Published in final edited form as:

IEEE Trans Med Imaging. 2010 February ; 29(2): 273–281. doi:10.1109/TMI.2009.2023787.

Nonlinear Regularization for Per Voxel Estimation of Magnetic Susceptibility Distributions from MRI Field Maps

Bryan Kressler [Student Member, IEEE],

Department of Biomedical Engineering, Cornell University, Ithaca, NY, 14853 USA.

Department of Radiology, Weill Cornell Medical College, New York, NY, 10021, USA.

Ludovic de Rochefort,

Department of Radiology, Weill Cornell Medical College, New York, NY, 10021, USA.

Tian Liu [Student Member, IEEE],

Department of Biomedical Engineering, Cornell University, Ithaca, NY, 14853, USA.

Department of Radiology, Weill Cornell Medical College, New York, NY, 10021, USA.

Pascal Spincemaille [Affiliate Member, IEEE],

Department of Radiology, Weill Cornell Medical College, New York, NY, 10021, USA.

Quan Jiang, and

Department of Neurology, Henry Ford Hospital, Detroit, MI, 48202, USA.

Yi Wang [Member, IEEE]

Department of Biomedical Engineering, Cornell University, Ithaca, NY, 14853, USA.

Department of Radiology, Weill Cornell Medical College, New York, NY, 10021, USA.

Abstract

Magnetic susceptibility is an important physical property of tissues, and can be used as a contrast mechanism in magnetic resonance imaging. Recently, targeting contrast agents by conjugation with signaling molecules and labeling stem cells with contrast agents have become feasible. These contrast agents are strongly paramagnetic, and the ability to quantify magnetic susceptibility could allow accurate measurement of signaling and cell localization. Presented here is a technique to estimate arbitrary magnetic susceptibility distributions by solving an ill-posed inversion problem from field maps obtained in an MRI scanner. Two regularization strategies are considered, conventional Tikhonov regularization, and a sparsity promoting nonlinear regularization using the ℓ_1 norm. Proof of concept is demonstrated using numerical simulations, phantoms, and in a stroke model mouse. Initial experience indicates that the nonlinear regularization better suppresses noise and streaking artifacts common in susceptibility estimation.

Index Terms

Magnetic susceptibility; Deconvolution; Magnetic resonance imaging; Quantification; Magnetic field; Inversion

I. INTRODUCTION

The ability to accurately estimate magnetic susceptibilities from MRI data has a number of potential applications. Many tissues possess susceptibilities that differ from their surrounds. For example, calcium possesses a negative susceptibility relative to water, so that accurate estimation of the susceptibility in calcified bone could provide a measure of bone mineral

density [1]. Being able to accurately measure bone tissue properties would be novel for MRI because calcified bone water has a very short transverse relaxation time (< 1 ms) [2], appearing as a dark region in the image.

Considerable interest has developed in tracking labeled cells with MRI [3]–[5]. Iron labeled cells typically appear as dark regions because of strong T_2^* effects, and iron quantity is difficult to assess from such images. If the iron concentration can be accurately quantified from estimated susceptibility distributions, it could allow accurate assessment of the number of labeled cells in a region. This is potentially useful for stem cell therapies, where susceptibility quantification could be used to monitor cell placement and division [6].

A number of previous techniques have addressed quantification of magnetic susceptibility from MR images [7]–[15]. Several of these techniques require geometrical assumptions about the susceptibility distribution of interest [7]–[10], or have been limited in that they assume regions of uniform susceptibility situated in a background with uniform susceptibility [11], [12]. Another technique allowed arbitrary susceptibility distributions [13] but required that the voxels analyzed be selected before quantification. These techniques yield good results because they use powerful forms of prior knowledge. However, the information provided by these priors must be accurate, or incorrect results will be produced. In many imaging situations, these priors may not be applicable. Another technique has been developed to estimate the susceptibility of every voxel without prior information, but requires that the object be imaged in multiple orientations [14]. Application of this technique in human imaging would be challenging because of the difficulty of rotating the subject in an MRI scanner.

Recent work has attempted to devise an inversion scheme to estimate the susceptibility of each voxel in an imaging volume [15]. Here we explore this technique in more detail, showing that the system being inverted is ill-conditioned, requiring the use of prior information to achieve a stable approximate inversion. We examine the application of standard Tikhonov regularization, as well as non-linear regularization based on the ℓ_1 norm, and show that in cases in which the susceptibility distribution being estimated is sparse, the ℓ_1 based regularization technique yields better background noise suppression and more consistent estimation of foreground susceptibilities.

A. Magnetic Susceptibility

A material with a magnetic susceptibility different from vacuum alters the magnetic field around it, resulting in a variation in resonance frequency which appears as a spatially varying phase in the image. In the forward problem of estimating the field variation given a susceptibility distribution, the susceptibility distribution is convolved with the response of a dipole, as

$$\Delta B_z(\mathbf{r}) = \frac{1}{4\pi} \int \chi(\mathbf{r}') \frac{3\cos^2\theta - 1}{|\mathbf{r}' - \mathbf{r}|^3} d^3\mathbf{r}' \quad (1)$$

where χ is the spatial susceptibility distribution, $\mathbf{r}$ is the spatial coordinate vector, and θ is the angle between $\mathbf{r}' - \mathbf{r}$ and the applied field [13]. ΔB_z is the relative field shift, given by

$$\Delta B_z(\mathbf{r}) = \frac{B_{meas}(\mathbf{r}) - B_0}{B_0} \quad (2)$$

where B_{meas} is the measured magnetic field and B_0 is the applied field. In Fourier space, this operation becomes point-wise multiplication with a kernel [16], [17]

$$\Delta B_z(\mathbf{r}) = \mathcal{F}^{-1} \left(\left(\frac{1}{3} - \frac{k_z^2}{k^2} \right) \mathcal{F}(\chi) \right) \quad (3)$$

where $k^2 = k_x^2 + k_y^2 + k_z^2$.

