Published in final edited form as:

IEEE Trans Med Imaging. 2010 May ; 29(5): 1156–1172. doi:10.1109/TMI.2010.2041787.

Basic Properties of SS-PARSE Parameter Estimates

Donald B. Twieg[Member, IEEE] and

Department of Biomedical Engineering, University of Alabama at Birmingham, Birmingham, AL 35294 USA

Stanley J. Reeves[Senior Member, IEEE]

Electrical and Computer Engineering Department, Auburn University, Auburn, AL 36849 USA

Donald B. Twieg: twieg@uab.edu

Abstract

Single shot parameter assessment by retrieval from signal encoding (SS-PARSE) is a recently introduced method to obtain quantitative parameter maps from a single-shot (typically 65 ms) magnetic resonance imaging (MRI) signal. Because it explicitly models local magnetization decay and phase evolution occurring during the signal 1) it can provide quantitative estimates of local transverse magnetization magnitude and phase, frequency, and relaxation rate and 2) it is free of geometric distortion or blurring due to field nonuniformities within the tissues. These properties promise to be advantageous in functional brain MRI (fMRI) and other dynamic imaging applications. In this paper, the basic phenomena underlying the performance of SS-PARSE in practice are discussed. Basic sources of bias errors in the parameter estimates are discussed, and performance of the method is characterized in terms of parameter estimates from simulation, experimental phantoms, and *in vivo* studies. Characteristics of the sum-of-square-error cost function and the iterative search algorithm are discussed, and their relative roles in determining estimation accuracy are described. Practical guidelines for use of the method are presented and discussed. *In vivo* parameter maps are also presented.

Index Terms

Error analysis; estimation; image generation; image reconstruction; magnetic resonance imaging (MRI)

I. Introduction

Most current methods of magnetic resonance imaging (MRI) reconstruction interpret raw MRI signal values as samples of the Fourier transform of the object. This is computationally convenient, but it neglects relaxation and off-resonance evolution in phase angle, both of which can occur to a significant extent during a typical MRI imaging signal. A commonly observed consequence of this “static object” assumption is geometric distortion or blurring errors in the presence of field inhomogeneities. The recently introduced parameter assessment by recovery from signal encoding (PARSE) methods [1] employ a more accurate model of the imaging signal which includes local relaxation and phase evolution as unknowns. This allows explicit estimation of local frequency and relaxation rates and at the same time eliminates off-resonance geometric defects. In this paper, we describe

Correspondence to: Donald B. Twieg, twieg@uab.edu.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

fundamental characteristics of this inverse-problem approach to MRI, formulated as a parameter estimation problem.

A single-shot version of PARSE (SS-PARSE) [1] has been developed to apply this approach to the single-shot imaging case, in which there is often a motivation to obtain quantitative maps of net transverse relaxation rate R_2^* (the inverse of net relaxation time T_2^*), initial transverse magnetization $M_\perp$ and local resonant frequency f . In studies of brain function by BOLD functional MRI (fMRI) [2], this promises several advantages over the more commonly used T_2^* -weighted single-shot methods, including quantitative mapping of relaxation rates and frequency, and freedom from the geometric errors commonly seen in conventional single-shot MRI when tissue susceptibility nonuniformities are present.

A number of other methods have been proposed for estimation of $M_\perp$, local frequency and R_2^* decay from single-shot data [3]–[7]. The simplest of these methods use a series of separate EPI or spiral readouts following a single excitation pulse, reconstructed in the conventional way to produce a series of images with different echo times TE_j , $j = 1, 2, \dots, J$. Temporal fits are then applied to pixel values over this series of images to obtain estimates of $M_\perp$ and R_2^* [3]–[5]. In these methods, signal decay and phase evolution due to frequency offset are not modeled *within* the signal for each image readout. Instead, the j th image frame is reconstructed as if all its data were acquired instantaneously at TE_j . Other methods model the data more accurately, including local decay and/or phase evolution occurring within the signal from each of the serial EPI or spiral signal readouts [6], [7]. Olafsson *et al.* [7] describe a method which jointly estimates R_2^* and frequency. It uses several simplifying assumptions to convert the problem to a series of linearized optimizations, which can be rapidly computed. Assuming constant $M_\perp$ over the series of frames in an fMRI series, it makes a linear approximation to changes in R_2^* and frequency based on a reference map. SS-PARSE differs from these methods in two respects. First, it uses data acquired along a trajectory NOT designed for a series of Fourier image reconstructions—instead, it uses a trajectory theoretically better suited to the estimation of local R_2^* and frequency, as discussed below. Second, it attempts to solve the general inverse problem [(8) in Section II] without simplifying assumptions which reduce computational burden but also require prior information from previous scans, and may limit the range of applications because of additional assumptions. For example, because it does not assume that $M_\perp$ or frequency is constant over a dynamic series of single-shot frames, SS-PARSE has the flexibility to monitor simultaneously changing $M_\perp$, R_2^* , and frequency during the series. This may be advantageous in fMRI, in dynamic susceptibility contrast (DSC) studies following injection of paramagnetic agents [8], [9], and in hybrid studies for simultaneous measurement of BOLD functional effects and tissue blood perfusion by arterial spin labeling (ASL) [10], [11]. In these cases both R_2^* and $M_\perp$ may vary from frame to frame, and their variations are of interest independently. In fMRI studies, SS-PARSE does not require an assumption of minimal motion between frames. The price SS-PARSE pays for solving the inverse problem independently for each time frame without time-saving assumptions is a longer computation time.

In addition to estimation of $M_\perp$, local frequency and R_2^* , PARSE methods can also be used to assess other dynamic changes in magnetization, such as changes deliberately imposed upon local magnetization phase or amplitude to confer sensitivity to local quantities such as velocity [12], [13], or, in principle, diffusion distributions. This paper describes some basic properties of SS-PARSE but does not directly address the estimation of these other parameters. However, the basic principles discussed here are expected to apply in estimation of other parameters. Similarly, SS-PARSE is a single-shot method, but it has a

straightforward extension to multiple-shot higher-resolution applications and to parallel multiple-receiver acquisition. Both cases supply additional data so that they can support estimation of greater numbers of parameters (realizable as greater spatial resolution or additional parameters characterizing the local signal). Multiple-shot and parallel acquisition methods, as well as spin-echo SS-PARSE (which estimates an additional local parameter, R_2) are currently under development.

The PARSE method estimates parameters from MRI data by matching a parametric signal model to the measured experimental signal $s(t)$. The best parameter estimates are defined as those which produce the signal values that best match the measured experimental signal in the least squares (LS) sense. Because the signal equation is nonlinear in the parameters, there is not a generally applicable closed-form solution for the LS optimal parameters, and in practice an iterative search algorithm is used to find the LS minimum.

Two factors limit estimation accuracy: 1) the accuracy of the LS solution for the parameter estimates and 2) the performance of the search algorithm in finding the LS solution. The first issue is addressed here by considering the shape and variability of the LS cost function and demonstrating experimental and simulated results. Obviously, the location of the LS cost function minimum determines the LS solution in parameter space, while the shape (the partial derivatives) of the LS cost function near the minimum determines the computational effort needed to approach the LS optimum to a given degree of accuracy. Also of interest are the presence or absence of secondary minima in the cost function. If the solution is to be found by an iterative search, conditions leading to a secondary minimum should be avoided.

Overall, the goal of this paper is to demonstrate the practical issues involved in applying SS-PARSE and to discuss some basic concepts potentially useful for understanding and improving the methodology and designing future applications.

Following a brief presentation of the theory of the method, we describe the properties of the least-square solution in terms of the cost function behavior. Basic sources of bias errors and methods for minimizing them are discussed. Experimental phantom data demonstrate the ability of SS-PARSE to provide undistorted parameter maps even in the presence of very large background gradients and large off-resonance bandwidths. Initial *in vivo* data from a 4.7T primate system are presented and discussed.

II. Theory

Here we examine the process of quantitative MRI mapping, and MRI in general, as a basic encoding/decoding process similar to a classical information theory problem.

A. Spatial Encoding and Decoding

In MRI, a radiofrequency excitation pulse is applied to nuclear magnetization in a strong magnetic field, producing transverse magnetization. After the excitation pulse, the local M induces in the detector a radiofrequency signal known as a free induction decay (FID), with decaying amplitude and precessing phase angle exactly reflecting the decaying magnitude and angular precession of the local $M_{\perp}(\mathbf{x}, t)$. We refer to this as the *local signal*, from within a small local region associated with an imaging volume element (voxel). Our goal is to measure the frequency and decay rate of the local signals at all voxels within a preselected slab or volume of tissue.

To accomplish this, we must decode the local signals from the net signal produced by the detector. This involves first encoding the location data into the net signal by imposing spatial modulating functions, and then decoding each of the local signals from the net signal.

The n th sample of the net MRI signal for a spatially continuous object with complex-valued transverse magnetization $M_{\perp}(\mathbf{x}, t) \equiv M_{x'}(\mathbf{x}, t) + iM_{y'}(\mathbf{x}, t)$ is the integral over space

$$s(n\Delta t) = \int_V C(\mathbf{x}) M_{\perp}(\mathbf{x}, n\Delta t) \psi(\mathbf{x}, n\Delta t) d\mathbf{x} \quad (1)$$

where signal samples are acquired at times $n\Delta t$, $n = 0, 1, \dots, N-1$, and $\psi(\mathbf{x}, t) = \exp(i2\pi \mathbf{k}(t) \cdot \mathbf{x})$ is the spatial modulating function produced by time-varying linear gradient fields $\mathbf{G}(t)$. $\mathbf{k}(t)$ is the k -trajectory representing the spatial-frequency sampling path [14]

$$\mathbf{k}(t) \equiv \gamma \int_0^t \mathbf{G}(t') dt' \quad (2)$$

where γ is the cyclic gyromagnetic ratio (4258 Hz per Gauss). $C(\mathbf{x})$ is the sensitivity map for the receiver channel.

For convenience, the signal equation is written as a sum over discrete volume elements (with the approximation that C_j and $\psi_{j,n}$ are constant over the voxel volumes)

$$s_n = \sum_j C_j M_{\perp,j,n} \psi_{j,n} \quad (3)$$

where $C_j = C(\mathbf{x}_j)$ and

$$\psi_{j,n} = \exp(i2\pi \mathbf{k}_n \cdot \mathbf{x}_j). \quad (4)$$

In (3), the continuous spatial distribution $M_{\perp}(\mathbf{x}, t)$ at time $n\Delta t$ is represented by a discrete spatial array defined by

$$M_{\perp,j,n} = \int_V M_{\perp}(\mathbf{x}, n\Delta t) u(\mathbf{x} - \mathbf{x}_j) d\mathbf{x} \quad (5)$$

where $u(\mathbf{x})$ is a voxel function (defined as a boxcar function, for example), and the $\mathbf{x}_j$, $j = 1, 2, \dots, J$, are the discrete spatial grid locations at which estimates are made.

