

Magnetic resonance tissue density estimation using optimal SSFP pulse-sequence design

Anand, C.; Sotirov, R.; Terlaky, T.; Zheng, Z.

Published in:
Optimization and Engineering

Publication date:
2007

[Link to publication](#)

Citation for published version (APA):

Anand, C., Sotirov, R., Terlaky, T., & Zheng, Z. (2007). Magnetic resonance tissue density estimation using optimal SSFP pulse-sequence design. *Optimization and Engineering*, 8(2), 215-238.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright, please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Magnetic Resonance Tissue Quantification using Optimal bSSFP Pulse-Sequence Design

Christopher Kumar Anand, Renata Sotirov, Tamás Terlaky,
Zhuo Zheng

*Department of Computing and Software, McMaster University, Hamilton, Ontario,
Canada L8S 4K1*

Abstract. We propose a merit function for the expected contrast to noise ratio in tissue quantifications, and formulate a nonlinear, nonconvex semidefinite optimization problem to select locally-optimal balanced steady-state free precession (bSSFP) pulse-sequence design variables. The method could be applied to other pulse sequence types, arbitrary numbers of tissues, and numbers of images. To solve the problem we use a mixture of a grid search to get good starting points, and a sequential, semidefinite, trust-region method, where the subproblems contain only linear and semidefinite constraints. We give the results of numerical experiments for the case of three tissues and three, four or six images, in which we observe a better increase in contrast to noise than would be obtained by averaging the results of repeated experiments. As an illustration, we show how the pulse sequences designed numerically could be applied to the problem of quantifying intraluminal lipid deposits in the carotid artery.

Keywords: magnetic resonance imaging, balanced steady-state free precession, Dixon method, semidefinite programming, trust-region algorithm

1. Introduction

Magnetic Resonance Imaging (MRI) is widely used in diagnosis, treatment monitoring and research to generate information both about structure and function of tissues. Most clinical imaging focuses on qualitative imaging (Haacke et al., 1999, Ch. 15), but quantitative applications do exist (Haacke et al., 1999, Ch. 21), especially in functional imaging, be it flow measurements in arteries or brain activity (Haacke et al., 1999, Ch. 24). In this paper, we focus on quantitative imaging of structure. Quantification of different tissue components is commonly performed in clinical MR spectroscopy (MRS) (Salibi and Brown, 1998), which offers very detailed information about tissue concentrations, but has either restricted spatial resolution, or extremely long acquisition times; and via the Dixon method (Dixon, 1984), in which concentrations of fat and water are calculated by taking linear combinations of images that were acquired by using pulse-sequence design variables chosen to create prescribed phase relationships between fat, water and (less com-

monly) silicone. Many pulse sequences provide a measure of fat-water separation, which may be referred to as quantification. We will reserve “quantification” for the case where the reported quantities depend only on tissue quantity and not other tissue parameters (resonance offset, relaxation) described below. Sequences which are described in the literature as “ T_1 – or T_2 –weighted” are not quantitative in our strict sense, because pixel intensity in these images depends on relaxation rates in addition to tissue quantity.

After explaining the Dixon method, we show how it can be generalized by allowing more general phase relationships, and we develop a measure of the efficacy of a given set of pulse-sequence design variables for tissue quantification. We carry this out for balanced Steady-State Free Precession (bSSFP) (Carr, 1958; Oppelt et al., 1986) imaging, but other sequence types can be similarly modeled. Having such a model opens up the possibility of choosing pulse-sequence design variables to minimize the expected noise in the computed tissue densities, by solving a nonlinear, semidefinite optimization problem (Ben Tal and Nemirovski, 2001). We present here an iterative method for solving this problem numerically. Our algorithm converges to a local optimum.

Note that although our approach can be developed into a clinical replacement for some applications of MRS, it is not closely related to MRS, and we have not considered the application of our models and method of optimization to the maximization of information in MRS signals.

The paper is organized as follows. In Section 2, we present several advantages of our method of quantification. In Section 3, we sketch the theory of bSSFP signal generation, which we need in Section 4 to formulate the optimization problem. In Section 5, we present an algorithm for solving the problem, and in Section 6 we give numerical results to support the adequacy of this strategy. In Section 7, we present an illustration of a possible clinical application of these methods for carotid artery imaging. In the final section, we address the many possibilities for applying our method, generalizing the model to other sequence types and including sources of systematic error, and further developing the theory and practice of the solution to the optimization problem.

2. Advantages of the Proposed Method of Quantification

In cases where distinct tissue types are well separated physically, and large in extent relative to image resolution, quantitative tissue volume

has been successfully extracted from qualitative images by calculations based on manual and automatic contour estimation (Jardim and Figueiredo, 2003). These methods of tissue separation are based on selectively suppressing tissue from undesired components, and are sensitive to main field inhomogeneity and other conditions, which make them unsuitable for quantification. However, our method has no such restrictions.

An alternative to MR imaging is MR spectroscopy, which can be used to quantify multiple molecules in regions of interest. Unfortunately, it is extremely limited in spatial and temporal resolution, and it is at least an order of magnitude slower to acquire than imaging methods, such as the method we propose.

Our method improves on conventional Dixon methods (Glover, 1991; Huang et al., 2004; Reeder et al., 2004; Vasanawala et al., 2000) since it

- does not dictate a particular configuration of phases (*e.g.*, in and out of phase) and look for pulse-sequence design variables to match this configuration,
- takes all tissue parameters into account (relaxation constants, as well as resonance offset),
- can be used to quantify any number of tissues, and
- increases contrast-to-noise ratio as a function of images acquired faster than signal averaging—and allows greater latitude in trading off imaging time versus contrast to noise.

These advantages are the direct results of formulating the selection of pulse-sequence design variables as an optimization problem, including the formulation of an objective function which measures contrast to noise.

Since our model takes all tissue parameters into account, we expect to be able to separate tissue types which could not have been separated by conventional Dixon methods, but we have only verified this numerically in the present work, and plan to undertake further work to make this technique robust enough for clinical application. Most importantly, we have started to incorporate calibration and correction for unavoidable variation in the main magnetic field (called B_0 inhomogeneity estimation in the MR literature), with the same aim as (Sutton et al., 2003), but with different numerical techniques.

3. Theory of MRI

There are many types of MR imaging sequences, to which our methods apply. We only formulate the design problem for tissue optimization for bSSFP, introduced in (Oppelt et al., 1986). We generally follow the notation of (Hargreaves et al., 2003). To make the paper self-contained, we sketch the basic theory needed to formulate this problem, both for mathematicians who want some understanding of the source of the problem, and for imaging experts. Readers looking for a more details should consult (Haacke et al., 1999). References to specific sections are contained within the text.

3.1. MR SIGNAL GENERATION AND THE BLOCH EQUATION

Magnetic Resonance Imaging and Spectroscopy measure the aggregate magnetic field generated by the spins of one or more nuclei in the object. Proton (hydrogen) imaging is by far the most common. Because the number of protons is very large, we model them as a density, using continuous or discontinuous functions. The most common functions come from a regular division of the imaging volume into equal cubes or more general rectangular prisms, called voxels. Approximating functions can be either piecewise constant, or sums of delta functions centered on each voxel. To each voxel, then, we can assign a vector $M \in \mathbb{R}^3$ representing the net magnetic moment of all the protons in that voxel. For each type of tissue, the magnitude of M is proportional to the quantity of protons. Since we are interested in quantifying different tissue types, our voxel model includes magnetization vectors for tissues of different types. This tissue model is an approximation, with all real tissues comprising different cell types, with multiple molecular environments. How detailed to make the model depends on the application: for example, some applications may require separate “tissue” components for intracellular and extracellular water for each tissue, while an automated fat volume estimator may work with only two components. The mathematics does not require that components correspond to tissues, but we do so to simplify the presentation.

In the rest of this paper, we will simplify the exposition, by considering a single voxel with fixed tissue concentrations, *i.e.* we ignore flow across the voxel boundary. The evolution in the magnetization in one voxel is independent of the magnetization in neighboring voxels, hence, we can consider the signal from each voxel separately.