B. Ill-conditioned Deconvolution

It has been suggested that the susceptibility distribution can be recovered by multiplying by the inverse of the kernel [18], [19]

$$\chi(\mathbf{r}) = \mathcal{F}^{-1} \left(\left(\frac{1}{3} - \frac{k_z^2}{k^2} \right)^{-1} \mathcal{F}(\Delta B) \right) \quad (4)$$

However, this problem is ill-posed because of the zeros in the k-space filter (i.e. points where $k^2 = 3k_z^2$, the “magic angle”). Haacke, et. al. [19] note that once this problem is discretized it is no longer ill-posed because the sample points can be chosen avoid the zeros of the filter. This is consistent with the mathematical literature because the matrix is nonsingular and therefore invertible. However, the resulting discrete problem will be ill-conditioned [20], resulting in severe noise amplification.

The discrete forward system can be written as

$$\mathbf{C}\chi = \psi \quad (5)$$

where χ is the vector representation of the discretized susceptibility distribution, ψ is the vector representation of the sampled field map, and $\mathbf{C}$ is the matrix representation of convolution with the dipole response. Here $\mathbf{C}$ is the matrix that must be inverted to estimate the susceptibility distribution given a measured field map.

In analyzing error propagation in discrete inverse problems, a useful quantity is the condition number $\kappa(\mathbf{C}) = |\lambda_{\max}|/|\lambda_{\min}|$, the ratio of the largest and smallest eigenvalues of the matrix being inverted. In the presence of noise, the actual measurement is given by $\hat{\psi} = \psi + \delta\psi$, where $\delta\psi$ is a vector representing the noise in the measured data. When inverting the system, the hope is that $\hat{\chi} = \chi + \delta\chi = \mathbf{C}^{-1}(\psi + \delta\psi)$ is close to the true χ , where $\delta\chi$ is the noise vector transformed to the space of the solution. To analyze noise propagation, $\mathbf{C}^{-1}\delta\psi$, the noise as it appears in the estimated solution, must be bounded. The normalized error in the estimated χ can be bounded as [21]

$$\frac{\|\delta\chi\|_2}{\|\chi\|_2} \leq \kappa(\mathbf{C}) \frac{\|\delta\psi\|_2}{\|\psi\|_2}. \quad (6)$$

A large condition number indicates severe noise propagation. If the convolution involved is taken to be circular, $\mathbf{C}$ becomes circulant, such that

$$\mathbf{C} = \mathcal{F}^H \mathbf{D} \mathcal{F} \quad (7)$$

where $\mathcal{F}$ represents a matrix whose rows are the Fourier basis vectors, and $\mathbf{D}$ is a diagonal matrix whose diagonal elements are the coefficients of the Fourier transform of the convolution kernel. Note that since an appropriately normalized Fourier matrix gives an orthonormal set of vectors, this gives a simple closed form for the eigendecomposition of $\mathbf{C}$. The sample points of the convolution kernel will in general be very close to the k-space zeros, resulting in a very large condition number. The simple inversion scheme works for

noise free datasets, but will quickly break down in the presence of even small amounts of noise.

Figure 1 demonstrates noise amplification for inversion of a simulated susceptibility distribution. Given the true susceptibility distribution (Fig. 1(a)), the true field shift map is computed (Fig. 1(b)) and a small amount of noise is added (Fig. 1(c)). The susceptibility distribution producing this field shift is estimated by inversion (Fig. 1(d)) and its field map is computed (Fig. 1(f)). Notice the field map induced by the estimated susceptibility distribution almost perfectly matches the noisy field map from the true susceptibility distribution (the residual is shown in Fig. 1(e)), but the noise is much more severe in the estimated susceptibility distribution than in the original field map.

II. REGULARIZATION

The discrete problem (5) can be formulated in terms of a least squares problem

$$\min_{\chi} \|C\chi - \psi\|_2^2 \quad (8)$$

where χ is the solution vector, C is the matrix representation of the convolution, ψ is the measured relative field map, and $\|y\|_2$ denotes the Euclidian ℓ_2 norm of y defined as

$\|y\|_2 = \sqrt{y^T y}$. An issue that arises in practice is that reliable relative field information may not be available at all points in an image, but a weighting term can be added to the least squares problem to compensate for nonuniform noise properties. Consider the diagonal matrix W which has elements derived from the inverse of the standard deviation of the estimated frequency. The least squares problem then becomes

$$\min_{\chi} \|W(C\chi - \psi)\|_2^2. \quad (9)$$

A. Tikhonov Regularization

A standard approach to applying prior knowledge to invert ill-conditioned systems is regularization. A common form of regularization is Tikhonov regularization. Given the standard least squares problem as defined in (9), the Tikhonov regularized version is given by

$$\min_{\chi} \|W(C\chi - \psi)\|_2^2 + \alpha^2 \|L\chi\|_2^2 \quad (10)$$

where α is a tunable regularization parameter and L is a regularization matrix. The solution of this system can be found by finding the zero of the gradient of the normal equation,

$$(C^T W^T W C + \alpha^2 L^T L) \chi = C^T W^T W \psi. \quad (11)$$

This can be solved using standard linear systems techniques.

An intuition into the action of the regularizer for deconvolution can be gained by examining the k-space formulation of the problem in the case where $L = W = I$, the identity matrix. Using the property that $\mathcal{F}^H \mathcal{F} = \mathcal{F} \mathcal{F}^H = I$, Tikhonov regularization in k-space becomes

$$(D^T D + \alpha^2 I) \mathcal{F} \chi = D^T \mathcal{F} \psi. \quad (12)$$

The diagonal matrix $(\mathbf{D}^T \mathbf{D} + \alpha^2 \mathbf{I})$ is trivially invertible, and $\mathcal{F}\chi$ is the Fourier transform of the vector being solved for. Here we see that the regularizer effectively adds a constant value onto each of the eigenvalues of $\mathbf{C}^T \mathbf{C}$. The condition number then becomes

$$k(\mathbf{C}^T \mathbf{C} + \alpha^2 \mathbf{I}) = \frac{\lambda_{\max}^2 + \alpha^2}{\lambda_{\min}^2 + \alpha^2} \quad (13)$$

which can yield a substantial improvement if $\alpha^2 \gg \lambda_{\min}^2$. In practice, the convolution is not circular, so the matrix $\mathbf{C}^T \mathbf{C}$ is not circulant. Rather, it is a Toeplitz matrix, so the Fourier decomposition does not yield the eigendecomposition. However, the eigenvalues of the Toeplitz matrix and the corresponding circulant matrix are asymptotically equivalent [21] and the inverse of the circulant matrix can be used as an effective preconditioner for inverting the Toeplitz matrix [22], thus this analysis yields some insight into the action of the regularizer.