The local magnetization is assumed to decay in magnitude with exponential rate R_2^* , and to evolve in phase at the constant frequency f_j so that the local signal at location j is

$$M_{\perp,j,n} = M_{\perp,j} \exp[-(R_2^* + i2\pi f_j)n\Delta t] \quad (6)$$

where $M_{\perp,j}$ is the complex magnetization at $t = 0$.

With (6) as the local signal, the net signal is the sum over all local signals, with the addition of an eddy current contribution

$$s_n = \sum_j \psi_{j,n} M_{\perp,j} \exp[-(R_2^* + i2\pi f_j)n\Delta t + i\Phi_n]. \quad (7)$$

Here Φ_n is a time-varying spatially-constant phase value representing the effective static field variations produced by eddy currents. (Φ_n is system- and sequence-dependent, but it

can be measured straightforwardly (see Section III), and is independent of the object being imaged, so it need be measured only once). Once Φ_n is measured and removed from the signal, the net signal becomes

$$s_n = \sum_j \psi_{j,n} M_{\perp j} \exp [-(R_{2j}^* + i2\pi f_j) n \Delta t]. \quad (8)$$

SS-PARSE simply estimates the unknown arrays $M_{\perp}$, R_{2j}^* , and f from the net signal, i.e., it solves the inverse problem (8).

It is worthwhile to consider the relationship of this inverse problem to Fourier-based MRI methods. If the $\psi_{j,n}$ functions are those of (4), and if their k -sampling locations fall on a rectilinear grid, then for an object with both R_{2j}^* and $f=0$ everywhere, then (8) becomes a linear equation representing the spatial discrete Fourier transform, in which case applying a simple 2-D or 3-D inverse FFT produces a complex image. (If sample locations in k -space are irregularly spaced, as in spiral acquisitions [15], [16], the samples may be interpolated to an equivalent rectilinear sample set before the FFT).

When this FT-based decoding scheme is applied to data from an object in which R_{2j}^* and f are not zero everywhere (i.e., an actual object), image errors occur. These are especially noticeable in single-shot imaging methods such as echo-planar imaging (EPI) [17], notably 1) the well-known geometric distortion or blurring due to off-resonance frequencies [18], [19], and 2) filtering effects due to differential T_2^* -weighting of spatial frequency components [20]. Despite these errors, the FT-based decoding scheme is attractive because it is very fast, capable of producing images in real time, as fast as they can be acquired.

In inverting the nonlinear signal (8) to produce estimates of R_{2j}^* and f as well as $M_{\perp}$ at $t=0$, SS-PARSE is much more computationally demanding. The method we have adopted searches for the set of parameter values which will [when substituted in (8)] produce a net signal best matching the observed noisy signal in the sense of minimizing the sum of square difference between the estimated and observed signals.

B. Least Squares Cost Function and Bias Errors

The observed signal has the form $z_n = s_n + e_n$, where s_n is the deterministic complex signal of (8) and e_n is complex measurement noise assumed to be zero-mean and independent Gaussian-distributed noise.

The parameter estimation process searches for the parameter values such that the model signal $\hat{s}(\mathbf{p})$ computed from the current parameter values best matches the actual observed signal. In the investigations reported here, a conjugate gradient search was used to iteratively adjust the 4J unknown parameter values (concatenated into the column vector $\mathbf{p}$) until the sum-square residual

$$Q(\mathbf{p}) = \sum_n |z_n - \hat{s}_n(\mathbf{p})|^2 \quad (9)$$

is minimized to a desired tolerance, yielding the least-square parameter estimate $\mathbf{p}_{LS}$.

The estimation error for parameter p_j is the difference between the actual value $p_{act,j}$ and the estimated value $p_{est,j}$. Bias errors correspond to a nonzero expected value of estimation error $p_{err,j}$, and random errors are characterized by the standard deviation or variance of $p_{err,j}$. To

detect activation-related changes in parameter R_2^* with optimal sensitivity, we should minimize the level of random estimation errors in R_2^* , because these form the noise background against which activation-related changes must be detected.

C. Relationship Between k-Space Sampling and the Spatial Evaluation Grid

In principle, parameter estimates can be computed for any set of locations within the field-of-view (FOV), but it is convenient to evaluate parameter maps on a square Cartesian grid of spatial locations for display in a conventional format. Because the parameter estimates are evaluated at a grid of spatial locations, we call the corresponding region of k -space the *evaluation k-band*. The evaluation k -band is the region in (k_x, k_y) surrounding the origin containing all spatial frequencies uniquely representable as impulsive samples on the spatial Cartesian grid of pixels. For an $N \times N$ spatial pixel grid with FOV, this corresponds to spatial frequencies $|k_x|, |k_y| < N/(2\text{FOV})$. In this study, a rosette k -trajectory [21] was used (see Section III-A). The k -space coordinates k_x and k_y sampled along the rosette determine the spatial encoding functions $\psi_n(\mathbf{x})$, $n = 1, \dots, N$, in (4). Because of the rosette's non-Cartesian sampling pattern, these functions are not orthonormal, but if the rosette does not contain large gaps (roughly, gaps $> 1/\text{FOV}$) they do form a frame for the spatial impulsive pixel grid (or for a set of continuous bandlimited spatial functions). In contrast, the information represented in the rosette k -space data are gathered from a disc centered at $k = 0$ (see Fig. 1). We refer to this as the *acquisition k-band*. In Fig. 1, three possible square evaluation grids are considered for the trajectory used here, giving square k -bands which are 1) inscribed within the acquisition k -band (52×52), 2) intermediate (64×64), and 3) circumscribed about the acquisition k -band (72×72). In all cases, there are either portions of the signal representing spatial frequencies not representable in the image array, or there are spatial frequencies appearing in the image array which are not sampled by the signal within the acquisition k -band. The consequences of these discrepancies are considered below.

D. Bias Errors Arising From Model Discrepancies

Because the estimation process relies on matching a forward model of the local signal with the actual signal, any discrepancies between the experimental signal source and the forward model signal (8) can lead to bias errors in the estimate. There are three significant potential discrepancies between the actual physical signal and the forward signal model used in these studies.

1. Local signal model—The simplest implementation has the signal model in impulsive form as in (8). An impulsive signal model does not model dephasing across the voxel, which occurs in static background gradients due to susceptibility variations within the tissue. This leads to periodic regularly spaced sidebands in k , which can produce *sideband leakage*, representing overlapping (aliased) signal contribution in the principal acquisition k -band. These overlapping sidebands are *not* present in the signal from the actual (continuous) object (see Fig. 2). Sideband leakage leads to small bias errors in all parameter estimates, as demonstrated below. An alternative implementation under development uses a local interpolated signal model defined as originating from a continuous spatial distribution specified by a smooth voxel function.
2. The evaluation k -band may not match the acquisition k -band. The evaluation band for the conventional square grid of spatial samples is a square in k -space, while the acquisition band is a disc. Where the evaluation band exceeds the acquisition band, spatial frequencies are present in the solution which are not sampled by the rosette's disc-shaped k -band. This leaves the possibility for high-spatial-frequency

“rogue” components to develop during the iterative search which are not represented in the signal. Similarly, if the acquisition k -band exceeds the evaluation band, the energy appearing in the signal at these k -locations will appear aliased into lower spatial frequencies in the estimate maps. The practical consequences in SS-PARSE, as well as corrective procedures, are discussed below.

3. Background gradients lead to signal drift in k -space (see Fig. 2). The object—more precisely, the transverse magnetization distribution—has a k, t distribution which in general spreads in k as time progresses. Consider the k, t signal distribution of an object with a linear background frequency gradient within a given spatial segment. De-phasing across this segment steadily increases over time, so that its k, t signal distribution effectively drifts away from the origin as time progresses and gradually leaves the sampled baseband. This creates two problems: 1) because the signal from the object leaves the acquisition baseband, for times beyond the exit time, the sampled signal lacks information reflecting the segment, unavoidably degrading the quality of parameter estimates from that region, and 2) the sidebands (mentioned above) drift in parallel with the principal band, so that portions of the modeled distribution leaving the baseband on one side will appear to enter from the opposite side at the same time. At these k, t -locations, the forward model signal will have no counterpart in the measured signal, so that attempts by the minimization algorithm to better match the two signals will introduce estimate errors.

III. Methods

In this section, we describe basic methods used in implementing SS-PARSE and methods used for assessing estimation errors described above.

A. k, t Trajectory; k - and Φ —Calibration

The rosette k -trajectory, $k(t) = k_f \cos(\omega_1 t) \exp(i\omega_2 t)$ (with complex $k = k_x + ik_y$) was designed with $\omega_1 = g_{\max}/k_f$ and $\omega_2 = (37/64) \omega_1$, maximum readout gradient $g_{\max} = 3.64$ gauss/cm and sampling within a k -space radius $k_f = 2.82 \text{ cm}^{-1}$ (A k -disc with this radius has the same area as a square $5 \text{ cm}^{-1} \times 5 \text{ cm}^{-1}$ cartesian grid, so the point spread function width and nominal spatial resolution are nearly equivalent to that of a 64×64 single-shot EPI for the same 12.8 cm FOV). A rosette k -trajectory was chosen because of its high sensitivity to the temporal evolution of magnetization phase and amplitude [21]. Because actual k -sampling locations deviate from nominal values due to nonideal gradient system performance and system-specific eddy-current effects, the effective k, t locations of the signal samples were determined experimentally by the method of Zhang *et al.* [22]. This procedure also determines $\Phi(t)$, a time-varying spatially constant field offset attributable largely to eddy currents. The signal sampling rate was 198 511.166 Hz, satisfying the Nyquist k -sampling criterion ($\Delta k = 1/\text{FOV}$) for a 12.8 cm FOV with the specified $g_{\max}$. Ramp-up and ramp-down sections were added at the beginning and end of the readout gradient waveforms to return gradient and k -values to zero. It should be emphasized that the performance of the method does not require that the k, t samples be at their nominal (i.e., designed) locations; it only requires that the (k, t) locations used in the estimation algorithm be based on their actual values, including effects of nonideal gradient system performance, eddy currents, etc. These actual values are determined at a single session by the calibration procedure and apply to all subsequent studies, at least until the gradient system hardware is altered or malfunctions.

