \mathbf{E}_{M'*M} \times \boldsymbol{\rho}_{M*N} , \tag{3.2}$$

where $\mathbf{E}$ is the combined phase-encoding and sensitivity-encoding operator, and M'

denotes the product of the number of coil elements, N_c, and the (reduced) number of

phase-encoding steps, M_R $(=M/R$, where R as the SENSE acceleration factor). While

defined in Ref. [66], $\mathbf{E}$ can be constructed by stacking the product of $\mathbf{PE}$ with the

sensitivity encoding matrix, $\mathbf{SE,}$ of each coil element, as

$$\mathbf{E} = \begin{bmatrix} \mathbf{PE}_{M'*M} \times \mathbf{SE}^1_{M*M} \\ \mathbf{PE}_{M'*M} \times \mathbf{SE}^2_{M*M} \\ \vdots \\ \mathbf{PE}_{M'*M} \times \mathbf{SE}^{Nc}_{M*M} \end{bmatrix} , \tag{3.3}$$

where 1, 2, ..., N_c index each coil element. Furthermore, as described in Ref. [66] for SNR optimization, "pre-whitening" can be done to both sides of Eq. (3.2) by multiplying $(\mathbf{L}^{-1} \otimes \mathbf{I})_{M'*M'}$, where $\mathbf{I}_{Nc*Nc}$ is obtained from a Cholesky decomposition of the noise covariance matrix [66]; $\mathbf{I}_{M_R*M_R}$ is an identity matrix; and $\otimes$ is the Kronecker operator [65].

3.2.1 SLAM localization with prior knowledge

For simplicity, Eq. (3.2) is used throughout to represent both conventional CSI and the pre-whitened SENSE CSI reconstruction, the latter differentiated by the "SENSE" label. Introducing an auxiliary matrix, $\mathbf{b}$, containing the spatial information defining the C compartments segmented from MRI, results in:

$$\mathbf{s}_{M'*N} = \mathbf{E}_{M'*M} \times \mathbf{b}_{M*M}^{-1} \times \mathbf{b}_{M*M} \times \boldsymbol{\rho}_{M*N} . \tag{3.4}$$

As described in Chapter 2, $\mathbf{b}$ is composed by adding "-1" elements into C columns of an identity matrix. Note that the first dimension of the $\boldsymbol{\rho}$ matrix carries ordered spatial information for all of the voxels. Accordingly, the location of each of the C columns corresponds to the first voxel of each of the C compartments. The "-1" elements are located in each of the C columns after the first voxel, and correspond to all the rest of the voxels in each compartment. These elements are used to eliminate hypothetically identical rows in the $\boldsymbol{\rho}$ matrix in accordance with the compartment model.

Assuming that the individual CSI spectra in each of the C compartments are identical, dimensional reduction of Eq. (3.4) then leads to:

$$\mathbf{s}_{M'*N} = \mathbf{E}_{M'*M} \times \mathbf{b}_{M*C}^{r} \times \boldsymbol{\rho}_{C*N}^{r} , \tag{3.5}$$

where $\boldsymbol{\rho}^r_{C*N}$ is obtained from retaining the C non-eliminated rows in $\boldsymbol{\rho}_{M*N}$, which

correspond to the spectra of the C first voxels in the C compartments, respectively. $\mathbf{b}^r_{M*C}$

is obtained by retaining the C columns in $\mathbf{b}^{-1}_{M*M}$ corresponding to the C non-eliminated

rows.

3.2.2 Algorithms for SLAM and SENSE SLAM reconstruction

Two algorithms are used to reconstruct SLAM or SENSE SLAM spectra. The

first is the same one described in Chapter 2:

$$\boldsymbol{\rho}^r_{C*N} = (\mathbf{E}_{M'*M} \times \mathbf{b}^r_{M*C})^+ \times \mathbf{s}_{M'*N}, \tag{3.6}$$

where "+" denotes the Moore-Penrose pseudo-inverse when $M' > C,$ or the inverse when

$M' = C$. The second, slightly different, algorithm is denoted with asterisks as SLAM* or

SENSE SLAM*:

$$\boldsymbol{\rho}^r_{C*N} = (\mathbf{b}^r_{M*C})^+ \times (\mathbf{E}_{M'*M})^+ \times \mathbf{s}_{M'*N}. \tag{3.7}$$

Both algorithms require that $M' \geq C$, which is easily fulfilled in practice, e.g. $C = 3$

for cardiac spectroscopy [50], or $C = 4$ or 5 for SLAM MRS of brain, as exemplified later.

With conventional (Eq. 3.6) SLAM, $M' \leq M$ always, and typically $M' << M$, which means

that $(\mathbf{E}_{M'*M})^+$ is generally under-determined. Conversely, for SENSE SLAM* (Eq. 3.7),

$M' = N_c M/R$ could easily exceed M making $(\mathbf{E}_{M'*M})^+$ over-determined (e.g., with a

combination of an $N_c = 32$ element coil and an acceleration factor $R = 16$). In any case,

numeric regularization is recommended, especially where SENSE reconstruction is

involved and SNR is low. Here, a truncated singular value decomposition (TSVD) [67]

method is utilized wherein values below, for example 2% of the maximum, are discarded

to ensure that the condition number [30] is not greater than 50. In practice, the level of

numeric regularization may be optimized for non-ideal/low SNR data, by increasing the level of numeric regularization until the results become stable.

3.2.3 The discrete spatial response function

The continuous spatial response function (SRF) has been used to analyze image-based spectroscopic localization methods [20, 50]. While this assumes that the continuously distributed spectra from the subject are "ground truth", the spatially continuous spectra are generally unknown for anything other than test phantoms. Here instead, and analogous to the fractional SLAM leakage/error minimization approach used in Chapter 2, a discrete form of the SRF (dSRF) is used which treats the spatially discrete CSI spectra as ground truth. The CSI spectra are measurable in practice, and CSI's widespread use for multi-voxel MRS renders it a standard whose compartmental average SRF can reasonably approximate the various compartments, albeit dependent on spatial resolution. The dSRF is especially suited for analysis of the SLAM reconstruction because prior knowledge is incorporated into its under-sampled CSI model. Note that unlike the continuous SRF, the dSRF is only defined at the discrete CSI sample points. Also note that a discrete point spread function was used in a similar way to characterize under-sampling interference between Fourier voxels in sparse MRI [61].

For (SENSE) SLAM,

$$\mathbf{dSRF}_{C*M} = (\mathbf{E}_{M'*M} \times \mathbf{b}^r_{M*C})^+ \times \mathbf{E}_{M'*M} \tag{3.8}$$

and for (SENSE) SLAM*,

$$\mathbf{dSRF}_{C*M} = (\mathbf{b}^r_{M*C})^+ \times (\mathbf{E}_{M'*M})^+ \times \mathbf{E}_{M'*M} . \tag{3.9}$$

The calculation of contributions from any specific spatial region, to the reconstructed (SENSE) SLAM spectrum from these dSRFs is shown in Appendix A.

3.2.4 Incorporating spatial and temporal inhomogeneity

In practice, the reconstructed CSI and SENSE CSI spectra are often further corrected for spatial and temporal B_0 and B_1 field inhomogeneities, for example, due to the use of surface coils [20] and/or gradient eddy currents that induce spatially- and/or time-dependent phase variations [68]. These can be accommodated by modifying Eq. (3.4) to incorporate a diagonal matrix **A** comprised of the spatially-dependent terms:

$$\mathbf{s}_{M'*N} = \mathbf{E}_{M'*M} \times \mathbf{A}^{-1}_{M*M} \times \mathbf{b}^{-1}_{M*M} \times \mathbf{b}_{M*M} \times \mathbf{A}_{M*M} \times \boldsymbol{\rho}_{M*N} . \tag{3.10}$$

The diagonal elements in **A** are correction factors, such as reciprocals of the receiver coil sensitivity profile, or conjugates of the phase variation at each voxel. The reconstruction for (SENSE) SLAM and (SENSE) SLAM* are, respectively:

$$\boldsymbol{\rho}^r_{C*N} = (\mathbf{E}_{M'*M} \times \mathbf{A}^{-1}_{M*M} \times \mathbf{b}^r_{M*C})^+ \times \mathbf{s}_{M'*N}, \tag{3.11}$$

$$\text{and} \quad \boldsymbol{\rho}^r_{C*N} = (\mathbf{b}^r_{M*C})^+ \times \mathbf{A}_{M*M} \times (\mathbf{E}_{M'*M})^+ \times \mathbf{s}_{M'*N} . \tag{3.12}$$

If temporal correction terms involving time-varying phases are required [68], the reconstruction in Eqs. (3.10–3.12) must be repeated at each time point. The ^{1}H brain MRS examples shown later utilize both spatial and temporal eddy current corrections [68, 69]. Note that such corrections can alter the dSRF, especially when **A** varies abruptly. Appendix B lists the dSRF formulae with **A** included.

3.2.5 Summary of (SENSE) SLAM implementation

The proactive SLAM and SENSE SLAM experimental protocols are as follows:

(i) Acquire a scout MRI;

(ii) Apply a subset of $M' \geq C$ phase-encoding gradients selected from the central k-space of a regular (SENSE) CSI sequence (or an alternative user-defined subset – see Chapter 2), and acquire the MRS data;

(iii) Register the regular (SENSE) CSI grid onto the corresponding MRI;

(iv) Segment the (SENSE) CSI voxels into the C desired anatomical compartments based on the MRI;

(v) Reconstruct the compartmental spectra using (SENSE) SLAM or the modified (SENSE) SLAM* algorithms (Eqs. 3.6 and 3.7).

For retroactive validation, (SENSE) SLAM is simply applied to the central k-space signals of previously acquired (SENSE) CSI data in Step (ii), and the resulting compartmental average spectra compared to those obtained by co-adding CSI spectra from the constituent voxels in the identical compartments. Acceleration factors (the value of M') were chosen by empirical testing of (SENSE) SLAM/SLAM* on conventional (SENSE) CSI data sets, to achieve acceptable agreement for the intended application. Expressions for the relative SNR of compartmental average (SENSE) CSI, (SENSE) SLAM and SLAM* spectra are presented in Appendix C.

3.3   Methods

All phantom and human studies were performed on a 3T Philips Achieva MRI/MRS system. All human studies were performed on consenting participants using protocols approved by our Institutional Review Board. Reconstruction was performed off-line on a personal laptop computer (quad-core, 1.87 GHz) using software developed in *Matlab* (Mathworks, Natick, MA). (SENSE) CSI, (SENSE) SLAM and SLAM* spectra are displayed as true compartmental averages, normalized to account for differences in compartmental volume, and without time-domain filtering, unless otherwise stated. Vertical scaling is the same for each subject, independent of method, but may vary between subjects depending on receiver gain. Compartments in the brain

were segmented on CSI grids to include any voxel with a volume that intersected a region-of-interest on the corresponding ^{1}H image by 50% or more.

3.3.1 ^{31}P MRS studies of phantoms

^{31}P MRS studies were performed with a single loop 14-cm ^{31}P transmit/receive surface coil, with two disk phantoms filled with 300mM H_3PO_4 and H_3PO_2 placed on top in skewed orientations. After multi-slice gradient refocused echo (GRE) scout ^{1}H MRI, a 3D ^{31}P CSI [4] sequence was applied with adiabatic-half-passage (AHP) excitation (repetition time, TR = 717 ms; field-of-view, FOV, X = Z = 160 mm, Y = 100 mm with X, Y, and Z as the horizontal, vertical, and longitudinal FOV extents in the magnet bore; spatial resolution per voxel, $\Delta X = \Delta Z = 16$ mm, $\Delta Y = 5$ mm; number of phase-encoding steps, PE = 10x10x20 = 2000; total scan time = 23.9 min). Proactive 3D SLAM was then performed with all parameters the same, except that only the 2x2x5 = 20 central-most k-space PEs were applied, reducing the scan time to 15.8 s (including two startup cycles). The CSI grid was co-registered with the scout MRI, three compartments segmented, and proactive SLAM spectra reconstructed. These were compared with the co-added, compartmental average CSI spectra from the constituent voxels in the identical compartments, as noted above.

3.3.2   Proactive ^{1}H MRS brain studies in healthy subjects

In vivo studies of the human brain in healthy volunteers first utilized a commercial 8-channel receive-only head coil operating in quadrature mode to avoid the complications of SENSE reconstruction. Scout MRI was provided by a spin-lattice relaxation (T_1)-weighted 3D turbo-GRE sequence (flip-angle FA = 8^0; echo-time TE = 3.2 ms; TR = 6.9 ms; FOV, Y = 256 mm, X = 228 mm, Z = 156 mm; resolution $\Delta X = \Delta Y$

= 1 mm, ΔZ = 1.2 mm; total scan time = 4.5 min). A spin-echo single-slice 2D ^{1}H CSI

sequence [47, 70, 71] was then applied (TR = 2 s; TE = 144 ms; FA = 90^0; FOV, Y = 200

mm, X = 180 mm; slice thickness, SL = ΔZ = 15 mm; resolution $\Delta X = \Delta Y$ = 10 mm;

number of PEs = 360; scan-time = 12 min) with eight outer volume saturation (OVS)

slabs [70] angled around the head to suppress the lipid signal from the scalp, along with a

"VAPOR" [72] pre-pulse for water suppression. This was followed by a 2D SLAM

sequence with identical parameters except that the number of PEs were reduced to 49

from central-most k-space, reducing the scan-time to 1.7 min (including two startup

cycles). The CSI and MRI grids were co-registered and C = 4 compartments segmented

to: (1) a small user-defined region in the brain; (2) the 'rest of the brain'; (3) the scalp;

and (4) a background compartment. Compartmental average CSI spectra were computed

for these four compartments for comparison with SLAM. SLAM was also implemented

using just the central 25 of the 49 PEs, for an effective acceleration factor of R = 14.

3.3.3 Retroactive ^{1}H MRS brain studies in cancer patients

MRI and ^{1}H CSI MRS data were obtained from 16 consecutively recruited

patients enrolled in an ongoing clinical research study, who had brain tumors with

discernible choline (Cho), total creatine (CR, unphosphorylated creatine, Cr, plus

phosphocreatine, PCr) or N-acetylaspartate (NAA). These data were acquired with a

commercial 32-element SENSE head coil. The acquisition protocol involved: (a) a

SENSE coil-sensitivity reference scan; (b) a multi-slice FLAIR [73, 74] sequence (TE =

120 ms; TR = 11 s; Inversion recovery delay, TI = 2.8 s; FOV, Y = 212 mm, X = 189

mm, Z = 132 mm; resolution, ΔY = 0.83 mm, ΔX = 1.04 mm; ΔZ = 2.2 mm; duration =

3.85 min); and (c) a three-slice SENSE-accelerated 2D spin-echo ^{1}H CSI sequence

(vertical and horizontal acceleration factors $R_Y = 3$, $R_X = 2$; TR = 2.5 s; TE = 144 ms; FA = 90^0; FOV, Y = 230 mm, X = 188 mm, SL = ΔZ = 13.2 mm; slice gap SG = 4.4 mm; ΔX = ΔY = 7 mm) [47, 53, 70, 71]. The SENSE acceleration reduced the number of phase encodes from 891 (= 33x27) to 154 (= 11x14). This was further reduced to 120 by omitting the corners of k-space. Eight lipid-suppression OVS slabs and dual-band water-lipid-suppression (HGDB) pre-pulses were used [75]. (d) An additional set of 27 CSI PEs were acquired without HGDB pulses to provide eddy-current compensation [68]. The total scan-time for (c) and (d) was 6.2 min. Eddy-current-corrected CSI spectra were reconstructed using the scanner's default settings.

Because the raw k-space data were not saved for these studies, retroactive SLAM and SLAM* reconstruction were performed using 1/6th of the k-space data obtained by inverse Fourier transforming the eddy-current-corrected CSI spectra. Five compartments were assigned for SLAM and SLAM*. These were: (1) tumor; (2) contralateral brain (since clinical MRS assessment is routinely based on differences between a suspected pathology and its contralateral region); (3) 'rest of the brain' compartment excluding the tumor and the contralateral compartments; (4) scalp; and (5) background. Magnitudes of peaks in SLAM, SLAM* and the compartmental average CSI spectra, were least-squares fitted to Gaussian models in the frequency domain using software written in *Matlab*, and Cho, CR and NAA peak areas calculated. The results were compared with the compartmental average CSI spectra obtained by co-adding the CSI spectra from constituent voxels in the identical compartments.

3.3.4 SENSE and proactive ^{1}H MRS brain studies

A further set of 8 *in vivo* brain tumor studies were acquired with the identical CSI protocol except that the raw *k*-space data was separately saved. As in the scanner's default reconstruction, the 120 water-suppressed SENSE CSI signals (SENSE factors R_Y = 3, R_X = 2) with the corners of *k*-space omitted, were zero-filled back to the 11x14 = 154 rectangular array. SENSE SLAM and SENSE SLAM* analysis was performed using the 25 central *k*-space PEs, and, as an extreme example, using only 1 PE (zero gradient, but 32 channel acquisition) to provide effective speed-up factors of 5 and 120, respectively. The 11x14 SENSE CSI *k*-space acquisition scheme with the corners omitted, and the *k*-space sampling for SENSE SLAM* with a further 5-fold and 120-fold acceleration factors are illustrated in Fig. 3.1.
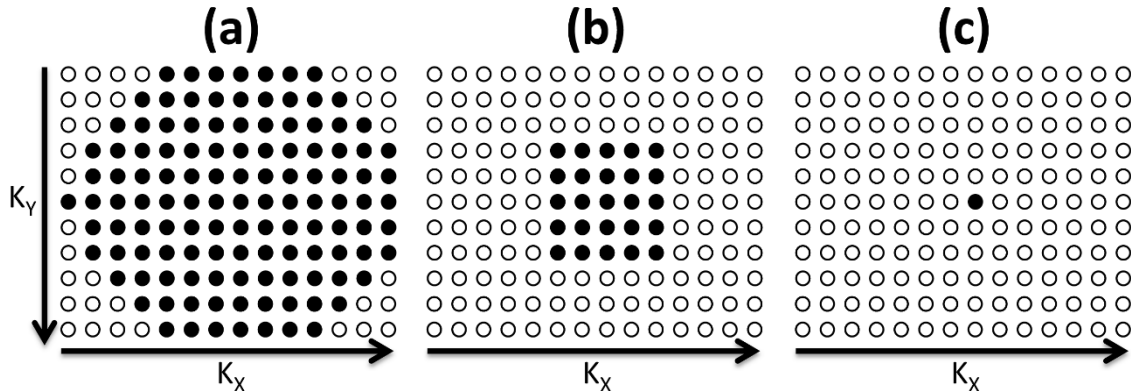
Figure 3.1: The *k*-space acquisition schemes for: (a) 11x14 SENSE CSI with corners omitted (120 steps); and SENSE SLAM/SLAM* acquisitions with further (b) 5-fold (25 steps total) and (c) 120-fold (1-step) acceleration factors, respectively. For simplicity, the 6-fold *k*-space undersampling afforded by SENSE, is omitted. Each filled circle denotes a k-space sample in the k_Y (or Y) and k_X (or X) directions.

In the last 5 of these patients, a proactive three-slice 2D SENSE SLAM sequence was also implemented using 25 PEs for a total scan-time of 1.1 min, including two startup cycles but no additional eddy-current acquisition. Excluding the latter, this corresponded to R = 5 compared to the SENSE CSI protocol above. SENSE SLAM

spectra were compared to compartmental average SENSE CSI spectra both without and with eddy current corrections. For eddy current compensation, the 27 non-water-suppressed PEs from the CSI eddy-current acquisition were reconstructed by SENSE CSI zero-filled to 154, from which the spatial and time-dependent phases were estimated and interpolated across the sample volume [68] after correcting for phase jumps [69]. The same spatial phases were then used for correcting SENSE CSI, SLAM and SLAM* spectra.

Finally, a five-slice 2D SENSE SLAM sequence was implemented pro-actively by extending TR to 3.4 s (duration = 1.5 min including two startup cycles) to provide whole brain coverage while maintaining all other parameters the same.

3.3.5   In vivo 3D ^{31}P human heart MRS

In vivo cardiac ^{31}P MRS studies were performed on healthy volunteers oriented prone on a 17 cm/11 cm diameter dual loop transmit and 8 cm single loop receive coil [29]. Turbo-GRE scout ^{1}H MRI was performed (FA = 30^0; TR = 4.3 ms; TE = 2.1 ms; FOV, Z = 450 mm, X = 450 mm, Y = 160 mm; resolution ΔZ = 2.34 mm, ΔX = 4.69 mm, ΔY = 5 mm), followed by a cardiac-gated pulse-and-acquire 3D ^{31}P CSI sequence (AHP excitation; TR = 2 s; FOV Z = 160 mm, X = 180 mm, Y = 160 mm; resolution $\Delta Z = \Delta X$ = 20 mm, ΔY = 10 mm; PEs = 1152; duration = 38.7 min including 8 startup cycles). Three compartments—chest, heart and background—were segmented from the co-registered images. SLAM spectra were reconstructed and compared to the compartmental average CSI spectra.

3.3.6   Effect of compartmental misregistration and sample heterogeneity

To investigate the robustness of the SENSE SLAM* and SENSE SLAM reconstruction to mis-registration and signal heterogeneity, the segmented compartment maps from 8 different MRS studies with 5 segmented compartments, were superimposed on a 9^{th} data set, resulting in gross misplacement of compartments relative to the anatomy. The resultant SENSE SLAM* and SENSE SLAM spectra from the mis-registered compartments that contain mixtures of various involved and uninvolved tissues, were compared with the average SENSE CSI spectra from same compartments.

3.3.7   Localization analysis

To evaluate differences in the dSRF of (SENSE) SLAM and (SENSE) SLAM*, a quantitative analysis was performed with 5 segmented compartments in a ^{1}H MRS data set acquired from a patient with a brain tumor, in accordance with Appendix A. For the dSRF computation and analysis, the acceleration factor of SENSE SLAM and SENSE SLAM* was 6 times faster than SENSE CSI (25 phase encodes vs. 154 phase encodes) without omitting the corners of k-space. To compare SLAM and SLAM* two acceleration factors, R = 6 and 36 vs. CSI, were chosen.

To quantify the potential contamination of the tumor dSRF from extra-compartmental signals, the integral of both the dSRF and its modulus, |dSRF| were calculated by summation over each of the five segmented compartments for: (i) SENSE SLAM*, (ii) SENSE SLAM, (iii) SLAM*, and (iv) SLAM*, both without (TSVD threshold =0%) and with numeric regularization at a 2% TSVD threshold. Note that while the summed dSRF benefits from the cancellation effects of an oscillating spatial response [50], the summed |dSRF| represents a worst possible scenario with no cancellation

whatsoever, requiring both positive and negative MRS signals having the same sign as the dSRF. Because numeric regularization has a broadening effect on the dSRF, we also computed the integrated dSRF and |dSRF| over a region expanded by ~20% in each direction about the tumor compartment, effectively making it 40% larger, analogous to the effect of a cosine-apodized point spread function [76].

Finally we compute the dSRF for a brain compartment in the extreme case of the 120-fold accelerated SENSE SLAM* spectrum acquired with only a single phase-encode.

3.4   Results

3.4.1 ^{31}P MRS studies of phantoms

Figure 3.2 shows the segmentation on a scout MRI of the disk phantom, and spectra representing the compartmental average 3D CSI and 3D SLAM reconstructions. Minimal difference is observed between the proactively acquired SLAM (red) and CSI spectra (blue), despite the fact that the SLAM spectra were acquired 100 times faster than the CSI spectra, excluding the 2 startup cycles. The H_3PO_4 phantom has a singlet at 2.9 ppm and the H_3PO_2 phantom has three peaks at 2.9 ppm, 13.5 ppm and 24.1 ppm. No spatial leakage from the 13.5 ppm and 24.1 ppm peaks of the H_3PO_2 phantom spectrum (Fig. 3.2b), into the H_3PO_4 spectrum (Fig. 3.2c) is discernible in either CSI or SLAM spectra, and spatial leakage into background (Fig. 3.2d) is also negligible.
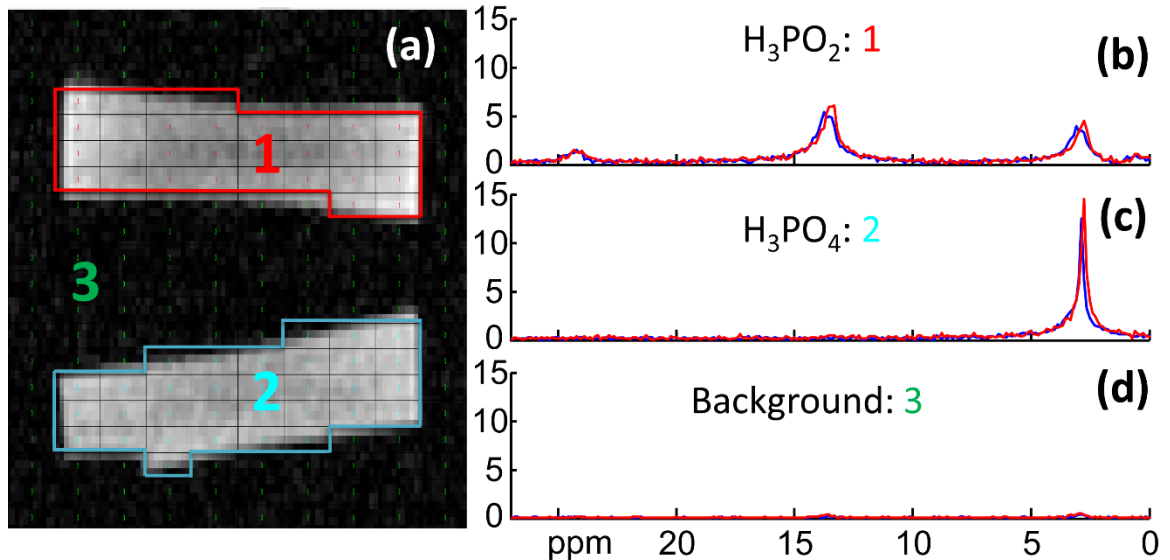
Figure 3.2: Comparison of 3D proactive ^{31}P SLAM spectra in a phantom. (a) Segmentation based on a co-registered phantom image. The CSI voxels are segmented into three compartments: 1, H_3PO_2; 2, H_3PO_4; 3, background. (b-d) Red SLAM spectra reconstructed from 20 phase encodes compared to the blue CSI compartmental average spectra reconstructed from 2000 phase encodes. The SLAM acceleration factor is R=100.
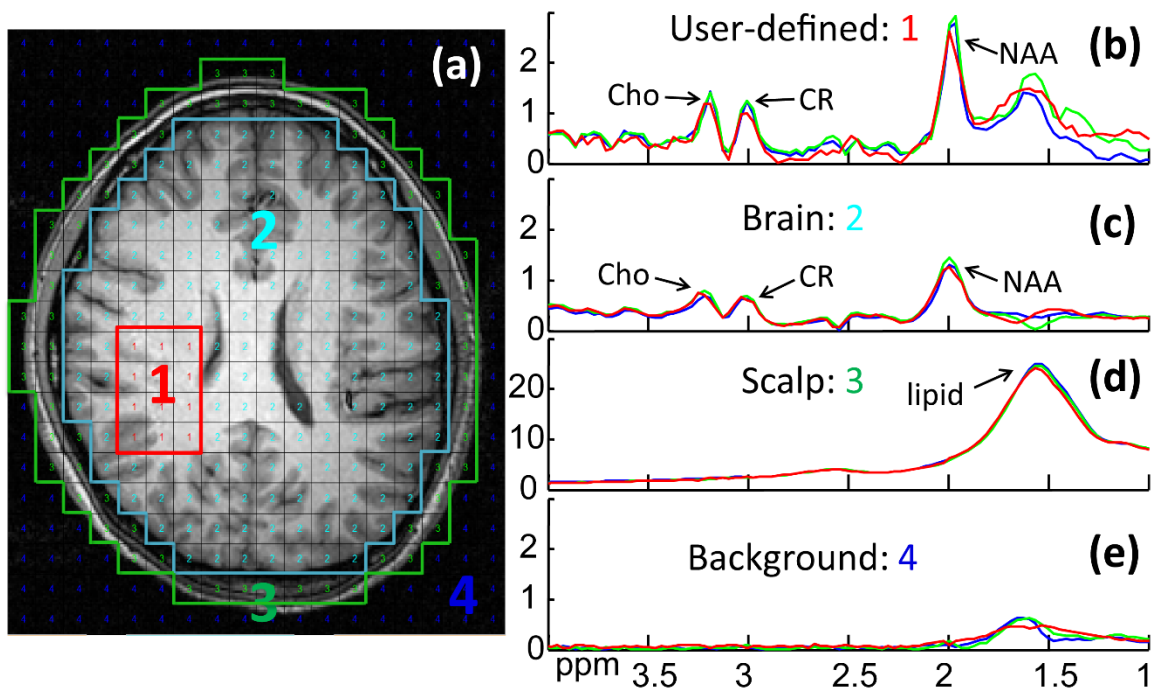
Figure 3.3: Retroactive and proactive SLAM in normal brain. (a) MRI of normal brain, overlaid with the CSI grid and segmented into four compartments post MRS acquisition: (1) user-defined area; (2) 'rest of the brain'; (3) scalp; and (4) background. (b-d) 2D CSI (blue) and 2D SLAM spectra (red and green) acquired from the four compartments. The red SLAM spectra were acquired in a proactive scan using data 14 times faster than CSI. The green SLAM spectra were reconstructed retroactively from a subset of 1/14th of the CSI data.