The solution to the minimization problem given by (10) is the image χ which comes close to satisfying the system governed by (9) and has the smallest ℓ_2 norm of vectors close to satisfying the system. This indicates that large values in the derived susceptibility distribution will be heavily penalized while small values will not be heavily penalized because the penalty is quadratic. Often this is not desirable because we expect some small regions of the image to have large susceptibilities.

B. ℓ_1 Regularization

The excessive attenuation associated with Tikhonov regularization can be avoided by using a norm other than the ℓ_2 norm. The ℓ_1 norm tends to promote sparsity in the values of the voxels, where sparsity means that only a small number of voxels will have non-zero value. The ℓ_1 norm has been widely used in the field of geophysics for deconvolution of seismogram data to find sparse changes in density of the earth's crust [23], and has recently received much attention in signal processing for sparse decomposition and approximation of signals (often referred to as basis pursuit) [24]. The ℓ_1 norm is also the driving force behind the compressed sensing techniques that have recently appeared in the information theory and MRI literature [25], [26].

Here we consider regularization using the ℓ_1 norm of the susceptibility distribution by solving a problem of the form

$$\min_{\chi} \|\mathbf{W}(\mathbf{C}\chi - \psi)\|_2^2 + \nu \|\chi\|_1 \quad (14)$$

where $\|\chi\|_1 = \sum_i |\chi_i|$. We hypothesize that this formulation will work well for estimating susceptibility distributions with a small number of voxels having large susceptibilities, e.g. imaging distributions of cells that have been labeled with iron oxide particles. This problem is convex and can be formulated as a second order cone program, for which a globally optimal solution can be efficiently computed [27].

C. Determination of Regularization Parameters

One question in regularized inversion techniques is how to choose an appropriate regularization parameter ν or α^2 . A choice with a simple statistical justification is to recognize that the term $\|\mathbf{W}(\mathbf{C}\chi - \psi)\|_2^2$ is closely related to the variance of the acquired data. Given knowledge of the noise variance of the estimated frequency map, a certain deviation from the acquired data is expected when the frequency map is computed from the estimated susceptibility distribution,

$$\mathbb{E}[\|W(C\chi - \psi)\|_2] = \frac{1}{\omega_{cf}} \sqrt{\sum_i W_i^2 \sigma^2(\omega_i)} \quad (15)$$

where $\mathbb{E}[x]$ denotes the expected value of x in a probabilistic sense and ω_{cf} is the center frequency of the magnet. Division by ω_{cf} corresponds to B_0 scaling to match the units of the residual determined by the solver. The regularization parameter can be chosen to match the squared residual to the noise variance of the data, indicating in essence the solution should be faithful to the acquired data up to the noise level.

III. FIELD MAP COMPUTATION

A. Frequency Fitting

A critical step in susceptibility estimation is generating a map of the spatially varying magnetic field induced by the susceptibility distribution of interest. Here, two schemes for this estimation are discussed, one in which multiple echos are acquired at different points after the RF pulse, and one in which scan time requirements permit only a single echo to be acquired.

1) Multiple Echoes—To estimate the frequency map, the phase of each voxel across multiple TEs is first unwrapped by removing jumps larger than $\pm\pi$. The phases are then fit using least squares regression, and the slope of the resulting fit line is used as the frequency of each voxel. In performing a least squares fit, it is desirable to weight each data point to have the same variance as all other points, scaling deviations from each acquired data point uniformly. The variance of a phase image can be estimated from the variance of the corresponding complex MRI image [28]. Assuming that the variances of the real and imaginary components of the image I are equal ($\sigma^2(\text{Re}\{I\}) = \sigma^2(\text{Im}\{I\})$), the standard deviation of the noise in the phase image can be estimated as

$$\sigma(\phi) = \frac{\sigma(\text{Re}\{I\})}{|I|} \quad (16)$$

on a voxel by voxel basis. From (16), the variance of the computed phase is proportional to the inverse squared image intensity. It can be assumed that the noise in each voxel is independent and identically distributed across time points. Using linear regression weighted by the image intensity, the resonance frequency ω of a voxel can be computed as

$$\omega = \frac{(\sum |I_i|^2)(\sum |I_i|^2 t_i \phi_i) - (\sum |I_i|^2 t_i)(\sum |I_i|^2 \phi_i)}{(\sum |I_i|^2)(\sum |I_i|^2 t_i^2) - (\sum |I_i|^2 t_i)^2} \quad (17)$$

where I_i and ϕ_i are the magnitude and phase of the voxel at the i^{th} time point, respectively, and t_i is the time of the i^{th} time point.

A measure that is useful for both determining the weights used in the estimation of the susceptibility images and for determining the regularization parameter is the variance of estimated frequencies. From the covariance matrix of the frequency fitting problem [21], the variance of the estimated frequency is

$$\sigma^2(\omega) = \frac{\sum \frac{1}{\sigma^2(\phi_i)}}{\left(\sum \frac{1}{\sigma^2(\phi_i)}\right)\left(\sum \frac{t_i^2}{\sigma^2(\phi_i)}\right) - \left(\sum \frac{t_i}{\sigma^2(\phi_i)}\right)^2} \quad (18)$$

for each voxel in the fit frequency map. The quadratic dependence of the denominator on echo time indicates that noise variance can be greatly decreased by choosing widely spaced echo times.