B. Experimental Acquisition Methods

Experimental data were acquired using a 4.7T 60-cm-vertical-bore Varian primate MRI system (Varian Inc., Palo Alto, CA), using a stripline resonator quadrature head coil (Insight

Neuroimaging, Worcester, MA). Transverse cross-sectional rosette signals (slice thickness = 2 cm) were taken from an 8-cm-diameter beaker with four 1.6-cm-diameter tubes containing agarose gel with either CuSO_4 or Sephadex beads, in varying concentrations to provide different R_2^* values.

Data from an anesthetized rhesus macaque were acquired using an IACUC-approved protocol, using the same experimental MRI parameters used for the experimental phantom. In the macaque study a fat saturation pulse was used to suppress fat signal, which is not modeled in the current PARSE forward signal model.

To provide “gold standard” frequency and R_2^* maps for comparison with SS-PARSE parameter estimates, standard spoiled gradient-echo images (GEMS, 128 matrix, slice thickness 2 mm, 12.8 cm FOV) of a transaxial slice of the 4-tube phantom were acquired at the same session as a series of rosette acquisitions. Separate acquisitions were made with 16 different echo times (TE = 5, 8, 10, 12, 15, 20, 25, 20, 25, 30, 35, 40, 45, 50, 55, 60, 65, and 70 ms). An R_2^* map and a frequency map were computed pixel-by-pixel from the complex image data. Simple linear fits to the log of the magnitude and the unwrapped phase angle at each frame were used for preliminary estimates of R_2^* and frequency. These were used as starting values for a more accurate nonlinear fit of the complex exponential $\exp(-(R_2^* + 2i\pi f)n\Delta t)$ to the pixel histories, giving the reference images shown in Fig. 13.

C. PLCG Search and Postprocessing

Parameter optimization was accomplished by means of a modified conjugate gradient search algorithm [23], which adjusts the parameter values in the forward signal model (8) to minimize the sum-square-error (SSE) cost (9). A progressive-length modification of the conjugate gradient algorithm (PLCG) is made to improve convergence over wider frequency bandwidths, as discussed in Section IV. The progressive lengths used in the PLCG algorithm are increasing integer multiples of the basic “swoop”—a single segment of the rosette passing across the diameter of the rosette, giving a single gradient echo. In the k -trajectory used in these studies, this swoop consists of 120 data points. Iteration at each length is continued until the fractional decrement in cost $\Delta Q/Q$ at a single iteration becomes smaller than a specified value, termed the *stopping criterion*. A typical series of progressive lengths in terms of numbers of swoops would be (2, 3, 5, 7, 10, 15, 22, 29, 38, 47, 58, 71, 84, 99, 107), with a stopping criterion of 0.01 for each of the first 15 lengths and 0.001 for the final length. Thus the first set of iterations minimizes the SSE cost over the first two swoops only, and continues until the fractional cost decrement $\Delta Q/Q = 0.01$; the second set of iterations is over the first three swoops, and continues until $\Delta Q/Q = 0.01$, etc.

Convergence to the correct parameter values at a given location depends upon its frequency being correctly estimated, and for frequencies outside a certain range, convergence may not be possible. To ensure convergence of the search at all locations, it is often necessary to apply a bias frequency to the FID signal [i.e., multiply the signal by $\exp(2i\pi f_{\text{bias}}t)$] before beginning the parameter optimization search. Because the FID signal is seen as a sum of individual local FIDs, each modulated by the temporal course of its local modulation function $\psi_{j,n}$ [as in (7), this simply artificially shifts the effective frequency at all locations by the same frequency f_{bias}]. A good choice of f_{bias} is one that places the center of the object’s frequency distribution near zero. This improves the chances the frequencies present in the entire object will fall within the required range for convergence. This range is indicated by the width of the broad cost function valley seen in Fig. 4(a)–(c). In practice, for cases of poor shimming or large susceptibility gradients when the frequency range of the object is unknown, the best f_{bias} can be chosen by trying a series of values and picking a value which produces parameter maps including the entire object.

The CG algorithm is known to have the “semi-convergence” property: i.e., in the presence of measurement noise, during the CG search the model signal converges monotonically toward the actual signal in the LS sense, while the parameter estimates begin to converge toward the minimum-norm solution in parameter space, but then begin to move away [24]. Thus, while searching at the last data length, the parameter maps are often observed to become ever rougher, especially under low signal-to-noise ratio (SNR) circumstances. We investigated this phenomenon using experimental 4-tube phantom data. To limit roughness, separate regularization terms were added to the cost function Q for the complex magnetization map ($M_{x'} + iM_{y'}$) as well as the exponential map $\mathbf{B} = \mathbf{R}_2^* + 2i\pi f$. With this addition, the cost function minimized is

$$Q = \sum_n |z_n - \widehat{s}(\mathbf{p})|^2 + \sum_j \left[\alpha_1 |\{\mathbf{H} * \mathbf{M}_\perp\}_j|^2 + \alpha_2 |\{\mathbf{H} * \mathbf{B}\}_j|^2 \right]$$

where α_1 and α_2 are regularization parameters, $\mathbf{M}$ and $\mathbf{B}$ are the spatial arrays corresponding to the magnetization map and exponential maps, $*$ denotes the spatial convolution operation, and $\mathbf{H}$ is the difference array used to gauge roughness

$$\mathbf{H} = \begin{pmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{pmatrix}.$$

The PLCG algorithm is computationally intensive. A typical initial estimation run for 64×64 -format parameter maps requires 1–4 min using MATLAB code, depending on the data and the settings of the PLCG. When processing a series of similar frames, as in fMRI applications, after the first frame is processed, subsequent frames can be processed more quickly by simply iterating at full length (T_a) only, using as initial parameter values the final estimates computed at the first run.

D. Simulated Phantom Studies

The parameter estimate maps produced by SS-PARSE represent the results of an iterative search to arrive at the LS solution, where the cost function reaches a minimum. The shape of the cost function determines how difficult it will be for the iterative search algorithm to find the LS solution. Additionally, the location of the cost function minimum establishes a limit to performance; if the LS solution does not coincide with the actual parameter values, then even if convergence to the LS solution is successful, there will be estimation errors. To gain some insight into the shape of the cost function in parameter space, we computed surface and contour plots representing the “partial” cost function, i.e., the effect on total cost when two of the parameters $M_\perp$, f , or R_2^* were varied at a specific pixel location with all parameters at other pixels held constant.

To investigate the relationship of the parameter convergence history to the LS estimate and to the actual parameter values for representative pixel locations, partial cost functions (i.e., cost function variations when only parameters of interest are changed and all others are held fixed at their actual values) were computed and plotted as contours along with convergence tracks for the parameters involved. Contour plots were computed from 100×100 arrays of values of M_{xy} and R_2^* and for 100×100 arrays of values of R_2^* and frequency. These simulations used a small analytic disc phantom with different levels of noise added to the

signal, tracking the course of the parameter estimates throughout the search. To prevent bias errors due to the evaluation k -band exceeding the acquisition k -band, an “inscribed” spatial evaluation grid (52×52) was used. To prevent bias errors due to aliasing from data outside the square evaluation band, the rosette data outside the evaluation band was not used in the estimation process. (An alternative to omitting this data would be to apply regularization at a suitable level on a higher-resolution grid.) The PLCG search was continued to a very small tolerance ($\Delta Q/Q = 10^{-8}$; 3205 iterations) to ensure the final estimates would be close to the LS minimum for the data set at hand.

To investigate the process of PLCG convergence for these locations, a series of estimations were performed using an analytic simulated phantom. Parameter estimates for three selected adjacent pixels were tracked throughout the 3205 iterations to observe the course of convergence of parameter values. All simulations and processing were performed using MATLAB (Mathworks, Natick, MA), running on a dedicated computer with eight AMD Opteron 880 dual-core processors with 64 GB of Random Access Memory (@Xi Computer Corporation, San Clemente, CA).

To demonstrate effects of the impulsive/continuous model discrepancy, analytically simulated test signals were generated, representing the signal from a simulated version of the 4-tube experimental phantom. Additional simulated signals were computed for a nonphysical spatially impulsive version of the 4-tube phantom. This impulsive simulated object matched the impulsive forward signal model, so the resulting estimated parameter maps should lack bias errors caused by the impulsive/continuous discrepancy.

To investigate the effects of relatively low $|M_{\perp}|$ values on convergence, a two-part disc phantom signal was simulated analytically, with a small central region having a lower intensity than the outside annulus of unit intensity. Relative disc amplitudes of 0.5, 0.2, 0.1, and 0.05 were tested. Noise levels evaluated (rms noise levels as fraction of maximum FID amplitude) were 0.001, 0.002, 0.005, 0.01, 0.02, 0.05. Equivalent image SNRs were 200, 100, 40, 20, 10, and 4.0, respectively.

In all simulations reported here, noise levels were specified in terms of fraction of maximum signal—i.e., the signal level at the earliest $k = 0$ crossing, the peak of the first gradient echo. To translate these measurement SNR values into terms of equivalent image SNR, we generated Cartesian k -space data for the object being simulated with the same FOV and matrix size, then added measurement noise with the same relative rms power as in the rosette data. (Image SNR was defined as the ratio of maximum image amplitude to the Gaussian real-valued standard deviation. The Gaussian standard deviation of pixel noise was estimated as (mean of the Rician-distributed pixels in a background region)/1.253 [25]). Simulation results include these equivalent image SNR values as an aid for interpreting noise values in familiar terms.

E. Comparing Encoding and Decoding Strategies

In a separate series of simulation studies, we compared encoding and decoding strategies.

An encoding method sometimes used for single-shot R_2^* and frequency estimation is multiple-echo EPI (MEPI), a series of repeated complete EPI readouts following a single excitation pulse [3]. These repeated EPI readouts are used to generate separate EPI images with different echo times TE. Relaxation rates R_2^* and frequencies can be estimated for each pixel by fitting complex exponentials to the values at that pixel through the image series. The MEPI acquisition used in our simulation was based on a simplified version of the EPI sequence used in our Varian system, with 64 echoes of 64 samples each, using a maximum readout gradient of 2.29 G/cm, with FOV 12.8 cm, and echo times TE = 20, 60, 100, 140 ms. For convenience in the simulation, the four EPI readouts were concatenated, omitting

the initial gradient dephasing and final gradient rephasing pulses required in the actual sequence. This slightly idealized version of MEPI is expected to give a slightly optimistic estimate of performance relative to the actual sequence, which has acquisition dead times during the de-phasing and rephasing gradients. Nominal resolution and FOV for MEPI and rosette were identical. Note that the MEPI sequence has a total 160 ms duration, significantly longer than the 65 ms readout of the rosette, giving the MEPI a nominal advantage in terms of acquisition time. We used analytically computed signals from phantoms with four discs with different relaxation rates ($R_2^* = 15, 25, 35$, and 45 s^{-1}) and frequencies, and we compared three sets of encoding and decoding strategies for estimating $M_\perp$, f , and R_2^* : 1) MEPI encoding, followed by 2DFT reconstruction of each of the four frames and a nonlinear exponential fit of (1) for the local FID, 2) MEPI encoding followed by PLCG search, and 3) a rosette acquisition for encoding (with the same 2.29 gauss/cm maximum readout gradient as the MEPI, with k-band circumscribing the square evaluation k-band), followed by PLCG search. In this simulation, the rosette had a k -radius circumscribing the band of the evaluation grid (see Fig. 1).