3.4.2 Proactive ^{1}H MRS brain studies in healthy subjects

Figure 3.3 compares both proactive and retroactive implementation of 2D SLAM with 2D CSI in a normal brain. The user-defined region inside the brain (Fig. 3.3a) was chosen to mimic a potentially "abnormal" compartment. Both retroactive (green) and proactive SLAM spectra (red) agree with the CSI spectrum from the compartment (blue), and are essentially coincident in the brain and scalp compartments (Fig. 3.3c,d). Here, SLAM is 14 times faster than CSI. The proactive SLAM spectrum from the "abnormal" compartment (Fig. 3.3b) differs slightly more from the CSI spectrum than the retroactive SLAM spectrum, likely reflecting scan-to-scan variations including SNR effects, fat/water suppression and/or physiologic motion.

Reconstruction time (solving Eq. 3.6 or 3.7) for matrix sizes up to 800x5 was ≤ 2 min on the laptop computer for all studies.

3.4.3 Retroactive ^{1}H MRS brain studies in cancer patients

Figure 3.4 is an example of retroactive implementation of SLAM (spectra b1-f1) and SLAM* (spectra b2-f2) in a patient with a low-grade astrocytoma (Fig. 3.4a). The CSI spectra were generated from 33x27 PEs while SLAM used only 1/6th (=13x11) of them. Both SLAM and SLAM* spectra (red) basically coincide with the compartment average CSI spectra (blue) over nearly the entire spectral range. SLAM* is slightly more consistent with CSI than SLAM in this example (e.g., Cho in b1 vs. b2, and NAA in d1 vs. d2). Importantly, the high Cho/CR peak area ratio in the tumor compartment is faithfully preserved by SLAM and SLAM* (spectra b1, b2), while the contralateral compartment does not show high Cho/CR (spectra c1, c2). Elevated Cho levels are common in tumors [45], and the Cho/CR ratio may be useful for grading them [77].
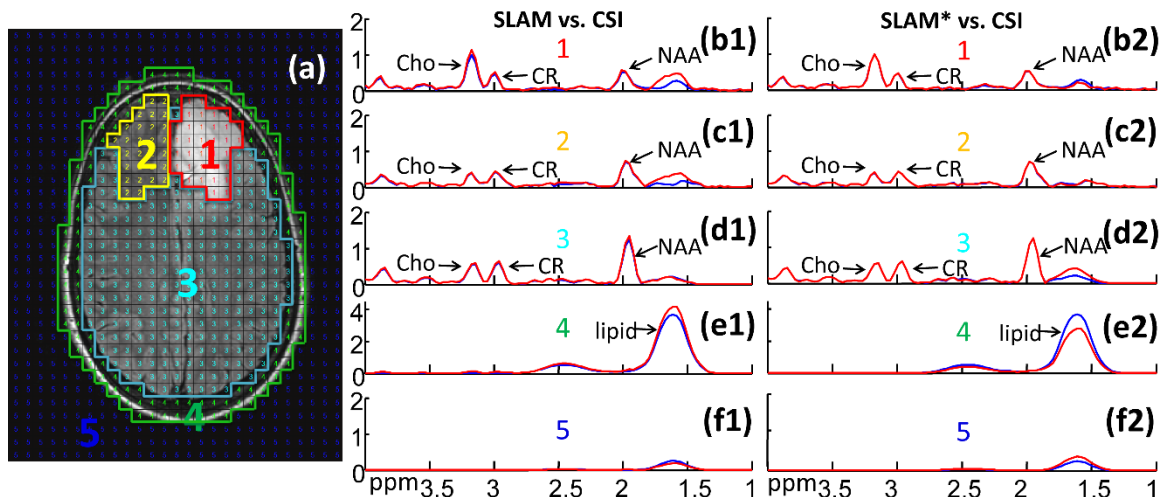
Figure 3.4: Comparison of SLAM and SLAM* brain spectra with CSI. (a) Image with co-registered CSI grid and segmented into five compartments: (1) tumor (low-grade astrocytoma); (2) contralateral brain; (3) 'rest of the brain'; (4) scalp; and (5) background. Spectra (b1-f2) show the CSI spectra (blue) from the corresponding compartments, along with SLAM (b1-f1) and SLAM* (b2-f2) spectra (red) reconstructed from $1/6^{th}$ of the CSI data for an acceleration factor R = 6.
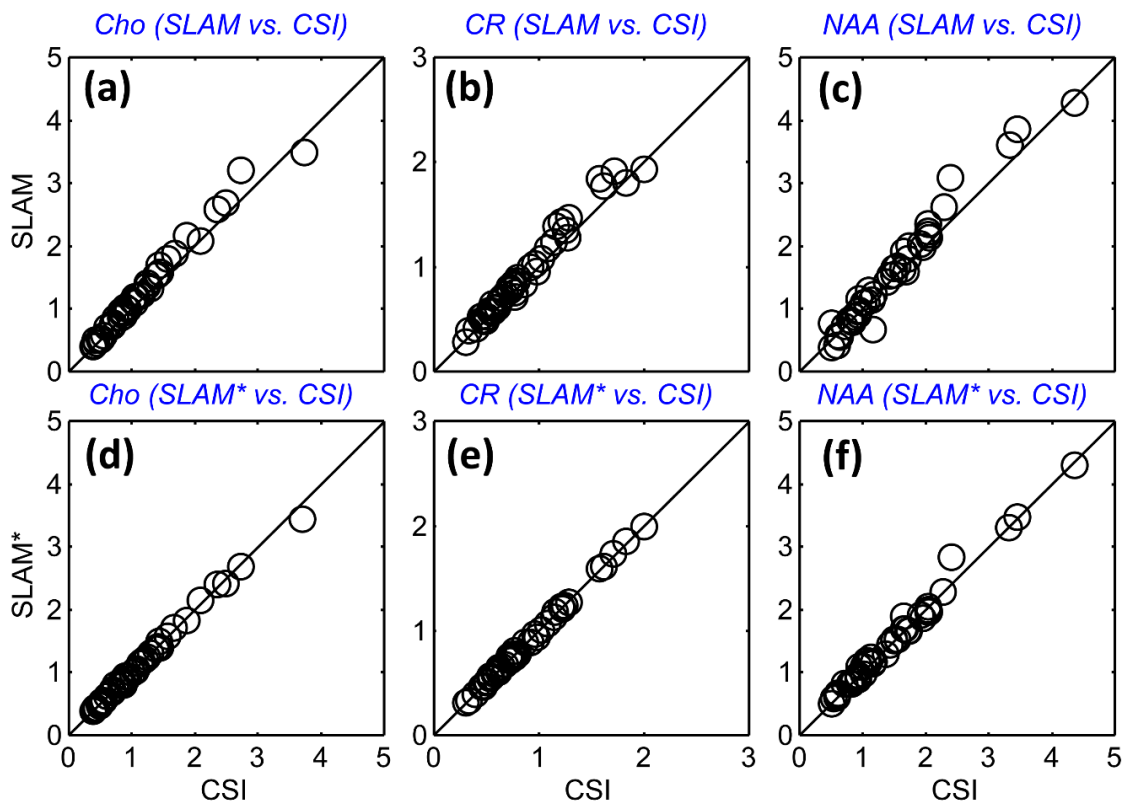
Figure 3.5: Quantitative comparison of SLAM and SLAM* with CSI data from the 16 patients. Cho, CR, and NAA levels (arbitrary units) as quantified in SLAM (a-c), and in SLAM* (d-f) spectra in tumor, contralateral brain and 'rest of the brain' compartments, as a function of those levels measured in the CSI spectra from the same compartments. The SLAM and SLAM* spectra were reconstructed from $1/6^{th}$ of the 2D CSI data for an effective R = 6. The correlation coefficients are R ≥ 0.98 for all cases.

70

Figure 3.5 shows the Cho, CR, and NAA levels quantified from CSI, SLAM and SLAM* spectra, reconstructed retroactively from tumor compartments in the 16 patient data sets. For the tumor, contralateral brain, and 'rest of the brain' compartments, both SLAM and SLAM* are quantitatively consistent with CSI, while the SLAM* reconstruction shows slightly less deviation than SLAM. The same consistency of results compared to CSI was observed for the Cho/Cr metabolite ratios (not shown), and Bland-Altman analysis for Cho, CR (Fig. 3.6) and NAA from the tumor and contralateral brain compartments did not reveal any significant systematic effects. The standard deviation (SD) of the ratio of SLAM to CSI levels of NAA, Cho and CR, were respectively 15%, 5.9% and 6.7%, as compared to 6.2%, 2.8% and 2.1% for SLAM*, reflecting perhaps the greater lipid contamination of NAA in the SLAM spectra.

3.4.4 SENSE and proactive ^{1}H MRS brain studies

Figure 3.7 demonstrates SENSE SLAM implementation on a patient with a large glioblastoma whose raw k-space data had been saved. Here the SENSE SLAM and SENSE SLAM* spectra (red) are effectively 5 times (spectra b1-f1 and b2-f2) and 120 times faster (spectra b3-f3) than the SENSE CSI data (blue), respectively. With SENSE, the modified SLAM* algorithm (spectra b2-f2) consistently outperforms the original SLAM algorithm (spectra b1-f1), in the ~1.7 ppm lipid region where strong, broad, lipid peaks tend to corrupt the NAA peaks. The SENSE SLAM* spectra (b2-d2) are relatively unaffected by the lipid peaks. Substantially less spatial leakage is also evident in the background compartment of SENSE SLAM* compared to SENSE SLAM (spectra f2 vs. f1). In the extreme case where only the central zero-gradient PE step was used and spatial-encoding depends entirely on SENSE, SENSE SLAM* produces surprisingly

good spectra (spectra b3-f3), while the earlier SLAM algorithm performed poorly (not shown). Based on Appendix C, for the tumor compartment (c3), SENSE SLAM* maintains 23% of the SNR of SENSE CSI. After correcting for the over two-orders-of-magnitude difference in acquisition time, the SNR per unit volume per unit time for SENSE SLAM* is about 3 times higher than SENSE CSI. Importantly again, the high Cho/CR characteristic of the tumor is preserved in both SLAM* and SLAM tumor spectra (spectra c1, c2, c3).
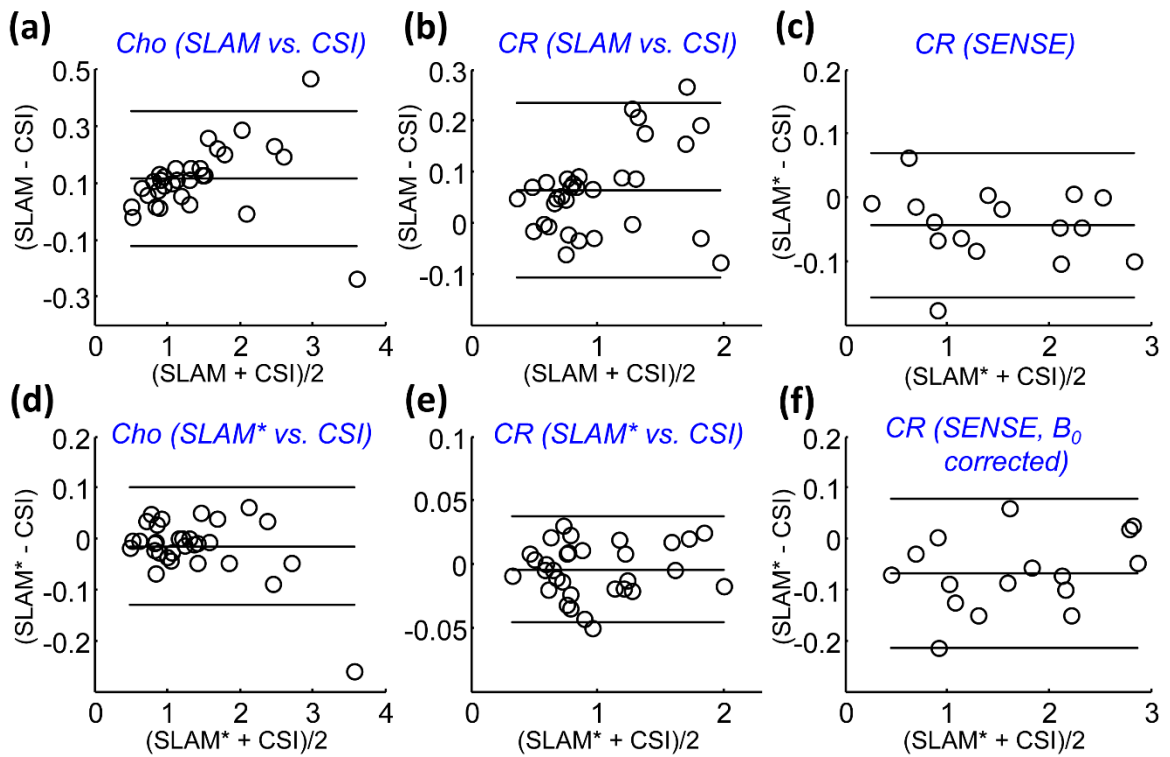
Figure 3.6: Bland-Altman plots for Cho (a, c) and CR (b, e) measured by SLAM (a, b) and SLAM* (d,e), as compared to CSI. Parts (c) and (e) show Bland-Altman plots for CR measured by SENSE SLAM* as compared to SENSE CSI, without (c), and with (f) B_0 corrections.
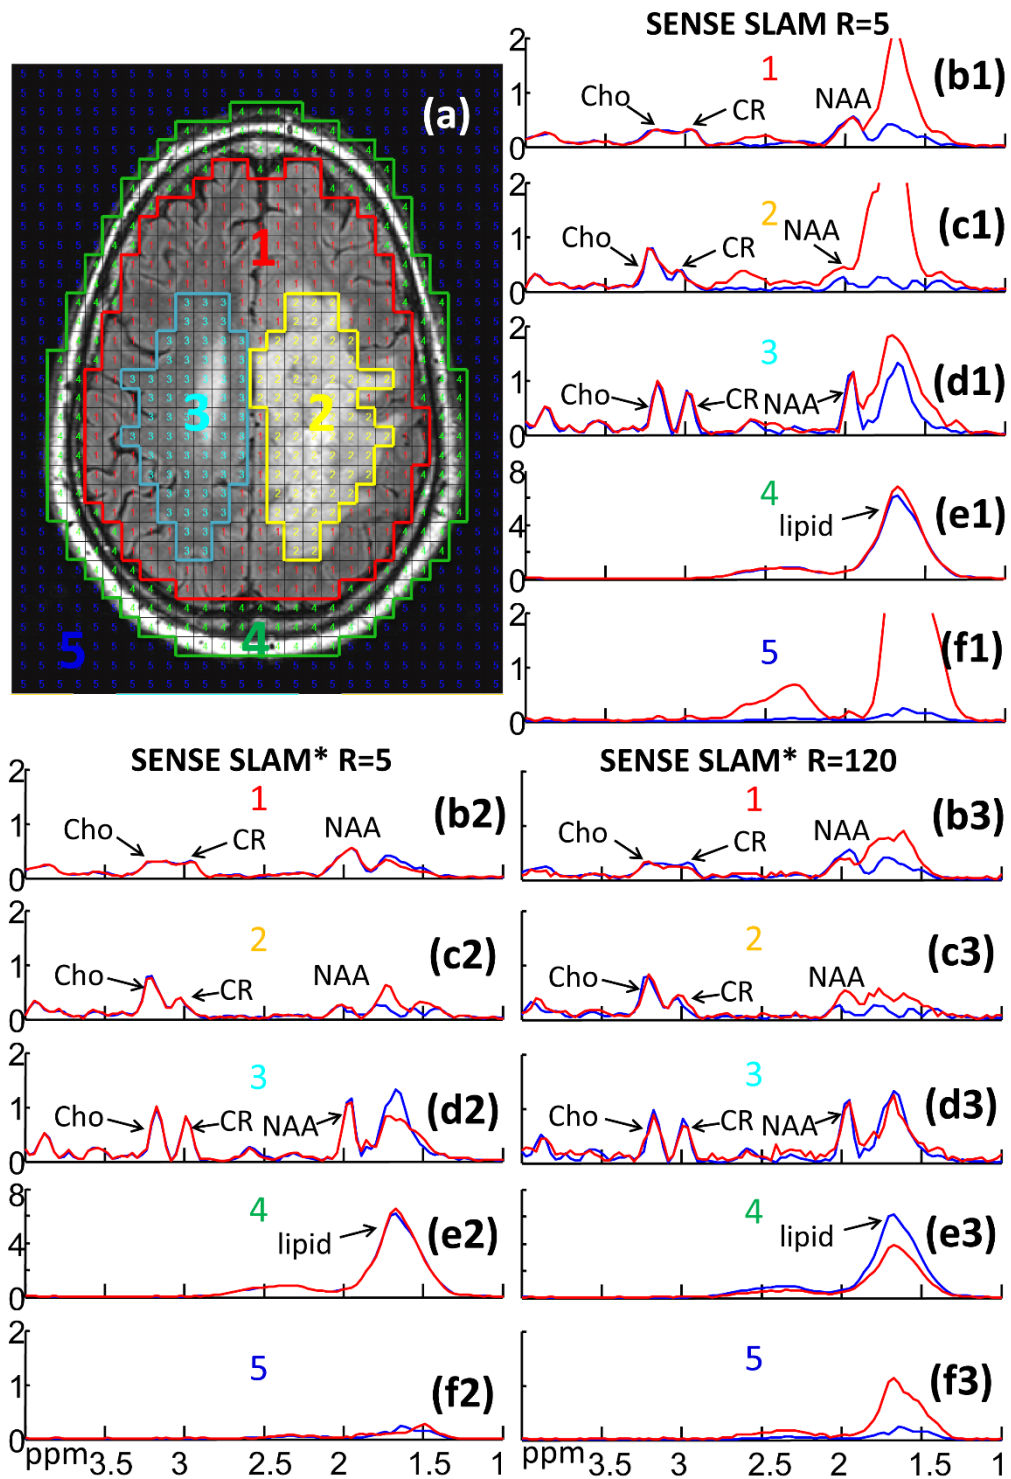
Figure 3.7: 2D retroactive SENSE SLAM and SENSE SLAM* results. (a) MRI showing segmentation of five compartments: (1) 'rest of the brain'; (2) tumor (a glioblastoma); (3) contralateral brain; (4) scalp; and (5) background. Spectra (b1-f3) are from the corresponding compartments with SENSE CSI spectra in blue for comparison. SENSE SLAM spectra (b1-f1) and SENSE SLAM* spectra (b2-f2) in red, were reconstructed with 1/5th of the SENSE CSI data, for an acceleration factor R=5. Spectra (b3-f3) were reconstructed with R=120 using SENSE SLAM*.

Figure 3.8 shows retroactive (green) and proactive (red) SENSE SLAM*, and SENSE CSI results (blue), without (spectra b1-f1) and with (spectra b2-f2) corrections for eddy-current induced time-varying spatial phases. Without eddy-current correction, the SENSE SLAM* used 25 PEs, while SENSE CSI used 120 phase encodes. All three methods used the same eddy-current CSI-based eddy-current acquisition to estimate the time/space-varying phase. A high level of coincidence is evident among the three sets of spectra. Eddy current correction noticeably improves spectral resolution (spectra d2 vs. d1) while reducing lipid contamination (spectra e2 vs. e1).

Figure 3.9 shows quantitative Cho, CR and NAA levels for SENSE CSI and SENSE SLAM* without and with eddy-current correction for tumor and contralateral brain compartments in the 8 patients. SENSE SLAM* reliably generates quantitatively indistinguishable results as SENSE CSI both without and with (b, d) eddy current correction, 5 times faster. Similar results were found for measurements of metabolite ratios derived from tumor and contra-lateral compartments (not shown), and Bland-Altman analysis for CR, Cho, and NAA measurements with SENSE did not reveal any significant systematic effects (Fig. 6c,f shows results for CR only).