An MR imager is basically a big magnet. The moment, M , precesses in the external field, $B \in \mathbb{R}^3$, according to

$$\frac{dM}{dt} = \gamma M \times B, \quad (1)$$

where γ is a physical constant (Haacke et al., 1999, Ch. 2). This is the same equation which governs a gyroscope suspended on a string. The magnetic field, B , is a function of the chemical environment, so it varies from one tissue to another (Liang and Lauterbur, 1999, Ch. 3). (In a real imager, manufacturing tolerances and other factors also influence B , but we don't discuss them in this paper.) This equation can be empirically validated for short time periods (microseconds), but for long time periods (milliseconds), quantum interactions must be taken into account. Quantum interactions cause the net magnetization to return to thermodynamic equilibrium $M_0 = (0, 0, \text{proton density})^T$ (Haacke et al., 1999, Ch. 5). Adding two terms to account for the empirically measured relaxation, we arrive at the Bloch Equation (Haacke et al., 1999, Ch. 4):

$$\frac{dM}{dt} = \gamma M \times B - \frac{1}{\tau_2} \begin{pmatrix} M_x \\ M_y \\ 0 \end{pmatrix} - \frac{1}{\tau_1} \begin{pmatrix} 0 \\ 0 \\ M_z - M_{0,z} \end{pmatrix}, \quad (2)$$

here $M(t) \in \mathbb{R}^3$ is the aggregate magnetization and τ_1, τ_2 are empirically-determined, tissue-dependent constants that represent the longitudinal relaxation time and transverse relaxation time, respectively. (In the MR literature, T_1 and T_2 is the more common notation, but this conflicts with the natural names for design variables to be introduced later.) Since (2) depends on material properties, one solves for the evolution of each tissue's magnetization separately. The measured signal is a linear combination of the signals from each tissue component.

In an MR imager, we use electro-magnets to transmit magnetic waves at different frequencies to perturb spins from equilibrium, M_0 , parallel to the main field, and measure the projection of the resulting precession on the transverse plane, which we interpret as a complex-value, $M_x + iM_y$. We will next explain how the tissue quantities can be recovered from these complex values. Then we will return to the Bloch Equation and model the generation of the signal for a particular pulse sequence, and explain how given such a model we can minimize the error in the resulting tissue quantification.

3.2. TISSUE QUANTIFICATION

For a sample voxel that consists of m tissue types with concentrations $(\rho_1, \rho_2, \dots, \rho_m) \in \mathbb{R}^m$, which is imaged n times by varying the pulse sequence, the resulting signals $(\iota_1, \iota_2, \dots, \iota_n) \in \mathbb{C}^n$, are given by

$$\begin{aligned}\iota_1 &= a_{11}\rho_1 + a_{12}\rho_2 + \dots + a_{1m}\rho_m, \\ \iota_2 &= a_{21}\rho_1 + a_{22}\rho_2 + \dots + a_{2m}\rho_m, \\ &\vdots \\ \iota_n &= a_{n1}\rho_1 + a_{n2}\rho_2 + \dots + a_{nm}\rho_m,\end{aligned}$$

where by convention $\iota_k = M_x(k) + \sqrt{-1}M_y(k)$ is the projection of the magnetization $M_{k\text{th tissue}}$ to the x - y plane, which is identified with the complex plane, and $a_{ij} \in \mathbb{C}$ gives the expected signal in image i of a unit quantity of tissue j . If $\text{rank}(S) = m$ (where S is considered as a real matrix of having twice the number of rows as the complex coefficient matrix $S^{\mathbb{C}} = (a_{ij})_{n \times m}$ by splitting the real and imaginary parts of a_{ij} into two adjacent rows), then we can invert this linear system to find the tissue concentrations. This is the basic idea of tissue quantification, the interesting part is how to choose pulse-sequence design variables so that S is well conditioned.

3.3. DIXON METHOD

All imaging methods in MR are based on understanding (and working with) the behavior of solutions to the Bloch equation (2). The Dixon Method, (Dixon, 1984), uses chemical differences between tissues which manifest themselves as different external field values, and hence different resonant frequencies. For tissues with different resonant frequencies, the signals will go in and out of phase periodically. Dixon observed that for fat and water, if one image is acquired while both are in phase and one while they have opposite phases, then

$$S^{\mathbb{C}} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad (3)$$

with respect to the natural bases for the complex vector space of pixel values in two images, and the real vector space of tissue quantities. Expressed with respect to the real basis of real and imaginary parts of image one, followed by real and imaginary parts of image two:

$$S = \begin{pmatrix} 1 & 1 \\ 0 & 0 \\ 1 & -1 \\ 0 & 0 \end{pmatrix}. \quad (4)$$

Note that addition and subtraction of images is all that is required to recover the original fat and water concentrations.

The difference between the effective excitation time and the effective measurement time is called the echo time. Dixon fixed the phase relationship by altering the echo time, but ignored differences caused by relaxation. When we formulate the model for bSSFP sequences, we will be concerned with repetition time, the time between successive excitation pulses. Repetition time always effects signal generation, but for spoiled pulses, it can be excluded from calculations of the complex phase of the signal, whereas for bSSFP, phase (and magnitude) are very sensitive to repetition times (Zur et al., 1988; Scheffler, 1999).

The Dixon method usually uses spoiled sequences, because (ignoring the calibration issue) relaxation parameters can be ignored. For bSSFP sequences, relaxation parameters cannot be ignored, and Hargreaves et al., (Hargreaves et al., 2003), adjust the pulse-sequence design variables of bSSFP sequences by proper selection of the sequence repetition time and the center frequency offset respectively in order to realize a 180° phase difference between fat/water signals. Other authors, *e.g.*, Vasanaawala et al., (Vasanaawala et al., 2000), have considered more general tissue separation problems, but always with similar simple structures for S^C .

Dixon introduced the principle of separating different tissue types by manipulating the phase relationships, which can be easily extended to the cases more than fat and water, as shown in our simulation section (see §7). The linear transformations S used in the Dixon method are simple by design, and therefore usually not written in matrix form. Not doing so obscured the fact that more general linear combinations of tissue densities can be used for tissue quantification. Recently, Rybicki *et al.*, (Rybicki, 2001), and Reeder *et al.*, (Reeder et al., 2004), have considered matrices containing square roots and arbitrary complex roots of unity, respectively. This complex matrix can be approximated by manipulating a single pulse design variable (echo time), when the resonant frequencies of the different tissues are well separated. To the best of our knowledge, we are the first to point to, and to exploit the generalization to arbitrary complex values. The introduction of a merit function, and a systematic computational method of optimizing this function over all possible choices of pulse-sequence design variables for the case of bSSFP pulses are the main contributions of this paper. Note that Reeder *et al.*, do use noise in the tissue separations as an objective, and discuss the merits of their proposed heuristic for choosing echo times. They find that in the best case, noise in computed tissue separations is reduced as would be expected from signal averaging alone. By

optimizing over a larger number of design variables, we observe better than the expected reduction from signal averaging alone (see §7).

3.4. BALANCED STEADY-STATE FREE PRECESSION

When a spin system is excited by a train of periodic radio frequency (RF) pulses with repetition time $T \ll \tau_2$, the spin system will reach a dynamic equilibrium, known as the steady state (Carr, 1958). There are quite a number of studies about the bSSFP signal and bSSFP pulse sequences in recent years, such as (Hanicke and Vogel, 2003; Scheffler, 2003; Hargreaves et al., 2001) because bSSFP has the advantage of yielding high signal in short scan time, high image resolution and good image contrast, etc. For this reason, we use bSSFP sequences to demonstrate our approach of optimally determining the pulse-sequence design parameters to maximize contrast to noise.

The physical system is well described by the Bloch equation (2), but for our purposes it is simpler to discretize the system by assuming the fixed simple pulse-sequence design shown in Figure 1 and integrating the Bloch equation in each section. The magnetization we measure is the unique fixed-point corresponding to the steady state of this dynamical system, which is a function of the design variables and parameters describing the physical properties of the tissues we wish to measure.

The dynamical system is the composition of several simple components, parameterized by the tissue parameters

$$(\tau_1, \tau_2, \kappa) \in \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}, \quad (5)$$

and the pulse-sequence design variables

$$(\alpha, f, T) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}_+, \quad (6)$$

where τ_1, τ_2 are decay times of the magnetization measured in milliseconds (ms), κ is the resonance offset of the tissue measured in Hz , \mathbb{R}_+ denotes the set of nonnegative real numbers, α is the flip angle, f is the angle of RF phase cycling, T is the echo time which we fix to be half the repetition time. Fixing the echo time to be midway between RF pulses corresponds to the most basic spin-warp and radial k-space trajectories (Haacke et al., 1999, Ch. 9). Other choices of readout trajectories would lead to other fixed ratios. Simultaneously optimizing RF pulse design and readout trajectories would lead to a more complicated optimization problem, which is beyond the scope of this paper.