It was expected that the rosette sampling trajectory might be more effective in acquiring information regarding the parameters $M_\perp$, f , and R_2^* than the MEPI trajectory, and this would be manifested by lower estimate variances. This was investigated by evaluating the Cramer-Rao Lower Bound (CRLB) [26], a theoretical lower bound for error variance in an unbiased parameter estimate for a signal with a given level of measurement noise. CRLB estimates were computed for a small disc object of constant amplitude with a range of relaxation times R_2^* . Because of the large matrix inversion required in CRLB computation, the sample object used was a disc with about 200 pixels, much smaller than the experimental phantom. Thus the CRLB values may not be accurate predictors of actual variances to be expected in the experimental data, but may be valuable as indicators of relative merit.

In principle, a decoding algorithm is optimal (i.e., an efficient estimator [26]) if it succeeds in extracting all of the information encoded into the signal by the encoding mechanism. In that case, the estimation errors for an unbiased estimator with a given measurement noise level should be consistent with the CRLB predictions of the error variance floor of the encoding mechanism.

To compare the three encoding and decoding strategies with one another and with relative CRLB limits, a set of 16 individual noisy signals (differing only in the noise realization) were simulated for low and high noise levels for both MEPI and rosette sampling. Each simulated FID was normalized to amplitude 1.0, and normally-distributed independent zero-mean complex noise with standard deviations of 0.0005 or 0.002 was added to each normalized complex FID signal. Equivalent image SNRs were 101.1 and 25.7, respectively.

IV. Results

A. Typical Estimate Characteristics: Experimental Phantom Studies

Spatial and temporal variability of parameter estimates were examined in a series of 30 sequentially acquired signals from the same slice level in the experimental 4-tube phantom. The only differences expected among these signals are different noise realizations and perhaps small instrumental variations. A 64×64 spatial evaluation grid was used. Fig. 3 depicts superimposed profiles across the $|M_\perp|$, R_2^* and f maps at row 35, passing through the center (water only) section of the phantom. The f maps exhibited spatially global frequency “jitter,” in that the local frequency estimates at all pixels changed by virtually the same amount from one frame to the next. Fig. 3(d) depicts the row-35 frequency profile after global frequency jitter has been removed. Note that 1) for all parameters, temporal

variability at the pixels within the phantom is seen to be smaller than the spatial variability, 2) the $|M|$ and R_2^* estimate maps show a significant spatial roughness, in spite of the fact that the values of these parameters are expected to be relatively uniform in the water pixels, 3) the f maps are spatially smoother, with frequency variations consistent with the effects of the nearby tubes, which have susceptibilities different from water and therefore perturb the local magnetic field, and 4) after removal of the global jitter from the frequency maps, temporal standard deviation (TSD) of the frequency estimates is quite small, ranging from 0.08 to 0.16 Hz. Roughness can be reduced by regularization, which was not used in this example.

Regions of interest were defined for the central water region of the phantom (32 pixels) and the upper left tube (42 pixels). Mean parameter values and mean TSD values were computed for the regions. Results are shown in Table I, along with parameter estimates from corresponding regions of the “gold standard” Gradient Echo data. Note that SS-PARSE frequency and R_2^* values are similar, and that the temporal standard deviations in all parameters are small.

B. SSE Cost Function and Convergence Behavior

Fig. 4 depicts the partial cost function behavior for a single pixel location, when $|M_\perp|$, R_2^* , and f are varied. The plots shown here represent SSE costs for noisy single-shot data from an analytically simulated 5-cm-diameter disc with $R_2^*=30\text{ s}^{-1}$ and offset frequency $f=10\text{ Hz}$ using the standard rosette k -trajectory. Cost functions were computed from three data lengths (5, 20, and 65 ms) to demonstrate the cost function behavior as the iterative search continues through the progressive lengths. Note that as data length increases, the contours become narrower in the frequency direction, and that small secondary minima (seen in Fig. 4(c) and (f) at low $|M_\perp|$ values at ca. 50–60 Hz frequency offsets) become spaced more closely in frequency. Fig. 4(j) and (k) show cost function surface and contours for $|M_\perp|$ and R_2^* . The noise level was chosen to represent moderate-to-high noise conditions, with data SNR ratio about 500, corresponding to equivalent image SNR of about 65. The contour plot for lower SNR data (data SNR about 200, image domain SNR about 26) in Fig. 4(l) suggests that with increasing noise levels, secondary minima become deeper with respect to the global minimum. However, note that the cost function for short data lengths appears well-behaved and free of secondary minima within a roughly $\pm 100\text{ Hz}$ bandwidth. The progression of cost function shapes from Fig. 4(a)–(g) illustrates the rationale for the progressive length search. The starting estimates of $|M_\perp|$ and f are both zero. During the initial search within the short-length cost function shown in Fig. 4(a), the estimate of $|M_\perp|$ moves toward larger values, and the estimate of f moves toward the actual f value. Thus as the search begins at the second data length, the $|M_\perp|$ estimate will be well away from the region where secondary minima occur, facilitating convergence to the local minimum. A single-shot image with SNR = 65 would typically be regarded as too noisy for use in BOLD fMRI. This “worst case” example was chosen to demonstrate that low SNR does not produce problematic secondary minima, at least in the partial cost functions investigated here.

The 2-D partial cost function shown here represents SSE values on a 2-D plane within a vastly larger-dimensional parameter space, so the cost functions of Fig. 4 represent only a very partial view of the SSE behavior. For convenience, these cost functions were computed under the assumption that parameter estimates at other pixels had converged to their actual values. To determine the effect of this approximation, we computed cost functions for full-length data with all other parameters perturbed in a Gaussian distribution from their actual value with rms deviation in f and R_2^* of 2 Hz and 2 s^{-1} respectively, and rms deviation of .01

in $|M_{\perp}|$. The resulting cost functions (not shown) were very similar, and did not introduce additional secondary minima in the $(|M_{\perp}|, f)$ plane.

Fig. 5 shows the R_2^* estimate histories for the three selected pixels during the early (iterations 1-250) and late (iterations 150-3205) portions of convergence. Starting values of all parameters were 0. Long-term convergence plots demonstrate that the parameters converge to stable values, but that the estimates are slightly different for the different pixels. These errors likely represent both noise propagated from measurement noise, and bias errors due to model mismatch—e.g., the Gibbs' ringing visible in Fig. 5(a). The early convergence plots [Fig. 5(e) and (f)] demonstrate that $|M_{\perp}|$ converges to near the final $|M_{\perp}|$ estimates by iteration 20, much earlier than the frequency estimates, which converge more gradually. The implication is that for pixels with sufficiently large $|M|$, during early stages of convergence in PLCG, while the cost function resembles that of Fig. 4(a) and (b) without secondary minima over a broad range of frequency, the current search position moves to larger $|M_{\perp}|$ values, well to the right in the $|M_{\perp}| - f$ plane. This leaves the search position in a region of the cost function without secondary minima during subsequent stages [Fig. 4(c)–(g)], so the search can converge to the primary minimum.

Fig. 6(a) demonstrates cost function contours near the SSE minimum for pixel 1 (R_2^* versus $|M_{\perp}|$), with actual parameter value (open circle), final estimate (square), and convergence path. Fig. 6(b) and (c) are plots of cost function contours for R_2^* versus $|M_{\perp}|$ for pixels 2 and 3. Fig. 6(d) shows convergence in frequency v $|M_{\perp}|$ for pixel 1. For Fig. 6(a)–(d), rms measurement noise levels were 0.001 (Equivalent image SNR = 168). Fig. 6(e) and (f) are representative plots of $R_2^* v \cdot |M_{\perp}|$ and $\hat{v} \cdot |M_{\perp}|$ for pixel 1 for rms noise = .01 (equivalent image SNR = 16.8). Contours represent increments of SSE cost function (as fraction of minimum value) of 1×10^{-4} , 2×10^{-4} , 5×10^{-4} , 1×10^{-3} , 2×10^{-3} . These results demonstrate the following.

1. Final estimates are very near the SSE cost function minimum in all cases, suggesting the PLCG has converged to the LS solution in this plane of the parameter space.
2. For each pixel and each parameter, the LS solution does not coincide with the actual value of the parameter. (This was true for all noise levels, even in the absence of noise, due to model mismatches).
3. SSE contours in a $|M_{\perp}|, R_2^*$ plane exhibit a diagonal trough, suggesting the minimization problem is somewhat poorly conditioned with respect to these parameters, and $|M_{\perp}|$ and R_2^* estimates can be expected to have correlated errors.
4. Within a realistic range of noise levels, increasing the level of measurement noise does not introduce roughness or secondary minima to the cost function in the region of the primary minimum, but it does lead to a shallower minimum, which may slow convergence. As noise levels are increased, the parameter estimate variances are expected to become larger—i.e., the expected distance of the LS solution from the actual value is expected to become greater.

An additional factor influencing local convergence is simply the relative strength of the local signal contributed by a given location. From an intuitive view, if the local signal is weak, its relative influence on the cost function will be weak. Even in the asymptotic limit near the LS minimum, where the cost function is expected to behave approximately quadratically [27], the gradient of the cost function along the parameter axes belonging to a low-level pixel will be very small, providing little incentive for descent along those directions. Thus the search may not converge at all along the parameter directions associated with low-intensity pixels.

A thorough analysis of this local nonconvergence is beyond the scope of this publication, but its basic behavior is demonstrated here in a series of simulated phantom studies.