Figure 3.10 demonstrates the feasibility of providing essentially whole-brain coverage for brain tumor patients using proactive SENSE SLAM* in ~1.5 min. In this case, the tumor was too large to permit coverage by the regular SENSE CSI protocol and a whole-brain SENSE CSI acquisition was precluded by scan-time, so no CSI data were available for comparison. Elevated Cho/CR is evident in the tumor compartment in the three adjacent slices (Fig. 3.10b-d).
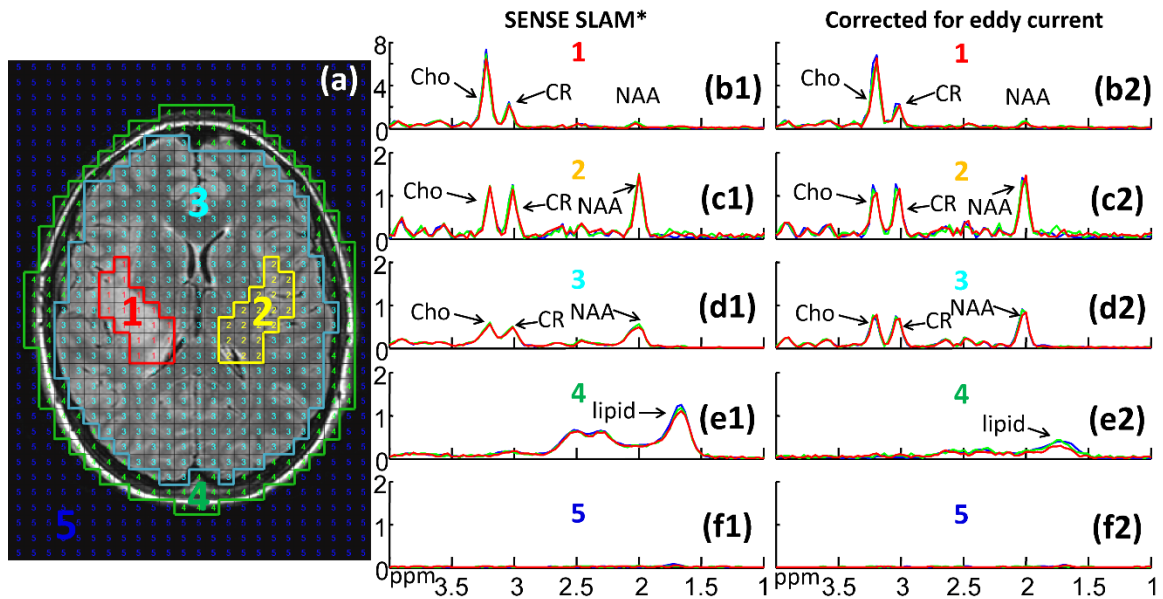
Figure 3.8: SENSE CSI and SLAM* spectra without and with eddy current corrections. (a) MRI depicting segmentation: (1) brain tumor; (2) contralateral brain; (3) 'rest of the brain'; (4) scalp; and (5) background. Spectra (b1-f2) are from the corresponding compartments without (b1-f1) and with (b2-f2) eddy current correction. Blue spectra are SENSE CSI. The SENSE SLAM* spectra are reconstructed from 1/5th of the CSI data set (green), or proactively in a separate scan with R=5 (red).
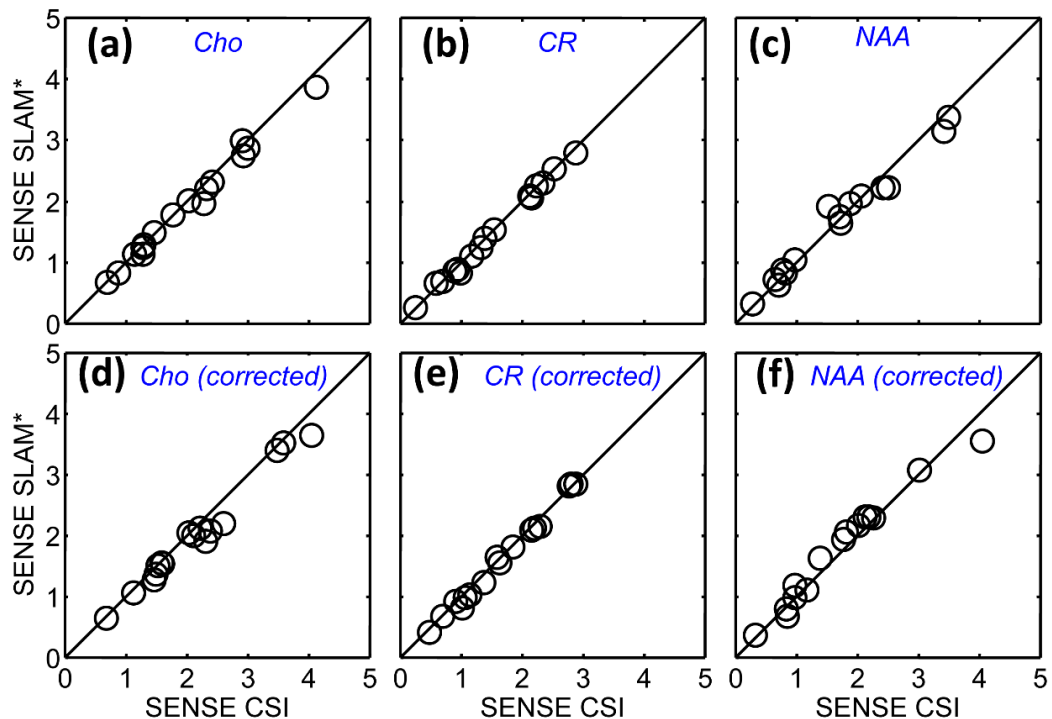
Figure 3.9: Quantitative comparison of SENSE SLAM* measures of Cho, CR, and NAA levels (arbitrary units) in tumor and contralateral brain with SENSE CSI measures from the same compartments, without (a-c) and with (d-f) eddy current corrections. The SENSE SLAM* spectra were reconstructed with 1/5th of the CSI data for R = 5. Correlation coefficients are $R \geq 0.98$ for all cases.
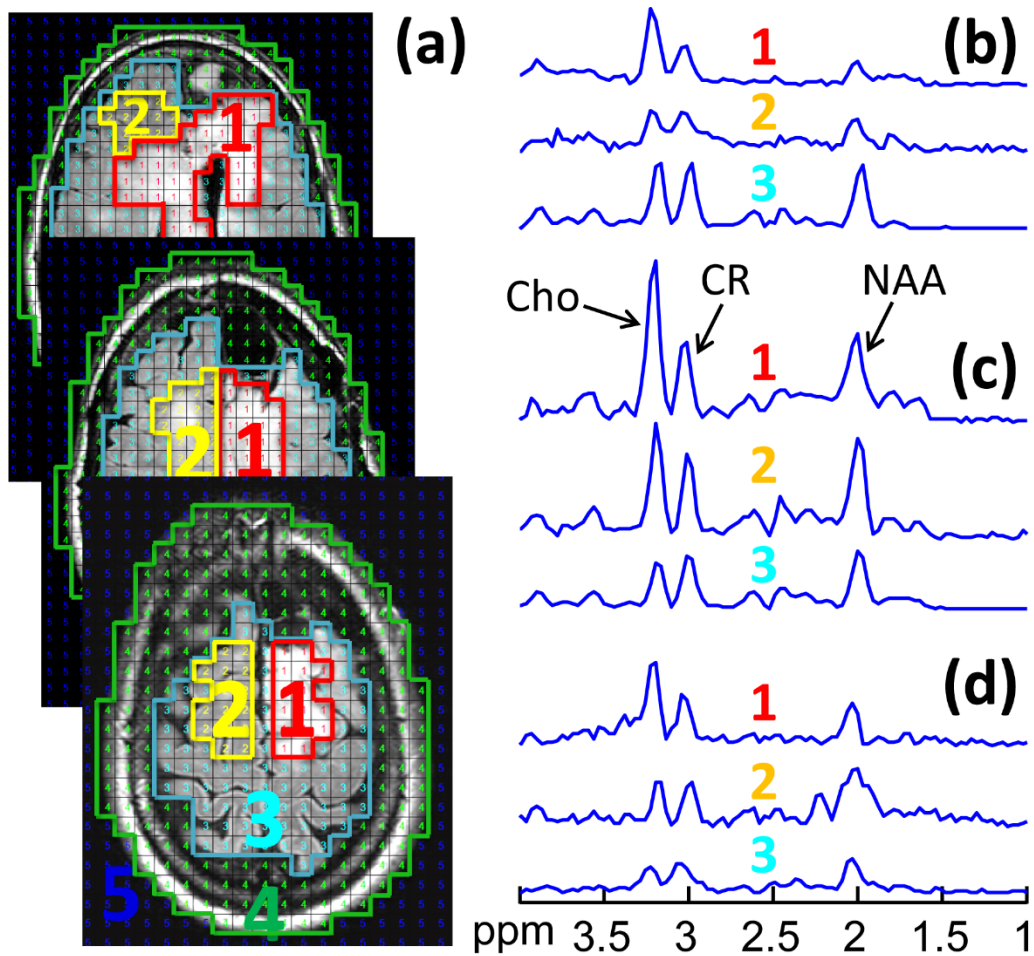
Figure 3.10: (a) Proactive SENSE SLAM* from the top three sections (separation, 17.6 mm) of a 5-slice data set, annotated to show: (1) tumor; (2) contralateral brain; (3) 'rest of the brain'; (4) scalp; and (5) background compartments. (b-d) SENSE SLAM* spectra of the corresponding compartments in each of the three slices.
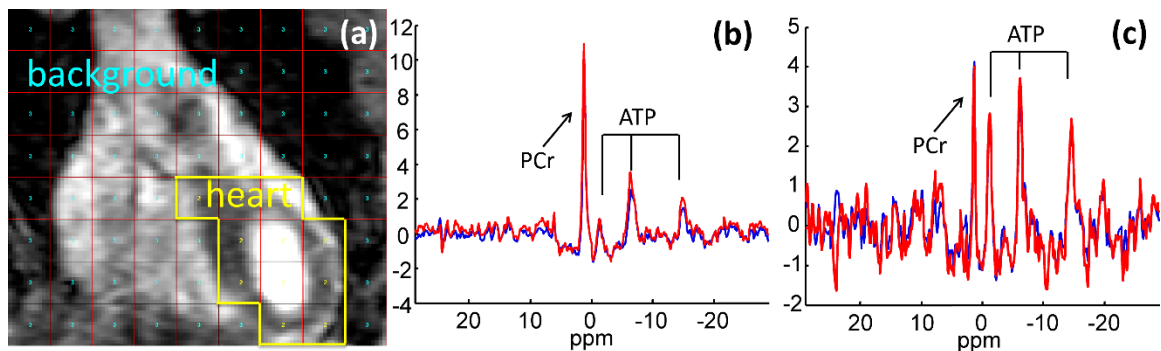
Figure 3.11: 3D SLAM in a human ^{31}P cardiac study. (a) Cardiac MRI showing segmented compartments: (1) chest; (2) heart; and (3) background. SLAM (red) and CSI (blue) spectra from (b) the chest and (c) the heart, with $R = 7$ and a 15 Hz exponential filter.

3.4.5   In vivo 3D ³¹P human heart MRS

Figure 3.11 illustrates the application of 3D SLAM reconstruction to a conventional thoracic ³¹P 3D CSI data set acquired in 38.7 min at 3T, which was reduced by SLAM to 5.5 min, representing an acceleration factor of 7.

3.4.6   Compartmental misregistration and sample heterogeneity

Figure 3.12 shows the sensitivity of SENSE SLAM* (spectra a2-h2) and SENSE SLAM (spectra a3-h3) to segmentation errors due to gross maladjustment in both size and location. Despite gross intra-compartmental anatomical inhomogeneity, SENSE SLAM* spectra are remarkably consistent with SENSE CSI from the same compartments in all cases, consistent with manageable, if not negligible signal bleed. SENSE SLAM reconstructed with Eq. (6), does show some discrepancies (e.g., compare spectra c3 and e3), and inferior fitting to the SENSE CSI spectra used as the standard, compared to SENSE SLAM*. Importantly, all SLAM and SLAM* spectra were reconstructed from the same 25 PEs from central k-space, with an acceleration factor of 5 vs. SENSE CSI.
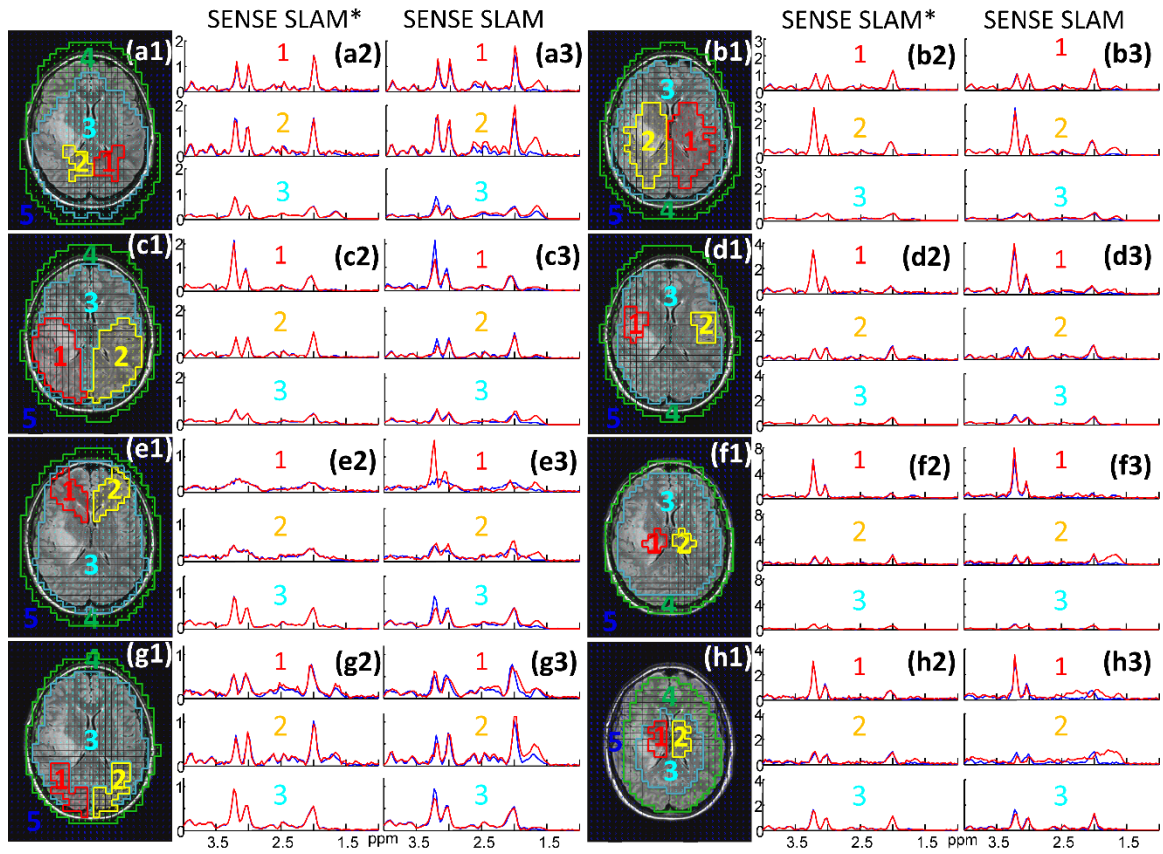
Figure 3.12: The effects of incorrect segmentation and inhomogeneity. (a1-h1) Brain MRI showing eight grossly maladjusted segmentations (Fig. 8a shows the correct segmentation). The corresponding SENSE SLAM* (a2-h2, red) and SENSE SLAM spectra (a3-h3, red) reconstructed from the maladjusted compartments are compared with the average SENSE CSI cspectra (blue) from the same compartments. For SENSE SLAM* and SENSE SLAM, R = 5 compared to SENSE CSI.
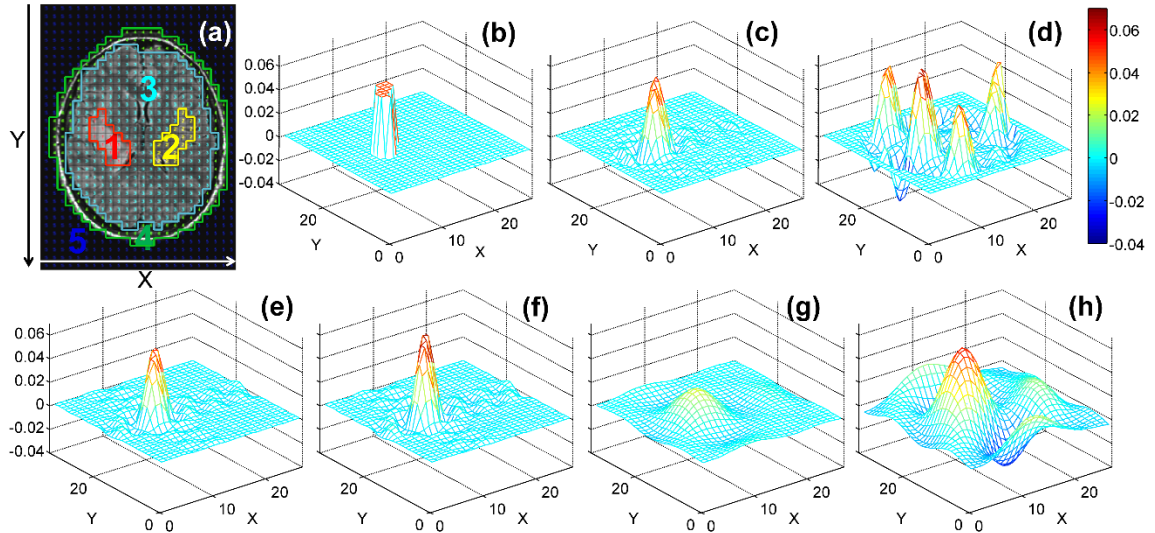
Figure 3.13: The computed tumor dSRF. (a) Brain MRI showing tumor segmentation(1, in red). The real part of dSRF is shown for the following reconstructions: (b) SENSE SLAM* (R=6 vs. SENSE CSI, without numeric regularization); (c) SENSE SLAM* (R = 6 vs. SENSE CSI, with numeric regularization and TSVD threshold ≤2% of the maximum); (d) SENSE SLAM (R=6 vs. SENSE CSI); (e) SLAM* (R=6 vs. CSI); (f) SLAM (R=6 vs. CSI); (g) SLAM* (R=36 vs. CSI); (h) SLAM (R=36 vs. CSI). Note that the dSRF is essentially not changed with or without the 2% TSVD in (d-h).
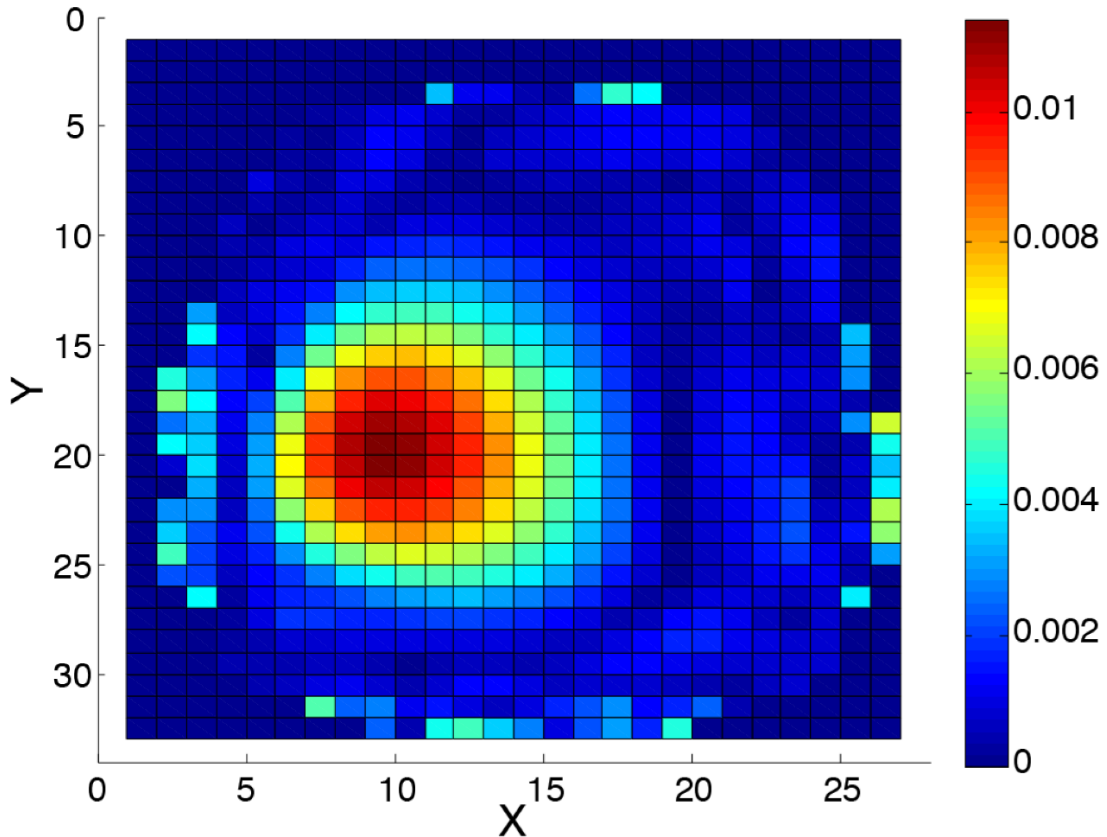
Figure 3.14: The computed dSRF for compartment #3 of Fig. 7 in the extreme case of a single phase-encode (intensity scale is arbitrary).

79

Table 3.1: Localization analysis based on dSRF's of the tumor compartment shown in Figure 3.11 for, SENSE SLAM* with 2% TSVD, and SENSE SLAM*, SENSE SLAM, SLAM* and SLAM without regularization, respectively. Summation of dSRF and absolute dSRF (|dSRF|) are computed over different spatial regions.

	Sum of	Tumor	Contra-lateral	'Rest of the brain'	Scalp	Back-ground
(a): SENSE SLAM* w/o regularization	dSRF	1	0	0	0	0
	|dSRF|	1	0	0	0	0
(b): SENSE SLAM* with 2% TSVD	dSRF	0.8304	0.0095 - 0.0014i	0.1678 + 0.0014i	-0.0023 - 0.0005i	-0.0054 - 0.004i
	|dSRF|	0.8304	0.0304	0.5671	0.1367	0.072
(c): SENSE SLAM	dSRF	1	0	0	0	0
	|dSRF|	1.0001	0.0873	2.3905	1.8984	0.8022
(d): SLAM*	dSRF	0.8180	0.0186	0.1804	-0.0208	0.0038
	|dSRF|	0.8180	0.034	0.5839	0.0948	0.1519
(e): SLAM	dSRF	1	0	0	0	0
	|dSRF|	1	0.0361	0.7876	0.1253	0.2011
Summation over ~20% extended tumor region						
(f): SENSE SLAM* w/o regularization	dSRF	1	0	0	0	0
	|dSRF|	1	0	0	0	0
(g): SENSE SLAM* with 2% TSVD	dSRF	0.9720 - 0.0015i	0.0095 - 0.0014i	0.0262 + 0.0028i	-0.0023 - 0.0005i	-0.0054 - 0.004i
	|dSRF|	0.9721	0.0304	0.4255	0.1367	0.072
(h): SLAM*	dSRF	0.962	0.0186	0.0364	-0.0208	0.0038
	|dSRF|	0.962	0.034	0.4399	0.0948	0.1519

3.4.7 Localization analysis

Figure 3.13 plots the dSRF's for a tumor compartment reconstructed from the

(SENSE) SLAM and (SENSE) SLAM* strategies, applied to the same anatomic

segmentation (Fig. 3.13a). The difference between the modified SENSE SLAM*

algorithm, Eq. (3.7), and SENSE SLAM employing Eq. (3.6) is best appreciated from

parts (c) and (d). SENSE SLAM* has a spatially restricted dSRF, whereas the dSRF of

SENSE SLAM is expansive with both positive and negative lobes.

The dSRF of SENSE SLAM* performed with Eq. (3.7) without numeric

regularization, is in theory perfect (Fig. 3.13b). This is because combining the 25 PEs

with the 32 separate acquisition channels creates $M' = 25 \times 32 = 800$ known k-space

signals. Although there are $M = 33 \times 27 = 891$ unknown image-space signals, application

of the sensitivity mask during SENSE reconstruction [6] effectively reduces the number

of unknowns from 891 to less than 800. Consequently, $(\mathbf{E}_{M'*M})^{+}$ in Eq. (3.9) is over-

determined, hence the perfect dSRF. With numeric regularization, which is typically

needed to ensure a well-conditioned matrix inversion when SNR is low, 'the perfect

dSRF' of Fig. 3.13(b) degrades slightly to the dSRF for SENSE SLAM* shown in Fig.

3.13(c). The dSRF remains spatially well-constrained, with edge definition improving as

the TSVD threshold is reduced, or as more PEs are added. On the other hand, as Fig.

3.13(d) shows, the dSRF of SENSE SLAM employing Eq. (3.6) does not closely

approximate the perfect case, even with all phase encodes (154 PEs in SENSE CSI)

included.

However, without SENSE reconstruction, the dSRF of SLAM* (Fig. 3.13e) and

SLAM (Fig. 3.13f) using 1/6th of the CSI PEs, are very similar to each other. They are

also similar to the dSRF of numerically regularized SENSE SLAM* (Fig. 3.13c), which

is 6 times faster than SENSE CSI and 36 times faster than CSI. However, the similarity

between SLAM* and SLAM dSRFs disappears as the number of PEs is reduced further,

as illustrated with only 1/36th of the CSI PEs in Figs. 3.13(g) and (h). Note that while

numeric regularization affects the dSRF of SENSE SLAM*, it barely alters that of SLAM,

SENSE SLAM or SLAM*. We found that numeric regularization with the 2% TSVD threshold was adequate for the data sets analyzed herein.

To quantify the dSRF's of the tumor compartment in Fig. 3.13, the summed dSRF and |dSRF| were calculated over the five segmented compartments (tumor and contralateral brain, 20 CSI voxels each; 'rest of the brain' compartment #3, 288 voxels; scalp, and background). Without numeric regularization, SENSE SLAM* achieved perfect reconstruction (Table 3.1, row a) consistent with its dSRF (Fig. 3.13b). Numeric regularization broadened the dSRF of SENSE SLAM*, resulting in non-zero sums of dSRF and |dSRF| from the surrounding compartments (Table 3.1, row b). On the other hand, SENSE SLAM always forces the integral of dSRF to zero over the other compartments. This is in part achieved by a spatially oscillating dSRF (Fig. 3.13d), reflected in the non-zero |dSRF| sum in Table 3.1 (row c), which explains the strong lipid contamination in the SENSE SLAM spectra (Fig. 3.7, spectra b1-f1). Unlike SLAM (Table 3.1, row e), SLAM* (row d) does not force the summed dSRF to zero, but SLAM* does reduce the |dSRF| sum compared to SLAM. This is consistent with the reduced lipid contamination noted in Fig. 3.7.

Both of the dSRF's in Fig. 3.13(c,e) are broadened compared to the perfect case (Fig. 3.13b). Table 3.1, rows f-h, show the effect of summing the dSRF and |dSRF| of the tumor compartment over a region expanded by ~20% in each direction (tumor extended from 20 to 28 voxels; same 20-voxel contralateral brain compartment; remaining brain region decreased to 280 voxels). Without regularization SENSE SLAM* is unchanged (Table 1, row f vs. row a). However, with numerical regularization the summed dSRF in the tumor increases from 0.83 to 0.97, while that in the rest of the brain decreases from

82

0.17 to 0.03 (row g). A similar result obtains for SLAM* (Table 3.1, row h). This means that nearly all of the contributions to the reconstructed SLAM* and SENSE SLAM* spectra derive from a region localized around the tumor. Thus, it is reasonable to interpret the numerically regularized dSRF of SLAM* or SENSE SLAM* as a broadening (~20% in each direction for the specific geometry shown in Fig. 3.13a) of the ideal case.

Finally, Fig. 3.14 depicts the dSRF in the extreme case of 120-fold accelerated SENSE SLAM* for compartment 3 of Fig. 3.7. Here, with only a single PE, localization is essentially entirely attributable to the SENSE coils in combination with the prior knowledge. This results in some signal rise around the periphery, close to the coils. Note that such extreme acceleration may not work well for other segmentation, data, or applications.

3.5   Discussion

Even though information about spatial heterogeneity may be lost, in human and animal studies in practice, metrics derived from global or compartmentally-averaged spectra are often used to index metabolic status for diagnostic or prognostic purposes. SLAM is an MRS localization method that provides average spectra from segmented anatomical compartments using a limited set of phase-encoded acquisitions. Because these spectra are reconstructed by linear algebra for which any subset of the conventional CSI PE set can be selected, choosing only those at central k-space where SNR is highest, generally results in the best SNR by avoiding SNR losses associated with co-adding conventional CSI spectra, post-acquisition [8, 50]. This means that SLAM can employ the same (central k-space) PE set for any given study protocol without requiring additional scanner-side prescription. The SNR gain compared to co-adding CSI spectra

then provides a substantial offset to the SNR loss that results from cutting the scan-time by dropping the high-order phase-encodes of the conventional CSI acquisition. Compared to 2D and 3D CSI whose PE sets typically require 10^2 to 10^3 acquisitions, the decimated PE acquisition sets possible with SLAM can produce very large reductions in scan-time indeed. Such reductions can easily enable inclusion or extension of MRS in a patient or animal MRI/MRS protocol that would otherwise be precluded by time constraints, as was the case for the brain tumor patient in Fig. 3.10 for whom conventional SENSE CSI could not otherwise be performed.

In Chapter 2, we demonstrated up to an 8-fold acceleration for 1D SLAM limited to a single receiver coil, when compared to an existing standard 16-PE 1D CSI protocol used for human cardiac ^{31}P MRS [49]. With a 4-fold speed-up, measurements of human cardiac energy metabolism obtained from 1D SLAM and 1D CSI were indistinguishable. In this chapter, we have expanded 1D single-coil SLAM in *five* major and novel ways.

First, we have extended the SLAM method to 2D and 3D CSI achieving acceleration factors of up to 100-fold (Fig. 3.2) and 120-fold (Fig 3.7). Even the lower acceleration factors of 5- or 7-fold reported for the human data (Figs 3.7 – 3.12) represent meaningful, enabling, scan-time advantages. Importantly, these acceleration factors are measured relative to CSI acquisitions that have already been truncated by dropping (22% of the) PEs from the corners of k-space. The 7-fold reduction for 3D ^{31}P cardiac CSI, for example (Fig. 3.11), slashes a barely tenable 39-min scan, which would be impractical for quantitative kinetic studies involving repeat acquisitions [27, 39, 41], to just 5½ min. This is short enough to include with a whole functional cardiac MRI study in a single patient exam. But even if the time-saving were not needed, averaging 7 SLAM

acquisitions, for example, could still deliver up to a $\sqrt{7}$-fold SNR benefit compared to a compartmental-average CSI spectrum acquired in the same time: averaging four 1D SLAM acquisitions was already shown to produce a $\sim\sqrt{4} = 2$-fold SNR gain for leg ^{31}P MRS [50], and a ~3-fold gain in SNR per unit volume and time was demonstrated in Fig. 3.7.

Second, we have combined SLAM with multi-receiver SENSE. State-of-the-art brain ^{1}H CSI often employs SENSE parallel imaging techniques to dramatically reduce the CSI PE set and hence scan-time. This raises questions of whether SLAM can do better than SENSE CSI, or even if, or how, SLAM could be combined with SENSE. The analysis and results from a 32-channel head-coil presented here show that, with conventional SENSE calibration, the signals from the individual receiver channels can actually be used in lieu of PEs in SLAM, for spatial encoding. Consequently, SENSE SLAM can provide significant additional scan-time acceleration on top of that provided by SENSE CSI. This was demonstrated with 5 to 6-fold acceleration factors in Figs. 3.7 – 3.10 and 3.12. The example of a 120-fold acceleration compared to SENSE CSI illustrates the extreme case of replacing *all but a single central k-space PE* with the SENSE inputs from the 32-channel coil (Figs. 3.7, 3.14). Indeed, in many cases adding SENSE results in a reconstruction that is many-fold over-determined for the number of compartments being solved. Note that while the choice of gradient PEs to be used in conjunction with the SENSE reconstruction is arbitrary, deteriorating error and SNR away from central *k*-space argue for always choosing the centermost *k*-space steps first.

Third, we introduced a modification to the SLAM algorithm—SLAM* (Eq. 3.7) that reduces errors due to the effects of signal heterogeneity within compartments [50].