The dynamical system (Hargreaves et al., 2001; Janes, 2000) has the form

$$M_{k+1} = \bar{A}M_k + b \quad (7)$$




Figure 1. bSSFP pulse sequence showing two pulse repetitions (a) and (b). Each pulse interval contains one readout interval (c) when data is collected, and one RF pulse (d). We model the RF pulse and data readout as occurring at single points of time, with the readout time equally spaced (by time T (e) and (f) from its bracketing RF pulses. The affine transformations associated with these events are R and Q acting at (d), and P and $C \cdot +D$ acting twice during each repetition at (e) and (f).

with steady state $M_{SS}(\alpha, f, T, \tau_1, \tau_2, \kappa)$ which has been calculated explicitly in (Freeman et al., 1971). We will build up the dynamical system from its components slightly differently from previous developments, because we want to optimize the signal at the middle point of the readout. This computation is not difficult, but establishes the notation we use, and makes the paper self-contained for readers without sufficient knowledge of MR physics.

The first component is the rotation of the magnetization vector caused by the RF pulse (consider the RF pulse to be a delta function)

$$R = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\alpha) & \sin(\alpha) \\ 0 & -\sin(\alpha) & \cos(\alpha) \end{pmatrix}. \quad (8)$$

Variations in bulk susceptibility of the different tissues cause different tissue to observe different magnetic fields B , which we capture by the parameter κ , this causes a rotation about the z axis:

$$P = \begin{pmatrix} \cos(\kappa T) & \sin(\kappa T) & 0 \\ -\sin(\kappa T) & \cos(\kappa T) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (9)$$

In designing the pulse sequence, we can change the axis of rotation of the RF pulse. This is referred to as RF pulse phase cycling, because

current generation scanners have the ability to rotate this angle by a fixed constant, which is usually an integral fraction of 2π . To simplify this analysis, it is easier to fix the RF phase, and pretend that the rest of the experiment rotates from one RF pulse to the next by

$$Q = \begin{pmatrix} \cos(f) & \sin(f) & 0 \\ -\sin(f) & \cos(f) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (10)$$

We collect the effect of relaxation into two components:

$$C = \begin{pmatrix} e^{-T/\tau_2} & 0 & 0 \\ 0 & e^{-T/\tau_2} & 0 \\ 0 & 0 & e^{-T/\tau_1} \end{pmatrix} \quad (11)$$

and

$$D = \begin{pmatrix} 0 \\ 0 \\ (1 - e^{-T/\tau_1}) \end{pmatrix}, \quad (12)$$

where the scaling of D assumes that M has been scaled so that the minimum energy state is $M_0 = (0, 0, \text{proton density})^T$. Therefore

$$M_{k+1} = P(C(RQP(CM_k + D)) + D).$$

If a steady state exists, it satisfies

$$AM_{\text{SS}} = b, \quad (13)$$

where $\bar{A} = PCRQPC$, $A = I - \bar{A}$ and $b = PCRQPD + PD$.

Since the rotations R , P , and Q have unit eigenvalues, and all the eigenvalues of the relaxation matrix C are less than one (and positive), the matrix A is invertible, and so the steady state exists.

3.5. IMAGING

Let m be the number of tissues, and n the number of experiments. We denote $u_l = (u_{1l}, u_{2l}, u_{3l}, u_{4l}, T_l) \in \mathbb{R}^5$, $l = 1, \dots, n$, $t_k = (\tau_{1k}, \tau_{2k}, \kappa_k) \in \mathbb{R}^3$, $k = 1, \dots, m$, where we represent the angles α_l and f_l by unit vectors $(u_{1l}, u_{2l}) = (\cos(\alpha_l), \sin(\alpha_l))$ and $(u_{3l}, u_{4l}) = (\cos(f_l), \sin(f_l))$. If $M_{\text{SS}}(u_l, t_k) \in \mathbb{R}^3$ is the steady-state magnetization corresponding to design variables u_l and tissue t_k , with components $M_{\text{SS},x}$, $M_{\text{SS},y}$, $M_{\text{SS},z}$, and ρ_k is the density of tissue k , then the measured signal from experiment l is

$$\left(\sum_{k=1}^m M_{\text{SS},x}(u_l, t_k)\rho_k, \sum_{k=1}^m M_{\text{SS},y}(u_l, t_k)\rho_k \right).$$

If we write the results of n experiments (*i.e.* images) as a $2n$ dimensional real vector, and the m dimensional tissue densities as an m dimensional real vector, the transformation from tissue densities to measurements is

$$S = \begin{pmatrix} M_{SS,x}(u_1, t_1) & \dots & M_{SS,x}(u_1, t_m) \\ M_{SS,y}(u_1, t_1) & \dots & M_{SS,y}(u_1, t_m) \\ \vdots & \vdots & \vdots \\ M_{SS,x}(u_n, t_1) & \dots & M_{SS,x}(u_n, t_m) \\ M_{SS,y}(u_n, t_1) & \dots & M_{SS,y}(u_n, t_m) \end{pmatrix}, \quad (14)$$

If S does not have full rank, we cannot reconstruct tissue densities from this set of images. If S does have full rank, the Moore-Penrose Pseudo-Inverse (Horn and Johnson, 1985) of S , $(S^T S)^{-1} S^T$, is an unbiased maximum likelihood estimator (Mardia et al., 1979, Thms. 6.2.1, 6.2.2) for the tissue densities as a function of the measured image data. In the next section, we will show how the pulse-sequence design variables can be chosen, not only to avoid singular transformations, but to maximize the quality of the computed tissue densities.

4. Optimal bSSFP Design for Tissue Segmentation

Our objective is to choose the pulse-sequence design variables such that the error in the reconstructed tissue densities is minimized. As is standard practice in MR imaging, we will assume that measurement noise is white, that is, independent and normally distributed (Haacke et al., 1999). Under this assumption the error in the reconstructed tissue densities will also be normally distributed (since the reconstruction is a linear transformation), but the error in different tissue components will not necessarily be independent. Barring additional information about the use of the tissue segmentation (*e.g.*, the method of diagnosis in which they will be employed), we assume that the objective is to minimize the worst-case error among different tissue densities. Even if in a particular application, we were only concerned with the quantification of one tissue, it is likely that radiologists would occasionally want to examine other tissue quantification images, so minimizing the worst-case error is a good conservative design criterion.

4.1. FORMULATION

Since the Moore-Penrose Pseudo-Inverse (Horn and Johnson, 1985) is a linear map from measured signals to quantity estimates, it follows that if $\epsilon \in \mathbb{R}^{2m}$ is a vector of measured noise, then the resulting errors in

the tissue densities are also normally, but not identically distributed, and given by $(S^T S)^{-1} S^T \epsilon$.

We can calculate the expected error in the tissue quantity estimates by using a singular-value decomposition $S = V^T D U$, where V and U are orthonormal and D is diagonal. Then we have

$$(S^T S)^{-1} S^T \epsilon = U^T D^{-1} V \epsilon. \quad (15)$$

Since each measurement noise is independent with distribution $N(0, \sigma)$, for any two rows V_i and V_j in the orthonormal matrix V , $V_i \epsilon$ and $V_j \epsilon$ are also independent with distribution $N(0, \sigma)$, which means that each component in $\epsilon' = V \epsilon$ is also independent with distribution $N(0, \sigma)$. Similarly, each component in $\epsilon'' = D^{-1} \epsilon'$ is independent with distribution $N(0, \lambda_i^{-1} \sigma)$, in which λ_i is the i th eigenvalue of D . When left multiplied by the matrix U^T , the resulting error $\epsilon''' = U^T \epsilon''$ also has independently distributed components, with the i th element having distribution $N(0, \sigma \sqrt{\sum_j (U_{i,j}^T)^2 \lambda_j^{-2}})$. To minimize the worst-case error, we want to minimize the largest value

$$\sigma \sqrt{\sum_j (U_{i,j}^T)^2 \lambda_j^{-2}},$$

for all values of i . This is equivalent to minimizing the maximal λ_i^{-2} , which is equivalent to maximizing the minimal λ_i . Since

$$S^T S = U^T D V V^T D U = U^T D^2 U,$$

and the eigenvalues of D^2 and $U^T D^2 U$ coincide because $U^T U = I$, the eigenvalues of D^2 are the same as the eigenvalues of $S^T S$. Thus our goal is to

$$\max_{\text{design variables}} \min\{\text{eigenvalues of } S^T S\}. \quad (16)$$

We formulate the eigenvalue optimization problem (16) using semi-definite inequalities as

$$\begin{aligned} & \max \lambda \\ & S^T S - \lambda I \succeq 0, \end{aligned}$$

where \succeq means positive semidefiniteness, and additional constraints for S and for the design variables describe the set of accessible design variables.