2) Single Echo—In some cases it is not possible to acquire multiple echoes due to constraints on total acquisition time. If the signal to noise ratio is sufficiently high, accurate field maps can be obtained from a single echo. In this situation, the phase is assumed to be 0 at the center of the RF pulse (i.e. at TE = 0). The equation for determining the frequency then becomes

$$\omega = \frac{\phi}{TE}. \quad (19)$$

This simple technique poses several challenges beyond limited SNR. One challenge is the phase at time 0 may not be 0, leading to errors in the estimated frequency from variations in phase across the image. These variations can result from velocity induced phase, eddy currents, B_1 inhomogeneity, coil phase variations, and other other system imperfections. In multi-echo imaging, this phase is compensated for in the fitting process which allows an arbitrary initial phase. In single echo imaging, removal of this phase becomes part of background field removal (Section III-B), as the phase arising from effects other than susceptibility is expected to be slowly varying.

The second challenge lies in phase wrapping due to large off-resonance frequencies. The minimum TE is limited by the gradient strength and the resolution and field of view of the image being acquired. A longer TE is often desirable to yield a better phase SNR. Given TE, the maximum deviation from the center frequency (in radians/sec) that can be resolved is

$$\omega_{max} = \frac{\pi}{TE}. \quad (20)$$

In multi-echo imaging, the maximum frequency is determined by the spacing between echoes, rather than the minimum TE, allowing the maximum measurable frequency to be increased by decreasing echo spacing. One way to address wrapping in single echo imaging is to employ a spatial unwrapping algorithm. This problem is difficult to solve for two or three dimensional images, and has been an area of active research for a number of years [28]–[34]. The single echo images shown in this work were unwrapped using the phase unwrapping algorithm in [34].

In the case of a single echo acquired at time TE after the RF pulse, the equation for the frequency fit is

$$\phi = \omega TE. \quad (21)$$

From the properties of Gaussian random variables, the corresponding variance is given by

$$\sigma^2(\omega) = \frac{\sigma^2(\phi)}{TE^2} \quad (22)$$

where $\sigma(\phi)$ is calculated as before.

B. Background Field Removal

In an MR scanner, several sources of frequency variation exist, aside from those induced by susceptibility distributions. One main source is main field inhomogeneity, which is often on

the order of the frequencies induced by the presence of MR compatible objects with different susceptibilities. Another source of frequency variation is chemical shift from spins bound in molecules other than H₂O. These effects must be removed in order to accurately estimate susceptibility distributions. This work does not address chemical shift since the spin systems consist entirely of hydrogen atoms bound to oxygen in water. However, the effect of chemical shift from fat could be removed using Dixon techniques, which yield fat and water images as well an inhomogeneity map without the effects of chemical shift [29], [35]. Two separate techniques were employed to remove background fields in different imaging situations.

1) Background Field Subtraction—One way to cope with background field inhomogeneity is to derive frequency maps relative to a reference phantom [36]. Thus, all susceptibility distributions are changes relative to the susceptibility distribution of the reference phantom. Given a frequency for a particular voxel in the reference phantom ω_{ref} and a frequency in the susceptibility distribution of interest ω_{obj} the relative frequency shift is computed as

$$\psi_{rel} = \frac{\omega_{obj} - \omega_{ref}}{\omega_{cf}} \quad (23)$$

where ψ_{rel} is the desired relative field shift and ω_{cf} is the center frequency of the magnet. The relative field shift is equal to the relative frequency shift. The reference phantom is a phantom identical to the phantom of interest, but with the susceptibility distribution of interest removed.

2) High Pass Filtering—Another technique to account for the effects of background field inhomogeneity is to high pass filter the frequency maps. The background field is expected to be smooth, so a high pass filter will remove these effects. Susceptibility differences at the boundary between the phantom and surrounding air are expected to induce only a slowly varying field away from the boundary, which is also removed by a high pass filter. However, the dipole responses induced by a sparse susceptibility distribution are dominated by high frequency components and will survive the high pass filter largely unaltered. Using the high pass filter technique allows susceptibility estimation in cases where a background inhomogeneity map is not available, and would facilitate use in in vivo applications.

For this work, a simple high pass filter was implemented. A sliding average of frequency was computed for a small spherical window across the entire image. The average was weighted by the squared magnitude of the signal intensity at each voxel and subtracted from the frequency of each voxel to yield an estimate of the susceptibility induced field inhomogeneity. This process can be described as

$$\psi_{rel}(x) = \frac{\omega_{obj}(x) - \frac{\sum_{y \in S_x} W^2(y) \omega_{obj}(y)}{\sum_{y \in S_x} W^2(y)}}{\omega_{cf}} \quad (24)$$

where S_x represents the spherical shell surrounding the voxel x , and $W(y)$ is a weight proportional to the inverse of the phase noise standard deviation at each voxel. $W(y)$ and $\omega_{obj}(y)$ were taken to be zero outside of the imaged volume.

IV. SOLVER IMPLEMENTATIONS

For reasonably sized images, the systems to be inverted become large, making computing a closed form inverse to (11) computationally intractable. For a $64 \times 64 \times 64$ voxel dataset,

the matrix $\mathbf{C}^T\mathbf{C}$ is 262144×262144 in size, contains over 68.7 billion elements, and would consume half a terabyte of memory if it were explicitly formed using an 8 byte double precision data type. Using iterative inversion techniques, it is not necessary to explicitly form or invert this system. Instead, the forward system can be repeatedly applied to a residual vector to provide updates to an estimated solution. In this work, we use the conjugate gradients algorithm for the ℓ_2 based regularization because of its rapid convergence properties. Conjugate gradients can be used to solve the system given by equation (11) because the matrix to be inverted is symmetric positive definite. For the ℓ_1 regularization scheme, the standard conjugate gradients algorithm cannot be used because the equation to be solved is not a quadratic form. In this case, we use a log barrier interior point method whose details are described in Appendix A.