Fig. 7 depicts a set of frequency estimates within a center disc for (b) 52×52 , (c) 64×64 , (d) 72×72 evaluation grids, for rms noise levels as fraction of maximum FID amplitude 0.001, 0.002, 0.005, 0.01, 0.02, 0.05, (Equivalent image SNR = 167, 85, 34, 17, 8.6, 3.6) for relative disc amplitudes of 0.5, 0.2, 0.1, 0.05. Convergence of local pixel parameters has failed if the frequency estimate is far from the actual value (8 Hz in this example). Note that convergence is achieved for intensity 0.5 even for relatively high noise level 0.02, while for lower intensity levels, convergence occurs only for low noise levels. Note in Fig. 7 the object appears within a faint disc seen in the background. This disc is simply a reduced FOV. To speed the estimation process, we estimate parameters only in pixels within a circular FOV known to contain all pixels producing a signal. When dealing with a small object such as this one, we often reduce the FOV accordingly to reduce the total number of parameters that need to be estimated. In short, the circle in the background simply represents the region within which estimates were made.

C. Optimality and Comparison of Encoding and Decoding Strategies

CRLB comparisons of R_2^* variance for MEPI sampling and for rosette sampling for R_2^* values of 15, 25, 35, and 45 s^{-1} are shown in Fig. 8(a) and (b) for two noise levels. Note the expected standard deviations are nearly the same for 15 s^{-1} but diverge increasingly with increasing R_2^* . CRLB-predicted s.d. values are proportional to levels of measurement noise, so the values shown in Fig. 8(b) are exactly 4 times those in Fig. 8(a) and are shown only for purposes of comparison with the actual s.d.s from simulated noisy signals shown in Fig. 8(d). Fig. 8(c) and (d) compare results at two noise levels for 1) MEPI encoding with conventional 2DFT followed by pixel-by-pixel nonlinear fit to a complex exponential, 2) “raw” (not transformed to the spatial domain by DFT) MEPI data analyzed by the PLCG estimation algorithm, and 3) rosette data processed by PLCG for the same four R_2^* values. Note that estimates from MEPI data from PLCG are lower than with conventional 2DFT and LS pixel-by-pixel fit, suggesting that the PLCG estimation is superior at extracting information on R_2^* from the encoded signal. The s.d.s for SS-PARSE are smaller, especially for faster decay rates, than for either MEPI-based approach. The simulation results appear qualitatively consistent with the CRLB predictions. In summary, these results suggest that for purposes of R_2^* estimation, the rosette encoding is superior to MEPI encoding, and that for MEPI encoding, the PLCG optimization is superior to 2DFT/pixel-wise fitting as a decoding method. This does not constitute a comprehensive analysis of encoding and decoding performance, but does suggest that there are significant differences in the basic performance levels of these different encoding and decoding approaches.

D. Bias Errors in Parameter Maps; Modeling Errors

Fig. 9 demonstrates the appearance of bias errors in basic parameter maps in a study using simulated data. Each column has the three parameter estimate maps, $|M_\perp|$ (top), R_2^* (center), and f (bottom). Columns (a), (b), and (c) were estimated on a 52×52 grid, and columns (d), (e), and (f) on a 72×72 grid. Columns (a) and (d) were estimated from signals computed from a physically unrealizable impulsive object, in which local signals originate only from the grid points. Signals for columns (b) and (e) were computed analytically from a spatially continuous object model, as in the other simulations in this paper. Signals for columns (c) and (f) were obtained experimentally from the actual 4-tube phantom using the 4.7T Varian system. The simulated signals had no measurement noise added to them. The frequency maps in the simulated tubes are much more uniform than the experimental case, simply because for convenience the tubes were simulated with a uniform frequency.

Estimate maps from the unrealizable impulsive phantom [Fig. 9(a) and (d)] are extremely smooth. In fact, if no measurement noise is added to it, the impulsive phantom signal gives estimates with vanishingly small errors. This result confirms that a PLCG search based on the impulsive signal model (8), with the correct parameter values, can match the signal exactly when the actual signal originates from an impulsive object. However, a continuous object, as in the simulation [Fig. 9(b) and (e)] or the experiment [Fig. 9(c) and (f)], produces a signal which cannot be fitted exactly by the impulsive signal model, no matter what parameters are used, leading to spatially variable bias errors. In terms of the SSE cost function, this represents a displacement of the cost function minimum from the actual parameter values (as demonstrated in the partial SSE cost function contour plots of Fig. 6). In the figure, this is visually more apparent in the R_2^* maps than in the $|M_\perp|$ maps. This basic modeling error was mentioned in the Theory section above.

Another source of bias errors discussed in the Theory section is the mismatch of k -acquisition and k -evaluation bands. These bias errors can be seen in Fig. 9(d)–(f) as high-frequency checkerboard-like components corresponding to the outer corners of the evaluation k -band in Fig. 1, which are present in the evaluation k -band for 72×72 , but are not represented in (and cannot be influenced by) the signal data from the acquired k -disc. These checkerboard-like bias errors are especially apparent in the estimates from the unreal impulsive object (d), which have k -mismatch bias errors, but do not have bias errors from the impulsive/continuous model discrepancy.

The 52×52 $|M_\perp|$ maps Fig. 9(b) and (c) include ringing effects not seen in the 72×72 maps Fig. 9(e) and (f). This may be due to aliasing of the high-spatial-frequency signals from portions of the acquisition k -disc outside the smaller evaluation band of the 52×52 maps.

E. Approaches to Reduce Roughness and Bias Errors

In experimental 4-tube phantom studies, we studied the evolution of roughness and checkerboard-like artifacts as iterations continued. For evaluation grids larger than 52×52 , with continued iteration checkerboard “rogue component” artifacts steadily increased in amplitude. Fig. 10(a)–(c) demonstrates this effect in a series of 64×64 estimate maps in which convergence was stopped at fractional decrements $\Delta Q/Q = 0.01, 0.001$, and 0.0001 , respectively. Increasing roughness and checkerboard artifacts are visible, especially in the R_2^* maps. This appears to be associated with the “semiconvergence” behavior of the CG algorithm [24].

In Fig. 10(d)–(f), iteration was stopped at $\Delta Q/Q = 0.001$, but different approaches to reduce roughness and checkerboard artifacts were used. In Fig. 10(d), a bandlimiting operation was used to remove out-of-acquisition-band components at each iteration of the PLCG algorithm; i.e., the parameter maps were bandlimited at each iteration to the evaluation band by applying a spatial-frequency window. In Fig. 10(e) and (f), the PLCG included regularization with parameters $\alpha_1 = \alpha_2 = 10$ and $\alpha_1 = \alpha_2 = 200$, respectively. In this modified PLCG algorithm, regularization alters the cost function during the optimization search, penalizing roughness in the parameter maps. Note the high-frequency checkerboard-like “rogue” components visible in Fig. 10(b) and (c) are effectively removed in Fig. 10(d)–(f).

F. Bias Errors and Local Nonconvergence

In experimental data, Fig. 11 demonstrates the effects of local frequency estimation error. Fig. 11(a)–(c) are $M_\perp$, frequency, and R_2^* estimate maps for a typical four-tube phantom study. As mentioned in the Methods section, to ensure convergence at all locations, it is often necessary to apply a bias frequency to the FID signal (i.e., simply twist it in phase

angle) to place its mean frequency near zero. This step improves the chances the frequencies present in the entire object will fall within the required range for convergence, so that all pixels within the phantom or subject cross-sectional image will converge properly. To demonstrate consequences of local nonconvergence, a poorly chosen offset frequency was applied to the FID to cause the PLCG algorithm to fail to converge to the proper frequency at some pixels. This can be seen along the upper right margin of the phantom, as seen in (e). As a result, the $M_{\perp}$ and R_2^* maps contain localized bias errors [Fig. 11(d) and (f)] both near (solid arrows) and remote from (open arrows) the location of the frequency error. In practice, non-convergent pixels can lead to ghost artifacts distributed through the parameter maps, with strength roughly proportional to the number and intensity of the nonconvergent pixels. This can be thought of as a spatial misassignment of signal energy due to nonconvergence at some locations.

G. Large Off-Resonance Frequencies and Large Background Gradients

Data were acquired from the experimental 4-tube phantom with severe de-shimming to investigate the problem of dealing with large background gradients and large off-resonance frequencies. A two-stage estimation approach was applied to the data: 1) a preliminary frequency map was obtained by PLCG, and a low-order 2-D expansion was fitted to it; 2) the local temporal basis functions were modified by adding linear phase corresponding to the rough local frequency map, and the PLCG algorithm was applied again using the modified basis functions. The resulting parameter maps and pseudo- T_2^* -weighted image are shown in Fig. 12(a)–(d). For comparison a gradient-echo EPI image of the deshimmied phantom is shown in Fig. 12(e) and a profile through the frequency map [Fig. 12(f)]. The EPI image is highly distorted and exhibits significant dropouts as well as N/2 ghosts. The SS-PARSE parameter maps are of poorer quality than the well-shimmed case of Fig. 10 but are free from distortion and N/2 ghosts. As measured by the final SS-PARSE frequency map, the frequency range was -158 Hz to $+108$ Hz with a maximum measured background gradient of 106 Hz/cm. It should be emphasized that though this was a multistage process involving computation of a preliminary frequency map, it was all performed using one 65-ms single-shot signal (in fact, only the first 23 ms of this signal were used in the analysis, as explained below). Thus this or a similar approach could be used in dynamic studies with true high temporal resolution (i.e., without the need to acquire a separate preliminary data set and without the need to assume the frequency map is constant over time).

H. Comparison With “Gold Standard” Gradient Echo Data

Figs. 13 and 14 demonstrate direct comparisons between the parameter maps from the multiple-TE “gold standard” acquisition (128×128 , fitted from the series of multiple-shot GEMS acquisitions) and the SS-PARSE results (64×64 , from a single 65-ms readout). Overall, the R_2^* and frequency parameter maps match well. Beyond the obvious coarser resolution of the 64×64 SS-PARSE maps, there are two notable differences in the R_2^* maps. First, the bright rings around the individual tubes in the SS-PARSE R_2^* map are more noticeable. These are thought to represent simply the larger effective R_2^* within the tube/water boundary voxels including the differing frequencies and R_2^* of the water and the doped agarose gel sharing the larger voxels of the 64×64 array. The frequency values seen in the frequency maps are quite similar. The GE frequency map has been masked to eliminate very large frequency values outside the phantom which would have dominated the scaling of the display. The SS-PARSE frequency map has not been masked. A comprehensive comparison of SS-PARSE maps with the gold standard maps will be presented in a separate publication. In Fig. 14, T_2^* -weighted GE images are compared with pseudo- T_2^* -weighted images computed as $|M_{\perp}| \exp(-R_2^* TE)$, where R_2^* and $|M_{\perp}|$ are the SS-PARSE estimated parameter

maps. Even though the 64×64 snapshot images are of coarser resolution than the 128×128 GE images, contrast is comparable, and small irregularities seen in the tubes are clearly visible in both GE and SS-PARSE maps.