The formulation of the constraints involves three angles that always occur as sine-cosine pairs, which define rotations around two of the coordinate axes in \mathbb{R}^3 . We can replace the sines and cosines that appear

in the rotation matrices R and Q with unit vectors u_{1l} and u_{2l} , and add the constraints

$$u_{1l}^2 + u_{2l}^2 = 1, \quad (17)$$

$$u_{3l}^2 + u_{4l}^2 = 1. \quad (18)$$

In our model we relax the constraints (17)-(18) to the convex quadratic constraints

$$u_{1l}^2 + u_{2l}^2 \leq 1, \quad (19)$$

$$u_{3l}^2 + u_{4l}^2 \leq 1. \quad (20)$$

The numerical computations justify such relaxation and verify our expectations; the convergence of the relaxed constraints to the boundary of the feasible set occurs during iterations of our solving algorithm. At the local optimum we obtain equalities in (19)-(20).

In addition to the angles, we have variables T_l , which control the repetition time of the pulse sequence used to collect an image. In practical applications, this variable is bounded below by hard physical constraints on the instrumentation (T_{\min}), and it is bounded above by practical limits on bSSFP image stability (T_{\max}), and limits on the patients ability to remain still (possibly including holding their breath). (These limits also depend on the particular model of imager, on the part of the body being imaged, and the field strength of the magnet.)

4.2. COMPLETE SYSTEM

Recall that m stands for the number of tissues and n stands for the number of experiments, and

$$A(u_l, t_k) = I - P(T_l, \kappa_k)C(T_l, \tau_{1k}, \tau_{2k})R(u_{1l}, u_{2l})Q(u_{3l}, u_{4l})P(T_l, \kappa_k)C(T_l, \tau_{1k}, \tau_{2k}),$$

$$b(u_l, t_k) = (P(T_l, \kappa_k)C(T_l, \tau_{1k}, \tau_{2k})R(u_{1l}, u_{2l})Q(u_{3l}, u_{4l}) + I)P(T_l, \kappa_k)D(T_l, \tau_{1k}),$$

and

$$S(u_1, \dots, u_n, t_1, \dots, t_m) = \begin{pmatrix} M_{SS,x}(u_1, t_1) & \dots & M_{SS,x}(u_1, t_m) \\ M_{SS,y}(u_1, t_1) & \dots & M_{SS,y}(u_1, t_m) \\ \vdots & \vdots & \vdots \\ M_{SS,x}(u_n, t_1) & \dots & M_{SS,x}(u_n, t_m) \\ M_{SS,y}(u_n, t_1) & \dots & M_{SS,y}(u_n, t_m) \end{pmatrix}. \quad (21)$$

Using this notation, we can now write the entire optimization model in the compact form:

$$\begin{aligned}
& \max \lambda \\
& \text{s.t. } S^T S - \lambda I \succeq 0 \\
& \quad A(u_l, t_k) M_{SS}(u_l, t_k) = b(u_l, t_k) \quad \forall l, k \\
& \quad u_{1l}^2 + u_{2l}^2 \leq 1 \quad \forall l \\
& \quad u_{3l}^2 + u_{4l}^2 \leq 1 \quad \forall l \\
& \quad T_l \in [T_{\min}, T_{\max}],
\end{aligned} \tag{NL-SDO}$$

where $l = 1, \dots, n$, $k = 1, \dots, m$. Note that in the nonlinear semidefinite optimization problem (NL-SDO), constraints $u^2 + v^2 \leq 1$ are second-order cone constraints (SOCO), see (Ben Tal and Nemirovski, 2001).

5. A Trust-Region Algorithm for NL-SDO

In this section, we explain our algorithm for solving (NL-SDO), where (T_1, \dots, T_n) and (t_1, \dots, t_m) are fixed parameters. We name the model with fixed T_l and t_k values (NL-SDO), (see page 14). The values of the tissue parameters (t_1, \dots, t_m) that are used in the algorithm for solving (NL-SDO), are representative values chosen from the literature, and the values of the design variables T_l , $l = 1, \dots, n$ are obtained by a grid search that is explained in §5.2.

Our aim is to solve the nonlinear semidefinite problem (NL-SDO) by solving a sequence of linear mixed semidefinite (SDO) and second-order cone (SOCO) trust-region subproblems. We chose to fix the values of T_l in the nonlinear problem because the T_l appear in multiple places in the constraints (within both real exponentials and sine and cosine) which come from the dynamical system (13), and we couldn't find a suitable change of coordinates (as we did in changing from angles to unit vectors) to reduce the nonlinearity. This decision was supported by subsequent visual exploration of the objective function, which strongly suggested that the objective is 'less convex' and has a larger Lipschitz constant in the T_l directions than the other variable directions. Rapid changes in the objective as a function of the variables T_l is consistent with the inverse relationship between the repeat time and the resonance frequency. More investigation is required to understand the differences in roles of the pulse-sequence design variables, methods such as automatic differentiation are needed to overcome the technical difficulty in treating the T_l variables—and in dealing with many more degrees of freedom, and more complicated pulse design problems. Given the focus

of the present paper on demonstrating the validity of the SDO formulation and gauging the difficulty of solution (on a clinically relevant, but relatively simple test problem), we felt it was reasonable to fix the values of T_l in the subproblems.

In this section we first derive a linear mixed SDO and SOCO trust-region subproblem (P_{lin}), (see page 17) obtained by linearizing (NL-SDO) with respect to $(u_{1l}, u_{2l}, u_{3l}, u_{4l})$, $l = 1, \dots, n$, and then explain our sequential SDO-trust-region algorithm based on (P_{lin}), (see page 17). Sequential semidefinite programming methods are also observed in (Freund and Jarre, 2003) and (Kruk and Wolkowicz, 1998).

5.1. LINEARIZATION AND THE SDO-TRUST REGION SUBPROBLEM

It is known that mixed SDO and SOCO problems can be solved efficiently with interior-point methods (IPMs), (see (Sturm, 2002)). Therefore, for the eigenvalue problem (NL-SDO) we introduce SDO-SOCO subproblems defined on a trust region. The semidefinite eigenvalue constraint

$$S^T S - \lambda I \succeq 0 \quad (22)$$

is not fitting in the form of linear conic optimization, therefore we substitute out the quadratic term as $X = S^T S$ and replace constraint (22) by the following two constraints

$$\begin{aligned} X &= S^T S \\ X - \lambda I &\succeq 0. \end{aligned}$$

The first constraint is nonlinear while the second one is a standard SDO constraint that can be straightforwardly implemented in any standard solver. By taking into consideration that T_l are *fixed parameters*, we rewrite (NL-SDO) problem as follows

$$\begin{aligned} \max \quad & \lambda \\ \text{s.t.} \quad & X - \lambda I \succeq 0 \\ & X = S^T S, \\ & A(u_l, t_k) M_{SS}(u_l, t_k) = b(u_l, t_k), \quad \forall l, k \\ & u_{1l}^2 + u_{2l}^2 \leq 1, \quad \forall l \\ & u_{3l}^2 + u_{4l}^2 \leq 1, \quad \forall l, \end{aligned} \quad (\text{P})$$

where $l = 1, \dots, n$, and $k = 1, \dots, m$.