Although the integral of the SRF computed with the prior SLAM reconstruction (Eq. 3.6) is forced to zero outside a desired segmented compartment, non-zero integrals can result when signal variations are large. While the summed dSRF magnitudes in Table 3.1 represent highly unlikely worst-case scenarios in which nothing at all cancels, the experimental studies suggest that in practice, incomplete cancellation primarily arises from the superficial scalp compartment. Here, variations in intense lipid signals are exacerbated by the OVS pulses, as well as by the field inhomogeneity at the scalp which is furthest from the magnet iso-center, resulting in lipid signals bleeding into the brain SLAM spectra. We showed that the spatially limited dSRF of the SLAM* reconstruction (Fig. 3.13) greatly ameliorates this effect. For this reason it is preferred, especially with SENSE. The cost of SLAM* is some broadening of the dSRF. The broadening decreases as the number of PEs or the compartment size increases, and also as the threshold for numeric regularization is reduced. The upshot, is that SLAM* and especially SENSE SLAM*, are relatively insensitive to inhomogeneity or gross maladjustment of the segmentation (Fig. 3.12), such that lipid contamination of SLAM* spectra differs little from the CSI standard (Figs. 3.4, 3.7, 3.8, 3.12). Even so, the dSRF is a function of the acceleration factor, the SENSE coil geometry, as well as the sample distribution and compartment segmentation, and will therefore likely require some tailoring to optimize results for a given study protocol.

Fourth, all of these techniques–2D, 3D, SLAM and SLAM*, with and without SENSE, have been reduced to practice *in vivo*, in human studies. Moreover, as the data in Figs. 3.5 and 3.9 attest, when these techniques are applied to [1]H MRS studies of patients with brain tumors and key [1]H MRS metabolites—Cho, CR, and NAA—are quantified

and compared with the average CSI spectra from the same compartments, the same results obtain. This, 5 times faster than SENSE CSI, which is already accelerated.

Fifth, we have shown that existing tools commonly applied to compensate CSI spectra for field inhomogeneities and eddy-currents, can conveniently be incorporated into SLAM on a voxel or compartment-by-compartment basis similar to CSI, without affecting the quantitative agreement between SLAM and CSI (Figs. 3.8 and 3.9). As usual, the treatment of spatial and temporal field inhomogeneities (Section 3.2.4), does require quantification of their behavior.

In Chapter 2, We reviewed SLAM in the context of earlier SLIM [10], GSLIM [56] and SLOOP [57] approaches noting their differences, as well as single voxel approaches which are limited to rectangular voxels that can't be altered post-acquisition [50]. SLIM, GSLIM and SLOOP could in principle offer ultra-high resolution and eliminate artifacts from signal bleed. However, this would require perfectly accurate compartmental segmentation and uniform compartments that may be difficult to realize in practice. Ultra-high resolution MRS, which would at least be needed to establish a standard for error quantification, may also be problematic. In contrast, SLAM accepts conventional CSI as the standard. It is therefore easily validated with existing CSI data sets, retro- or pro-actively acquired (Fig. 3.8, 3.10), and does not need ultra-high resolution MRS and accurate segmentation is not critical (Fig. 3.12).

In conclusion, SLAM is a new CSI-based MRS localization technique that offers huge reductions in scan-time and potentially large increases in volume coverage, while preserving SNR. The method does not even require scanner-side prescription of anatomic compartments. In our opinion, its reduction-to-practice in 2D, in 3D, with SENSE, with

eddy-current compensation, and with quantitative validation in ^{1}H MRS studies of 24 patients with brain tumors, should suffice to demonstrate that the technique is now ready for prime time.

Chapter 4 : Partial volume correction by grid shifting (PANGS)

4.1 Introduction

Magnetic resonance spectroscopy (MRS) is a valuable tool in the study of normal physiology and numerous diseases [28, 39, 44, 77-79]. Chemical shift imaging (CSI) [4], is a main-stay of multi-voxel MRS localization techniques, but its *in vivo* resolution is typically limited to 0.2–1 ml, primarily due to the low concentrations of the metabolites and moieties of interest (MOIs) and sensitivity. The low resolution renders CSI susceptible to partial volume errors (PVE) in signal intensity, due to tissue heterogeneity within CSI voxels, especially those near the edges of a subject being scanned.

In proton (^{1}H) brain CSI studies in particular, PVE from peripheral voxels often manifests as intense lipid signals from the scalp bleeding into adjacent brain voxels, which can obscure important metabolic information. A number of schemes for suppressing such lipid bleed artifacts, have been implemented. One common approach is to avoid lipid-rich regions altogether by using a PRESS [48] or STEAM [49] pre-localization sequence prior to CSI. Another is to saturate the unwanted signal areas using outer volume suppression (OVS) [70, 80, 81]. These methods are limited by the fact that the head's anatomy does not conform to the respective rectilinear localization gradients, leaving some voxels of interest either with significant residual PVE, or artificially attenuated by the suppression. Lipid suppression based on its longitudinal relaxation time (T_1) [82], for example using the T_1-null method [83, 84]; or those based on frequency-selective excitation [75, 85, 86] or signal refocusing [87, 88] are other options. These will also attenuate the MOIs by saturation and/or are susceptible to the effects of local field (B_0) inhomogeneity.

In any case, such methods can be combined with post-processing techniques, which are the focus of the present work. Typically, a spatial apodization (cosine) filter is applied to the *k*-space data prior to Fourier transformation to limit bleeding of intense peripheral lipid signals. Unfortunately this reduces the nominal spatial resolution by ~40% [76]. Another option is to extrapolate the *k*-space data based on the assumption that the lipid signal is strictly band-limited to the scalp, as identified in a separate MRI or lipid map extracted from the CSI data itself [89]. Dual-density weighted-average *k*-space acquisitions can also be used on data in which the lipid regions are defined: basically, high-resolution lipid data are acquired with low-averaging and low-resolution metabolite data are acquired at high-averaging [90, 91]. However, this strategy requires multiple averages and is unsuitable when operating with a single acquisition per phase-encode.

Here, a new method—*pa*rtial volume correctio*n* by *g*rid *s*hifting ('PANGS') is proposed to minimize PVE and signal bleed. The method iteratively shifts the spatial reconstruction coordinates in a specific region of image space, without affecting the *k*-space data or spatial resolution, to minimize PVE. PANGS is applied to ^{1}H brain CSI data from healthy volunteers and patients with brain tumors to reduce lipid contamination and resolve MOIs in spectra, and metabolite maps constructed therefrom. Measures of *N*-acetylaspartate (NAA), total creatine (CR), and choline (Cho) quantified in PANGS spectra are compared to those in CSI with and without cosine spatial filtering, to demonstrate lipid artifact suppression.

4.2   Theory

In CSI, spatial voxels are resolved by discrete Fourier transform (DFT) of the *k*-space signals[92, 93],

$$\mathbf{s}_{M*N} = \mathbf{PE}_{M*M} \times \boldsymbol{\rho}_{M*N} \, , \qquad\qquad (4.1)$$

where $\mathbf{s}$ is the data vector in k-space, $\mathbf{PE}$ is the Fourier phase-encoding operator in one

or more spatial dimensions, $\boldsymbol{\rho}$ is the spectral matrix vector in the spatial domain, M is the

number of phase-encoding steps and N is the number of temporal spectral data points.

$\mathbf{PE}$ is constructed as,

$$\mathbf{PE}_{M*M} = \begin{bmatrix} e^{i2\pi k_1 x_1/M} & e^{i2\pi k_1 x_2/M} & \cdots & {}_M/M \\ e^{i2\pi k_2 x_1/M} & e^{i2\pi k_2 x_2/M} & \cdots & {}_M/M \\ \vdots & \vdots & \ddots & \vdots \\ e^{i2\pi k_M x_1/M} & e^{i2\pi k_M x_2/M} & \cdots & x_M/M \end{bmatrix} \qquad (4.2)$$

where k_1, k_2, …, k_M refer to different phase-encoding gradient steps and x_1, x_2, …, x_M

refer to different reconstruction coordinates. k and x range from 0 to M-1, or from $-M/2$ to

$M/2$-1, depending on where the origin is chosen.

$\mathbf{PE}$ involves two conjugate variables that relate to the sampling grid in k-space

(k), and to the reconstruction grid in image space (x). For CSI, both the k-space and

image space grids are uniformly distributed, resulting in a uniform DFT for

reconstruction. When the DFT is uniform, the reconstruction coordinate is placed in the

geometric center of each voxel, with the implicit assumption that a point source in its

geometric center accurately represents the whole voxel (Fig. 4.1a). For voxels completely

filled with homogeneous tissue this assumption is true, and therefore the PVE is small

(~13%) [94, 95] and confined to adjacent voxels as shown in Fig. 4.1(b). In this case, the

non-zero bleed is due to intra-voxel dephasing from the CSI gradients [94]. When voxels

are only partially filled, such as in thin lipid layers in the scalp, the locations of their

centers-of-mass (COMs) differ significantly from their geometric centers (Fig. 4.1d),

causing much larger PVE that affects many voxels, as depicted in Fig. 4.1(e).

PANGS shifts the reconstruction coordinates (x) away from the geometric centers of designated voxels to match their COMs, resulting in a non-uniform (or non-integer x) DFT reconstruction. When the reconstruction coordinate is shifted to coincide with the voxel's COM, the PVE is dramatically reduced as shown in Fig. 4.1(f), whereas significant PVE is generated by shifting the reconstruction coordinate away from the COM (Fig. 4.1c).

4.2.1   Estimating the centers-of-mass (COMs)

The crucial step for PANGS is to determine the COM for each voxel, which are typically unknown. In practice, we estimate the COM based on prior knowledge. Specifically, to suppress lipid signal in ^{1}H brain CSI, a scalp region is defined from a co-registered anatomical image. Then, the reconstruction coordinates of the scalp voxels are iteratively shifted to minimize the lipid signals in the non-scalp region, without altering the raw k-space data.

Two strategies are utilized to expedite the iterative optimization. First, the COMs are quickly estimated by tracing the scalp on the co-registered MRI. These estimates are used as initial values for the iterations. Second, for each iteration, the reconstruction coordinates of only one or a small number of voxels are changed. Ordinarily, whenever any reconstruction coordinate is changed, the inverse of **PE** must be recomputed. However, the inefficiency of inverting a non-uniform DFT poses a major computational bottleneck for optimization. But if only a few reconstruction coordinates (corresponding to a few columns in **PE**) are changed in each iteration, the matrix inversion can be converted to a matrix multiplication [96] as

$$(\mathbf{PE} + \mathbf{\Delta PE})^{-1} = \mathbf{PE}^{-1} - \frac{1}{1 + trace(\mathbf{\Delta PE} \times \mathbf{PE}^{-1})} \mathbf{PE}^{-1} \times \mathbf{\Delta PE} \times \mathbf{PE}^{-1} \quad (4.3)$$

Here, **PE** and $\mathbf{PE}^{-1}$ are the encoding matrix and its inverse with non-uniform

reconstruction coordinates (*x*) from the previous iteration, and **ΔPE** is the updating

matrix of rank 1 (with only one non-zero column) describing the change of only one

reconstruction coordinate. The right side of Eq. (4.3) is repeated for each changed

coordinate. The computation time is significantly reduced with Eq. (4.3) because matrix

multiplication is more computationally efficient than inversion.
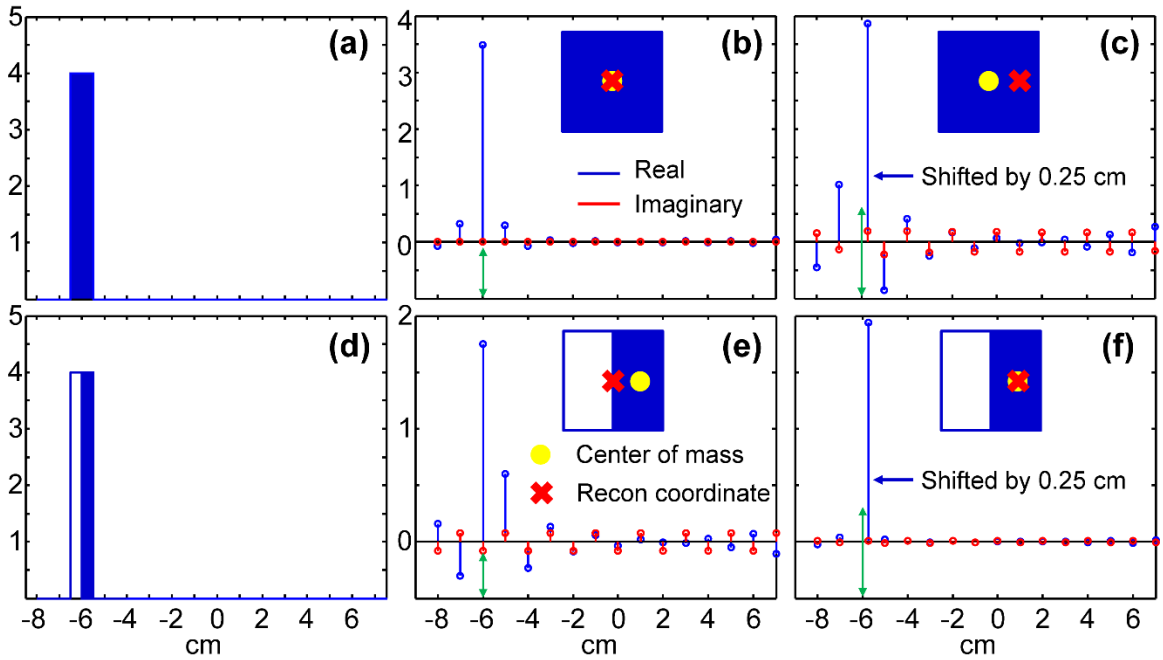
Figure 4.1: Simulated effect of shifting the reconstruction coordinates on a 16-step 1D simulation with 1 cm resolution. (a) 1 cm and (d) 0.5 cm thick uniform subjects in a 16 cm field-of-view. (b) and (e) show the single-point spectra reconstructed with standard CSI for (a) and (d), respectively. (c) and (f) show reconstructed spectra from PANGS with the third reconstruction coordinate shifted from -6 cm to -5.75 cm for (a) and (d), respectively. Minimum PVE results when the reconstruction coordinate is matched with the COM in (b) and (f). Blue and red circles denote the real and imaginary parts of the reconstructed spectra, respectively.

93

4.3   Methods

4.3.1   Processing and analysis

All processing and analysis was performed offline on a personal laptop computer (quad-core, 1.87 GHz).

4.3.2   PANGS optimization

For PANGS, raw k-space data acquired in ^{1}H CSI studies of the human brain was input as $\mathbf{s}$ to the left side of Eq. (4.1). Reconstructions were performed with the coordinates of each scalp voxel in $\mathbf{PE}$ (x) constrained to a range of $\pm 1/2$ of the voxel dimension from the voxel's geometric center. This generated many spectral data sets ($\boldsymbol{\rho}$) after solving Eq. (4.1). A 2Hz exponential filter and 4-fold zero-filling were applied to the time-domain data sets. Then, a metric of the lipid signal bleed was computed for each spectral data set by integrating the magnitude spectra from all non-scalp voxels over the spectral range of the lipid signal. The best reconstruction grid and spectral data set was taken as that exhibiting the minimum lipid bleed metric. The constrained optimization routine, '*fmincon*', of *Matlab 8.2* (Mathworks, Natick, MA) was deployed to iteratively adjust the reconstruction coordinates for scalp voxels and automatically select the best combinations of voxel locations to achieve a 'global' optimum. To expedite optimization, the starting values for *fmincon* were the estimated COMs based on the intensities of the high-resolution anatomical images, and Eq. (4.3) was incorporated into the objective function used in *fmincon*. Relative to these starting points, the reconstruction coordinates were further constrained to within $\pm 1/4$ of the voxel size using the 'Active-Set' minimization algorithm [97, 98] (with termination tolerance on independent variables = 0.01; maximum of 100 iterations; all other *fmincon* parameters were default values).

4.3.3 Spatial response functions

Spatial response functions (SRFs) [31, 50] were computed to evaluate the mechanisms giving rise to PVE, and its potential reduction by PANGS. The SRF of a single voxel spectrum at -4 cm was calculated (Fig. 4.1), and the effects of spatial cosine filtering (weighting, k-space center/edges = 1/0) investigated.

4.3.4 Quantitative results and metabolite maps

To investigate PANGS' ability to suppress PVE, PANGS spectra from cortical regions were compared to those from CSI with and without cosine filters. For quantitative comparisons, eight regions-of-interest (ROIs) each comprised of 4 CSI voxels, were selected from the central field-of-view (FOV) of the brain, where corruption of the CSI scans by PVE is expected to be minimal. Average ROI spectra from PANGS and CSI were fitted in the frequency domain using the 'circle-fit' method [99] and identical input parameters. Metabolite maps from CSI and PANGS were computed by integrating absolute spectra over 1.9–2.1 ppm for NAA, 2.9–3.1 ppm for CR and 3.1–3.3 ppm for Cho MOIs. For display, the integrated spectral areas from the scalp and background were set to zero, and the brain metabolite maps were linearly interpolated to four-times the originally acquired spatial resolution in each direction.

Because the spectral range of the lipid resonances can be broad due to poor shimming and imperfect OVS suppression, the effect of using different lipid frequency ranges on PANGS was investigated by comparing metabolite maps reconstructed with the lipid bleed metric minimized over 0.5–1.75 ppm to those minimized over 0.5–3 ppm. The larger lipid range for PANGS minimization is preferable for more stable outcomes, but can raise concerns about the influence of the NAA signal on the outcome.

4.3.5 Global PANGS vs. row-by-row PANGS

To further expedite optimization, a compromised 'row-by-row' PANGS approach was devised and the results compared to the 'global' optimization approach above. The central difficulty for iterative minimization is that the number of possible COM locations grows exponentially with the number of scalp voxels. For example, if there are 70 voxels in the scalp region and each voxel can have 10 trial locations, the total number of combinations is 10^{70}—a daunting prospect to evaluate exhaustively. Although the *fmincon* function can deal with this quite well, it is rather slow. However, due to the separability property of the Fourier transform, if a row of voxels from the x (or y) direction of a 2D CSI data set is selected, any reconstruction coordinate shift in that row along the x (or y) direction will only affect the spectra of that row. This allows PANGS to be implemented in a row-by-row manner instead of the global approach above. The number of possible trials in each row can be dramatically reduced, e.g. 4 scalp voxels each with 10 possible locations in a row, results in 10^4 combinations. The *fmincon* function (with maximum iterations reduced to 50) was still used for this row-by-row PANGS implementation, even though an exhaustive search was now practical.

For the 20x18 phase-encoded human data, row-by-row PANGS was iterated 38 times. The compromise in the row-by-row PANGS compared to the global PANGS optimization, is that once a row is optimized it won't be revisited. This reduces the likelihood of finding global optimal coordinate combinations, but dramatically reduces computation time (typically, by ~ 20-fold).

4.3.6   Human studies

All human studies were approved by the Johns Hopkins Institutional Review Board and were performed on the 3T Philips (Best, Netherlands) Achieva system. Informed consent was obtained from all participants.

[1]H MRI and CSI data from three healthy volunteers were acquired with a head coil configured in quadrature mode. A single-slice spin-echo CSI sequence [47, 70] (repetition time, TR = 2 s; echo time, TE = 144 ms; in-plane resolution = 1x1 cm; FOV = 20x18 cm; slice thickness, ST = 1.5 cm; and 12.1 min scan duration with two startup cycles) was used with OVS lipid suppression [70] and VAPOR [72] water suppression pre-pulses. Anatomic images were acquired with a 3D T_1-weighted MP-RAGE [100] sequence (flip-angle, FA = 8°; TR = 6.9 ms; TE = 3.2 ms; resolution = 1x1x1.2 mm; FOV = 256x228x156 mm; scan duration = 4.5 min; inversion time, TI = 813 ms). The CSI voxel grid was co-registered with anatomical images to identify the scalp.

Three patients with brain tumors were studied using the scanner's 32-receive-channel head coil. The single-slice spin-echo [1]H CSI sequence (same parameters as above except: FOV = 22x18 cm; ST = 13.2 mm; scan duration = 13.3 min with two startup cycles) was employed with OVS and VAPOR pre-pulses. A multi-slice FLAIR [73, 74] sequence was used to provide anatomical images (TR = 11 s; TE= 120 ms; TI = 2.8 s; refocusing angle = 120°; resolution = 0.83x1.04x2.2 mm; FOV= 212x189x132 mm; scan duration = 3.9 min; SENSE [6] acceleration factor = 1.5). For CSI with and without cosine filters, the multi-channel patient data was processed in the same way as the quadrature-coil studies of healthy volunteers, except that spectra from individual channels were combined using a root-of-the-sum-of-the-squares method [101, 102]. For

PANGS, all individual channels shared the same reconstruction coordinates and results were combined using the same root-of-the-sum-of-the-squares method.

4.4 Results

4.4.1 Simulated SRF

Fig. 4.2 compares the SRF of CSI with that after PANGS is applied. The SRF of the spectrum of a MOI at -4 cm has a zero-crossing at -6 cm, where we posit that an intense scalp lipid resonance resides (Fig. 4.2a). If the voxel at -6 cm voxel contains uniformly distributed lipid tissue, it can reasonably be approximated by a point source at -6 cm. Since the SRF of the -4 cm spectrum has a zero-crossing at -6 cm, the MOI spectrum at -4 cm will be minimally affected by signals from the -6 cm voxel, whose contributions at -4 cm are as depicted in Fig. 4.1(b). However, if the -6 cm lipid voxel is not well represented by a point source at -6 cm but rather by a point source at -5.75 cm as in Fig. 4.1(d), then the effect of PANGS is to shift the reconstruction coordinate of the zero-crossing of the SRF from -6 cm to the COM at -5.75 cm (Fig. 4.2b). As a consequence, the -6cm lipid signals now have negligible bleed at the other zero-crossings where the reconstruction grid is sampled (Fig. 4.1f), including at the location of the MOI in the -4 cm spectrum (Fig. 4.2b). Cosine spatial apodization, on the other hand, disrupts the zero-crossings of the CSI SRF (Fig. 4.2c), which constrains the spatial extent of the SRF at the cost of a broadened spatial linewidth.
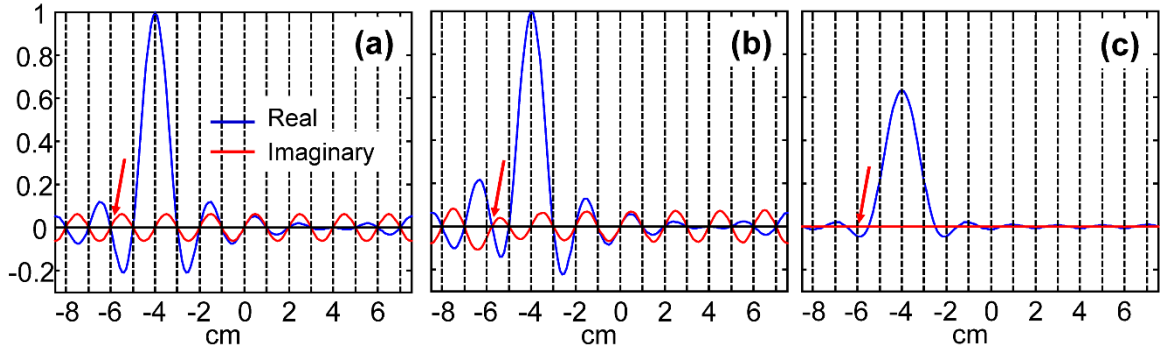
Figure 4.2: Spatial response functions (SRFs) of (a) CSI, and (b) PANGS with a 0.25 cm shift, for a voxel at -4 cm. Arrows denote the SRF amplitude at the sampled location of the subject in Fig. 1(d). (c) The cosine-filtered SRF corresponding to (a). Blue and red lines denote the real and imaginary parts, respectively. In practice the voxels are discretely sampled at the vertical dashed lines, except for the -5.75 cm point in (b) (arrow).
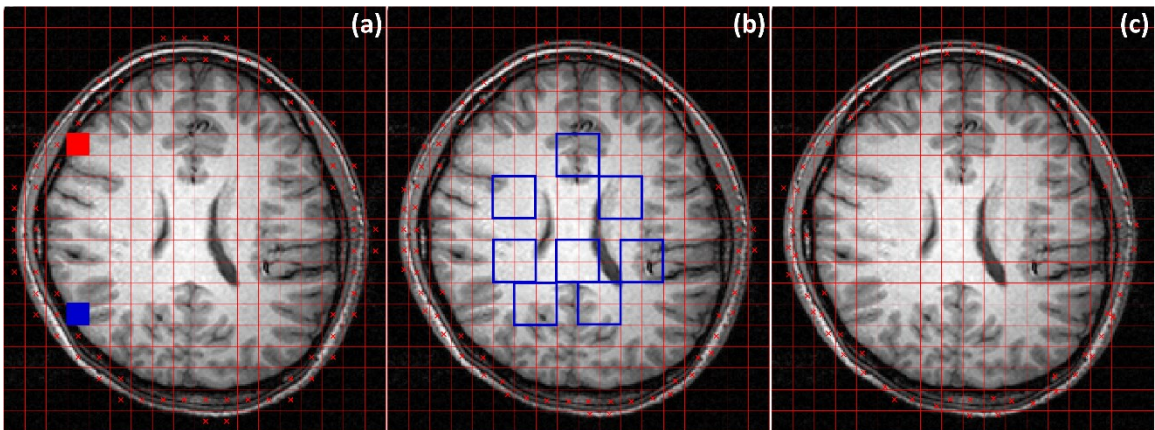
Figure 4.3: CSI grids of voxels co-registered with anatomical MRI and labelled with scalp voxels (x = reconstruction coordinates). (a) The reconstruction coordinates are coincident with the geometric centers of each voxel as in conventional CSI. Spectra from voxels shaded red and blue are displayed in Fig. 4.4. (b) Estimated COMs based on the co-registered high-resolution anatomical MRI, used as starting coordinates for PANGS. Blue outlines denote 8 ROIs quantified for Table 4.1. (c) Final PANGS coordinates that minimize the lipid signals in non-scalp regions.

4.4.2   Studies of healthy subjects

For conventional CSI, the reconstruction coordinates were at the geometric centers (Fig. 4.3a). Inspection reveals that the tissue COMs of most of the voxels in the scalp did not coincide with their geometric centers. Estimated COMs based on anatomical images are shown in Fig. 4.3(b). After minimizing the lipid metric value over 0.5–1.75 ppm, the reconstruction coordinates of the scalp voxels no longer reside at the geometric centers as shown in Fig. 4.3(c), but the non-scalp voxels are unmoved.

Fig. 4.4 shows that PANGS spectra (central column; Fig. 4.4b, e) have significantly reduced baseline contamination and artifacts (red arrows) compared to CSI, without (left; Fig. 4.4a, d) and with (right; Fig. 4.4c, f) cosine filtering. The CSI spectra without cosine filtering are contaminated by PVE due to scalp lipid signals adjoining these cortical voxels. The cosine-filtered CSI spectra are also corrupted by the filter's spatial broadening effect.