I. In Vivo Results

The *in vivo* SS-PARSE maps shown in Fig. 15 depict clear tissue contrast in the R_2^* and frequency maps, more readily seen in the pseudo- T_2^* -weighted image. Though the accuracy of SS-PARSE R_2^* maps has been assessed in phantom studies (in [28], with more complete results in a separate forthcoming publication), the accuracy of *in vivo* R_2^* measurements remains to be evaluated thoroughly in animal and human studies. There are a few cloudy artifacts in the R_2^* and $|M_\perp|$ maps, some corresponding to very rapidly decaying ($R_2^* > 60 \text{ s}^{-1}$) components. These are thought to represent artifacts associated with failure of convergence in the extracranial tissues, as in the artifacts of Fig. 11. Gray matter/white matter/CSF tissue contrast is apparent in the R_2^* map and in the pseudo- T_2^* -weighted image.

V. Discussion

The goal of this paper has been to describe basic properties of the SS-PARSE method and to characterize key features of its practical performance when the PLCG algorithm is used to solve it. Unfortunately, it is difficult to characterize the solution to this large nonlinear problem in comprehensive terms theoretically. Consequently, characterizations of SS-PARSE behavior presented here have been mostly empirical.

A. Convergence Issues

In simulations, experimental phantom and *in vivo* studies, we have demonstrated successful convergence, even under conditions of adverse off-resonance and background gradients. However, parameters for pixel locations with low signal magnitude and/or large frequency offsets are more likely to be misestimated, due to failure of the search algorithm to converge to the LS solution along those parameter-space axes. It was demonstrated that local nonconvergence can lead to bias errors near and far from the misestimated pixel. Severity of these cloudy bias artifacts appears to be related to the amount of signal misestimated.

It is possible to define a “convergence envelope,” a curved surface enclosing the range of parameters $|M_\perp|$, R_2^* , and f within which convergence is guaranteed. A detailed investigation of this convergence envelope, and the effects on the envelope of PLCG control parameters (e.g., final convergence tolerance, degree of regularization) will be presented in another report. The convergence envelope is established by the relative gradient slopes within the LS cost function and is thus a matter of the conditioning of the problem. It is likely that the convergence envelope can be enlarged for a given data set by incorporating appropriate preconditioning or adaptive conditioning measures in the search algorithm. Such algorithmic improvements are under investigation.

In the simple example of Fig. 7, it was seen that local nonconvergence may occur in pixels with low relative values of $|M|$, regardless of local SNR. This is thought to represent a conditioning problem that may be amenable to preconditioning measures we have not explored. However, it should be emphasized that in typical brain cross sections, the $|M|$ values are relatively uniform, simply because proton density in normal tissues is relatively uniform. Thus, in many applications this behavior may not present a practical problem even with the current algorithm.

It is reasonable to ask how the PLCG can converge to a unique solution when during its initial stages the inverse problem is underdetermined. Because the LS solution exists even for underdetermined problems, there will be an LS solution, though probably an inaccurate one, even for a short data length. Even though the LS solutions for initial abbreviated data lengths can be expected to include aliasing errors due to undersampling of k -space, these errors are observed empirically to have little if any impact on the final solution [this can be observed in Fig. 5(e) and (f)]. The reason for this may be that the iterative search innovations $\Delta \mathbf{p}$ for initial iterations at any data length are driven by the largest discrepancies between the experimental signal and the model signal. Thus they are weighted heavily toward low spatial frequencies, because the density of rosette samples is high near the k -space origin, and also because the signal values are typically much larger at lower spatial frequencies. Furthermore, typical frequency maps tend to be relatively smooth over space—i.e., they are dominated by low spatial frequencies. Thus even though initial data lengths are short and the overall problem is initially underdetermined, the frequency map tends to converge methodically throughout the PLCG process, with lower spatial frequency components converging first and higher spatial frequency components converging later. Thus aliasing errors appearing in these maps at intermediate stages of the iteration will tend to be corrected during the final iterations at full data length.

One might suspect the large-dimensioned nonlinear Least Squares problem to be solved here would be characterized by a cost function with multiple secondary minima, and that it would be quite difficult in practice to prevent the optimization search from being trapped in one of these secondary minima. We investigated here the shapes of partial SSE cost functions and the search tracks taken by the progressive length conjugate gradient (PLCG) algorithm given noisy signals. For pixels with adequate magnetization magnitude and within a broad frequency bandwidth, our results suggest that even though secondary minima exist, the PLCG search exploits the properties of the SSE cost function to avoid them and converges to the primary minimum. Furthermore, the behavior of the cost function near the global minimum appears to remain concave and without discernible roughness in contour plots even when measurement noise levels are higher than would be acceptable in MRI practice.

B. Bias Errors

Even if the search algorithm behaves perfectly, the SS-PARSE parameter estimates may include significant bias error, because the LS solution is inherently biased by modeling errors. It was shown that these modeling errors include 1) the discrepancy between the spatially continuous source of the radio-frequency signal and the spatially impulsive source of the forward signal model and 2) the mismatch between the disc-shaped acquisition k -band and the square evaluation k -band. “Leakage” errors from the use of the impulsive model are typically small: for a 10-cm-diameter disc object in a 19.2 cm FOV with 64×64 image matrix, the energy of leakage errors is only 0.8% of the signal energy in the baseband. These same leakage errors occur in standard 2-D MRI any time the 2-D FFT alone is used to reconstruct images. Nonetheless, an interpolated forward signal model is under investigation [29] to reduce the leakage errors. Bias errors from problem (b) include aliasing of acquired signal from outside the evaluation band, and spurious “rogue components” introduced during the iterative search. A simple approach to eliminate aliasing of acquired signal outside the evaluation band, and at the same time prevent “rogue components” from arising, is to use a spatial sampling grid with an evaluation k -band that is inscribed in the acquisition k -band, and then simply discard the acquired data and the modeled signal data outside this square band, performing the search on the basis of the data that is inside the square. In theory, this has the disadvantage that it discards legitimate information, but in terms of signal energy, the useful information lost is negligible (for a 10-cm-diameter disc object, only about 2% of the signal energy is lost). The amount of information sacrificed can be

reduced somewhat by using a hexagonal evaluation grid [30]. This has a hexagonal evaluation k -band, which fills the acquisition k -band more completely, but requires interpolation computations to form Cartesian-grid parameter maps for display. Another approach is to use a larger image matrix, with evaluation band entirely containing the acquisition band. This reduces leakage errors, and eliminates the aliasing of acquired signal outside the evaluation band (because the evaluation band now includes the entire acquisition band). With this option, it is necessary to use either bandlimiting during the iterative search, or roughness-penalizing regularization, to limit growth of “rogue components” during the search. This tactic also increases the number of parameters to be estimated by a factor of two beyond the number used for the inscribed evaluation band, potentially significantly increasing computation time.

Bias errors from several sources can be effectively reduced by including roughness-penalizing regularization in the PLCG algorithm or by preweighting the signal in k to provide a roll-off at higher spatial frequencies. However, regularization and pre-weighting also smooth the maps, somewhat reducing effective spatial resolution. Although it is likely that regularization will play an important role in achieving the best results in SS-PARSE, we did not address the issue of regularization comprehensively in this study because we wanted to concentrate first on the more basic issues of PARSE optimization. Regularization in this work has been used to reduce roughness, but on an empirical basis.

It should be emphasized that even though effects of regularization are well-understood in the context of conventional MRI using non-Cartesian k -space sampling patterns, and in linearized models such as that of [7], a thorough study of the effects of regularization for the general dynamic signal model used here have not yet been investigated.

C. Temporal Variability of R_2^* Estimates

The temporal standard deviations illustrated in Fig. 3 and listed in Table I demonstrate that for a well-shimmed, relatively uniform-intensity object, SS-PARSE estimates of the local parameters are consistent and have low temporal variability.

Random measurement noise propagates into estimation errors by displacing the position of the SSE minimum from the position representing the true parameter values in parameter space. The prediction of estimate error given by the Cramer-Rao Lower Bound can be applied to discrete models, even in nonlinear problems [27], though the CRLB cannot give a completely accurate prediction of variance errors here, because 1) the actual signal source is spatially continuous rather than discrete and 2) the CRLB cannot be evaluated practically for the large number of free parameters in a typical image FOV and must be evaluated only for smaller-dimensioned “toy” problems as was used in deriving the results shown in Fig. 9.

Nonetheless, the “toy problem” CRLB predictions for R_2^* variability for MEPI and rosette trajectories are consistent with the PLCG results from simulated data, suggesting the CRLB and analyses based on the information matrix may be useful in comparing the effectiveness of encoding by different k , t -sampling patterns.

D. Background Gradients

Background field gradients due to local tissue susceptibility variations are problematic in conventional single-shot imaging methods, typically producing distortion or blurring. Our experimental study using a misshimmed phantom confirmed the absence of distortion in SS-PARSE parameter maps, even with large background gradients producing substantial distortion in EPI. The fact that the rosette samples near $k = 0$ earlier and much more frequently than most single-shot methods is likely an advantage for recovering useful information on these problem areas. Nonetheless, large background gradients limit the

useful signal duration in SS-PARSE (as well as other single-shot methods). Specifically, it was shown in Fig. 2 that an object (or a section of an object) with a background gradient produces a signal in k , t -space which drifts in k . If the background gradient is very strong, such that the signal from the object reaches the k -radius within which data are acquired, the data beyond this time should be discarded in the PLCG analysis. Attempting to fit to the signal past this exit time leads to a rapidly increasing discrepancy between the model signal (based on an impulsive model with periodic structure in k -space) and the actual signal (derived from a continuous distribution in space, which does not have periodic k -space structure), producing artifacts due to the reentering model signal indicated by the arrow in Fig. 2. In the experimental study shown in Fig. 12, for example, the background gradients were very strong, and portions of the signal in k , t reached the acquired k -radius in about 23 ms. For this reason, only 23 ms of data were used in the estimation. This necessary temporal data truncation may be responsible for the bias errors (i.e., artifacts) seen in the parameter maps of Fig. 12. Future developments may greatly reduce this problem. Continuous-voxel signal models may be developed to reduce or eliminate model errors occurring when the acquisition time exceeds the k -exit time, though in this case continued acquisition would add no more information regarding the problem region. Furthermore, higher-resolution multi-shot and modified single-shot rosette acquisition schemes with expanded (even possibly undersampled) k -space coverage may be useful in acquiring more signal (and making better parameter estimates) from regions with large frequency gradients.