The forward system can be quickly applied to a vector because it consists primarily of a convolution. While straightforward multiplication of the matrix with an image that is $M \times N \times P$ points would require $O(M^2N^2P^2)$ operations, performing convolution using a fast Fourier transform requires only $O(MNP \log(MNP))$ operations. Care must be taken in performing the convolution using a discrete Fourier transform. If the DFT of an $M \times N \times P$ image is multiplied by the DFT of an $Q \times R \times S$ kernel, the resulting image has been circularly convolved with the kernel, such that the kernel wraps from one side of the image to another (indexing is modulo (M, N, P)). To avoid this, both the image and the kernel must be padded with zeros to size at least $(M+Q-1) \times (N+R-1) \times (P+S-1)$ before the DFT is computed, and in the resulting image, the appropriate $M \times N \times P$ sized region is selected. In the case of generating a phase map from a susceptibility distribution, the image is $M \times N \times P$ and the convolution kernel is $2M \times 2N \times 2P$, so the DFT must be of size at least $(3M-1) \times (3N-1) \times (3P-1)$.

To generate the dipole response function for the deconvolution, a $2M \times 2N \times 2P$ k-space response was calculated from (3). This response was then transformed back to image space, and zero padded to size $3M \times 3N \times 3P$. This image space response was then transformed back to k-space to allow efficient convolution using Fourier transforms.

V. EXPERIMENTS & RESULTS

A. Numerical Phantom

To evaluate the feasibility of the technique in situations with sparse susceptibility distributions, a numerical phantom consisting of a $32 \times 32 \times 16$ image with 2048 randomly selected indexes of known susceptibility was generated. The values for the susceptibilities at these points were chosen from a uniform random distribution between -16 ppm and 16 ppm. To simulate T_2^* effects causing signal loss near regions of high susceptibility, the weights (image intensities) assigned to each voxel with non-zero susceptibility and its 26-connected neighbors were 0.4, and the weights of all other voxels were set to 1. The susceptibility distribution was padded with zeros to $96 \times 96 \times 48$, and Fourier transformed. In k-space, the transform was multiplied by the Fourier transform of the response of a dipole for a $64 \times 64 \times 32$ image zero-padded to $96 \times 96 \times 48$. The inverse Fourier transform was computed, and the appropriate $32 \times 32 \times 16$ region was selected as the difference field map in image space. Zero mean Gaussian noise with a standard deviation of 0.02 ppm was added to the relative field map before inversion. The noise standard deviation was divided by the weight at each voxel to simulate the noise increase that is encountered with reduced signal intensity.

Susceptibilities were estimated from the noisy field map and the weight image, using both regularization techniques and a number of regularization parameters. The results are shown in Fig. 2. Figures 2(a) and 2(b) show L-curves for the two regularization techniques. Notice that neither regularization technique produces the characteristic L shape, making the L-curve

method for estimating the best regularization parameter uninformative [37]. Figures 2(c) and 2(d) show plots of the slope of a line fit to the estimated vs. true susceptibilities as the regularization parameter is varied. In both cases, overregularization produces a slope near zero, and decreasing the regularization parameter results in a smooth transition between slopes of zero and one. Figure 2(e) shows the true image, the noisy field map, and the images reconstructed by each inversion technique. Notice that for regularization parameters too large, the susceptibility distribution is excessively attenuated, while for regularization parameters too small, noise becomes more apparent.

The best regularization parameter for this data set was determined based on the noise variance of the field map. The images corresponding to these regularization parameters are outlined in black in Fig. 2. These parameters visually appear to provide a good trade off between fidelity to the true image and noise suppression. Notice in Figs. 2(a) and 2(b) these parameters occur on the flat portions of the L-curves, which is not expected in the conventional L-curve analysis [37]. In Figs. 2(c) and 2(d), the optimal regularization parameter occurs where the slope of the fit line has leveled off for the ℓ_1 regularization, while for the ℓ_2 regularization it occurs while the slope is significantly less than 1.

B. Phantom Images

A convenient way to generate a sparse susceptibility distribution using iron oxide is to print a dot pattern using a conventional laser printer. It has previously been noted that labels produced using a laser printer will produce strong signal dephasing artifacts [38], [39] because some toner manufacturers employ styrene coated magnetite (Fe_3O_4) as the ink in their toner cartridges, a compound that has a strong susceptibility. By varying the dot size over a range, dots with different total quantities of iron can be produced. It was found that a single pixel printed at 300 dpi is sufficient to produce a strong dephasing and measurable susceptibility variation. A phantom consisting of dots of sizes {2, 4, 6, 9, 12, 16, 20, 25, 30} was printed using an HP LaserJet 4050 N laser printer. For this experiment, it was assumed the susceptibilities of the paper and the styrene coating on the toner were negligible, a reasonable assumption given the magnitude of the dephasing produced by the toner. It was also assumed the total quantity of iron varies linearly with the number of pixels in each dot. In the resulting images, it was found the sheet of paper on which the dots were printed produced a decrease in signal magnitude, but did not produce a noticeable field variation.

The printer dot phantom was suspended in a water filled container using an acrylic frame located far from the dots, resulting in a nearly homogeneous field over the region where the dots were located. The phantom was imaged in a 1.5T GE Excite MR scanner using an interleaved multi-echo gradient echo pulse sequence and a single channel extremity coil. Pulse sequence parameters were matrix size = $96 \times 96 \times 24$, voxel size 1 mm isotropic, flip angle = 30° , receiver bandwidth = 62.5 KHz, 16 echos, $\text{TE}_1 = 1.8$ ms, $\Delta\text{TE} = 0.5$ ms, $\text{TR} = 50$ ms. The magnetic field was oriented along the readout direction. For susceptibility estimation, a $48 \times 48 \times 16$ voxel region in the center of the phantom containing the printed dots but excluding the acrylic frame was selected. The field inhomogeneity map was filtered using a kernel with a diameter of 11 voxels.

Results for the printed dots phantom are shown in Fig. 3. While the susceptibilities are not perfectly localized to single voxels, the total magnetization estimated for each voxel is expected to match the iron content of each dot. Thus, the integral of susceptibility over a region surrounding each dot is plotted against dot size in Fig. 3(a). Figure 3(b) shows the field inhomogeneity map, while Figs. 3(c) and 3(d) show the susceptibility maps estimated using the ℓ_2 and ℓ_1 regularizations. The regularization parameters used for display and analysis were determined from the noise variance criteria. Both techniques result in

approximately linear results, with the ℓ_1 regularization producing a higher slope. This is expected based on the results of the numerical phantoms.