Table 4.1 presents a comparison of NAA, CR, and Cho peak areas quantified in the eight centrally-located voxels outlined in blue in Fig. 4.3(b), in unfiltered CSI, PANGS and cosine-filtered CSI. MOI peak areas from PANGS and filtered CSI did not differ significantly from non-filtered CSI (P > 0.2, in all cases). The mean MOI peak area measured in PANGS spectra differed by < 1.5% from filtered and non-filtered CSI for all metabolites. However, the fitted spectral linewidths of spatially-filtered CSI spectra were significantly broader than in CSI without spatial filtering (P < 0.03, all MOIs), which in turn did not differ from PANGS (< 1%; P > 0.2, all MOIs). The small (≤ 6%) increase in linewidth of filtered vs. non-filtered CSI likely reflects the filter's broadening effect compared to non-filtered CSI and PANGS.
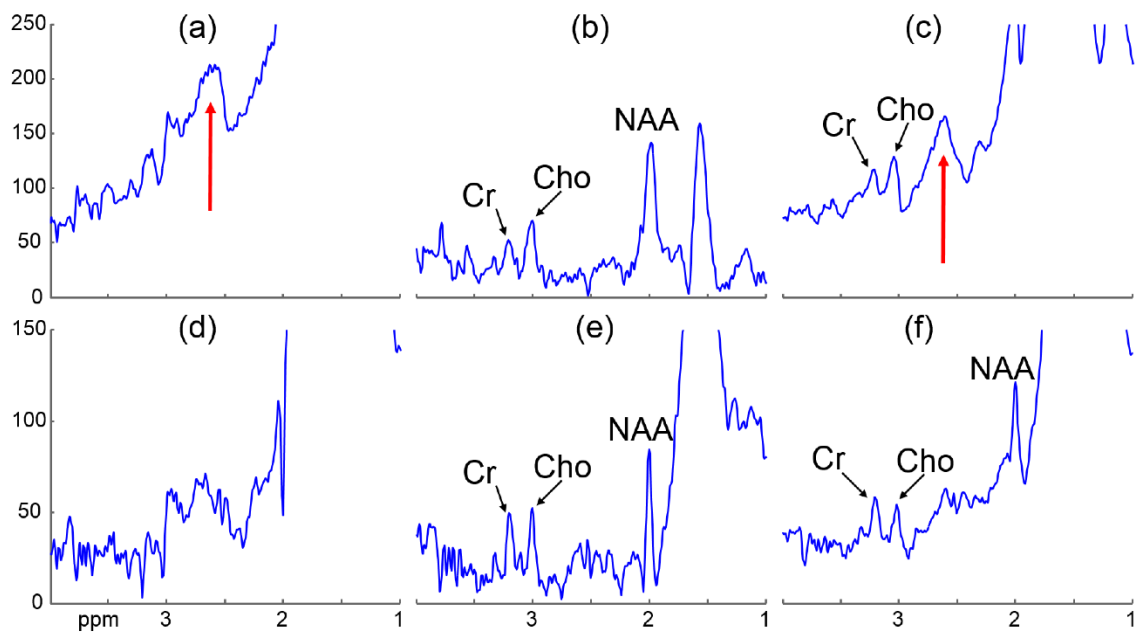
Figure 4.4: Magnitude spectra from CSI and PANGS from (a-c) the solid blue voxel, and (d-f) the red voxel in Fig. 4.3(a). Parts (a) and (d) are from CSI without a cosine filter. Parts (b) and (e) are the PANGS spectra. Parts (c) and (f) are from CSI with a cosine filter. Red arrows in (a) and (c) indicate an artifact peak. No baseline corrections were applied.

Table 4.1: Metabolite peak areas (arbitrary units) and linewidths (Hz) quantified in non-filtered CSI, PANGS and filtered CSI from the 8 selected ROIs outlined in blue in Fig. 4.3(b).

CSI w/o cosine	PANGS	CSI w/ cosine	CSI w/o cosine	PANGS	CSI w/ cosine
NAA peak areas (a.u.)			**NAA peak linewidth (Hz)**		
4931	4866	4803	7.38	7.31	7.58
5007	5128	5013	5.98	6.11	6.26
4930	4956	5055	6.17	6.19	6.87
5071	5015	4993	7.62	7.56	7.90
4700	4872	4826	10.48	10.77	10.46
5628	5593	5373	7.64	7.60	8.22
3910	4023	4275	5.09	5.19	5.72
4443	4521	4648	6.77	6.86	7.30
Error[a]	1.03%	1.33%		0.78%	6.22%
p-value[b]	0.21	0.54		0.24	0.003
CR peak areas (a.u.)			**CR peak linewidth (Hz)**		
1930	1995	1968	6.95	7.13	7.19
2060	2099	2191	6.68	6.81	7.15
2730	2688	2638	9.25	9.14	9.43
2423	2443	2357	8.19	8.23	8.48
2809	2794	2585	11.19	11.19	10.88
1816	1841	1916	7.28	7.35	8.06
2558	2436	2572	6.84	6.54	7.23
2599	2571	2556	8.40	8.35	8.81
Error[a]	-0.05%	-0.16%		0.00%	4.31%
p-value[b]	0.74	0.67		0.95	0.027
Cho peak areas (a.u.)			**Cho peak linewidth (Hz)**		
1627	1667	1667	6.19	6.30	6.58
1737	1749	1736	5.01	5.05	5.25
2023	2009	1928	5.68	5.65	5.89
2263	2272	2359	6.35	6.36	7.01
2001	1998	2062	8.07	8.07	8.47
2140	2159	2091	6.99	7.03	7.39
1996	1869	1880	6.78	6.47	6.69
1762	1744	1796	6.04	6.02	6.31
Error[a]	-0.47%	-0.13%		-0.27%	4.87%
p-value[b]	0.59	0.90		0.66	0.005

[a] ROI-wise difference from non-filtered CSI in percentage and averaged for 8 ROIs.
[b] paired t-test vs. non-filtered CSI.
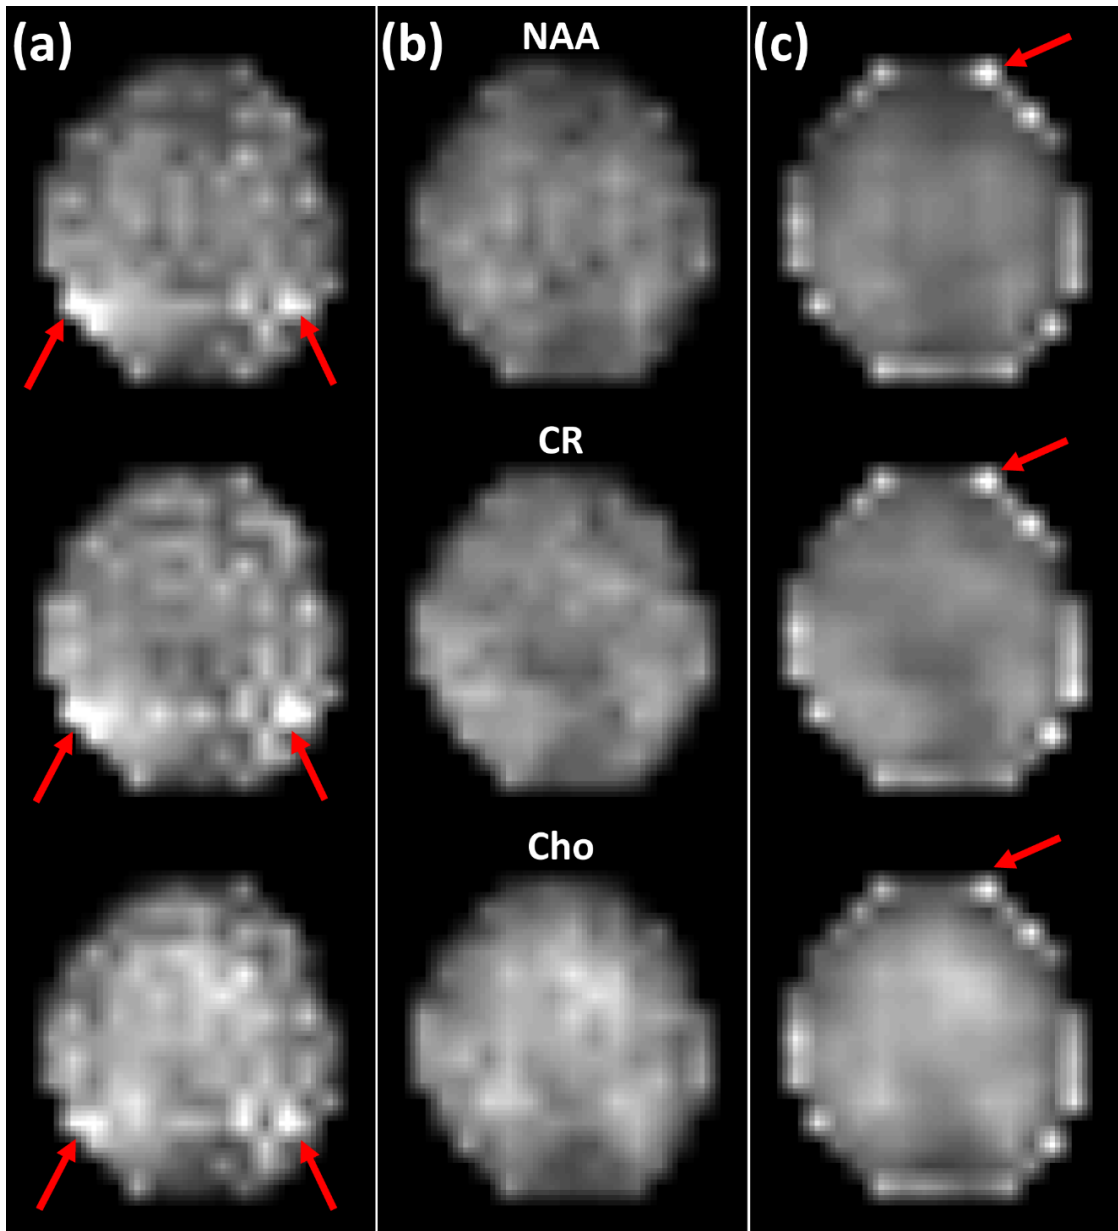
Figure 4.5: Metabolite maps of NAA, CR and Cho (top to bottom) from a healthy volunteer obtained with regular CSI and PANGS. (a) CSI without a cosine filter; (b) PANGS with lipid signals minimized over 0.5–1.75 ppm; and (c) CSI with a cosine filter. Image contrast thresholds are set identically. Red arrows indicate hyper-intense lipid artifacts.

Moreover, metabolite maps created from the PANGS data (Fig. 4.5b) were essentially free from hyper-intense lipid bleed artifacts seen in CSI (red arrows, Figs. 4.5a, c). In addition, bright rim artifacts (Fig. 4.5c) from the reduced spatial resolution of cosine-filtered maps, are absent from the PANGS images. Note that the fmincon routine is not guaranteed to yield a globally optimal result even though the lipid signal metric value monotonically decreases during the optimization process.

4.4.3   Global PANGS vs. row-by-row PANGS

Figs. 4.6 shows the results of applying the global PANGS optimization to generate metabolite maps minimizing lipid signals over 0.5–1.75 ppm (Fig. 4.6a) and 0.5–3 ppm (Fig. 4.6b). There are no significant differences between (a) and (b). This means that NAA signal contributions do not appreciably affect PANGS lipid signal minimization, even when the NAA and lipid spectral frequency ranges overlap. In addition, little difference is evident between the row-by-row PANGS algorithm (Fig. 4.6c) and the more-exhaustive global PANGS algorithm (Fig. 4.6a). However, the final lipid signal metric from row-by-row PANGS is about 15% higher than the global PANGS result, which represents a trade-off for a ~16-fold reduction in computation time (1.2 min vs. 18.7 min) in this implementation.

4.4.4   Brain tumor patients

Fig. 4.7(a) shows a FLAIR MRI from a patient following surgery for a high-grade tumor (arrow, top), with the PANGS optimized scalp reconstruction coordinates overlaid (bottom). Fig. 4.7(c) shows NAA, CR and Cho images reconstructed via row-by-row PANGS for comparison with CSI without (Fig. 4.7b) and with (Fig. 4.7d) spatial apodization. Here, the lipid frequency range for minimization was extended to 0.5–3 ppm

104

to deal with a more challenging deficiency in OVS lipid suppression. The computation time for row-by-row PANGS was 58.3 minutes due to the large amount of data from 32 channels. However, global PANGS would have taken hours, which might not be considered practical, even if it did provide better results than row-by-row PANGS. The hyper-intense lipid artifacts in the regular CSI scan (Fig. 4.7b, top) are eliminated in the PANGS NAA map (Fig. 4.7c, top), and artifacts in the cortical region of the cosine-filtered image (Fig. 4.7d. top; red arrow) are attenuated. The spatial resolution of cosine filtered CSI is worse than both non-filtered CSI and PANGS, which can be seen from the blurrier definition and increased size of the lateral ventricle and surgical lesion in the CR and Cho images (green arrows, Fig. 4.7d).
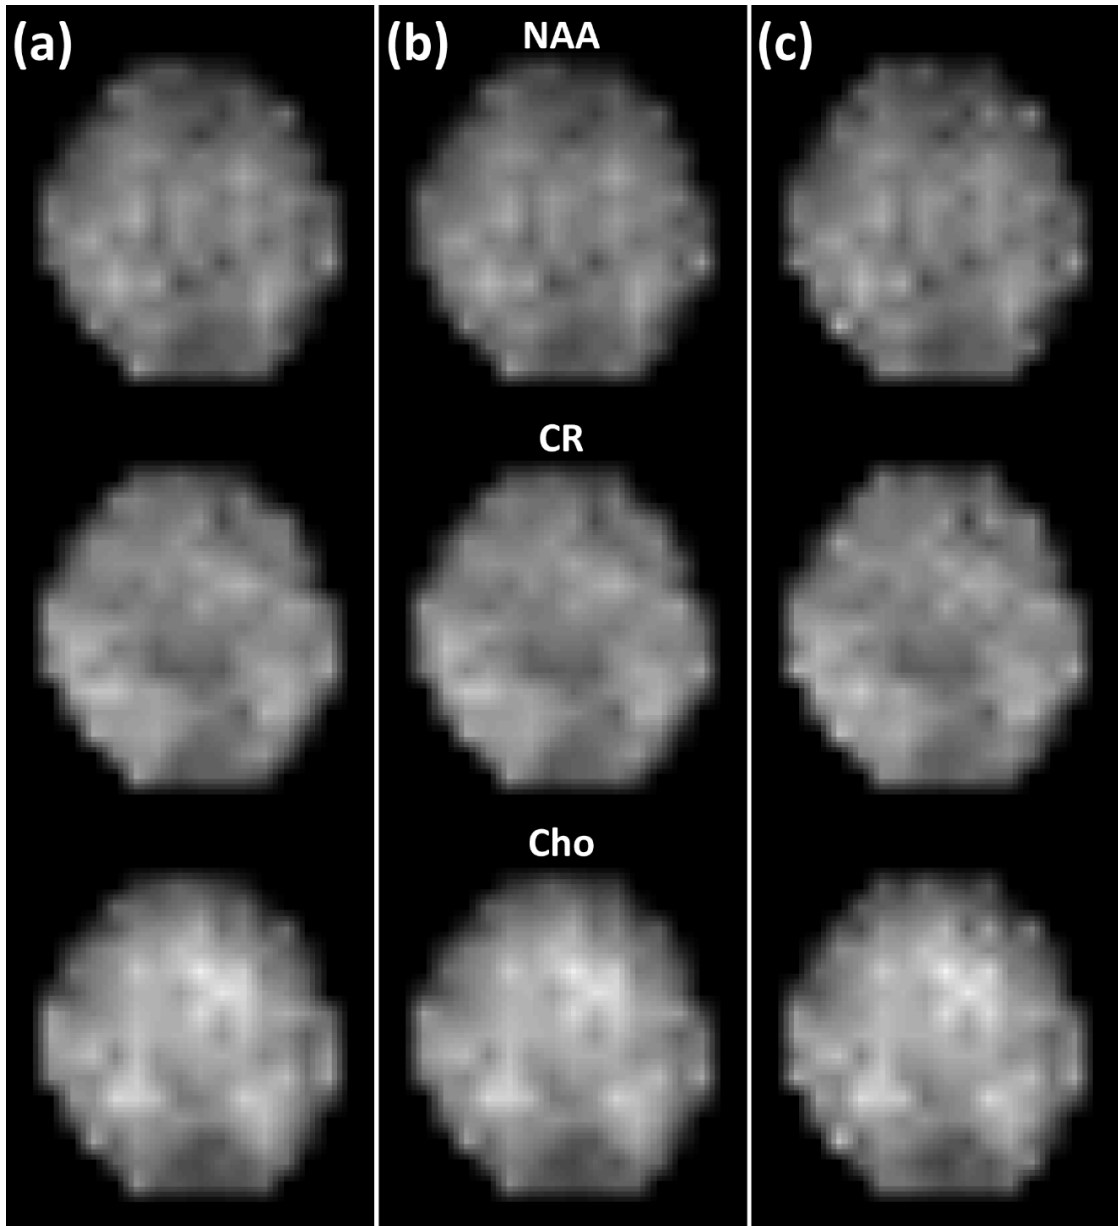
Figure 4.6: NAA, CR and Cho metabolite maps (top to bottom) obtained from a healthy subject and reconstructed with: (a) global PANGS with lipid signals minimized over 0.5–1.75 ppm; (b) global PANGS and lipid minimized over 0.5–3 ppm; and (c) row-by-row PANGS and lipid minimized over 0.5–1.75 ppm. Reconstruction times were 18.7, 31.9, and 1.2 min respectively. Image contrast thresholds are set identically.
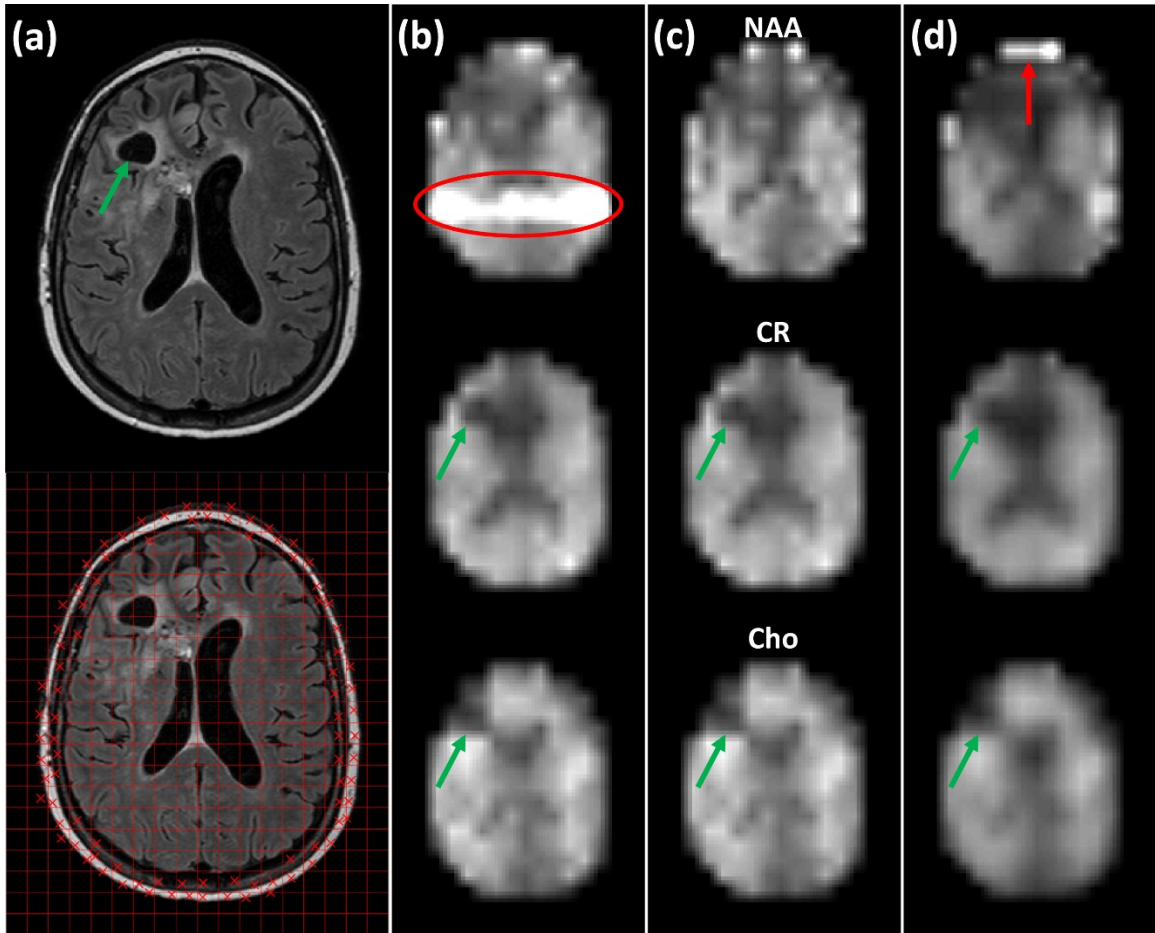
Figure 4.7: (a) FLAIR image from a patient with a resected glioblastoma, and overlaid with optimized PANGS scalp coordinates (x, bottom). Parts (b-d) are corresponding NAA, CR and Cho metabolite maps (top to bottom) reconstructed from: (b) CSI without a cosine filter, (c) PANGS and (d) CSI with a cosine filter. Image contrast thresholds are set identically. Regions outside the brain (including scalp) are masked from these maps. Hyper-intense lipid artifacts are indicated in red in (b) and (d). Green arrows indicate the surgical resection.

4.5   Discussion

In brain ^{1}H CSI, the scalp, being peripheral, is most susceptible to the effects of B_0 as well as RF-field inhomogeneity. The upshot is that multiple techniques are typically required to suppress the scalp's intense lipid resonances, which nevertheless often fail for nearby voxels that may include pathologies. Here, we introduced a new post-processing method, PANGS, which when combined with conventional OVS, provided excellent lipid suppression compared to standard non-filtered and filtered CSI, and reduced lipid artifacts in MOI maps (Figs. 4.5 and 4.7). Moreover, PANGS generated quantitatively the same fitted metabolite areas as conventional CSI in central voxels that were unaffected by lipid artifacts (Table 4.1), while improving the detection and resolution of metabolites in cortical regions (Fig. 4.4). The use of nonuniform DFT affords PANGS this unique ability to simultaneously correct PVEs while maintaining the spatial resolution of the original acquisition.

Spatial apodization is ubiquitous in CSI. It can effectively constrain the bleed due to PVE, and has a side-benefit of increasing the signal-to-noise ratio (SNR) [103], but the downside of reducing spatial resolution is largely ignored. Its main value for ^{1}H brain CSI is the suppression of scalp lipids, which primarily affect NAA. However, lipid suppression by spatial filtering may be unnecessary for MOIs spectrally distant from the lipid, such as Cho and CR (Fig. 4.7b), which apodization may needlessly compromise. The SNR benefit of spatial apodization derives from voxel averaging: the cosine filtering applied here is equivalent to shifting the FOV by a 1/2-voxel and averaging the result with non-shifted voxels [104]. Even if the broadening of the spatial resolution in cosine-filtered CSI is of no concern, PANGS can still offer improved spectral resolution in

cortical voxels (Fig. 4.4, Table 4.1). The spectral broadening of cosine-filtered CSI is attributable to the combined effects of spatial broadening and B_0-inhomogeneity including variations in phase.

Compared to k-space extrapolation [89], image-space PANGS does not assume that the lipid signals are limited to a spatial band in the scalp or are specifically reduced elsewhere, for example, it doesn't assume that brain MOI signal strengths are $\leq 1/6^{th}$ that of the scalp lipid [89]. Here, with OVS used in conjunction with PANGS, the lipid signals may not have been 6-fold higher than brain MOIs at some locations, which could have undermined such a k-space extrapolation. An accurate definition of the scalp is in fact not critical to PANGS. The only consequence of excluding a few lipid-rich voxels from the scalp compartment, is that the PVE bleed artifact from those voxels won't be corrected. If the scalp lipids are completely suppressed, the result of PANGS is identical to non-filtered CSI.

Lastly, the speed of the iterative optimization process could benefit many-fold from parallel computing techniques. Also, PANGS could have other applications as well, such as reducing ringing artifacts in MRI, where a possible objective function could involve minimizing the background signal energy. For now, we posit that PANGS is a flexible post-processing technique that can significantly reduce PVE bleed artifacts in metabolite maps and spectra, and improve metabolite detection in cortical regions without altering k-space data or compromising spatial resolution.

Chapter 5 : Accelerated and motion-corrected MR endoscopy[1]

5.1   Introduction

Unlike conventional MRI with external detector coils, intravascular (IV) MRI receives the signal using tiny internal detectors [105, 106]. At field strengths of 3T, these can afford high signal-to-noise ratios (SNR) that permit 80-300 μm resolution imaging of vessel walls and associated pathology, as fast as several frames per second (fps) [107, 108]. With suitably modified sensitivity profiles, transmit/receive IV MRI probes can also provide high-resolution imaging from the probe's point-of-view [109], analogous to optical endoscopy, IV Ultrasound, and Optical Coherence Tomography. Nevertheless, while excelling in soft-tissue contrast, 3T IV MRI still lags behind these other modalities in speed.

Factors contributing to IV MRI's slowness include the physics of spatially encoding the MRI data, and the delays required to accommodate magnetic resonance (MR) relaxation. MRI uses linear magnetic field gradients that encode the object in the spatial frequency domain (k-space). The encoding scheme is ordinarily subject to the Nyquist criterion which requires that sufficient k-space be filled to avoid aliasing errors upon Fourier reconstruction. Satisfying Nyquist contributes directly to the long acquisition times. MRI scan-times can be reduced using recent 'compressed sensing' (CS) reconstruction techniques [61, 110] which permit image reconstruction with a limited subset of k-space. Successful CS implementation requires 'sparse' data with sufficient SNR, and has been used in several conventional MRI settings including neurological, cardiac and dynamic applications [61, 110-112]. However, to date CS has not been

---

[1] The work in this chapter was primarily conducted by Dr. Shashank S. Hegde. My role was primarily associated with developing / implementing compressed sensing (CS) for the IVMRI probes.

applied to IV MRI whose intrinsically high local SNR, spatial sparsity and need for speed, make it an ideal application.

Slow scan speeds which are endemic to ultra-high resolution MRI, also render IVMRI susceptible to the effects of physiological motion—blood flow and respiration, as well as motion during probe advancement. Conventional cardiac- and breath-gating techniques reduce motion-sensitivity, but typically restrict speed and scan-time, depending on the period of the corresponding motion. While ungated acquisitions could allow frame-rates limited only by the SNR (or spatial resolution), their implementation would require other strategies for motion compensation. In past, several strategies for correcting conventional projection-reconstruction (PR) images for motion have employed data consistency criteria based on the zeroth and/or higher-order moments of each image projection [113-115]. These methods are easily affected by intensity variations and can fail in high-resolution IVMRI where differences in the order of a millimeter can translate to several tens of pixels.

In this work we first address the speed limitation by developing CS reconstruction techniques for under-sampled IVMRI. We use these to demonstrate effective frame-rate acceleration factors of up to eight-fold. Second, we present a motion-correction method for ungated PR IVMRI. The method reduces sensitivity to motion using data consistency criteria based on the intrinsic amplitude and phase properties of the IV MRI detector's sensitivity profile. The correction algorithm is applied at the repetition period (TR) of each projection prior to reconstruction, and is therefore amenable to real-time application to the data stream, including MRI endoscopy [109]. We show high-resolution results in fruit, human vessel specimens, and a rabbit aorta *in vivo*.