E. Conclusions

SS-PARSE is a single-shot example of a potentially larger group of PARSE MRI methods using parameter estimation based on inverse solution of a parameterized signal model. Immediate extensions of the SS-PARSE method (currently under investigation) include 1) higher-resolution multiple-shot PARSE, 2) spin-echo SS-PARSE for simultaneous frequency, $|M_{\perp}|$, R_2 and R_2' estimation, and 3) parallel acquisition SS-PARSE, which should permit higher spatial-resolution imaging and in principle, smaller estimation errors. In principle, PARSE allows the flexibility to estimate any intra-signal temporal variations which can be modeled parametrically in the signal model. For example, intra-signal information is encoded experimentally in phase velocity encoding and directional diffusion measurements. Recently, a velocity-estimating version of SS-PARSE has been demonstrated [3], [4].

For any potential PARSE method to be feasible, sufficient information to support the estimation must be encoded into the signal or signals. If no *a priori* information is introduced (as was the case in these studies of SS-PARSE), a basic requirement is that the degrees of freedom in the data (usually, the number of data samples) are not fewer than the degrees of freedom in the unknown signal parameters. A more complete evaluation of a method's basic robustness and its accuracy limits requires analysis of the conditioning of the inverse problem.

We have found the PLCG algorithm to be a convenient approach to obtain the LS solution in the signal domain. Other algorithmic approaches to obtain the signal-domain LS solution are possible and may have some advantages, though the PLCG appears to successfully avoid secondary minima in the cost function under typical circumstances. We have described here some basic properties of the LS signal-domain solution, which will affect the performance of any optimization algorithm searching for the cost function minimum, whether it be PLCG or not.

Clearly, there are several practical applications in which estimating local magnetization magnitude and net relaxation rate R_2^* have potential clinical value. The ability to estimate

local frequency precisely in a single shot also has a number of potential applications. It may provide a clear signature of draining veins, permitting them to be excluded from functional MRI maps. Measurement of rapid but very small local changes in frequency may also be useful in neural current MRI (nc-MRI) [31], a method with the potential to provide functional MRI maps at very high temporal resolution, more directly demonstrating neural activity. Extensions of precise frequency measurement to multiple-shot PARSE methods with higher spatial resolution may provide a useful tissue contrast in susceptibility imaging, especially at higher field strength [32].

Further development and study of SS-PARSE is called for, but the results at this relatively early stage are encouraging. It offers significant practical advantages as an acquisition technique for BOLD fMRI and other applications which monitor dynamic changes in transverse relaxation rates. At a more basic level, the analysis of the behavior of SS-PARSE presented here demonstrates the advantages of choosing a (k, t) trajectory based on its ability to encode the desired information rather than its suitability for subsequent discrete Fourier transformation to an image format.

This suggests a general observation. From the standpoint of optimal estimation, it is clear that the prevailing paradigm for MRI methods based on a static signal model and an FT-based decoding algorithm is a compromise that does not provide the best possible results in many applications. It seems more likely that the goal of estimating local NMR signals, a goal at the heart of all MRI techniques, would be better served in many (if not most) cases by 1) generating signals in a manner providing optimal encoding of the desired information, and then 2) using well-chosen parameter estimation algorithms to efficiently estimate the parameters to be measured. As computational power becomes continually cheaper and more readily available, the computational cost of such optimal or near-optimal estimation-based methods becomes an increasingly poorer argument against their adoption.

Finally, it should be emphasized that SS-PARSE and other PARSE methods can be used in any modern conventional MRI system without hardware modifications.

Acknowledgments

This work was supported in part by the National Institutes of Health under Grant R21/R33 EB003292, and the Alabama Eyesight Foundation.

References

1. Twieg DB. Parsing local signal evolution directly from a single-shot MRI signal: A new approach for fMRI. *Magn Reson Med*. Nov; 2003 50(5):1043–52. [PubMed: 14587015]
2. Ogawa S, Menon RS, Tank DW, Kim SG, Merkle H, Ellerman JM, Ugurbil K. Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging: A comparison of signal characteristics with a biophysical model. *Biophys J*. 1993; 64(3):803–812. [PubMed: 8386018]
3. Speck O, Hennig J. Functional imaging by 10- and T2*-parameter mapping using multi-image EPI. *Magn Reson Med*. Aug; 1998 40(2):243–8. [PubMed: 9702706]
4. Posse S, Wiese S, Gembris D, Mathiak K, Kessler C, Grosse-Ruyken ML, Elghawaghi B, Richards T, Dager SR, Kiselev VG. Enhancement of BOLD-contrast sensitivity by single-shot multi-echo functional MR imaging. *Magn Reson Med*. 1999; 42:87–97. [PubMed: 10398954]
5. Barth M, Metzler A, Klarhoefer M, Roll S, Moser E, Leibfritz D. Functional MRI of the human motor cortex using single-shot, multiple gradient-echo spiral imaging. *Magn Reson Imag*. 1999; 17:1239–43.
6. Sutton BP, Noll DC, Fessler JA. Dynamic field map estimation using a spiral-in/spiral-out acquisition. *Magn Reson Med*. 2004; 51(6):1194–204. [PubMed: 15170840]

7. Olafsson VT, Noll DC, Fessler JA. Fast joint reconstruction of dynamic R_2^* and field maps in functional MRI. *IEEE Trans Med Imag.* Sep; 2008 27(9):1177–1188.
8. Duyn JH, van Gelderen P, Talagala L, Koretsky A, de Zwart JA. Technological advances in MRI measurement of brain perfusion. *J Magn Reson Imag.* Dec; 2005 22(6):751–3.
9. Wintermark M, Sesay M, Barbier E, Borbely K, Dillon WP, Eastwood JD, Glenn TC, Grandin CB, Pedraza S, Soustiel JF, Nariai T, Zaharchuk G, Caille JM, Dousset V, Yonas H. Comparative overview of brain perfusion imaging techniques. *J Neuroariol.* Dec; 2005 32(5):294–314.
10. Williams DS, Detre JA, Leigh JS, Koretsky AP. Magnetic resonance imaging of perfusion using spin inversion of arterial water. *Proc Nat Acad Sci USA.* Jan; 1992 89(1):212–6. [PubMed: 1729691]
11. Williams DS. Quantitative perfusion imaging using arterial spin labeling. *Methods Mol Med.* 2006; 124:151–173. [PubMed: 16506421]
12. Zuo J, Walsh EG, Deutsch G, Twieg DB. Rapid mapping of flow velocity using a new parse method. *Magn Reson Med.* 2006; 55:147–152. [PubMed: 16315204]
13. Zuo J, Bolding M, Twieg DB. Validation of velocity-SS-PARSE for single-shot flow measurement. *Magn Reson Imag.* Apr; 2007 25(3):335–40.
14. Twieg DB. The k -trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods. *Med Phys.* 1983; 10:610–621. [PubMed: 6646065]
15. Ahn CB, Kim JH, Cho ZH. High-speed spiral-scan echo planar NMR imaging-I. *IEEE Trans Med Imag.* Mar; 1986 5(1):2–7.
16. Block KT, Frahm J. Spiral imaging: A critical appraisal. *J Magn Reson Imag.* Jun; 2005 21(6): 657–68.
17. Mansfield P. Multi-planar image formation using NMR spin echoes. *J Phys C.* 1977; 10:L55–8.
18. Farzaneh F, Riederer S, Pelc N. Analysis of T2 limitations and off-resonance effects on spatial resolution and artifacts in echo-planar imaging. *Magn Reson Med.* 1990; 14:123–139. [PubMed: 2352469]
19. Jezzard P, Clare S. Sources of distortion in functional MRI data. *Human Brain Mapp.* 1999; 8(2–3):80–5.
20. Haacke, EM.; Brown, RW.; Thompson, MR.; Venkatesan, RV. *Magnetic Resonance Imaging: Physical Principles and Sequence Design.* New York: Wiley-Liss; 1999. p. 285–287.
21. Noll DC. Multishot rosette trajectories for spectrally selective MR imaging. *IEEE Trans Med Imag.* Aug; 1997 16(4):372–377.
22. Zhang YT, Hetherington HP, Stokely EM, Mason GF, Twieg DB. A novel k -space trajectory measurement technique. *Magn Reson Med.* 1998; 39:999–1004. [PubMed: 9621924]
23. Press, WH.; Teukolsky, SA.; Vetterling, WT.; Flannery, BP. *Numerical Recipes in C. 2.* New York: Cambridge Univ. Press; 1992.
24. Bertero, M.; Boccacci, P. *Introduction to Inverse Problems in Imaging.* Philadelphia, PA: IOP; 1998. p. 164
25. Haacke, EM.; Brown, RW.; Thompson, MR.; Venkatesan, RV. *Magnetic Resonance Imaging: Physical Principles and Sequence Design.* New York: Wiley-Liss; 1999. p. 340
26. Melsa, JL.; Cohn, DL. *Decision and Estimation Theory.* New York: McGraw-Hill; 1978. p. 232–233.
27. Seber, GAF.; Wild, CJ. *Nonlinear Regression.* Hoboken, NJ: Wiley; 2003.
28. Deshpande, H.; Twieg, DB. Temporal variability of R_2^* estimates in single-shot methods. *Proc. ISMRM Annu. Sci. Meeting; Seattle, WA.* May 2006; p. 3375
29. Tang W, Reeves SJ, Twieg DB. Fast joint estimation of local magnitude, decay, and frequency from single-shot MRI. *Proc SPIE.* 2007; 6498:649818.
30. Erhardt JC. MR data acquisition and reconstruction using efficient sampling schemes. *IEEE Trans Med Imag.* Sep; 1990 9(3):305–9.
31. Harel H, Ugurbil K, Uludag K, Yacoub E. Frontiers of brain mapping using MRI. *J Magn Reson Imag.* 2006; 23:945–957.

32. Duyn JH, van Gelderen P, Li TQ, de Zwart JA, Koretsky AP, Fukunaga M. High-field MRI of brain cortical substructure based on signal phase. *Proc Nat Acad Sci USA*. Jul; 2007 104(28): 11796–801. [PubMed: 17586684]

Fig. 1. Schematic k -space plot. The rosette k -trajectory used in much of this work extends to radius $k = 2.82 \text{ cm}^{-1}$. A portion of the rosette is shown here to indicate the k -radius of its coverage. The square frequency bands (evaluation bands) for three possible evaluation grids are shown: inscribed (52×52 spatial grid), intermediate (64×64), and circumscribed (72×72).

Fig. 2.