To quantify the performance of the proposed technique, phantom scans were performed using varying amounts of gadopentetate dimeglumine to induce different susceptibility distributions. Gd has a susceptibility of approximately 326 ppm/M at standard temperature [40], [41], so it is possible to estimate the relative susceptibility change that should be produced by varying the concentration of Gd in a phantom.

A phantom was constructed using a Petri dish and seven small segments of soda straw. The straw segments were glued vertically to the bottom of the Petri dish, forming short vertical cylinders. The Petri dish and straws were first filled with water, and imaged using a five inch surface coil to provide a reference image. The water was then removed from each straw and replaced by { 1%, 2%, 3%, 4%, 5%, 6%, 7% } dilutions of a 0.5 M Gd solution (Magnevist, Berlex). The scan was then repeated with identical shimming, positioning, and pulse sequence parameters. Pulse sequence parameters were matrix size = $128 \times 128 \times 16$, voxel size $0.75 \times 0.75 \times 1$ mm, flip angle = 30° , receiver bandwidth = 62.5 KHz, 5 echos, $TE_1 = 1.6$ ms, $\Delta TE = 0.5$ ms, $TR = 50$ ms. The magnetic field was oriented along the phase encode direction.

In this phantom, there is a significant portion of the image that does not contain any signal. Attempting to fit a phase or frequency distribution in this area is clearly meaningless. Furthermore, the estimates of the noise variance in the phase maps (16) are only accurate for large SNR values. At points with a true signal intensity near zero, the image intensity is determined predominately by noise, so that scaling the image noise by the image intensity produces erroneous results. To alleviate this problem, the weights were thresholded at three times the noise standard deviation.

To assess the effect of varying SNR on the estimated susceptibility distributions, two reconstructions were performed. The first used 5 echoes to estimate the inhomogeneity maps, while the second used only the first 3 echoes. Images and fit lines from the reconstruction using 5 echoes are shown in Fig. 4 and results from the reconstruction using 3 echoes are shown in Fig. 5. The expected value of the residual is 7.46 ppm for the 5 echo case and 15.76 ppm for the 3 echo case, indicating that the noise standard deviation increases by a factor of greater than two when reducing the acquisition from 5 echos to 3 echos.

In the 5 echo reconstruction, the ℓ_2 norm produces an estimated susceptibility distribution with noticeable streaking artifacts (Fig. 4(d)). The ℓ_1 norm produces an estimated susceptibility distribution with some streaking artifacts (Fig. 4(e)), but these are greatly diminished compared to those from the ℓ_2 norm reconstruction. To evaluate the numerical performance of the inversion techniques, a region of interest spanning multiple slices and contained entirely within each tube was selected, and the estimated susceptibility was averaged across each region of interest. The results are shown in Fig. 4(a). Here, the ℓ_2 distribution produces a fit with a slope near 1, while the ℓ_1 norm overestimates the susceptibilities, producing a slope of 1.44. This is in contrast to the numerical simulations where the ℓ_1 based regularization produced a slope near 1.

In the 3 echo reconstruction, the reconstructed images are visually similar to those from the 5 echo reconstruction. However, the streaking artifacts in the background of the ℓ_1 regularized inversion are reduced from those in the 5 echo reconstruction. This is because the 3 echo reconstruction has a greater noise variance, making higher regularization parameters necessary. In this case, the ℓ_2 regularized reconstruction significantly underestimates the susceptibilities, producing a slope of 0.65. The ℓ_1 norm continues to

overestimate susceptibility, producing a slope of 1.43, almost the same as for the 5 echo case. This seems to indicate that while the ℓ_1 regularization overestimates the susceptibilities, it is more consistent in the slope it produces than the ℓ_2 regularization. This also indicates it is necessary to calibrate the susceptibility inversion technique if absolute values of the susceptibilities are of interest.

It is possible to achieve slopes of 1 between the estimated and actual susceptibility values for both the ℓ_2 and ℓ_1 reconstructions by appropriate choice of the regularization parameters. Shown in Fig. 6 are the resulting images when the regularization parameters are tuned to achieve a slope of nearly 1 for both the 5 and 3 echo acquisitions of the Gd tube phantom. For the ℓ_2 regularization with 5 echos (Fig. 6(a)), the slope was 1.02. For the ℓ_2 regularization with 3 echos (Fig. 6(c)), the slope was 1.03. The regularization parameters, slopes, and residuals for the estimated susceptibility distributions determined using the expected value of the residual and by matching the slope are summarized in Table I. The residuals for the cases with slopes of 1 are less than predicted by the method in Section II-C, and the underregularization is evident in the increased streaking and noise in the estimated susceptibility distributions. The ℓ_1 regularization applied to the 5 echo data set (Fig. 6(b)) achieves a slope of 1.08. For the 3 echo dataset, the ℓ_1 regularization (Fig. 6(d)) produces a slope of 1.07. The residuals for the ℓ_1 regularization are larger than those predicted by the method in Section II-C. Applying a large regularization parameter removes most streaking in the background but suppresses the tubes with low susceptibility values. While the results in Fig. 6 have slopes of 1 between estimated and true susceptibility values, determining the appropriate regularization parameter if the susceptibilities are not known *a priori* is an unsolved problem.

C. Mouse Brain Studies

To demonstrate the feasibility of the technique in a biological sample, an excised mouse brain containing superparamagnetic iron oxide particles was imaged in a Bruker 7 T small animal scanner. In this case, a single image was acquired, and phase unwrapping and high pass filtering were performed. The phase unwrapping and filtering algorithms were those described in [34]. The imaged resolution was 93.75 μm isotropic, and the image was acquired using an echo time of 10 ms. **[More information about mice from Quan Jiang].**

Susceptibility distributions estimated for the mouse brain are shown in Fig. 7. Figure 7(a) shows a magnitude map of the mouse brain, with reduced signal intensity due to T_2^* effects indicating the location of iron oxide particles. Figure 7(b) shows the relative field map, where it can be seen that the phase unwrapping and filtering break down near the boundaries of the brain, but produce good results near the center of the brain where the background susceptibility is homogeneous. A limited region ($48 \times 48 \times 32$ voxels) of the brain where most of the T_2^* effects were observed was selected for analysis using the susceptibility inversion technique. The estimated susceptibility distribution using the ℓ_2 regularization is shown in Fig. 7(c), and the estimated susceptibility distribution using the ℓ_1 regularization is shown in Fig. 7(d). Again, the regularization parameter used in the displayed susceptibility maps was estimated based on the noise variance of the field map.