5.2   Theory

5.2.1 Accelerating image acquisition

In two-dimensional (2D) Cartesian MRI, the data are acquired to fill a 2D rectilinear grid in k-space and the images are reconstructed by 2D Fourier Transformation (FT). On the other hand, the PR method fills k-space along radial spokes, and the FT of each radial spoke in k-space corresponds to a projection in image space through the object at an angle perpendicular to the direction of the spoke. This direct relationship permits image reconstruction in either domain, by regridding the radial k-space data onto a rectangular grid followed by 2D FT using a Non-Uniform fast FT (NUFFT) algorithm, or using filtered back-projection techniques in image space, as in Computed Tomography. The two methods are equivalent, so we limit the description to the former approach, implemented here with an open-source NUFFT algorithm [116].

For alias-free reconstruction, both Cartesian and PR acquisition schemes ordinarily require sufficient k-space samples to satisfy the Nyquist criterion. Under-sampling in Cartesian MRI results in artefacts that appear as replicates of the original image displaced in the under-sampled dimension. Under-sampling artefacts in PR acquisitions are manifest as hyper-intense radial lines or 'streaking'. CS reconstruction ameliorates these artefacts by tailoring the under-sampling scheme so that the artefacts appear 'noise-like', and uses an iterative reconstruction that incorporates penalties that suppress the noise in the final image. In Cartesian acquisitions, the central k-space is fully-sampled but outer k-space is randomly sampled [61]. For PR, k-space is sampled with a uniform angular distribution [110].

For CS, image reconstruction is set up as a constrained minimization problem, where a functional, $\Phi(\mathbf{x})$, is minimized.

$$\Phi(\mathbf{x}) = \|A\mathbf{x} - \mathbf{y}\|_2 + \sum_n \lambda_n R_n(\mathbf{x}) \tag{5.1}$$

Here $\mathbf{y}$ is the measured k-space complex data from all (PR) spokes or (Cartesian) lines stacked to form a vector, $\mathbf{x}$ is the image vector to be estimated, A evaluates the image's FT in undersampled k-space, the ℓ vector norm $\|\mathbf{z}\|_2 = (\sum_i |z_i|^2)^{1/2}$, and there are n penalty functions $R(\mathbf{x})$ with weighting factors λ to constrain the resultant image based on *a priori* knowledge.

Suitable penalty functions that minimize the coefficients of the underlying image in a sparse transform space such as the Wavelet Transform (WT) or spatial finite-differences (minimizing Total Variation, TV), have been described earlier [61, 110]. The WT is a multi-scale representation of the image with coarse- and fine-scale wavelet coefficients representing low- and high-resolution image components, respectively. The TV constraint assumes that the underlying image consists of areas with constant (or mildly varying) intensity. The penalty functions using these transforms are

$$R_{\text{WT}}(\mathbf{x}) = \|\Psi(\mathbf{x})\|_1 \tag{5.2}$$

$$R_{\text{TV}}(\mathbf{x}) = \sum_i |D_X(x_i)| + |D_Y(x_i)| \tag{5.3}$$

where Ψ is the WT and the ℓ vector norm $\|\mathbf{z}\|_1 = \sum_i |z_i|$, D_X, D_Y denote the derivatives in X and Y direction, respectively.

An additional penalty function was designed based on the field profile of IVMRI probes. For a loopless antenna oriented parallel to the main field (z-axis), the trans-axial RF-field, B_1, can be approximated by that of a long conductor with current I:

113

$$B_\phi = \frac{\mu I}{2\pi r} \tag{5.4}$$

where ϕ is the azimuthal angle, r is the radial distance from the probe, and μ is the permeability of the medium. Even when the probe is not parallel to the z-axis, $B_\phi \propto 1/r$ still holds (because MRI is only sensitive to the transverse B_1 component) causing hyper-intensity at the probe location (Figs. 5.1a, b). Undersampling leads to streaking artefacts with spokes emanating from the probe location (Fig. 5.1c). The penalty function smears the variation between the spokes (thereby reducing 'spokal variation' - SV) for a more uniform image. We divide the image pixels into concentric annular sectors, compute the difference of complex pixel-sums between azimuthally adjacent sectors and sum the absolute difference over all sectors (Fig. 5.1d).

$$SV(u,v) = \sum x(u,v+1) - \sum x(u,v) \tag{5.5}$$

$$R_{SV}(\mathbf{x}) = \sum_{u,v} |SV(u,v)| \tag{5.6}$$

where $x(u,v)$ denotes the pixels in the u^{th} annulus and v^{th} sector. With the penalty functions of Eqs. (5.2), (5.3) and (5.6), the functional in Eq. (5.1) becomes

$$\Phi(\mathbf{x}) = \|A\mathbf{x} - \mathbf{y}\|_2 + [\lambda_{WT} R_{WT}(\mathbf{x})] + [\lambda_{TV} R_{TV}(\mathbf{x})] + [\lambda_{SV} R_{SV}(\mathbf{x})] \tag{5.7}$$

This functional was minimized using iterative CS algorithms as noted below (IRGNTV [110, 117], GRASP [118], SPARSE-MRI [61]). These algorithms involved searching for minima using Gauss-Newton and/or conjugate gradient methods by sequentially varying each of the three λ terms in Eq. (5.7). Optimum values of λ were selected by comparing sample datasets from the undersampled and the fully-sampled datasets using the Structural Similarity index (SSIM [119]),

$$SSIM(A,B) = \frac{4\mu_A \mu_B \sigma_{AB}}{(\mu_A^2 + \mu_B^2)(\sigma_A^2 + \sigma_B^2)} \quad (5.8)$$

where A and B are the images being compared, σ_{AB} is their cross-correlation, μ and σ are the mean and standard deviation of the pixel intensities (if $A=B$, SSIM = 1). Images are also visually inspected for any loss in tissue structure.

5.2.2 Motion Correction

In-plane motion

Since the IVMRI probe is embedded within the body, in-plane body motion could cause the location of the probe to vary within the imaging plane. However, if the probe location could be determined in each projection, then shifting or aligning each projection at the probe location should ameliorate the effect of the motion.

To detect the probe in each projection, we see from Eq. (5.4) and Figs. (5.1a, b) that: (i) the field strength falls rapidly with distance r from the probe; and (ii) the detection phase varies azimuthally around the probe and reverses direction at the probe in every projection. In addition, (iii) the probe itself is metallic and does not contribute any signal. Thus, each projection in image-space has a volcano-like signal intensity maximum close to the probe, a dip or 'crater' in intensity at the probe location (Fig. 5.2a, top), and a phase reversal at the exact location of the probe (Fig. 5.2a, bottom). We found that a combination of both amplitude and phase detection provided the most robust means of detecting the probe. The center-of-mass (COM) of each projection served as the starting point. The probe's approximate location was found by amplitude detection, which was also used as the starting point for the phase detection algorithm.
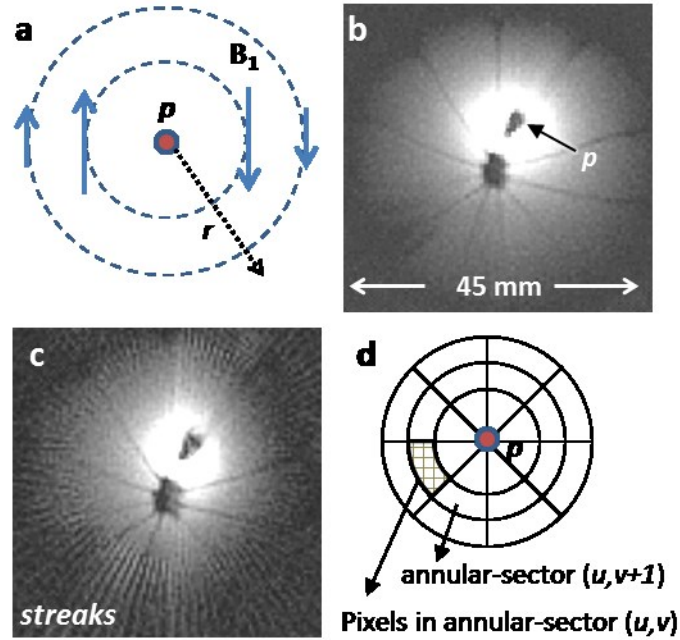
Figure 5.1: (a) The transverse field of a loopless antenna detector p showing decreasing B_1 with r and azimuthal variation in phase. (b) Typical IVMRI image of an orange obtained from probe p, showing intense bright region closest to the probe (c) Four-fold under-sampling of the image in (b) and conventional FT reconstruction showing streaking. The r^{-1} intensity filter has not been applied to (b) or (c). (d) Template used to calculate the Spokal Variation (SV) from adjacent annular sectors as explained in the text. The center of the template is placed on p.
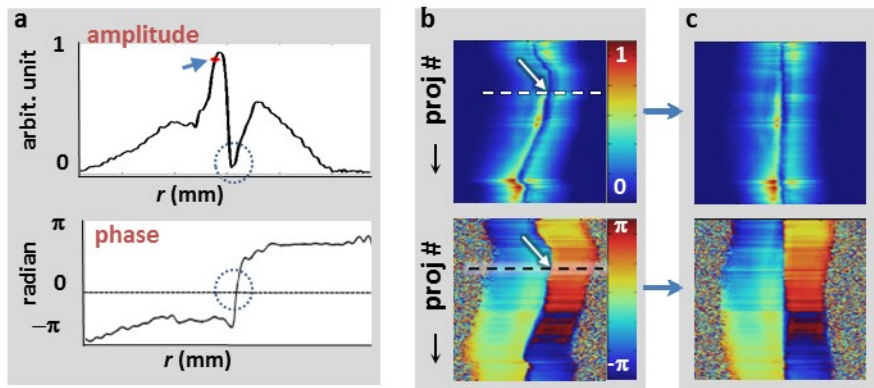
Figure 5.2: (a) Typical experimental projection amplitude (top) shows a 'crater' at the probe location (circled). The center-of-mass of the projection is also shown (arrowhead), but differs from the probe-location. The phase (bottom) of each projection shows a distinct phase reversal at the probe location (circled). (b) A stack of amplitudes (top) and phases (bottom) of all azimuthal angles reveals the probe location p (arrow) in each projection. The sample projection shown in (a) is highlighted using dashes. (c) Motion correction consists of re-aligning every azimuthal projection on p. Reconstruction of the aligned projections obtains an image with the probe as the center of the FOV.

116

Given the known physical dimensions of the probe and the image resolution, the width in pixels of the crater was determined and then used as *a priori* knowledge in a peak-fitting algorithm to find the crater. Note that the COM often differs from the probe location due to skewed intensity profiles (Fig. 5.2a, red diamond). The zero-crossing point of the phase-reversal was found using a simple level detection algorithm. Once the probe's location was detected, each projection was shifted to the center of the image field-of-view (FOV), and all projections were aligned based on the location of the probe (Figs. 5.2b, c). The image was then reconstructed with the probe at its center.

A further refinement to this method was implemented at high-resolution, where the shifting process was observed to be sensitive to intra-pixel jitter, which contributed to residual radial streaking. This involved comparing the detected probe location from one projection to its location in previously acquired projections, and determining the pixel-shift in image-space. If large positional deviations from the preceding, or an average of both preceding and succeeding projections were detected, a corresponding phase-shift in k-space was computed from the inverse FT of the pixel-shift (i.e. if $f(x) \xleftrightarrow{\text{FT}} F(\omega)$ then $f(x-n) \xleftrightarrow{\text{FT}} e^{-j\omega n} F(\omega)$). This 'corrected' phase-shift was then applied to the motion-corrupted projection in k-space, prior to reconstruction. In this way, only those few projections that were corrupted by motion were processed and subject to potential jitter introduced by the detection algorithm. Images reconstructed using this refinement do not necessarily have the probe at the center of the FOV, but because the probe was already located in preceding steps, a simple image translation was applied to shift it to the FOV center post-reconstruction as described previously [108].

Through-plane motion

Through-plane motion, especially from respiration, is a common source of artefacts *in vivo*, in general. Conventionally, it is overcome by navigator echoes placed on the lung-diaphragm interface [120]. In the absence of external navigators, we observe that when stacked, radial k-space projections exhibit a pattern of aberrant jumps (Fig. 5.3a) that correlate with breathing (~2s for a rabbit). These projections serve as 'internal navigators' that signified motion-corrupted projections that could be discarded, as in traditional navigator-gating. Because in PR-MRI successive projections are highly correlated, the aberrant projections were replaced by an average of the preceding and succeeding projections (Fig. 5.3b).

Undersampled Motion-Corrected reconstruction

Since the motion-correction methods outlined above act on each projection, they can also be applied to radially under-sampled data sets, as illustrated in Fig. 5.4. The motion-corrected undersampled data was iteratively reconstructed to produce an effectively faster motion-corrected image. These were compared both visually and using the SSIM (Eq. 5.8) *in vitro*, and visually *in vivo* where motion-free images were unavailable for comparison.
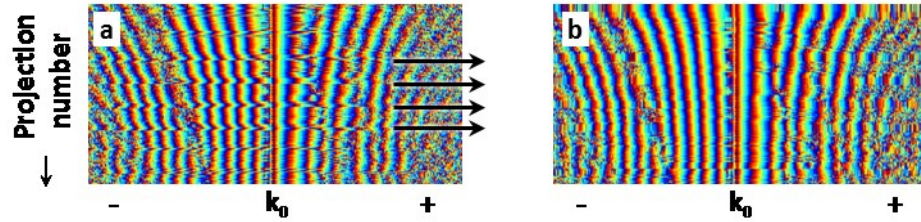
Figure 5.3: (a) Periodic aberrations (arrows) evident in a stack of *in vivo* k-space projection data (phase). The center sample k_0 and the positive and negative portions have been marked. Similar aberrations exist in the magnitude of k-space data but are less visible to the naked eye. (b) Aberrant projections are averaged using preceding and successive rows exhibiting a much smoother appearance. Reconstruction from (b) yields images with reduced streaking.
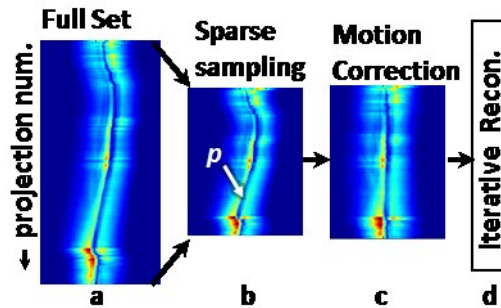
Figure 5.4: Acceleration and motion correction method showing (a) azimuthal projection amplitudes. (b) The projections are undersampled to a reduced dataset. (c) Motion correction consists of detecting probe p in each projection (signal void) and re-aligning every azimuthal projection on p and/or eliminating projections with excessive motion. (d) Images from the sparsely sampled motion-corrected projections are reconstructed using iterative techniques.

5.3   Methods

CS and motion-corrected radial PR-MRI was performed in a Philips 3T *Achieva* whole-body MRI scanner using internal loopless antennae [105] for signal reception and the scanner's body-coil for excitation [107]. The loopless-antenna was either a 2.75 mm outer-diameter (OD) 40cm-long semi-rigid copper coaxial cable with the inner conductor extended 42 mm to form the whip, or a 0.8 mm OD biocompatible super-elastic nitinol coaxial cable with a 42 mm whip suitable for IVMRI applications. Experiments employing Cartesian encoding and a 5-turn transmit/receive loop IV MRI endoscope [108, 109] are also reported. The loop antenna was 2.3 mm diameter at its widest, tuned with a 91 pF micro-capacitor, and connected to a 0.8 mm nitinol cable. The cable was connected to a single channel transmit/receive interface with a switchable PIN diode to decouple the receiver during conventional MRI, as described previously [108].

Data are presented from fruit to test undersampling and motion effects, human vessel specimens *in vitro*, and rabbits *in vivo*. The first study was performed with the semi-rigid loopless antenna in an orange that was manually shaken (± 3 mm in-plane, randomly) during MRI. The *in vitro* blood vessel study was performed with the semi-rigid loopless antenna in human iliac artery specimens immersed in saline. To mimic physiologically relevant motion, the container was placed on the abdomen of a free-breathing volunteer during MRI: in-plane motion due to respiration was ± 4 mm. *In vivo* studies approved by our Institutional Animal Care and Use Committee, were performed on 3 healthy New Zealand white rabbits. The rabbits were sedated with intramuscular acepromazine (1 mg/kg) and ketamine (40 mg/kg), induced with intravenous sodium thiopental, and intubated to maintain an open airway. The loopless antenna was advanced

into the descending aorta to the renal bifurcation via a femoral incision, as confirmed by

contrast-enhanced X-ray C-arm CT prior to transfer to MRI. The loop antenna was

inserted in the descending aorta via a surgical cut-down just inferior to the renal arteries.

Probes were fixed in place with surgical ties.
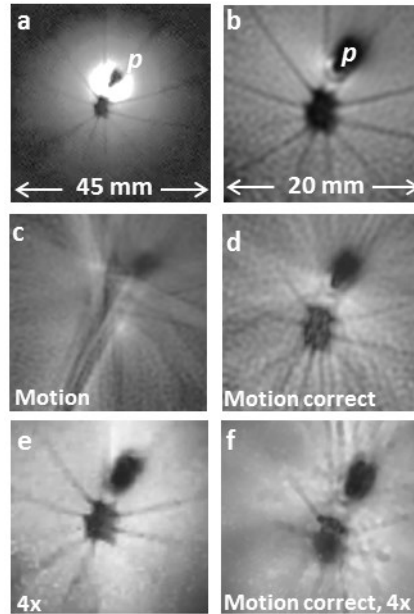
Figure 5.5: (a) MRI of an orange from internal probe p without the r^{-1} intensity filter and (b) its zoomed version with an r^{-1} intensity filter. (c) MRI of the same orange shaken $\pm$ 3 mm showing debilitating motion artifacts. (d) Motion correction substantially restores the image. (e) Fruit morphology revealed in (b) above is retained in a four-fold under-sampled radial CS reconstruction. (f) Four-fold acceleration and motion correction applied to (c) reveals fruit morphology but with some loss of detail in the periphery of the septa.

Scout MRI was performed to locate the probe, prescribe the imaging plane, and to

'volume shim' the vicinity of the internal probe. IVMRI was performed with and without

cardiac gating using radial and/or Cartesian gradient echoes (spoiled or balanced). For

radial acquisitions, azimuthal angles were incremented sequentially with alternating

direction reversals ('even/odd flyback'). All processing was performed off-line on 'raw'

k-space data using scripts written in MATLAB® (Mathworks, Nattick USA). IV images

were weighted by the receiver sensitivity using an r^{-1} image intensity filter centered on

the probe location (Fig. 5.5a, b) [105, 107], except where noted. Since the probe location was already determined by the motion correction algorithm, this step was relatively trivial.

All reconstruction was performed on a standard DELL XPS L502X laptop with an Intel Core i7-2760QM 2.40GHz processor, 8 GB RAM running Windows 7 Professional 64-bit OS. Freely available CS implementations (IRGNTV [117], GRASP [118], SPARSE-MRI [61]) were adapted for reconstruction. Even and odd k-space trajectories were handled separately for motion correction and were recombined for reconstruction, except where noted. Motion correction was done using either projection shifting alone (phantom and *in vitro* data), or combined with the averaging of aberrant k-space projections (*in vivo* data). The scan parameters, SSIM indices, CS software and reconstruction penalties are summarized in Table 5.1. For calculating SV in Eq. (5.4), the annular rings were confined to 1.5-3.5 cm from the probe.

5.4   Results

Motion of the shaken orange obliterated virtually all structure in the PR image from the loopless antenna, except for signal void at the probe location (Fig. 5.5c). Motion-correction on the same data set (Fig. 5.5d) substantially restored structures including the fruit's central void, and septa between the segments. A four-fold under-sampled radial CS reconstruction of the stationary fruit (Fig. 5.5e) retains essentially all of the structures of the original but without the streaking artefacts evident in a conventional undersampled NUFFT reconstruction (Fig. 5.1c). Application of the combined algorithm comprising motion-correction applied to the 4-fold under-sampled dataset (Fig. 5.5f) reveals essentially the same morphology except for some lack of

resolution of the septa near the periphery, as the fully-sampled motion corrected image

(Fig. 5.5d), but a dramatic improvement over the original corrupted image (Fig. 5.5c).
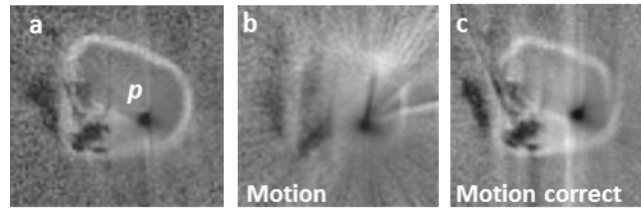
Figure 5.6: (a) IVMRI of a human iliac specimen *in vitro* (b) IVMRI from the sample mounted on the abdomen of a volunteer to simulate physiological motion showing major motion artifacts. (c) Motion correction removes streaking, revealing the specimen's underlying structure.
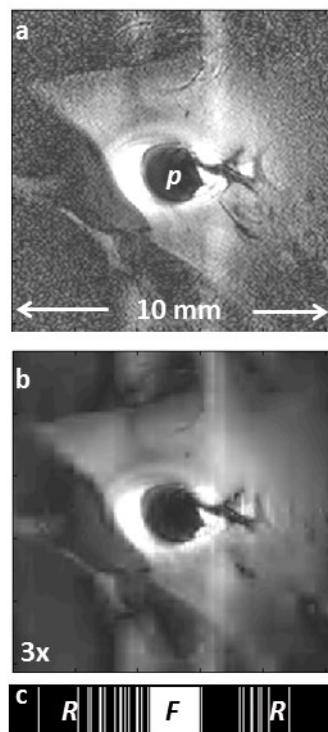
Figure 5.7: (a) Regular Cartesian MRI endoscopy of a rabbit aorta in vivo adapted from prior work [ref. 4] (b) Three-fold under-sampling yields a very similar image after CS reconstruction. (c) The under-sampling scheme depicts the fully (*F*) and randomly sampled (*R*) portions of corresponding k-space.

The effect of physiological motion of an iliac vessel specimen in a saline tank placed on the abdomen of a free-breathing volunteer is debilitating with most structure lost in this example (Fig. 5.6b vs a). However, motion correction restores nearly all of the lost structure (Fig. 5.6c).
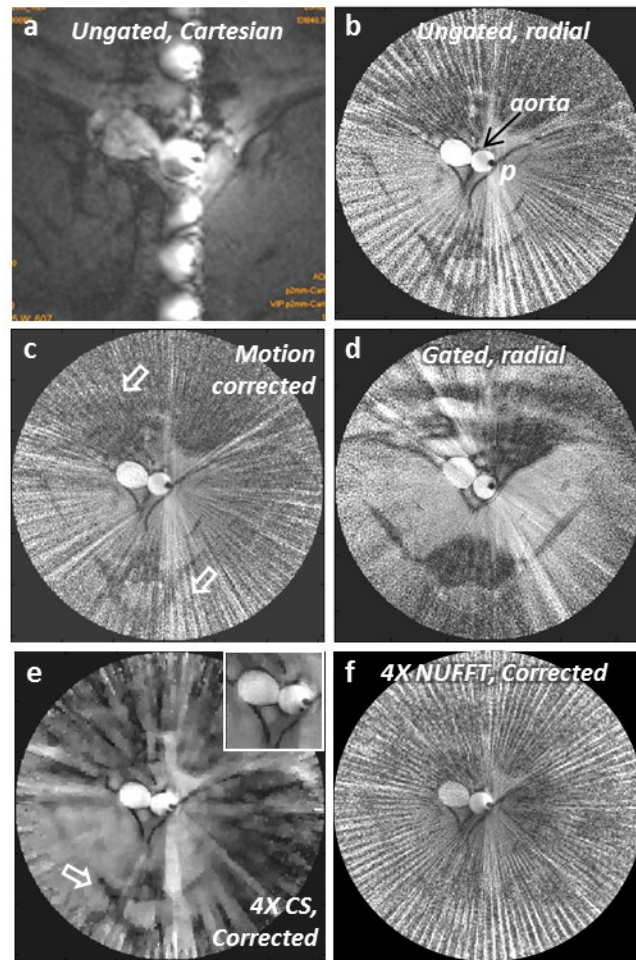
Figure 5.8: An *in vivo* rabbit aorta image acquired with an IV probe *p* and without any cardiac gating showing (a) intense ghost artefacts in Cartesian encoding and (b) streaking in radial encoding (with odd spokes) arising from physiological motion. (c) Motion-correction reduces the streaking especially at arrow locations. (d) Cardiac-gated *in vivo* IV MR image at approximately the same location reveals very similar tissue structure as (c). (e) Four-folds speedup of (c) retains overall morphology including aorta (inset). The tissue structure (arrow) is retained whereas in (f) the conventional FT reconstruction of the same dataset as (e) is swamped by streaking.

The effect of sparse sampling applied to *in vivo* Cartesian data to provide an effective three-fold speedup of loop-coil IV MRI endoscopy data acquired with conventional cardiac-gated is shown in Fig. 5.7. The CS reconstruction was performed with central k-space fully-sampled, and the edges randomly under-sampled (Fig. 5.7c, bar-graph).

Without cardiac or any other gating, *in vivo* images acquired with the loopless antenna at 200 μm resolution in the rabbit aorta can exhibit intense fold-over and ray artefacts for both Cartesian- and radially-encoded acquisitions respectively (Figs. 5.8a, b). Motion correction of the radial data set ameliorates the ray artefacts (Fig. 5.8c), yielding an image that is comparable to a cardiac-gated image acquired separately (Fig. 5.8d) but in 70% of the scan-time for full k-space coverage. CS reconstruction applied to 1/4th of the motion-corrected data, although somewhat patchy, retains overall structure (Fig. 5.8e) whereas conventional NUFFT reconstruction of the same undersampled data is overwhelmed by streaking artefacts (Fig. 5.8f). Zooming in on the aorta, we see that radial compressed sensing with only 1/4th of the original data (Fig. 5.9b) provides excellent rendition vs. the original (Fig. 5.9a), with streaking artefact reduced by the motion correction (Fig. 5.9c). CS reconstruction applied to 1/8th of the motion-corrected data provides an effectively eight-fold faster, motion-suppressed *in vivo* image (Fig. 5.9d). Ungated acquisition affords imaging speeds of up to 2 fps (Fig. 5.9e) and CS reconstruction using 1/2 the data (effectively 4 fps) faithfully renders the aorta but with some loss of detail in low-SNR regions (Fig. 5.9f).