Effect of background gradient on signal in k -space. In an object with a linear background gradient, transverse magnetization $M_{\perp}$ becomes linearly dephased as time progresses, corresponding to ever higher spatial frequencies, as indicated (a) in real-part images of $M_{\perp}$ at different times during the signal. Progressively tighter linear dephasing in time seen in (a) corresponds to linear drift of the spatial frequency distribution of the object (bright feature) seen in the instantaneous Fourier transforms shown directly below in (b). (c) Perspective plot of signal amplitude in k -space along earlier and later sections of the rosette k -trajectory, corresponding to drift seen in (b). The discrete impulsive model represented in the $N \times N$ array used in both PARSE and conventional MRI has sidebands in k -space. (d) The actual detected signal, which originates from a continuous object, does not have sidebands, so a sideband signal does not intrude from above as it does in (c). (e) A central cross-section in k , t space indicating the drift of the discrete object signal, including the drifting of the aliased object (indicated by arrow) at the extreme upper right. The discrepancy between the experimental signal (d) and the modeled signal (c) introduces estimation errors, which become severe if later sections of the signal are included in the analysis.

Fig. 3.

Comparison of bias errors and random errors (temporal standard deviation or TSD) in the parameter maps from a profile through row 35 the water section of the four-tube phantom, (a) Superposition of the individual $M_{\perp}$ map from sequential acquisitions 3 through 30. (b), (c) Superimposed row 35 of the R_2^* maps and frequency maps, respectively. (d) Frequency profile after removal of global frequency jitter from the serial frames. Temporal variability in all maps is smaller than spatial variability. Note $M_{\perp}$ and R_2^* estimates have rougher profiles than frequency, representing larger bias errors for those parameters. Roughness can be reduced by regularization, which was not used in this example.

Fig. 4.

Behavior of sum square error (SSE) cost function for parameters $|M|$ and R_2^* at a single pixel in 5-cm-diameter disc, for three progressive FID data lengths. Actual frequency = 10 Hz, $R_2^* = 30 \text{ s}^{-1}$, $|M| = 1.0$ (arbitrary units). Cost function surface plots and contour plots for 5 ms (a)–(c), 20 ms (d)–(f), and 65 ms (g)–(i). (c), (f), and (i) are contour plots over limited range to demonstrate secondary minima at low $|M|$ and decreasing frequency offsets with increasing length in (f) and (i). By beginning the CG search over very short FID durations, and continuing the search while gradually extending the data length, convergence can be obtained over a broad bandwidth. (j), (k). SSE surface plot for parameters $|M| \cdot v \cdot R_2^*$ for the same pixel as represented in the $|M| \cdot v \cdot f$ plots in (g) and (h). For (a)–(k), data SNR = 500, equivalent to image-domain SNR = 65. (l) SSE limited-range contour plot for 65 ms data with SNR = 200, equivalent to image-domain SNR = 26, demonstrates increased depth of secondary minima with increasing noise.

Fig. 5.

Late and early convergence histories for pixels in simulated phantom, showing convergence of R_2^* estimates to biased values: (a) $|M|$ map with three tracked pixels marked blue, black, and red, (b) R_2^* estimate histories during late stages of convergence to very small stopping tolerance (from iteration 150 to iteration 3205), (c) $|M_\perp|$ estimate histories, and (d) frequency estimate histories during the same period. Signal SNR = 1000, image equivalent SNR ≈ 340 . Iterations were stopped at $\Delta Q/Q = 10^{-8}$. (e) and (f) Early convergence histories (iterations 1–250) for $|M_\perp|$ and frequency, demonstrating $|M_\perp|$ typically converges earlier than f .

Fig. 6.

Convergence of parameter estimates for three representative pixel locations in simulation with noise. (a)–(c) Partial cost function contours for R_2^* versus magnetization magnitude, with actual parameter value (open circle), final estimate (square) with convergence path (blue) indicated by segments between successive intervals of 50 iterations, for test pixels 1–3, respectively. (d) Frequency versus magnetization magnitude for pixel 1. Measurement noise level for a – d = 0.001 (equivalent image SNR = 168). (e), (f) Plots of R_2^* versus $|M_\perp|$ and frequency versus $|M_\perp|$, respectively, for pixel 1 with rms noise = 0.01 (equivalent image SNR = 16.8). Scales in (e) and (f) are expanded by a factor of 3 relative to (a)–(d). Note PLCG converges to near cost function (SSE) minimum in all cases, and SSE minima are displaced from the respective actual values (by model errors and by noise). Contours represent increments of cost function Q (as fraction of minimum value): $\Delta Q/Q = 1 \times 10^{-4}, 2 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}, 2 \times 10^{-3}$.

Fig. 7.

Influence of relative pixel amplitude and noise level on convergence in noisy simulated phantom. Frequency convergence has occurred if frequency estimate is near actual value of 8 Hz. Frequency estimates within center disc for (b) 52×52 , (c) 64×64 , (d) 72×72 evaluations versus rms noise level as fraction of maximum FID amplitude, for relative center disc amplitudes of 0.5 (solid curve), 0.2 (long dash), 0.1 (dot dash), 0.05 (short dash). Convergence of PLCG fails for low relative intensities at high noise levels. Note that local SNR is not the sole determinant of the failure to converge. An additional influence is the ratio of the local signal amplitude to other signal amplitudes present in the object.

Fig. 8.

Comparison of temporal standard deviations (TSD) of R_2^* estimates, (a) TSD predicted by CRLB for the PLCG solution given MEPI data (blue bars) and rosette data (red bars) for R_2^* values of 15, 25, 35, and 45 s⁻¹. Data SNR = 2000 (equivalent image SNR = 101). (b) TSD predicted for data SNR = 500 (equivalent image SNR = 26). (c) TSD for R_2^* estimates observed in Monte Carlo simulation for MEPI with pixel-wise fit (blue bars), PLCG estimation from “raw” MEPI data (green bars), and SS-PARSE (red bars), for R_2^* values of 15, 25, 35, and 45 s⁻¹. Signal SNR = 2000 (image SNR = 101). (d) Corresponding TSD observed for signal SNR = 500 (equivalent image SNR = 26).

Fig. 9.

Comparison of estimate maps with different array sizes and signal sources: Each column has $|M_{\perp}|$ map at top, R_2^* map in middle, and frequency map at bottom. (a)–(c) are 52×52 maps (evaluation band is inscribed in acquisition band), and (d)–(f) are 72×72 maps (evaluation band circumscribes the acquisition band). (a) and (d) are estimated from noiseless simulated data from an artificial spatially impulsive object, (b) and (e) from noiseless simulated spatially continuous object, and (c) and (f) are from experimental data. Characteristics of 64×64 maps are intermediate between these. Note roughness and ringing artifacts in 72×72 maps, especially noticeable in the R_2^* maps.

Fig. 10.

Comparison of 64^2 estimate maps from experimental phantom with different convergence tolerances: (a) 10^{-2} , (b) 10^{-3} , and (c) 10^{-4} . Each column has $|M_{\perp}|$ map at top, R_2^* map in middle, and frequency map at bottom. Cases (d)–(f) used tolerance 10^{-3} , but different approaches were used to reduce roughness and checkerboard artifacts. (d) Used iteration-by-iteration bandlimiting to remove spatial frequencies in the estimate outside the acquisition band. (e) Regularization factors $a_1 = a_2 = 10$, (f) regularization factors $a_1 = a_2 = 200$. Note that checkerboard-like artifacts appear in the R_2^* and frequency maps (white arrows) at the smallest convergence tolerance.

Fig. 11.

Effects of local frequency misestimation. Poor choice of initial frequency search value for maps in lower row leads to bias errors (artifacts). (a), (b), and (c) are M_{xy} , frequency, and R_2^* estimate maps for a typical four-tube phantom study. When a poorly chosen offset frequency was applied to the FID before the search, the PLCG failed to converge to the proper frequency along the upper right margin of the phantom, as seen in (e). As a result, both the M_{xy} and R_2^* maps contain cloudy artifacts (bias errors) seen in magnetization magnitude map (d), both near (solid arrows) and remote from (dashed arrows) the location of the frequency error. Artifacts outside the phantom are inconsequential, but indicate similar low-level artifacts may be present in the interior of the phantom as well.

Fig. 12.

Intentionally misshimmed phantom. SS-PARSE maps (a) $|M|$, (b) frequency, (c) R_2^* , (d) pseudo- T_2^* -weighted image (TE = 25 ms). (e) T_2^* -weighted EPI image (TE = 28ms) of same cross section, and (f) frequency profile across row 35 of phantom. Note severe distortion, dropout, and $N/2$ ghosts in EPI image.

Fig. 13.

Comparison of single-shot (SS-PARSE) parameter maps on right with parameter maps computed by nonlinear least-squares fit to a series of standard 128×128 T_2^* -weighted complex gradient echo (GE) images, on left, (a) 128×128 R_2^* map computed from GE images, (b) 64×64 R_2^* map by SS-PARSE, (c) 128×128 frequency map computed from GE images (masked to exclude meaningless exterior values), and (d) 64×64 frequency map computed by SS-PARSE.

Fig. 14.

Comparison of standard 128×128 T_2^* -weighted gradient echo images (a) $TE = 20\text{ms}$, (c) $TE = 40\text{ ms}$) with SS-PARSE pseudo- T_2^* -weighted images [(b) and (d), respectively] at same echo times, computed from $|M_{\perp}|$ and R_2^* estimate maps.

Fig. 15.

SS-PARSE transaxial parameter maps from awake macaque: (a) $|M|$, (b) frequency, (c) R_2^* , (d) pseudo- T_2^* -weighted image ($TE = 25$ ms) computed from $|M|$ and R_2^* maps. Frequency and R_2^* maps are masked to suppress artifactual features outside the head, which would have dominated the image scaling, obscuring the meaningful detail within the head. Note high-level artifactual features outside head in the $|M|$ map have large R_2^* values, decay rapidly and are largely absent in the 25-ms pseudo- T_2^* -weighted map (d).

TABLE I

Means and Temporal Standard Deviations (TSD) of Parameter Estimates for Experimental Phantom Study From Regions of Interest in Water and Sephadex/gel Tube

	GEMS		SS-PARSE	
parameter	Water	Tube	Water Mean $\pm$ TSD ^I	Tube Mean $\pm$ TSD
M _{xy} (arb. units)	1.48	1.36	$9.73 \times 10^{-4} \pm 2.11 \times 10^{-6}$	$7.88 \times 10^{-4} \pm 2.57 \times 10^{-6}$
R ₂ * (sec ⁻¹)	0.27	18.87	2.33 ± 0.27	14.86 ± 0.59
f(Hz)	-17.07	-5.69	-22.68 ± 0.22	-5.71 ± 0.28

^ITSD = mean of temporal standard deviations for pixels within the region of interest.