In this case, validation of the numerical values obtained for the susceptibilities is not available. However, the regions of high susceptibility match closely with the signal voids visible in the magnitude images. The ℓ_1 regularization continues to produce better suppression of small susceptibility values than the ℓ_2 regularization. Both regularization techniques benefit from a regularization parameter larger than that indicated by the noise variance criteria. This is likely because the field map contains variations that are structured and therefore not due to noise, but also not produced by the injected iron oxide particles.

These result in voxels that do not contain large quantities of iron oxide particles having non-zero susceptibilities. Increasing the regularization parameter tends to suppress these voxels, while preserving the values of the iron oxide containing voxels.

VI. DISCUSSION

In this work, regularization techniques were developed to estimate sparse magnetic susceptibility distributions from MRI data. The technique was validated in numerical and physical phantoms, and initial feasibility was shown in an animal study. In the numerical phantom, the ℓ_1 regularization produced very good reconstructions of the susceptibility distribution while the ℓ_2 regularization significantly underestimated the susceptibility values. Using a phantom with known susceptibility, neither technique produced susceptibility maps with estimated susceptibility matching true susceptibility across varying noise level, but the ℓ_1 regularization technique provided better suppression of streaking artifacts at the magic angle and more consistent slope across varying noise levels than the conventional ℓ_2 based regularization. Both regularization terms produced satisfactory susceptibility images in an excised mouse brain. In all cases, the ℓ_1 regularization provided better suppression of background noise.

Several sources of error are present in the system implemented. A method for obtaining the best regularization parameter is an open question. In this application, it was found that seeking a residual matching its expected value based on noise variance in the field map was a useful criteria, while the L-curve method could not be applied. However, initial experience indicates that regularization parameters larger than indicated by the noise variance produce visually cleaner results. In part this may reflect visual preference for over regularized images because of reduced noise. It may also reflect sources of error in the system model not caused by noise.

One major source of error is phase wrapping near regions of strong susceptibility variation. The mapping of frequencies is not bijective in that frequency variations outside of the range $\left[\frac{\pi}{\Delta TE}, \frac{\pi}{\Delta TE}\right)$ are mapped into the interval $\left[\frac{\pi}{\Delta TE}, \frac{\pi}{\Delta TE}\right)$. Thus, large frequency variations are not accurately represented. This can be alleviated by reducing ΔTE , but this comes at a cost in phase SNR ((18) and (22)). In many cases, phase wrapping occurs in regions with signal voids, and these frequencies receive low weight in the inversion process. However, they will still have some adverse effect on the estimated susceptibility maps. Regrettably, these inaccuracies occur near the boundaries between regions with different susceptibilities, which should be the locations that provide the most information about the susceptibilities.

Another source of error comes from sampling and digitization effects. The model employed here assumes the susceptibility of each voxel is concentrated on a delta function centered in that voxel, and the phase of each voxel is the phase sampled on a delta function centered in that voxel. However, the susceptibility may not be centered in the voxel, or may be distributed throughout the voxel. Furthermore, the phase value in each voxel does not represent the phase at the center of that voxel. Instead, the signal is integrated over the sampling function [17], and the phase is the phase of this integrated signal.

Geometric distortion resulting from the magnetic field variations induced by the susceptibility distribution of interest is another source of error. The field variation results in local variations in precession frequency, which cause signal to be shifted from its true location along the readout direction [8]. The spatial shift can be estimated from the bandwidth per voxel, given by the receiver bandwidth divided by the number of voxels in the readout direction. For this work, it was assumed the field shifts are small relative to the

imaging bandwidth, resulting in minimal spatial distortion. However, this distortion is present and may adversely impact the estimation of susceptibilities. In the future, some mechanism to correct for this distortion should be incorporated to allow more accurate estimation and localization of susceptibility distributions.

Finally, the ℓ_1 regularization tends to overestimate the susceptibilities in the Gd doped tube phantom. Underestimation is expected in the ℓ_2 case because the magnitudes of the susceptibility values are directly penalized, but this overestimation in the case of the ℓ_1 norm is somewhat unexpected. Overestimation may result from the the ℓ_1 norm attempting to explain the field variation using as few non-zero components as possible. The regularization could be erroneously attributing field variations due to noise or other points with low susceptibilities (e.g. partial voxels along boundaries) to the points with large susceptibility. This could cause some overestimation of the susceptibility values. The overestimation could also result from the sampling effects noted earlier. Currently, the technique must be calibrated for the imaging situation if the absolute susceptibilities are of interest. However, near linearity is achieved, so relative susceptibility values are meaningful. By choosing the appropriate value for the regularization parameter ν , the ℓ_1 regularization can produce slopes of less than or equal to one, so a better method to choose the regularization parameter may at least partially rectify this situation.

Several improvements to the work presented here are necessary to make the technique applicable to biological imaging situations. Although phase filtering can remove much of the phase variation that is not due to the susceptibility distribution of interest, some variation will remain at the boundaries between air and tissue. Generalization of the ℓ_1 regularization term to total variation regularization may help address this issue. Total variation has been widely used in image denoising, and promotes sparsity of edges rather than sparsity of the estimated distribution, which may be a more realistic assumption. Another issue that must be addressed is removal of chemical shift artifacts, which give rise to resonance frequency variations that are not due to magnetic susceptibility. Existing fat/water separation techniques can yield a field map independent of chemical shift due to fat, and it may be possible to use these phase maps for susceptibility estimation. A final improvement that is necessary to allow accurate quantification of magnetic susceptibility is a method to determine the regularization parameter. In addition to the L-curve and expected residual techniques examined here, other methods to estimate a good choice of regularization parameter such as generalized cross-validation [42] exist, and should be examined in the context of susceptibility inversion. An improved method to determine the regularization parameter may lead to a slope closer to unity between the estimated and true susceptibility values.