CS reconstruction times depend significantly on the image resolution, FOV and iterative reconstruction parameters, ranging from 2 to 25 minutes, whereas the

corresponding FFT and NUFFT based reconstruction was in the order of milliseconds

(Table 5.1). The SSIM of all the corrected and/or undersampled CS reconstructed images

was $\geq$ 0.4 (Table 5.1).
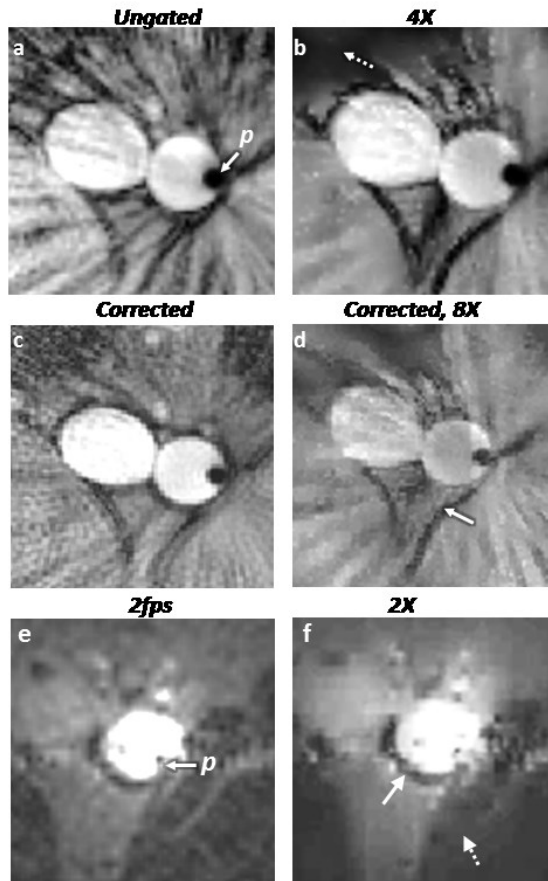
Figure 5.9: (a) Ungated radial IVMRI of the aorta shows streaking near the aorta and surrounding tissue (b) Tissue morphology is retained in a 4-fold under-sampled CS reconstruction albeit with some loss of signal in regions of low SNR (dashed arrow) (c) Motion-correction applied to (a) reduces streaking everywhere. (d) Eight-fold under-sampling of (c) retaining only odd spokes, retains morphology of the aorta and surrounding tissue (solid arrow), although the intensity of the blood signal is reduced. (e) IVMR images from a rabbit aorta in vivo, acquired at 2fps without any cardiac gating. (f) Two-fold undersampled CS reconstruction retains overall morphology (aorta, solid arrow), but some detail is lost in low SNR regions (dashed arrow). The r^{-1} intensity filter has not been applied to (e) or (f).

126

Table 5.1: Acquisition parameters for all images (acquisition time, T_{acq}); similarity indices (SSIM) computed with respect to the image in parentheses (e.g. Fig. 5.9f in last row has an SSIM = 0.82 compared to Fig. 5.9e); CS software used (IRGNTV, SPARSE-MRI or GRASP; reconstruction parameters, λ where applicable; and reconstruction times (T_{rec}) using conventional (NUFFT) or CS reconstruction.

Fig.	Image details
5.1b	2D radial GRE; 200 spokes ; 250 µm in-plane resolution; TR/TE = 15/6 ms. T_{acq} = 3s
5.1c	Same as 5.1b, but with 50 spokes. NUFFT T_{rec} ~ 100 ms
5.5a	Same as 5.1b
5.5b	5.5a zoomed
5.5c	SSIM = 0.28 (5.5b). Acquisition parameters same as 5.5a above.
5.5d	SSIM = 0.56 (5.5b). NUFFT T_{rec} ~ 100 ms
5.5e	SSIM = 0.52 (5.5b). CS: IRGNTV. λ_{WT} = 0; λ_{TV} = 2, 0.2, 0.02, 0.05 … (reduced every iteration); λ_{SV} = 0. T_{rec} = 4 mins.
5f	SSIM = 0.5 (5.5b). CS: IRGNTV, λ same as Fig. 5.5e above. T_{rec} = 17mins.
5.6a	2D radial GRE; 0.3x0.3x2mm^3 voxel; TR/TE=150/6 ms, 1000 spokes, 20mm FOV shown.
5.6b	SSIM=0.1 (5.6a). Acquisition parameters same as 5.6a above.
5.6c	SSIM = 0.4 (5.6a). NUFFT T_{rec} ~ 100 ms
5.7a	Cardiac gated 3D Cartesian GRE; TR/TE=250/12 ms; in-plane resolution 80 µm; T_{acq} = 3.1 min/5 contiguous slices
5.7b	SSIM = 0.53 (5.7a). CS: SPARSE MRI. λ_{WT} = 0.09; λ_{TV} = 0.01; λ_{SV} = 0. T_{rec} = 2 mins.
5.8a	Ungated, 0.2x0.2x5 mm^3 voxel, 2D Cartesian Balanced GRE, FA 90°, TR/TE 700/16 ms, 25 mm FOV. T_{acq} = 1.5 mins
5.8b	Ungated, 0.2x0.2x5 mm^3 voxel, 200 spokes balanced radial 2D GRE, FA 90°, TR/TE = 700/6 ms, 50 mm FOV shown. Effective T_{acq} = 2.5 mins
5.8c	NUFFT T_{rec} ~ 100 ms
5.8d	Cardiac gated, 0.2x0.2x5 mm^3 voxel, 400 spokes spoiled 2D radial GRE, FA 40°, TE 6 ms, 40 mm FOV. T_{acq} = 7 mins
5.8e	CS: GRASP. λ_{WT} = 0; λ_{TV} = 0.1; λ_{SV} = 0.0005. T_{rec} = 7 mins.
5.8f	NUFFT T_{rec} ~ 100 ms
5.9a	Ungated, 0.2x0.2x5 mm^3 voxel, 400 spokes (even+odd) balanced radial 2D GRE , FA 90°, TR/TE = 700/6 ms, 15 mm FOV shown. T_{acq} = 5 mins

5.9b	CS: IRGNTV, λ same as Fig. 5.5e above. T_{rec} = 10 mins.
5.9c	NUFFT T_{rec} ~ 500 ms
5.9d	50 spokes (only odd spokes) used. CS: IRGNTV. λ same as Fig. 5.5e above. T_{rec} = 25 mins.
5.9e	Ungated, 0.3x0.3x5 mm^3 voxel, Cartesian Balanced 2D GRE, FA 90°, TR/TE 9/4 ms, 15 mm FOV. T_{acq} = 0.5 secs.
5.9f	SSIM=0.82 (5.9e). CS: SPARSE MRI. λ_{WT} = 0.09; λ_{TV} = 0.01; λ_{SV} = 0. T_{rec} = 3 mins.

5.5   Discussion

This chapter presented strategies for accelerating image acquisition and suppressing motion artefacts in IVMRI, and demonstrated results in fruit for testing motion effects, vessel specimens *in vitro* (mounted on a volunteer to simulate physiological motion), and in rabbit aorta *in vivo*. CS reconstruction has been widely studied for MRI applications in the past, and our current implementation is based on freely available open-source CS implementations. We extended previously used penalties with a new IVMRI specific constraint (SV). Additional smoothing functions may be incorporated [110] to avoid some of the CS associated 'patchiness'. CS reconstruction takes significantly longer than conventional FT (minutes vs. milliseconds) and further software and hardware optimization (employing graphics processing units [121]) is required for practical use. As in any undersampling scheme, there is an SNR loss, and undersampling factors and reconstruction parameters have to be adjusted depending on the application. As far as we are aware, this is the first application of CS reconstruction to IVMRI applications [122].

Motion correction strategies based on PR have been used in the past in conventional MRI settings [113-115]. These methods use correction techniques based on

the zeroth-, first- and higher order moments of each projection, which can yield erroneous results when applied to projections having skewed intensity profiles (Fig. 5.2a). Our new method uses explicit probe localization based on the amplitude and phase profiles of IVMRI detectors. Moreover, in high-resolution IVMRI applications, there are dynamic changes in projections such that corrections based on the central k-space sample or constant spin-density assumptions [115] did not work well in our studies. All of our acquisition employed gradient echoes, which was the standard radial sequence provided by the manufacturer. Alternatively, radial spin-echo (SE) sequences could provide improved immunity to inhomogeneity (T_2*) effects and suppress signals from flowing blood for vessel wall imaging.

Note that motion correction using projection shifting and averaging of aberrant k-space projections do not account for the following sources of motion. First, probe motion with respect to the body could cause position-dependent amplitude or phase variations. In practice however, we observed that the tensile properties of the super-elastic *in vivo* nitinol probe pushed it against the vessel wall and provided a degree of immunity from blood-flow induced motion. In any case, this motion should have affected the cardiac-gated images (Fig. 5.8d) which were relatively free of such artefacts. Second, object rotation can cause a mismatch between the nominal acquisition angle and the actual projection. To correct this in high-resolution images, one could potentially model the motion based on fast low-resolution acquisitions. Third, pulsatile blood flow causes the vessel to expand and contract during image acquisition, affecting the surrounding tissue to a lesser extent. We explored a correction method wherein each projection was dilated

(compressed) to account for the radial tissue expansion (contraction) but found it difficult to reliably model the non-linear changes.

Nevertheless, motion correction separately applied to each projection as presented here, affords a sliding window reconstruction wherein the latest projection replaces the oldest, to provide an image-stream that updates at the TR rate. When the azimuthal angle increment in PR is small, the cross-correlation between successive projections can be used for probe detection, but can introduce a delay between acquisition and motion-correction. This will be larger when non-sequential angle increments are employed. Using the amplitude and phase profiles alone, automated detection was achieved for phantom and *in vitro* data, but some manual processing was required for *in vivo* data when spurious phase wraps occurred. Transmit/receive loop probes as in MR endoscopy [109] do not exhibit phase variations so that probe detection would have to rely solely on the near-field radial sensitivity profile.

We conclude that 3T IVMRI detectors are well-suited to compressed sensing and motion correction strategies based on their intrinsically radial and sparsely-localized sensitivity profiles, and high SNR. The benefits are faster free-breathing IVMRI with reduced motion sensitivity, while retaining the high-resolution image information.

Chapter 6 : Conclusion and Future Work

6.1 Conclusion

Be it spectroscopy or endoscopy, the key idea demonstrated here is to use prior knowledge for faster and better results. This dissertation has shown some applications of this idea, which has resulted in several new algorithms, journal papers, and patent applications. Conclusions and future directions are summarized in this chapter.

First, the proposed SLAM method was validated against the gold standard, CSI, in 1D ^{31}P cardiac and leg studies. SLAM yielded the same quantitative results four times faster in 24 cardiac patients and healthy subjects, compared to CSI. SLAM was further extended with fractional phase-encoding gradients (fSLAM) that optimized SNR and/or minimized both inter- and intra-compartmental contamination. Both SLAM and fSLAM were integrated into our 3T Philips Achieva scanner and demonstrated the potential to replace the current protocol used to perform ^{31}P MR spectroscopy in our lab. The work has resulted in a journal publication [50] and a patent application.

Second, the SLAM method was further extended to ^{1}H brain spectroscopy applications, where 2D and 3D imaging are typically used in conjunction with multi-element receive coils (parallel imaging). One major modification was that SLAM was adapted into SLAM* to improve its compatibility with the parallel imaging technique, SENSE. The new SENSE SLAM* method was shown to be highly robust to intra-compartment and inter-compartment non-uniformity. In addition, eddy current correction was added to the SLAM method. These new advances were all reduced to practice on patient studies and showed highly promising results with a dramatic acceleration advantage. This work also resulted in a journal publication [93].

131

In my opinion, these two projects resulted in a comprehensive tool for MRS and MRI processing, which made a third project possible. The proposed PANGS method was designed to reduce artifacts in the standard CSI technique, but not to accelerate the scan any more than in the two earlier projects. The CSI method places each reconstruction coordinate in the geometric center of the corresponding voxel, which often results in strong lipid artifacts in ^{1}H brain CSI. PANGS minimized this artifact by matching the reconstruction coordinates with the centers of mass without compromising spatial resolution. The PANGS method resulted in a paper being considered for publication.

Fourth, the last project on MR endoscopy was in collaboration with Dr. Shashank S. Hedge, who made major contributions to the work, including both motion correction and radial compressed sensing. My contribution involved the Cartesian compressed sensing part. The visually sparse nature of intravascular endoscopic images made them a good candidate for applying compressed sensing. The null signal from the endoscopic probe serves as a distinct marker for identifying its location for subsequent motion correction. We have summarized this work in a submitted journal paper manuscript and filed a patent application on it.

6.2   Future work

We are testing the feasibility of combining the SLAM method as applied to 1D ^{31}P MRS described in Chapter 2, with the "TRiST" technique [123] for measuring cardiac creatine kinase (CK) reaction rates using saturation transfer MRS methods. Adenosine triphosphate (ATP) is the most important energy source in the human body, which can be generated through the following CK reaction from Phosphocreatine (PCr) and ADP.

$$\text{PCr} + \text{adenosine diphosphate (ADP)} \xrightleftharpoons[k_f]{k_r} \text{ATP} + \text{creatine (Cr)}$$
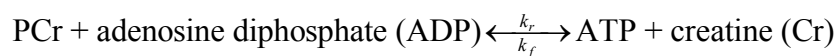
where k_f and k_r are the pseudo first-order forward and reverse reaction rates, respectively.

The forward CK ATP flux in heart failure patients was found to be reduced significantly compared to that in normal subjects [41]. One major limitation of the "TRiST" and other saturation transfer methods is that they require ~40 min or more, which may be intolerable for some patients and also may make the methods susceptible to motion artifacts. SLAM could reduce the scan time significantly if the CSI method used for TRiST were replaced by SLAM. But, one potential concern is that the SLAM TRiST results may differ from CSI TRiST results due to various errors that may accumulate as the saturation transfer methods involve multiple acquisitions.
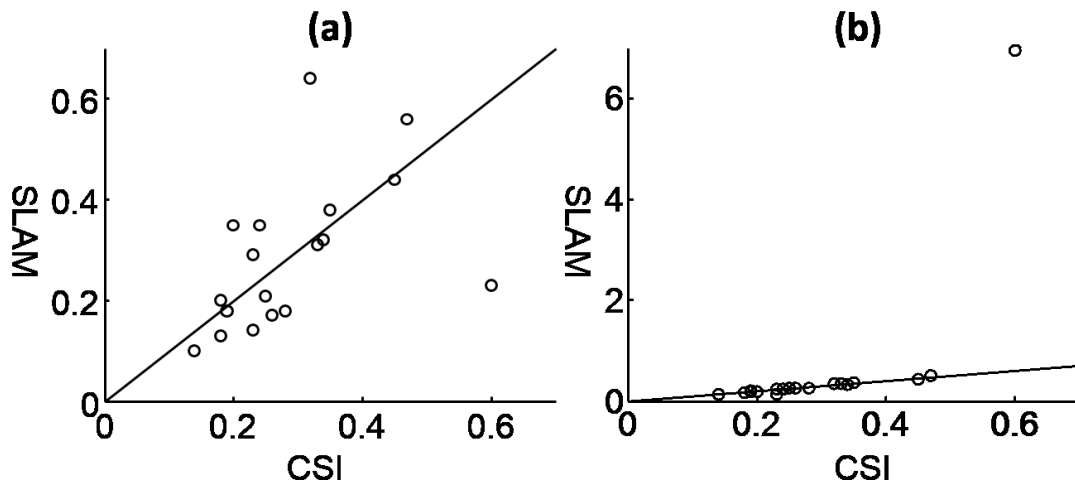
Figure 6.1: Pseudo-first-order forward rates, k_f, measured from 16-step CSI, compared with results from 4-step SLAM (a) and 8-step SLAM (b). R^2 values are 0.19 (a) and 0.24 (b). The solid line is the identity line.
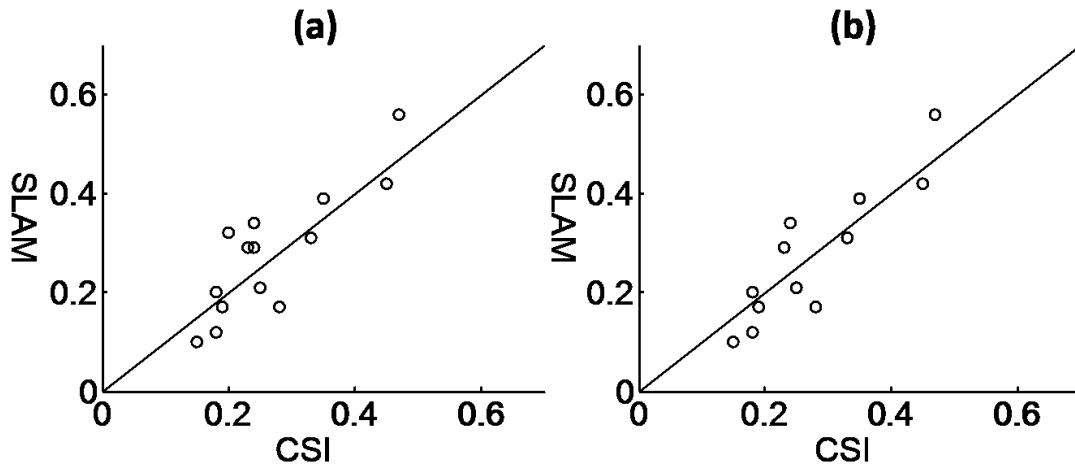
Figure 6.2: Forward rates, k_f, measured from 16-step CSI compared with those from 4-step SLAM after applying the proposed outlier rejection algorithm by setting the threshold to 40% (a) and 20% (b). Three and five outliers are discarded in (a) and (b), respectively. R^2 values are 0.71 (a) and 0.79 (b). The solid line is the identity line.
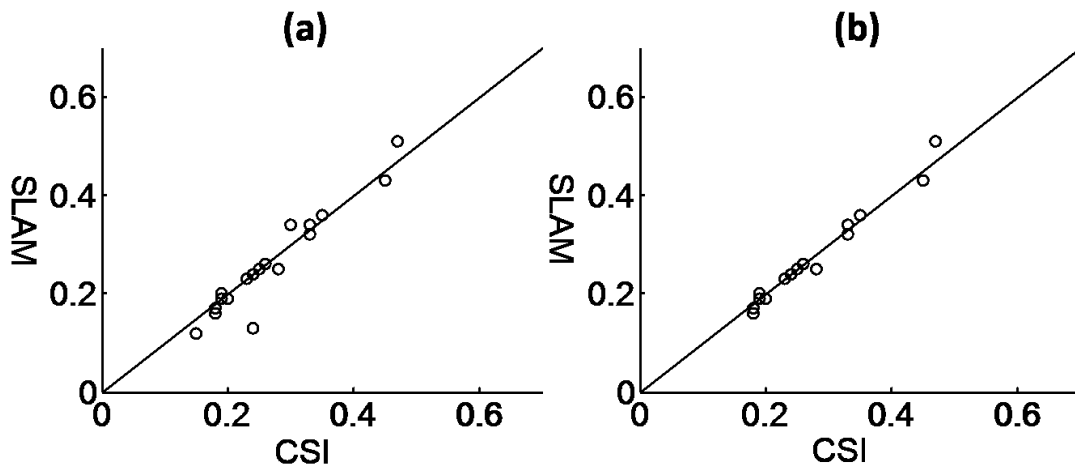
Figure 6.3: Forward rates measured from 16-step CSI vs. results from 8-step SLAM after applying the outlier rejection algorithm with the threshold set to 40% (a) and 20% (b). One and four outliers are discarded in (a) and (b), respectively. R^2 values are 0.90 (a) and 0.97 (b). The solid line is the identity line.

In an initial attempt, SLAM TRiST demonstrated significant scatter in k_f values as compared to CSI as a standard, as shown in Fig 6.1. However, the scatter was primarily due to a few outliers, which conceivably arose from motion artifacts or variations in position during the course of the studies. An outlier rejection method, based solely on the reduced SLAM dataset, was therefore proposed as follows. Step 1: Move the whole field of view by five different offsets 0 cm, 0.1 cm, -0.1 cm, 0.2 cm, and -0.2 cm, which are relatively small compared to the 1 cm nominal 1D CSI slice resolution. This offset can be achieved by applying a linear phase in Fourier space. Step 2: Apply the SLAM TRiST method for each offset and record the five calculated reaction rates. SLAM TRiST can be implemented with a user-defined number of phase encoding steps, such as 4 and 8. Step 3: Compute the mean and standard deviation of the five rates. Step 4: Identify the subject as an outlier if the ratio between standard deviation and mean is greater than a user-defined threshold, such as 40% or 20%.

In Fig. 6.2, the 4-step SLAM TRiST results for k_f are compared to 16-step CSI TRiST results after discarding outliers according to the proposed outlier rejection algorithm. For two different thresholds, 40% (Fig. 6.2a) and 20% (Fig. 6.2b), the consistency of k_f values measured by SLAM TRiST with respect to CSI TRiST was significantly improved compared to the results depicted in Fig. 6.1(a). Similarly, after applying the outlier rejection method, the consistency between 8-step SLAM TRiST and 16-step CSI TRiST k_f values, shown in Fig. 6.3, was significantly improved vs. the results shown in Fig. 6.1(b). Notably, the 8-step SLAM TRiST k_f measurements (Fig. 6.3) show less deviation from the identity line compared to 4-step SLAM TRiST (Fig. 6.2).

The SLAM method is currently being translated to quantitative MRI applications. Specifically, we are adapting SENSE SLAM for Chemical Exchange Saturation Transfer (CEST) imaging. The PANGS method can be further extended to incorporate SENSE for more potential uses, as well.

Appendix A: Localization analysis for (SENSE) SLAM

The reconstructed (SENSE) SLAM spectrum for the i[th] compartment of a set of C compartments can be determined from the dSRF as:

$$\boldsymbol{\rho}^{r}_{C*N}(i,:) = \sum_{j=1}^{M} \mathbf{dSRF}_{C*M}(i,j) \times \boldsymbol{\rho}_{M*N}(j,:),$$ (A.1)

where $\boldsymbol{\rho}^{r}_{C*N}(i,:)$ denotes that only the i[th] ($1 \le i \le C$) row corresponding to the i[th] compartment is selected from the (SENSE) SLAM spectral matrix, $\boldsymbol{\rho}^{r}_{C*N}$; $\mathbf{dSRF}_{C*M}(i,j)$ is the element at the i[th] row and j[th] column of $\mathbf{dSRF}_{C*M}$; j ($1 \le j \le M)$ is an index for each of all M spatial voxels; and $\boldsymbol{\rho}_{M*N}(j,:)$ is the j[th] spectrum in the (SENSE) CSI spectral matrix, $\boldsymbol{\rho}_{M*N}$.

If the spatial voxels of $\boldsymbol{\rho}_{M*N}$ in a specific region are grouped to form a user-defined region k, their contribution to the i[th] SLAM spectrum from the compartment of interest, $\boldsymbol{\rho}^{r}_{C*N}(i,:)$, is:

$$L_{k \to i} = \sum_{j \in region \ k} \mathbf{dSRF}_{C*M}(i,j) \times \left[\overline{\boldsymbol{\rho}^{k}} + \Delta\boldsymbol{\rho}^{k}_{M*N}(j,:) \right],$$ (A.2)

where $\overline{\boldsymbol{\rho}^{k}}$ is the average (SENSE) CSI spectrum for region k; and

$\Delta\boldsymbol{\rho}^{k}_{M*N}(j,:) = \boldsymbol{\rho}_{M*N}(j,:) - \overline{\boldsymbol{\rho}^{k}}$ is the deviation of each individual CSI voxel spectrum in the k[th] region from its regional mean.

A further examination of Eq. (A.2) leads to:

$$
\begin{aligned}
L_{k \to i} &= \overline{\boldsymbol{\rho}^{k}} \times \sum_{j \in region \ k} \mathbf{dSRF}_{C*M}(i,j) + \sum_{j \in region \ k} \mathbf{dSRF}_{C*M}(i,j) \times \Delta\boldsymbol{\rho}^{k}_{M*N}(j,:) \\
&\le \overline{\boldsymbol{\rho}^{k}} \times \sum_{j \in region \ k} \mathbf{dSRF}_{C*M}(i,j) + \sum_{j \in region \ k} \left| \mathbf{dSRF}_{C*M}(i,j) \right| \times \left| \Delta\boldsymbol{\rho}^{k}_{M*N}(j,:) \right|
\end{aligned}
$$ (A.3)

Eq. [A.3] relates the SLAM spectrum reconstructed for the i^{th} compartment, to the contributions from the mean signal of each region and the inhomogeneities, modulated by the dSRF over the specified regions. The inequality in |dSRF| sets an upper bound on contributions due to regional inhomogeneity.

Appendix B: **The dSRF for (SENSE) SLAM incorporating inhomogeneity effects**

The dSRF for (SENSE) SLAM reconstruction (Eq. 3.11), after incorporating spatial alternation terms, is:

$$\mathbf{dSRF}_{C*M} = (\mathbf{E}_{M'*M} \times \mathbf{A}^{-1}_{M*M} \times \mathbf{b}^r_{M*C})^+ \times \mathbf{E}_{M'*M} \times \mathbf{A}^{-1}_{M*M} . \tag{B.1}$$

Similarly, for (SENSE) SLAM* (Eq. 3.12):

$$\mathbf{dSRF}_{C*M} = (\mathbf{b}^r_{M*C})^+ \times \mathbf{A}_{M*M} \times (\mathbf{E}_{M'*M})^+ \times \mathbf{E}_{M'*M} \times \mathbf{A}^{-1}_{M*M} . \tag{B.2}$$

For time-varying spatial corrections, the dSRF should be computed separately for each time point.

Appendix C: **The SNR of (SENSE) SLAM**

The compartmental average (SENSE) CSI spectrum for the i^{th} compartment, ρ^i_{csi}, obtained from the "pre-whitened" Eq. (3.2) is:

$$\rho^i_{csi} = 1/\mathrm{L} \times \sum_{j \in compartment\ i} \mathbf{E}^+_{M*M_{csi}}(j,:) \times \mathbf{s}_{M_{csi}*N} , \tag{C.1}$$

where L is the number of voxels in the i^{th} compartment, j is a voxel index in the i^{th} compartment, $\mathbf{E}^+_{M*M_{csi}}$ is the (SENSE) CSI unfolding matrix, and M_{csi} is the number of acquired (SENSE) CSI signals from all coil elements. Assuming the SD of the noise in

the pre-whitened signal matrix, $\mathbf{s}_{M_{csi}*N}$ is σ, the noise in the i[th] compartmental average (SENSE) CSI spectrum is,

$$SD_{csi}^{i} = \sqrt{\sum_{k=1}^{M_{csi}} \left\| \left[1/L \times \sum_{j \in compartment\ i} \mathbf{E}_{M*M_{csi}}^{+}(j,k) \right] \right\|^2 \times \sigma^2}\ , \qquad (C.2)$$

where $k\ (1 \le k \le M_{csi})$ is an index for each acquired signal.

Similarly, the noise for the i[th] compartmental (SENSE) SLAM spectrum is,

$$SD_{SLAM}^{i} = \sqrt{\sum_{k=1}^{M'} \left| \left(E_{M'*M} \times b_{M*C}^{r} \right)^{+} (i,k) \right|^2 \times \sigma^2} \qquad (C.3)$$

and

$$SD_{SLAM*}^{i} = \sqrt{\sum_{k=1}^{M'} \left| \left[\left(b_{M*C}^{r} \right)^{+} \times \left(E_{M'*M} \right)^{+} \right](i,k) \right|^2 \times \sigma^2} \qquad (C.4)$$

for (SENSE) SLAM*. Here M' is the total number of k-space signals used in (SENSE) SLAM/SLAM*, and $(i,\ k)$ denotes the corresponding matrix element.

The relative SNRs for a given total scan-time obtain by comparing Eqs. [C.2 – C.4], assuming the compartment signal is constant.

References

[1] D.G. Nishimura, Principles of magnetic resonance imaging, Stanford University, 1996.

[2] P.C. Lauterbur, Image formation by induced local interactions: examples employing nuclear magnetic resonance, Nature, 242 (1973) 190-191.

[3] W. Edelstein, J. Hutchison, G. Johnson, T. Redpath, Spin warp NMR imaging and applications to human whole-body imaging, Physics in medicine and biology, 25 (1980) 751.

[4] T. Brown, B. Kincaid, K. Ugurbil, NMR chemical shift imaging in three dimensions, Proceedings of the National Academy of Sciences, 79 (1982) 3523-3526.