Now, we develop a *sequential semidefinite trust-region based algorithm* for solving (NL-SDO). At each iteration we define a linear SDO-SOCO trust-region subproblem by linearizing the nonlinear constraint

around the current point and restrict the movement to a certain trust region. Our next step is to linearize with respect to $(u_{1l}, u_{2l}, u_{3l}, u_{4l})$, $l = 1, \dots, n$ the constraint

$$X = S^T S, \quad (23)$$

To simplify the notation, we write $X = X(u_1, \dots, u_n, t_1, \dots, t_m)$. Each element in matrix S corresponds to x or y component of a steady-state magnetization that depends on tissue parameters t_k and design variables u_l , (see 21). Therefore, for the linearization of the nonlinear constraint (23) we explore the nonlinear constraints

$$A(u_l, t_k) M_{SS}(u_l, t_k) = b(u_l, t_k), \quad l = 1, \dots, n, \quad k = 1, \dots, m, \quad (24)$$

and remove them from the derived conic subproblem. It is easy to see that

$$\begin{aligned} X_{pq} &= \sum_{l=1}^n S_{lp} S_{lq} \\ &= \sum_{l=1}^n (M_{SS,x}(u_l, t_p) M_{SS,x}(u_l, t_q) + M_{SS,y}(u_l, t_p) M_{SS,y}(u_l, t_q)), \end{aligned} \quad (25)$$

and

$$\begin{aligned} \frac{\partial X_{pq}}{\partial u_{jl}} &= \frac{\partial M_{SS,x}(u_l, t_p)}{\partial u_{jl}} M_{SS,x}(u_l, t_q) + M_{SS,x}(u_l, t_p) \frac{\partial M_{SS,x}(u_l, t_q)}{\partial u_{jl}} \\ &\quad + \frac{\partial M_{SS,y}(u_l, t_p)}{\partial u_{jl}} M_{SS,y}(u_l, t_q) + M_{SS,y}(u_l, t_p) \frac{\partial M_{SS,y}(u_l, t_q)}{\partial u_{jl}} \end{aligned} \quad (26)$$

where $p, q = 1, \dots, m$, $j = 1, \dots, 4$, $l = 1, \dots, n$. In point (u_l, t_k) the values of the steady-state magnetization and the corresponding partial derivatives are computed from (24), and

$$A(u_l, t_k) \frac{\partial M_{SS,x}(u_l, t_k)}{\partial u_{jl}} = \frac{\partial b(u_l, t_k)}{\partial u_{jl}} - \frac{\partial A(u_l, t_k)}{\partial (u_l)_j} M_{SS,x}(u_l, t_k), \quad (27)$$

respectively, where $k = 1, \dots, m$, $M_{SS,x}(u_l, t_k) = (M_{SS}(u_l, t_k))_1$, and $M_{SS,y}(u_l, t_k) = (M_{SS}(u_l, t_k))_2$. Since $A(u_l, t_k)$ is nonsingular (see page 8), we can always solve for $M_{SS}(u_l, t_k)$ and $\partial M_{SS}(u_l, t_k) / \partial u_{jl}$ in (24) and (27). These values are then used for computing (25) and (26) that are explored further in linearization.

Let $h_l = (h_{1l}, h_{2l}, h_{3l}, h_{4l}, h_{5l}) \in \mathbb{R}^5$ denote the displacement in u_l for $l = 1, \dots, n$. Since in problem (P) parameters T_l are fixed, it follows that $h_{5l} = 0$ for all l . Now, from (25)–(27) we derive the first order approximation of the nonlinear constraint (23), *i.e.*

$$\tilde{X}(u_1 + h_1, \dots, u_n + h_n, t_1, \dots, t_m)_{p,q} = X_{pq} + \sum_{l=1}^n h_l^T \nabla_{u_l} X_{pq}. \quad (28)$$

Finally, by using (28) we derive the following problem that is obtained by linearizing the nonlinear, non-SOCO, non-SDO constraints of (P) around the point (u_1, \dots, u_n) , where T_l are fixed, and with respect to the trust-region radius Δ :

$$\begin{aligned} \max \quad & \lambda \\ \text{s.t.} \quad & X_{\text{lin}} - \lambda I \succeq 0 \\ & (X_{\text{lin}})_{pq} = \tilde{X}(u_1 + h_1, \dots, u_n + h_n, t_1, \dots, t_m)_{pq}, \quad \forall p, q \quad (\text{P}_{\text{lin}}) \\ & (u_{1l} + h_{1l})^2 + (u_{2l} + h_{2l})^2 \leq 1, \quad \forall l \\ & (u_{3l} + h_{3l})^2 + (u_{4l} + h_{4l})^2 \leq 1, \quad \forall l \\ & \sum_{l=1}^n \sum_{j=1}^4 h_{jl}^2 \leq \Delta^2, \end{aligned}$$

where $p, q = 1, \dots, m$, $l = 1, \dots, n$, and $\tilde{X}(u_1 + h_1, \dots, u_n + h_n, t_1, \dots, t_m)_{p,q}$ is computed as in (28). Note that the trust-region constraint is a second-order cone constraint, and therefore the optimization problem (P_{lin}) is a *linear* mixed SDO-SOCO problem. Observe that we do not need to linearize the constraints (24) for S , because the new S can be effectively computed as a function of the variables u_l , i.e. reconstructed from the optimal solution of (P_{lin}) , that by the definition of $X = S^T S$ naturally imply positive semidefiniteness of the new matrix X . If S were *only* defined by the constraints, we would have to linearize those constraints and project the solution of the linear subproblem onto the constraint manifold. Direct computation reduces the size of the linear subproblem and eliminates the projection step, see Section 6 for further details.

Problem (P_{lin}) is the *trust-region subproblem* in the algorithm that is described in the following subsection. We use the optimization software SeDuMi (Sturm, 1999) for solving these subproblems.

5.2. THE ALGORITHM

Here we describe our trust-region based algorithm that solves (P), (see page 15). Problem (P) is a nonconvex problem and the algorithm converges to a local optimum. We ensure that the point computed by the algorithm is at least a local optimum of (P), by numerically verifying

the Karush-Kuhn-Tucker conditions at that point. The algorithm is iterative and maintains at each iteration a feasible current point and the solution of the linearized model in it. It also maintains the smallest eigenvalue of $S^T S$, where S is computed at the current point.

Grid search to fix T_l and the Starting Point. Although we can start the algorithm from any feasible point of (NL-SDO), (see page 14), a grid search is developed for finding a “good” initial point and to find good T_l values, that are fixed for the rest of the algorithm. Note that in model (NL-SDO), $T_l, l = 1, \dots, n$ are variables. We form a grid with respect to $(\alpha_l, f_l, T_l), l = 1, \dots, n$, and compute S at each point of that grid. The “best” point obtained by the grid search, is the one with the property that the smallest eigenvalue of $S^T S$ is the largest, among smallest eigenvalues of the matrices $S^T S$ at different grid points. The values for T_l from the “best” grid point are now fixed parameters in the algorithm for solving (NL-SDO). The values for $(\alpha_l, f_l), l = 1, \dots, n$ at the “best” grid point, are the starting values of the algorithm. Our numerical experiments justify this choice of starting point (see Section 6).

Note that the quality of the “best” initial point in the described way depends on the density of the grid. Namely, the denser the grid, the “better” starting point is computed. However, the cost of the grid search increases exponentially with the number of the experiments n in the model and/or the grid density. In Section 6, where our numerical results are presented, we discuss the trade off between the time consumed in the grid search and the number of the experiments in the model.

We denote now the initial point obtained from the grid search, by $(u_1^0, \dots, u_n^0, t_1, \dots, t_m)$, and the smallest eigenvalue of $(S^0)^T S^0$ by λ^0 . Note that the input values T_l^0 are fixed during the optimization algorithm. We specify the *initial trust-region radius*, e.g., $\Delta^0 := 1$.

General step. To describe a general step of the algorithm, we assume to have a *current feasible point* $(u_1, \dots, u_n, t_1, \dots, t_m)$. Let the value of the objective function in the current point be $\bar{\lambda}$, which is the smallest eigenvalue of $S^T S$. The minimization of (P_{lin}) , (see page 17) around the current point with respect to the corresponding trust-region radius Δ gives a *new candidate point* $(u_1 + h_1^*, \dots, u_n + h_n^*)$, and the corresponding value of the objective function λ^* . Note that the new candidate point will satisfy (24) due to the construction of the algorithm, i.e. by solving (24) for $M_{SS}(u_l, t_k)$ at each iterate. Therefore the feasibility of each iterate is ensured. More precisely, the steps in the algorithm are

$$h^* \rightarrow u + h^* \rightarrow (\text{from(24)}) \quad M(u + h^*, t) \rightarrow (\text{from } (P_{\text{lin}})) \quad \lambda^*.$$

We denote by $\hat{\lambda}^*$ the smallest eigenvalue of $S^T S$, where S is computed at the new candidate point. Then we compute the *trust-region ratio* ϱ based on the information of the new candidate point, *i.e.*

$$\varrho = \frac{\bar{\lambda} - \hat{\lambda}^*}{\bar{\lambda} - \lambda^*}. \quad (29)$$

Note that the closer the value of ϱ is to one, the better approximation by the linearized model is realized. If a sufficient reduction at the objective function is obtained at the candidate point, then that point is accepted as the next iterate and the trust-region radius is expanded or kept the same, as specified by (30). A negative or very small ϱ indicates a poor approximation, and therefore the point is rejected and the trust-region radius is reduced. More precisely, we update the *trust-region radius* Δ , (see *e.g.*, (Berstekas, 1995; Conn et al., 2000)), in the following way:

$$\Delta^* = \begin{cases} c_1 \Delta, & \text{if } \varrho < r_1 \\ c_2 \Delta, & \text{if } \varrho > r_2 \\ \Delta, & \text{otherwise.} \end{cases} \quad (30)$$

In our computations we set $c_1 = 0.25$, $c_2 = 2$ and $r_1 = 0.2$, $r_2 = 0.95$. The choice of these parameters is made after extensive testing and benchmarking.

Stopping criteria. Several conditions are used for terminating the algorithm. When the trust region becomes very small (*e.g.*, $\leq 10^{-8}$), or when after a pre-specified number of iterations there is no significant improvement in the objective, we stop the algorithm. We also stop the algorithm if the total number of iterations reaches 25.

Finally, we present our sequential SDO-trust-region algorithm for solving problem (P), by the pseudo code in Figure 5.2.

6. Numerical Results

In this section, we present the results of numerical tests of our optimal pulse-sequence design algorithm, for the model (NL-SDO), (see page 14), as computed by our Algorithm I. We have implemented the algorithm in MATLAB for three tissues $m = 3$, and small numbers of experiments $n \leq 6$. Performance was measured on a PC (Pentium 4, processor 2.66 GHz).

In all our numerical computations, we use tissue parameters representative of fat, blood and muscle, as given in Table I. We selected these tissues because they are familiar to non-experts, and allow us to illustrate the utility of quantification at higher resolution in a very simple numerical model.

Algorithm I

Input:tissue parameters: $t_k = (\tau_{1k}, \tau_{2k}, \kappa_k)$, $k = 1, \dots, m$;repetitive time parameters: T_l^0 , $l = 1, \dots, n$;initial design variables: (α_l^0, f_l^0) , $l = 1, \dots, n$;compute: $u_l^0 = (\cos(\alpha_l^0), \sin(\alpha_l^0), \cos(f_l^0), \sin(f_l^0), T_l^0)$, $l = 1, \dots, n$;initial trust-region radius: $\Delta^0 = 1$;input parameters: $r_1 = 0.25, r_2 = 2, c_1 = 0.2, c_2 = 0.95$;**begin** $\Delta \leftarrow \Delta^0$; $u_l \leftarrow u_l^0$, $l = 1, \dots, n$;**while** one of the stopping criteria is satisfied (see page 19) solve (P_{lin}) , (see page 17) \rightsquigarrow optimal value λ^*, h_l^* , $l = 1, \dots, m$; new candidate point: $u_l^* = u_l + h_l^*$, $l = 1, \dots, n$; compute ρ from (29); **if** $\rho < r_1$ reduce Δ ; **else** update Δ according to (30) $u_l \leftarrow u_l^*$, $l = 1, \dots, n$; **end****end**return (α_l, f_l) , $l = 1, \dots, n$ and $\hat{\lambda}^*$;**end***Figure 2.* Sequential, trust-region, second-order conic algorithm.

The number of design variables in the bSSFP model for tissue segmentation with n experiments is $3n$, and these variables are (α_l, f_l, T_l) , $l = 1, \dots, n$. Let $\alpha := (\alpha_1, \dots, \alpha_n)^T$, $f := (f_1, \dots, f_n)^T$, and $T := (T_1, \dots, T_n)^T$. Note that in (NL-SDO) we introduce the substitutions $u_{1l} = \cos(\alpha_l)$, $u_{2l} = \sin(\alpha_l)$, $u_{3l} = \cos(f_l)$, and $u_{4l} = \sin(f_l)$. Algorithm I is initialized by the input values (α^0, f^0, T^0) that are obtained by the grid search (see §5.2, page 18), then we compute $(u_{1l}, u_{2l}, u_{3l}, u_{4l})$,

Table I. Tissue parameters for fat, blood and muscle.

	fat	blood	muscle
τ_1	250	1200	900
τ_2	60	50	50
κ	0.6π	0	-0.1π

Table II. Design variables for $n = 3$ obtained from grid search and computed by Algorithm I.

	α_l^0	f_l^0	T_l^0	α_l^*	f_l^*
1	20	180	3.0	17.71	190.48
2	20	90	4.5	15.00	97.92
3	30	180	5.0	23.63	180.51

$l = 1, \dots, n$, and finally Algorithm I returns the optimal values (α^*, f^*) . Our computational results for $n = 3, 4, 6$ are given in the sequel.

- *Case $n = 3$.* First, we give our numerical results for the model in which the number of tissues is equal to the number of experiments, *i.e.* when $m = n = 3$. In our experiments, the grid on which we search for the “best” initial point is: $\alpha_l = \{10, 20, \dots, 50\}$, $f_l = \{0, 30, \dots, 180\}$, $l = 1, 2, 3$, and $T_1 = \{3, 3.1, \dots, 3.9\}$, $T_2 = \{4, 4.1, \dots, 4.9\}$, $T_3 = \{5, 5.1, \dots, 5.9\}$. The search for the “best” point on this grid requires 7934 sec ≈ 2.2 hours.

In Table II, we give the values (α^0, f^0, T^0) that are obtained from grid search, and the optimal values (α^*, f^*) that are computed by Algorithm I. The smallest eigenvalue of $S^T S$, where S is computed for (α^0, f^0, T^0) is $\lambda^0 = 0.02448$, and the value of the objective function in (NL-SDO) for (α^*, f^*, T^*) is $\lambda^* = 0.02834$ (see also Table V). For obtaining that value, Algorithm I needs 2.21 seconds.

Convergence area. In order to find the radius of convergence of our algorithm around the “best” computed initial point, we did the following experiment. We fix $T^0 = (3, 4.5, 5)$ and form the grid around (α^0, f^0) . The computational results show that *for any* initial point taken on the grid (α^+, f^+, T^0) , where $\alpha_l^+ = \alpha_l^0 \pm d_l^\alpha$, $f_l^+ = f_l^0 \pm d_l^f$, $d_l^\alpha, d_l^f = 1, \dots, 7$, $l = 1, 2, 3$, Algorithm I converges to the same (α^*, f^*) . Our results also show that if $d_l^\alpha, d_l^f = 1, \dots, 9$, then Algorithm I con-

Table III. Design variables for $n = 4$ obtained from grid search and computed by Algorithm I.

	α_l^0	f_l^0	T_l^0	α_l^*	f_l^*
1	20	180	3.15	20.86	184.05
2	20	180	4.20	20.72	190.59
3	20	60	4.65	13.45	63.60
4	20	150	5.25	20.62	142.22

verges from 98% of points taken on that grid towards (α^*, f^*) , and if $d_l^\alpha, d_l^f = 1, \dots, 12$, then it converges to (α^*, f^*) for 93 % of the initial points on the grid. Hence, for the fixed T^0 our nonconvex problem (NL-SDO) is unimodal on a certain region around our “best” initial point. However, since our problem (NL-SDO) is nonconvex we can find points outside of the described region of convergence that converge to other local optimum. For instance, if we take the input values $\alpha^0 = (30, 10, 30)^T$, $f^0 = (180, 90, 150)^T$, then $\lambda^0 = 0.01796$ and Algorithm I converges to the local optimum $\lambda^* = 0.02839$, where $\alpha = (20.0621, 14.6718, 23.8554)^T$, $f = (203.1231, 97.4414, 179.4882)^T$. Observe that this solution is slightly better than the one generated by Algorithm I.

- *Case $n = 4$.* We give now the numerical results for the optimization problem (NL-SDO) in which $n = 4$. The grid was chosen with respect to the design variables as $\alpha_l = \{10, 20, \dots, 50\}$, $f_l = \{0, 30, \dots, 180\}$, $l = 1, \dots, 4$, and $T_1 = \{3, 3.15, \dots, 3.70\}$, $T_2 = \{3.75, 3.90, \dots, 4.45\}$, $T_3 = \{4.50, 4.65, \dots, 5.20\}$, $T_4 = \{5.25, 5.40, \dots, 5.95\}$. Our code for heuristic needs 55.2611 hours (!) to find the “best” initial values (α^0, f^0, T^0) , but Algorithm I needed only 3.10 seconds for computing (α^*, f^*, T^*) (see Table III) from the initial point. Note the dramatic increase in CPU time for the grid search computations, with respect to the number of experiments in the model. However, the CPU time for performing Algorithm I does not increase significantly with the increase of n in the model. The size of the cone program depends only on the number of tissues m , and the small increase in the CPU time is due to the cost for forming (28). In Table V are given values for λ^0 and λ^* .