[5] A. Macovski, Noise in MRI, Magnetic Resonance in Medicine, 36 (1996) 494-497.

[6] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, P. Boesiger, SENSE: sensitivity encoding for fast MRI, Magnetic Resonance in Medicine, 42 (1999) 952-962.

[7] M. Firbank, A. Coulthard, R. Harrison, E. Williams, A comparison of two methods for measuring the signal to noise ratio on MR images, Physics in medicine and biology, 44 (1999) N261.

[8] P.A. Bottomley, C.J. Hardy, Strategies and Protocols for Clinical ^{31}P Research in the Heart and Brain, Phil. Trans. R. Soc. Lond. A, 333 (1990) 531-544.

[9] T.R. Brown, B.M. Kincaid, K. Uguibil, NMR chemical shift imaging in three dimensions, Proc. Natl Acad Sci USA, 79 (1982) 3523-3526.

[10] X. Hu, D.N. Levin, P.C. Lauterbur, T. Spraggins, SLIM: spectral localization by imaging, Magnetic Resonance in Medicine, 8 (1988) 314-322.

[11] Z.P. Liang, P.C. Lauterbur, A generalized series approach to MR spectrocopic imaging, IEEE Transactions Medical Imaging, 10 (1991) 132-137.

[12] M. von Kienlin, R. Mejia, Spectral localization with optimal pointspread function, Journal of Magnetic Resonance, 94 (1991) 268-287.

[13] Z.P. Liang, F. Boada, R. Constable, E. Haacke, P. Lauterbur, M. Smith, Constrained reconstruction methods in MR imaging, Rev Magn Reson Med, 4 (1992) 67-185.

[14] A. Bashir, D.A. Yablonskiy, Natural linewidth chemical shift imaging (NL-CSI), Magnetic Resonance in Medicine, 56 (2006) 7-18.

[15] I. Khalidov, D. Van De Ville, M. Jacob, F. Lazeyras, M. Unser, BSLIM: Spectral Localization by Imaging With Explicit B_0 Field Inhomogeneity Compensation, Medical Imaging, IEEE Transactions on, 26 (2007) 990-1000.

[16] M. Jacob, X. Zhu, A. Ebel, N. Schuff, Z.P. Liang, Improved model-based magnetic resonance spectroscopic imaging, Medical Imaging, IEEE Transactions on, 26 (2007) 1305-1318.

[17] L. An, S. Warach, J. Shen, Spectral localization by imaging using multielement receiver coils, Magnetic Resonance in Medicine, (2011).

[18] Z. Dong, J.H. Hwang, Lipid signal extraction by SLIM: Application to 1H MR spectroscopic imaging of human calf muscles, Magnetic Resonance in Medicine, 55 (2006) 1447-1453.

[19] J.A. Kmiecik, C.D. Gregory, Z.P. Liang, P.C. Lauterbur, M.J. Dawson, Lactate quantitation in a gerbil brain stroke model by GSLIM of multiple-quantum-filtered signals, Journal of Magnetic Resonance Imaging, 9 (1999) 539-543.

[20] R. Loffler, R. Sauter, H. Kolem, A. Haase, M. von Kienlin, Localized spectroscopy from anatomically matched compartments: improved sensitivity and localization for cardiac [31]P MRS in humans, Journal of Magnetic Resonance, 134 (1998) 287-299.

[21] M. Meininger, W. Landschütz, M. Beer, T. Seyfarth, M. Horn, T. Pabst, A. Haase, D. Hahn, S. Neubauer, M. von Kienlin, Concentrations of human cardiac phosphorus metabolites determined by SLOOP [31]P NMR spectroscopy, Magnetic Resonance in Medicine, 41 (1999) 657-663.

[22] M. von Kienlin, M. Beer, A. Greiser, D. Hahn, K. Harre, H. Kostler, W. Landschutz, T. Pabst, J. Sandstede, S. Neubauer, Advances in human cardiac [31]P-MR spectroscopy: SLOOP and clinical applications, Journal of Magnetic Resonance Imaging, 13 (2001) 521-527.

[23] M. Beer, T. Seyfarth, J. Sandstede, W. Landschutz, C. Lipke, H. Kostler, M. von Kienlin, K. Harre, D. Hahn, S. Neubauer, Absolute concentrations of high-energy phosphate metabolites in normal, hypertrophied, and failing human myocardium measured noninvasively with [31]P-SLOOP magnetic resonance spectroscopy, Journal of the American College of Cardiology, 40 (2002) 1267-1274.

[24] R.G. Weiss, P.A. Bottomley, C.J. Hardy, G. Gerstenblith, Regional Myocardial Metabolism of High-Energy Phosphates during Isometric Exercise in Patients with Coronary Artery Disease, N Engl J Med, 323 (1990) 1593-1600.

[25] M.A. Conway, P.A. Bottomley, R. Ouwerkerk, G.K. Radda, B. Rajagopalan, Mitral regurgitation: Impaired systolic function, eccentric hypertrophy, and increased severity are linked to lower phosphocreatine/ATP ratios in humans, Circulation, 97 (1998) 1716-1723.

[26] R.G. Weiss, G. Gerstenblith, P.A. Bottomley, ATP flux through creatine kinase in the normal, stressed, and failing human heart, Proc Natl Acad Sci USA, 102 (2005) 808-813.

[27] C.S. Smith, P.A. Bottomley, S.P. Schulman, G. Gerstenblith, R.G. Weiss, Altered Creatine Kinase Adenosine Triphosphate Kinetics in Failing Hypertrophied Human Myocardium, Circulation, 114 (2006) 1151-1158.

[28] P.A. Bottomley, NMR Spectroscopy of the Human Heart, in: R.K. Harris, R.E. Wasylishen (Eds.) Encyclopedia of Magnetic Resonance, John Wiley: Chichester, 2009.

[29] A.-M. El-Sharkawy, M. Schär, R. Ouwerkerk, R.G. Weiss, P.A. Bottomley, Quantitative cardiac [31]P spectroscopy at 3 Tesla using adiabatic pulses, Magnetic Resonance in Medicine, 61 (2009) 785-795.

[30] R.A. Horn, C.R. Johnson, Matrix Analysis, Cambridge University Press, 1990.

[31] H.R. Brooker, T.H. Mareci, J. Mao, Selective Fourier transform localization, Magnetic Resonance in Medicine, 5 (1987) 417-433.

[32] R.E. Gabr, R. Ouwerkerk, P.A. Bottomely, Quantifying in vivo MR spectra with circles, Journal of Magnetic Resonance, 179 (2006) 152-163.

[33] P.A. Bottomley, Spatial localization in NMR spectroscopy in vivo, Annal NY Acad Sci, 508 (1987) 333-348.

[34] J. Frahm, H. Bruhn, M.L. Gyngell, K.D. Merboldt, W. Hanicke, R. Sauter, Localized high-resolution proton NMR spectroscopy using stimulated echoes: Initial applications to human brain in vivo, Magnetic Resonance in Medicine, 9 (1989) 79-93.

[35] R. Ordidge, A. Connelly, J. Lohman, Image-selected *in Vivo* spectroscopy (ISIS). A new technique for spatially selective NMR spectroscopy, Journal of Magnetic Resonance (1969), 66 (1986) 283-294.

[36] L.P. Panych, L. Zhao, R.V. Mulkern, PSF-choice: A novel MRI method for shaping point-spread functions in phase-encoding dimensions, Magnetic Resonance in Medicine, 54 (2005) 159-168.

[37] P.A. Bottomley, C.J. Hardy, P.B. Roemer, R.G. Weiss, Problems and expediencies in human ^{31}P spectroscopy. The definition of localized volumes, dealing with saturation and the technique-dependence of quantification, NMR in Biomedicine, 2 (1989) 284-289.

[38] C.J. Hardy, R.G. Weiss, P.A. Bottomley, G. Gerstenblith, Altered myocardial high-energy phosphate metabolites in patients with dilated cardiomyopathy, American Heart Journal, 122 (1991) 795-801.

[39] R.G. Weiss, P.A. Bottomley, C.J. Hardy, G. Gerstenblith, Regional myocardial metabolism of high-energy phosphates during isometric exercise in patients with coronary artery disease, New England Journal of Medicine, 323 (1990) 1593-1600.

[40] P.A. Bottomley, R.G. Weiss, Non-invasive magnetic-resonance detection of creatine depletion in non-viable infarcted myocardium, The Lancet, 351 (1998) 714-718.

[41] R.G. Weiss, G. Gerstenblith, P.A. Bottomley, ATP flux through creatine kinase in the normal, stressed, and failing human heart, Proceedings of the National Academy of Sciences of the United States of America, 102 (2005) 808-813.

[42] J. Kurhanewicz, D.B. Vigneron, H. Hricak, P. Narayan, P. Carroll, S.J. Nelson, Three-dimensional H-1 MR spectroscopic imaging of the in situ human prostate with high (0.24-0.7-cm^3) spatial resolution, Radiology, 198 (1996) 795-805.

[43] S.J. Nelson, D.B. Vigneron, W.P. Dillon, Serial evaluation of patients with brain tumors using volume MRI and 3D ^{1}H MRSI, NMR in Biomedicine, 12 (1999) 123-138.

[44] J. Scheidler, H. Hricak, D.B. Vigneron, K.Y. Kyle, D.L. Sokolov, L.R. Huang, C.J. Zaloudek, S.J. Nelson, P.R. Carroll, J. Kurhanewicz, Prostate Cancer: Localization with Three-dimensional Proton MR Spectroscopic Imaging—Clinicopathologic Study, Radiology, 213 (1999) 473-480.

[45] C. Dowling, A.W. Bollen, S.M. Noworolski, M.W. McDermott, N.M. Barbaro, M.R. Day, R.G. Henry, S.M. Chang, W.P. Dillon, S.J. Nelson, Preoperative proton MR spectroscopic imaging of brain tumors: correlation with histopathologic analysis of resection specimens, American Journal of Neuroradiology, 22 (2001) 604-612.

[46] P.B. Barker, J.H. Gillard, P. Van Zijl, B.J. Soher, D.F. Hanley, A.M. Agildere, S.M. Oppenheimer, R.N. Bryan, Acute stroke: evaluation with serial proton MR spectroscopic imaging, Radiology, 192 (1994) 723-732.

[47] P.B. Barker, D.D. Lin, In vivo proton MR spectroscopy of the human brain, Progress in Nuclear Magnetic Resonance Spectroscopy, 49 (2006) 99-128.

[48] P.A. Bottomley, Spatial localization in NMR spectroscopy in vivo, Annals of the New York Academy of Sciences, 508 (1987) 333-348.

[49] J. Frahm, K.-D. Merboldt, W. Hänicke, Localized proton spectroscopy using stimulated echoes, Journal of Magnetic Resonance (1969), 72 (1987) 502-508.

[50] Y. Zhang, R.E. Gabr, M. Schär, R.G. Weiss, P.A. Bottomley, Magnetic resonance Spectroscopy with Linear Algebraic Modeling (SLAM) for higher speed and sensitivity, Journal of Magnetic Resonance, (2012) 66-76.

[51] S. Posse, G. Tedeschi, R. Risinger, R. Ogg, D.L. Bihan, High Speed [1]H Spectroscopic Imaging in Human Brain by Echo Planar Spatial-Spectral Encoding, Magnetic resonance in medicine, 33 (1995) 34-40.

[52] E. Adalsteinsson, P. Irarrazabal, S. Topp, C. Meyer, A. Macovski, D.M. Spielman, Volumetric spectroscopic imaging with spiral-based *k*-space trajectories, Magnetic resonance in medicine, 39 (1998) 889-898.

[53] U. Dydak, M. Weiger, K.P. Pruessmann, D. Meier, P. Boesiger, Sensitivity-encoded spectroscopic imaging, Magnetic Resonance in Medicine, 46 (2001) 713-722.

[54] M.A. Griswold, P.M. Jakob, R.M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, A. Haase, Generalized autocalibrating partially parallel acquisitions (GRAPPA), Magnetic Resonance in Medicine, 47 (2002) 1202-1210.

[55] F. Breuer, D. Ebel, J. Ruff, M. Blaimer, N. Seiberlich, M. Griswold, P. Jakob, Parallel 2D and 3D spectroscopic imaging using GRAPPA, in: Proc. Intl. Soc. Mag. Reson. Med, 2006, pp. 3653.

[56] Z.-P. Liang, P. Lauterbur, A generalized series approach to MR spectroscopic imaging, IEEE Transactions Medical Imaging, 10 (1991) 132-137.

[57] M. von Kienlin, R. Mejia, Spectral localization with optimal pointspread function, Journal of Magnetic Resonance (1969), 94 (1991) 268-287.

[58] L. An, S. Warach, J. Shen, Spectral localization by imaging using multielement receiver coils, Magnetic Resonance in Medicine, 66 (2011) 1-10.

[59] D.L. Donoho, Compressed sensing, Information Theory, IEEE Transactions on, 52 (2006) 1289-1306.

[60] E.J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, Information Theory, IEEE Transactions on, 52 (2006) 489-509.

[61] M. Lustig, D. Donoho, J.M. Pauly, Sparse MRI: The application of compressed sensing for rapid MR imaging, Magnetic Resonance in Medicine, 58 (2007) 1182-1195.

[62] S. Hu, M. Lustig, A. Balakrishnan, P.E. Larson, R. Bok, J. Kurhanewicz, S.J. Nelson, A. Goga, J.M. Pauly, D.B. Vigneron, 3D compressed sensing for highly accelerated hyperpolarized [13]C MRSI with in vivo applications to transgenic mouse models of cancer, Magnetic Resonance in Medicine, 63 (2009) 312-321.

[63] M. von Kienlin, M. Beer, A. Greiser, D. Hahn, K. Harre, H. Köstler, W. Landschütz, T. Pabst, J. Sandstede, S. Neubauer, Advances in human cardiac [31]P-MR spectroscopy: SLOOP and clinical applications, Journal of Magnetic Resonance Imaging, 13 (2001) 521-527.

[64] O. Geier, A.M. Weng, A. Toepell, D. Hahn, M. Spindler, M. Beer, H. Köstler, Acquisition-weighted chemical shift imaging improves SLOOP quantification of human cardiac phosphorus metabolites, Zeitschrift für Medizinische Physik, 49 (2013) 1-6.

[65] C.F.V. Loan, The ubiquitous Kronecker product, Journal of Computational and Applied Mathematics, 123 (2000) 85-100.

[66] K.P. Pruessmann, M. Weiger, P. Börnert, P. Boesiger, Advances in sensitivity encoding with arbitrary k-space trajectories, Magnetic Resonance in Medicine, 46 (2001) 638-651.

[67] W.S. Hoge, D.H. Brooks, B. Madore, W.E. Kyriakos, A tour of accelerated parallel MR imaging from a linear systems perspective, Concepts in Magnetic Resonance Part A, 27 (2005) 17-37.

[68] J.R. Roebuck, D.O. Hearshen, M. O'Donnell, T. Raidy, Correction of phase effects produced by eddy currents in solvent suppressed [1]H-CSI, Magnetic Resonance in Medicine, 30 (1993) 277-282.

[69] A. Simonetti, W. Melssen, M. Van der Graaf, A. Heerschap, L. Buydens, Automated correction of unwanted phase jumps in reference signals which corrupt MRSI spectra after eddy current correction, Journal of Magnetic Resonance, 159 (2002) 151-157.

[70] J.H. Duyn, J. Gillen, G. Sobering, P. Van Zijl, C. Moonen, Multisection proton MR spectroscopic imaging of the brain, Radiology, 188 (1993) 277-282.

[71] H. Zhu, P.B. Barker, MR spectroscopy and spectroscopic imaging of the brain, Methods in molecular biology (Clifton, NJ), 711 (2011) 203-226.

[72] I. Tkac, Z. Starcuk, I. Choi, R. Gruetter, In Vivo [1]H NMR Spectroscopy of Rat Brain at 1 ms Echo Time, Magnetic Resonance in Medicine, 41 (1999) 649-656.

[73] J.V. Hajnal, B. De Coene, P.D. Lewis, C.J. Baudouin, F.M. Cowan, J.M. Pennock, I.R. Young, G.M. Bydder, High signal regions in normal white matter shown by heavily T2-weighted CSF nulled IR sequences, Journal of Computer Assisted Tomography, 16 (1992) 506-513.

[74] J.V. Hajnal, D.J. Bryant, L. Kasuboski, P.M. Pattany, B. De Coene, P.D. Lewis, J.M. Pennock, A. Oatridge, I.R. Young, G.M. Bydder, Use of fluid attenuated inversion recovery (FLAIR) pulse sequences in MRI of the brain, Journal of Computer Assisted Tomography, 16 (1992) 841-844.

[75] H. Zhu, R. Ouwerkerk, P.B. Barker, Dual-band water and lipid suppression for MR spectroscopic imaging at 3 Tesla, Magnetic Resonance in Medicine, 63 (2010) 1486-1492.

[76] R.A. De Graaf, In vivo NMR spectroscopy: principles and techniques, Wiley-Interscience, 2008.

[77] F. Howe, S. Barton, S. Cudlip, M. Stubbs, D. Saunders, M. Murphy, P. Wilkins, K. Opstad, V. Doyle, M. McLean, Metabolic profiles of human brain tumors using quantitative in vivo [1]H magnetic resonance spectroscopy, Magnetic Resonance in Medicine, 49 (2003) 223-232.

[78] J.R. Roebuck, K.M. Cecil, M.D. Schnall, R.E. Lenkinski, Human breast lesions: characterization with proton MR spectroscopy, Radiology, 209 (1998) 269-275.

[79] S.J. Nelson, Analysis of volume MRI and MR spectroscopic imaging data for the evaluation of patients with brain tumors, Magnetic Resonance in Medicine, 46 (2001) 228-239.

[80] T.K.C. Tran, D.B. Vigneron, N. Sailasuta, J. Tropp, P. Le Roux, J. Kurhanewicz, S. Nelson, R. Hurd, Very selective suppression pulses for clinical MRSI studies of brain and prostate cancer, Magnetic Resonance in Medicine, 43 (2000) 23-33.

[81] Y. Luo, R. De Graaf, L. DelaBarre, A. Tannus, M. Garwood, BISTRO: An outer-volume suppression method that tolerates RF field inhomogeneity, Magnetic Resonance in Medicine, 45 (2001) 1095-1102.

[82] P.A. Bottomley, T.H. Foster, R.E. Argersinger, L.M. Pfeifer, A review of normal tissue hydrogen NMR relaxation times and relaxation mechanisms from 1–100 MHz: dependence on tissue type, NMR frequency, temperature, species, excision, and age, Medical Physics, 11 (1984) 425.

[83] G. Bydder, I. Young, MR imaging: clinical use of the inversion recovery sequence, Journal of Computer Assisted Tomography, 9 (1985) 659-675.

[84] A. Ebel, V. Govindaraju, A.A. Maudsley, Comparison of inversion recovery preparation schemes for lipid suppression in [1]H MRSI of human brain, Magnetic Resonance in Medicine, 49 (2003) 903-908.

[85] M.A. Smith, J. Gillen, M.T. McMahon, P.B. Barker, X. Golay, Simultaneous water and lipid suppression for in vivo brain spectroscopy in humans, Magnetic Resonance in Medicine, 54 (2005) 691-696.

[86] M. Gu, D.M. Spielman, B1 and T1 insensitive water and lipid suppression using optimized multiple frequency-selective preparation pulses for whole-brain 1H spectroscopic imaging at 3T, Magnetic Resonance in Medicine, 61 (2009) 462-466.

[87] M. Mescher, A. Tannus, M. Johnson, M. Garwood, Solvent suppression using selective echo dephasing, Journal of Magnetic Resonance, Series A, 123 (1996) 226-229.

[88] J. Star-Lack, S.J. Nelson, J. Kurhanewicz, L. Huang, D.B. Vigneron, Improved water and lipid suppression for 3D PRESS CSI using RF band selective inversion with gradient dephasing (BASING), Magnetic Resonance in Medicine, 38 (1997) 311-321.

[89] C.I. Haupt, N. Schuff, M.W. Weiner, A.A. Maudsley, Removal of lipid artifacts in [1]H spectroscopic imaging by data extrapolation, Magnetic Resonance in Medicine, 35 (1996) 678-687.

[90] X. Hu, M. Patel, K. Ugurbil, A new strategy for spectroscopic imaging, Journal of Magnetic Resonance, Series B, 103 (1994) 30-38.

[91] B. Bilgic, B. Gagoski, T. Kok, E. Adalsteinsson, Lipid suppression in CSI with spatial priors and highly undersampled peripheral k-space, Magnetic Resonance in Medicine, (2012).

[92] Y. Zhang, R.E. Gabr, M. Schär, R.G. Weiss, P.A. Bottomley, Magnetic resonance Spectroscopy with Linear Algebraic Modeling (SLAM) for higher speed and sensitivity, Journal of Magnetic Resonance, 218 (2012) 66-76.

[93] Y. Zhang, R.E. Gabr, J. Zhou, R.G. Weiss, P.A. Bottomley, Highly-accelerated quantitative 2D and 3D localized spectroscopy with linear algebraic modeling (SLAM) and sensitivity encoding, Journal of Magnetic Resonance, 237 (2013) 125-138.

[94] A. Maudsley, Sensitivity in Fourier imaging, Journal of Magnetic Resonance (1969), 68 (1986) 363-366.

[95] Z. Wang, L. Bolinger, V.H. Subramanian, J.S. Leigh, Errors of Fourier chemical-shift imaging and their corrections, Journal of Magnetic Resonance (1969), 92 (1991) 64-72.

[96] K.S. Miller, On the inverse of the sum of matrices, Mathematics Magazine, 54 (1981) 67-72.

[97] M.J. Powell, A fast algorithm for nonlinearly constrained optimization calculations, in: Numerical Analysis, Springer: Berlin Heidelberg, 1978, pp. 144-157.

[98] M.J. Powell, The convergence of variable metric methods for non-linearly constrained optimization calculations, Nonlinear Programming, 3 (1978).

[99] R.E. Gabr, R. Ouwerkerk, P.A. Bottomley, Quantifying in vivo MR spectra with circles, Journal of Magnetic Resonance, 179 (2006) 152-163.

[100] J.P. Mugler, J.R. Brookeman, Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE), Magnetic Resonance in Medicine, 15 (1990) 152-157.

[101] P. Roemer, W. Edelstein, C. Hayes, S. Souza, O. Mueller, The NMR phased array, Magnetic Resonance in Medicine, 16 (1990) 192-225.

[102] M.A. Bernstein, K.F. King, X.J. Zhou, Handbook of MRI pulse sequences, Elsevier, 2004.

[103] D.L. Parker, G.T. Gullberg, P.R. Frederick, Gibbs artifact removal in magnetic resonance imaging, Medical Physics, 14 (1987) 640.

[104] E.M. Haacke, R.W. Brown, M.R. Thompson, R. Venkatesan, Magnetic resonance imaging: physical principles and sequence design, Wiley-Liss New York, 1999.

[105] O. Ocali, E. Atalar, Intravascular magnetic resonance imaging using a loopless catheter antenna, Magnetic resonance in medicine, 37 (1997) 112-118.

[106] E. Atalar, P.A. Bottomley, O. Ocali, L.C. Correia, M.D. Kelemen, J.A. Lima, E.A. Zerhouni, High resolution intravascular MRI and MRS by using a catheter receiver coil, Magnetic resonance in medicine, 36 (1996) 596-605.

[107] A.-M.M. El-Sharkawy, D. Qian, P.A. Bottomley, The performance of interventional loopless MRI antennae at higher magnetic field strengths, Medical Physics, 35 (2008) 1995-2006.

[108] S. Sathyanarayana, M. Schär, D.L. Kraitchman, P.A. Bottomley, Towards real-time intravascular endoscopic magnetic resonance imaging, JACC: Cardiovascular Imaging, 3 (2010) 1158-1165.

[109] S. Sathyanarayana, P.A. Bottomley, MRI endoscopy using intrinsically localized probes, Medical physics, 36 (2009) 908-919.

[110] K.T. Block, M. Uecker, J. Frahm, Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint, Magnetic Resonance in Medicine, 57 (2007) 1086-1098.

[111] U. Gamper, P. Boesiger, S. Kozerke, Compressed sensing in dynamic MRI, Magnetic Resonance in Medicine, 59 (2008) 365-373.

[112] M. Usman, D. Atkinson, F. Odille, C. Kolbitsch, G. Vaillant, T. Schaeffter, P.G. Batchelor, C. Prieto, Motion corrected compressed sensing for free‐breathing dynamic cardiac MRI, Magnetic Resonance in Medicine, 70 (2013) 504-516.

[113] G.H. Glover, D.C. Noll, Consistent projection reconstruction (CPR) techniques for MRI, Magnetic resonance in medicine, 29 (1993) 345-351.

[114] G. Glover, J. Pauly, Projection reconstruction techniques for reduction of motion effects in MRI, Magnetic resonance in medicine, 28 (1992) 275-289.

[115] A. Shankaranarayanan, M. Wendt, J.S. Lewin, J.L. Duerk, Two-step navigatorless correction algorithm for radial k‐space MRI acquisitions, Magnetic resonance in medicine, 45 (2001) 277-288.

[116] J.A. Fessler, B.P. Sutton, Nonuniform fast Fourier transforms using min-max interpolation, (2003).

[117] F. Knoll, C. Clason, K. Bredies, M. Uecker, R. Stollberger, Parallel imaging with nonlinear reconstruction using variational penalties, Magnetic Resonance in Medicine, 67 (2012) 34-41.

[118] Feng L, Chandarana H, Xu J, Block T, Sodickson DK, O. R, k-t radial SPARSE-SENSE: Combination of compressed sensing and parallel imaging with golden angle radial sampling for highly accelerated volumetric dynamic MRI, in: Proceedings of the International Society for Magnetic Resonance in Medicine, 2012, pp. 81.

[119] Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: from error visibility to structural similarity, Image Processing, IEEE Transactions on, 13 (2004) 600-612.

[120] R.L. Ehman, J. Felmlee, Adaptive technique for high-definition MR imaging of moving structures, Radiology, 173 (1989) 255-263.

[121] M. Uecker, S. Zhang, D. Voit, A. Karaus, K.D. Merboldt, J. Frahm, Real‐time MRI at a resolution of 20 ms, NMR in Biomedicine, 23 (2010) 986-994.

[122] S. Hege, Y. Zhang, P. Bottomley, Accelerated, motion-corrected high-resolution intravascular MRI at 3T, in: Proc. Intl. Soc. Mag. Reson. Med, 2013.

[123] M. Schär, A.M.M. El‐Sharkawy, R.G. Weiss, P.A. Bottomley, Triple repetition time saturation transfer (TRiST) ^{31}P spectroscopy for measuring human creatine kinase reaction kinetics, Magnetic Resonance in Medicine, 63 (2010) 1493-1501.

Curriculum Vita

Yi Zhang was born in the historic city of Anyang, China in 1988. He received his bachelor degree with honors in Electrical Engineering from Zhejiang University, China in 2009. He won the first prize in China Undergraduate Mathematical Contest in Modeling in 2007. He started his PhD training under the guidance of Dr. Paul A. Bottomley in the Department of Electrical and Computer Engineering at Johns Hopkins University in 2009. He is currently a student member of the International Society for Magnetic Resonance in Medicine (ISMRM) and the Radiological Society of North America. He received the Magna Cum Laude Merit Award from ISMRM in 2012 and 2013. His current research primarily involves localization, reconstruction and processing techniques in Magnetic Resonance Spectroscopy and Magnetic Resonance Imaging.