Convergence area. The numerical results indicate that for fixed $T^0 = (3.15, 4.20, 4.65, 5.25)^T$, the radius of convergence concerning (α^0, f^0) towards the local optimum is somewhere between 4 and 6. If we consider the points in the 8 dimensional “cube” of radius 6 around (α^0, f^0) then

Table IV. Design variables for $n = 6$ obtained from heuristic and computed by Algorithm I.

	α_l^0	f_l^0	T_l^0	α_l^*	f_l^*
1	30	60	3.00	11.76	59.22
2	30	180	3.50	24.77	201.23
3	30	180	4.30	21.86	223.30
4	30	60	4.50	16.29	83.72
5	30	180	5.00	20.82	226.79
6	50	180	5.95	25.91	172.03

Table V. λ^0 and λ^* values for $n = 3, 4, 6$, obtained by grid search and Algorithm I, respectively.

n	3	4	6
λ^0	0.02648	0.03765	0.04682
λ^*	0.02834	0.04108	0.06193

87 % points inside that cube converge towards the same local optimum, and for the “cube” of radius 7 only 81 % of the points.

- *Case $n = 6$.* Finally, we provide some computational results for the optimization problem (NL-SDO), where $m = 3$ and $n = 6$. We search for the “best” starting point on the grid: $\alpha_l = \{10, 30, 50\}$, $f_l = \{0, 60, 120, 180\}$, $l = 1, \dots, 6$, and $T_1 = \{3, 3.15, \dots, 3.45\}$, $T_2 = \{3.50, 3.65, \dots, 3.95\}$, $T_3 = \{4.00, 4.15, \dots, 4.45\}$, $T_4 = \{4.50, 4.65, \dots, 4.95\}$, $T_5 = \{5.00, 5.15, \dots, 5.45\}$ and $T_6 = \{5.50, 5.65, \dots, 6.00\}$. The program for the grid search was running for 1685 hours, and Algorithm I needed only 5.61 seconds for computing the local optimum given in Table V.

Since the problem of minimizing the errors in the reconstructed tissue densities is equivalent to solving the maximization problem (NL-SDO), and a value of the objective function in (NL-SDO) increases almost linearly with the number of experiments (see Table V), we conclude that adding multiple images considerably minimizes the noise in the calculated tissue concentrations.

7. Numerical Simulation and Validation

In this section we present results of a simplified numerical simulation of a carotid-artery cross-section, with three aims:

1. to validate the solver against the signal generation model,
2. to illustrate the kind of clinical applications, and the relative performance versus MR spectroscopy, and
3. to measure the advantage of removing restrictions on phase relationships.

Intraluminal lipids, *i.e.* fat deposits inside arteries, are an indication of arterial disease. For example, some strokes are caused by the rupturing of large lipid deposits in the carotid artery, which carries blood from the heart to the brain. Figure 3 depicts a cross-section of an idealized carotid artery. Large arteries are essentially composed of a tube of muscle (which contracts in sympathy with the heart to increase blood flow), containing (flowing) blood and surrounded by fat (and other) tissues. For simplicity, we can ignore the surrounding tissue which is not fat. In Figure 3, we also show a layer of lipid deposited uniformly on the inside of the vessel wall. The exact distribution of lipids is not important, since we have not chosen a resolution sufficient to determine the exact distribution. Superimposed on the vessel structure is a 1mm grid representing the voxel size for bSSFP tissue quantification, whereas the entire figure represents a single voxel for MR spectroscopy. So no matter how we line up this 1cm² voxel, we cannot separate the intraluminal lipids from the extraluminal lipids. Furthermore, a spectroscopic exam would take on the order of minutes, whereas the individual bSSFP images at this resolution would only require 256x6ms, or about 1.5 seconds, plus some overhead for instrument setup.

In order to demonstrate the advantage of generalizing the Dixon method by allowing arbitrary phase relationships between tissues (as expressed in the S matrix), we compare our 3-experiment with a 3-experiment result using conventional phase relationships:

$$S^{\mathbb{C}} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & 1 \\ 1 & -1 & -1 \end{pmatrix}. \quad (31)$$

To be comparable, we use bSSFP sequences for the conventional Dixon case, with pulse-sequence design variables in the same ranges. We used a grid search to find design values which result in an $S^{\mathbb{C}}$ of the form (31), after an overall complex scaling, and we allow a variation of 5 per




Figure 3. The cross-section of idealized carotid artery

cent in each entry. We allowed the variation in each entry to make the grid search reasonable. To not disadvantage the conventional case, we also use the Moore-Penrose Pseudo-Inverse for this case. The resulting three sets of pulse-sequence design variables are listed in Table VI. Note that the resulting sequences use very small flip angles, which would be difficult to work with in practice. We didn't want to add a constraint to the model with fixed phase relationships to prevent this, because that would adversely effect it in comparison with our generalized model.

The advantage of optimal sequence design over the design based on the conventional Dixon phase relationships is clear in a numerical simulation of the carotid artery simulation described by Figure 3. Figure 4 shows gray-scale images reconstructed by applying the Moore-Penrose Pseudo-Inverse to simulated bSSFP images, for different sequence designs, but with identical noise, to render noise effects comparable. The final row shows the results of the conventional Dixon phase relationships. The middle rows are the result of our algorithm, using only the grid search, and the grid search followed by the iterative method. The first row is the ideal case with zero noise. Note that without noise, all reconstruction methods and sequence designs produce the same results.

Table VII shows the relative performance of optimal pulse-sequence designs using our method for the cases of 3, 4 and 6 experiments. In




Figure 4. Simulation of tissue quantification: each column represents a different tissue, as labeled; ideal (zero noise) tissue densities are shown in row a), densities reconstructed from data collected with the optimal pulse-sequence design are in row b), row c) shows the densities from the pulse sequence found by grid search, and row d) illustrates the tissue quantification based on Dixon method. All values are displayed using the same gray scale.

Table VI. pulse-sequence design variables for Dixon method, where α and f are in degrees and T is in ms

α	f	T
4.0	45.0	1.2
1.0	21.0	4.9
1.0	86.5	5.0

Table VII. Numerical results for tissue quantification based on 1000 experiments, measured in mm^2 .

	3-Ex	4-Ex	6-Ex	Zero-noise
Actual intraluminal lipid	5.4000	5.4000	5.4000	5.4000
Estimated intraluminal lipid	5.4028	5.4016	5.3987	5.4000
Standard deviation of the estimate	0.0735	0.0650	0.0467	0.0000
Expected std. deviation from averaging	0.0735	0.0637	0.0520	0.0000
Expected std. deviation from objective	0.0735	0.0610	0.0497	0.0000

all cases, we use the sum of the fat components of all of the voxels marked with a Σ in Figure 3 to estimate the total intraluminal lipids in this cross-section. The estimated intraluminal lipid is close to the actual value in the simulations, for all numbers of experiments, but more importantly, the standard deviation of the lipid quantification decreases as a function of the number of experiments, and it does so faster than one would obtain by simply averaging the three experiments to reduce noise in the source images. The standard deviations in the simulation are also consistent with the deviations predicted by the objective function.

8. Conclusions and Future Work

We have shown that Dixon methods can be generalized by removing assumptions about the best complex phase relationships between tissue types across multiple images. The resulting model represents what clinicians and researchers want most: the reduction of error in the calculated tissue quantities. Although the resulting problem has semidefinite, and highly-nonlinear constraints, we have developed a method of solving it combining grid search and sequential SDO. Since we do not obtain conventional Dixon phase relationships for the generated signals, we can conclude that conventional Dixon methods are not optimal. More significantly, we have provided computational evidence that our generalized Dixon method produces more accurate results than signal averaging (see Table VII).

Identifying tissue distribution in-vivo has many applications in diagnostic imaging, treatment monitoring, and biological research. In many cases, information about tissue composition may be known from MR spectroscopy, but acquisition times are too long to apply spectroscopy to clinical diagnosis. Since MR does not involve ionizing radiation,

it is attractive for population-scale screening, and after developing appropriate sequences using our optimization method, it may be cost-effective to do so, *e.g.*, for quantifying liver fat. In a similar way, the reduced imaging time requirements for bSSFP-based quantification will make longitudinal studies feasible, *e.g.*, brain development studies in infants based on white and gray matter quantification. We have clinical partners for one such application and are planning to conduct research in this direction.

We have formulated the parameter selection problem in terms of bSSFP imaging. Since different sequence types are sensitive to different tissue properties, one would expect better quantification by looking at mixed imaging, using some bSSFP images and some other imaging types, *e.g.*, fast spin echo, spoiled sequences, gradient-reversed fast imaging with steady-state precession, inversion recovery sequences, see (Haacke et al., 1999). There is no fundamental obstacle to formulating mixed imaging parameter selection, but the software complexity will increase by having a nonconvex mixed continuous and discrete optimization problem.

In most imaging environments, technicians are free to adjust pulse-sequence design variables to suit the subject (including the subject's ability to remain still). Although our sequential trust-region SDO-SOCO algorithm is very fast, as expected from a modern interior point code solving a problem of this size, the initial grid search, whose cost grows exponentially with the number of experiments, would probably be too slow to use in this setting. We are exploring several heuristics to find good starting points quickly, even if we are less likely to find the global optimum. We are optimistic about such heuristics. We will also investigate including the variables T_i in the subproblems, as a possible way of increasing the basins of convergence (and thus reducing the cardinality of the grid). If these methods are not sufficient, we may choose to pre-compute good starting points for ranges of parameters to hide the cost of the grid search from the end user.

The most challenging area for future investigation is the adaptation of our method to situations in which parameters change between patients and within patients. Not all tissue parameters are accurately known, and some are known to vary as a function of pH and hydration. For our methods to be applicable in these situations, we will have to develop integrated methods of estimating these parameters. As a special case of this problem, we will need to incorporate field inhomogeneity into the tissue quantification process, and design sequence parameters which will guarantee a minimum contrast to noise ratio for a range of field inhomogeneity. The three-point Dixon technique (Glover, 1991), solves this problem by collecting two in-phase images at different echo

times. These images can be compared to generate a field map, which can be used to correct for inhomogeneity effects in the source images, before the linear transformation is applied. A similar approach would work to correct for inhomogeneity in our situation, but we would have to recalculate the linear transformation for each pixel. We would also have to calibrate RF penetration.

Although these issues are challenging, we believe that they are all solvable by adapting known techniques. We will also pursue extensions to the existing model to integrate estimation of these parameters into a larger inverse problem. Our success in modeling and solving the design problem for optimal pulse-sequence design variables, gives us confidence to tackle these larger problems.

9. Acknowledgments

This research work was supported by NSERC Discovery grants for all authors, the Canada Research Chair program and a MITACS grant for the third author, and a CFI grant for computing equipment.

We thank anonymous referees for their comments and suggestions to improve an earlier version of the paper.

References

- A. Albert, "Conditions for positive and nonnegative definiteness in terms of pseudoinverses," *SIAM Journal on Applied Math.*, vol. 17, pp. 434–440, 1969.
- A. Ben Tal and A. Nemirovski, *Lectures on Modern Convex Optimization. Analysis, Algorithms and Engineering Applications*, Volume 1 of MPS/ SIAM Series on Optimization, SIAM, Philadelphia, U.S.A., 2001.
- D.P. Bertsekas, *Nonlinear Programming*. Athena Scientific, U.S.A., 1995.
- H.Y. Carr, "Steady-state free precession in nuclear magnetic resonance," *Physics Review*, vol. 112, pp. 1693–1701, 1958
- A.R. Conn, N.I.M. Goud, and P.L. Toint, *Trust Region Methods*. MPS–SIAM Series on Optimization, 2000.
- W.T. Dixon, "Simple proton spectroscopic imaging," *Radiology*, vol. 153, pp. 189–194, 1984.
- R. Freeman and H.D.W. Hill, "Phase and Intensity Anomalities in Fourier Transform NMR," *Journal of Magnetic Resonance*, vol. 4, pp. 366–383, 1971.
- R.W. Freund and F. Jarre, "A sensitivity analysis and a convergence result for a sequential semidefinite programming method", Numerical Analysis Manuscript No. 03-4-09, Bell Laboratories, Murray Hill, New Jersey, 2003.
- G. Glover, "Multipoint Dixon technique for water and fat proton and susceptibility imaging," *Journal of Magnetic Resonance Imaging*, vol. 1, pp. 521–530, 1991.
- E.M. Haacke, R.W. Brown, M.R. Thompson, and R. Venkatesan, *Magnetic Resonance Imaging, Physical Principles and Sequence Design*. Wiley-Liss, 1999.

- W. Hanicke and H.U. Vogel, "An analytical solution for the SSFP signal in MRI," *Magnetic Resonance in Medicine*, vol. 49, pp. 771–775, 2003.
- B.A. Hargreaves, S.S. Vasanawala, J.M. Pauly, and D.G. Nishimura, "Characterization and reduction of the transient response in steady-state MR imaging," *Magnetic Resonance in Medicine*, vol. 46, pp. 149–158, 2001.
- B.A. Hargreaves, S.S. Vasanawala, K.S. Nayak, J.H. Brittain, B.S. Hu, and D.G. Nishimura, "Fat-suppressed steady-state free precession imaging using phase detection," *Magnetic Resonance in Medicine*, vol. 50, pp. 210–213, 2003.
- R.A. Horn and C.R. Johnson, *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1985.
- T.Y. Huang, H.W. Chung, F.N. Wang, C.W. Ko, and C.Y. Chen, "Fat and water separation in balanced steady-state free precession using the Dixon method," *Magnetic Resonance in Medicine*, vol. 51, pp. 243–247, 2004.
- S.V.B. Jardim and M.A.T. Figueiredo, "Automatic contour estimation in fetal ultrasound images," *IEEE International Conference on Image Processing-ICIP, Barcelona, Spain*, (2), pp. 1065–1068, 2003.
- E.T. Jaynes, "Matrix Treatment of Nuclear Induction," *Phys. Rev.*, vol. 98(4), pp. 1099–1105, 1955.
- S. Kruk and H. Wolkowicz, "SQ²P, Sequential Quadratic Constrained Quadratic Programming", in *Advances in Nonlinear Programming*, Kluwer Academic, pp. 177–204, 1998.
- Z.P. Liang and P.C. Lauterbur, *Principles of Magnetic Resonance Imaging*, IEEE Press Series in Biomedical Engineering, 1999.
- K.V. Mardia, J.T. Kent, and J.M. Bibby, *Multivariate Analysis*, Academic Press, 1979.
- A. Oppelt, R. Graumann, H. Barfuß, H. Fischer, W. Hartl, W. Schajor, "FISP—a new fast MRI sequence," *Electromedica (English Edition)*, vol. 54, pp. 15–18, 1986.
- S.B. Reeder, Z.F. Wen, H.Z. Yu, A.R. Pineda, G.E. Gold, M. Markl, and N.J. Pelc, "Multicoil Dixon chemical species separation with an iterative least-squares estimation method," *Magnetic Resonance in Medicine*, vol. 51, pp. 35–45, 2004.
- F.J. Rybicki, R.V. Mulkern, R.L. Robertson, C.D. Robson, T. Chung, and J. Ma, "Fast three-point Dixon MR imaging of the retrobulbar space with low-resolution images for phase correction: comparison with fast spin-echo inversion recovery imaging," *American Journal of Neuroradiology*, vol. 22, pp. 179–1802, 2001.
- N. Salibi and M.A. Brown, *Clinical MR Spectroscopy: First Principles*. Wiley-Liss, New York, 1998.
- K. Scheffler, "A pictorial description of steady states in fast magnetic resonance imaging," *Concept. Magnetic Res.*, vol. 11, pp. 291–304, 1999.
- K. Scheffler, "On the transient phase of balanced SSFP sequences," *Magnetic Resonance in Medicine*, vol. 49, pp. 781–783, 2003.
- J.F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11–12, pp. 625–653, 1999.
- J.F. Sturm, "Implementation of interior point methods for mixed semidefinite and second order cone optimization problems," *Optimization Methods and Software*, vol. 17(6), pp. 1105–1154, 2002.
- B.P. Sutton, D.C. Noll and J.A. Fessler, "Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities," *IEEE Trans. Medical Imaging*, vol. 22(2), pp. 178–188, 2003.
- S.S. Vasanawala, J.M. Pauly, and D.G. Nishimura, "Linear combination steady-state free precession MRI," *Magnetic Resonance in Medicine*, vol. 43, pp. 82–90, 2000.

- Y. Zur, S. Stokar and P. Bendel, "An analysis of fast imaging sequences with steady-state transverse magnetization refocusing," *Magnetic Resonance in Medicine*, vol. 6, pp. 175–193, 1988.