In the future, other regularization techniques should be investigated for magnetic susceptibility estimation. Direct minimization of the ℓ_0 semi-norm might yield improved suppression of background noise because the ℓ_0 semi-norm is optimal for promoting sparsity. This problem is NP-hard, but recently efficient methods based on the cross entropy method [43] and homotopy continuation [44] have been developed to approximately minimized the ℓ_0 norm. The $\ell_{1,1}$ norm should also be investigated as a regularization parameter [45]. Although the ℓ_1 norm promotes sparsity more strongly, the $\ell_{1,1}$ norm is everywhere differentiable, simplifying the solver by eliminating the need to formulate an equivalent constrained problem (Appendix A). This could greatly improve the efficiency of susceptibility mapping by eliminating the outer most iterations of the solver, while paying only a small penalty in the effectiveness of the regularization.

VII. CONCLUSION

In this work, ℓ_1 and ℓ_2 norm based regularization techniques were developed for magnetic susceptibility estimation from MRI field maps, and were evaluated using numerical simulations, phantom images, and data from an animal study. A criteria for selecting the regularization parameter based on noise variance was evaluated and shown to produce satisfactory results. The ℓ_1 regularization provided better suppression of streaking artifacts and background noise than the ℓ_2 regularization, producing better visual image quality. Furthermore, the ℓ_1 regularization generated more consistent susceptibility values across varying noise levels, as demonstrated in phantom images. These initial results show the feasibility of magnetic susceptibility estimation without the use of strong prior knowledge, and demonstrate the advantages of nonlinear sparsity promoting regularization strategies.

APPENDIX A

LOG BARRIER METHOD FOR ℓ_1 REGULARIZATION

The ℓ_1 norm regularization problem is non-linear and does not have a well defined gradient around zero. The difficulty with the gradient can be remedied by formulating an equivalent constrained problem

$$\begin{aligned} \min_{\chi, u} & \|W(C\chi - \psi)\|_2^2 + \nu \sum u_i \\ \text{s.t. } & \chi_i - u_i \leq 0, \quad i=1 \dots m \\ & -\chi_i - u_i \leq 0, \quad i=1 \dots m \end{aligned} \quad (25)$$

where m is the number of voxels in the image. This constrained problem can then be reformulated as an unconstrained problem using a logarithmic barrier

$$\min_{\chi, u} \|W(C\chi - \psi)\|_2^2 + \nu \sum u_i - \frac{1}{\tau} \sum \log(u_i - \chi_i) - \frac{1}{\tau} \sum \log(u_i + \chi_i) \quad (26)$$

The function $-\frac{1}{\tau} \log(x)$ is referred to as the log barrier, and has a finite value when its argument x is positive and an infinite value when the argument is negative (or exactly zero), enforcing the condition that $|\chi_i| < u_i$ so that $\sum u_i$ is an upper bound on the ℓ_1 norm of χ . The parameter τ governs how well the log barrier approximates a true barrier (zero when the constraint is satisfied and infinite when the constraint is violated). Larger values of τ yield a better approximation to the constraint, but are more difficult to minimize because the function becomes less smooth.

The systems given by (25) and (26) are convex, so that the minimum can be found to finding the point where the gradient is zero. The approach used here is iterative search using Newton's Method [27], [46]. Newton's method operates by locally approximating the function being minimized using a quadratic function, and then minimizing the quadratic approximation by finding the zero of its gradient. This requires both the gradient and the Hessian of the system being solved. Concatenating the vectors χ and u , and defining

$G(\chi, u) = \|W(C\chi - \psi)\|_2^2 + \nu \sum u_i - \frac{1}{\tau} \sum \log(u_i - \chi_i) - \frac{1}{\tau} \sum \log(u_i + \chi_i)$, the gradient of (26) is given by

$$\nabla G(\chi, u) = \begin{bmatrix} 2C^T W^T r - \frac{1}{\tau} f_{u_1}^{-1} + \frac{1}{\tau} f_{u_2}^{-1} \\ \nu + \frac{1}{\tau} f_{u_1}^{-1} + \frac{1}{\tau} f_{u_2}^{-1} \end{bmatrix} \quad (27)$$

and the Hessian is given by

$$\nabla^2 G(\chi, u) = \begin{bmatrix} 2C^T W^T W C + \frac{1}{\tau} \Sigma_{11} & \frac{1}{\tau} \Sigma_{12} \\ \frac{1}{\tau} \Sigma_{12} & \frac{1}{\tau} \Sigma_{11} \end{bmatrix} \quad (28)$$

where $f_{u1} = \chi - u$, $f_{u2} = -\chi - u$, $\Sigma_{11} = \text{diag}(f_{u1})^{-2} + \text{diag}(f_{u2})^{-2}$, $\Sigma_{12} = -\text{diag}(f_{u1})^{-2} + \text{diag}(f_{u2})^{-2}$, and $r = W(C\chi - \psi)$. Here $\text{diag}(x)$ is the diagonal matrix with diagonal elements composed of the elements from x , and x^{-1} is taken to be the component-wise inverse of the elements of x . The Newton system is then

$$\begin{bmatrix} 2\tau C^T W^T W C + \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{11} \end{bmatrix} \begin{bmatrix} \Delta\chi \\ \Delta u \end{bmatrix} = - \begin{bmatrix} 2\tau C^T W^T r - f_{u1}^{-1} + f_{u2}^{-1} \\ \tau v + f_{u1}^{-1} + f_{u2}^{-1} \end{bmatrix} = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \quad (29)$$

giving the step direction $[\Delta\chi^T \Delta u^T]^T$. The system is symmetric positive definite, so it can be solved using conjugate gradients. To speed the convergence of the conjugate gradient solver, the variables Δu can be eliminated using the Schur complement [27] of the Hessian, giving the reduced system

$$(2\tau C^T W^T W C + \Sigma_{11} - \Sigma_{12}^2 \Sigma_{11}^{-1}) \Delta\chi = w_1 - \Sigma_{12} \Sigma_{11}^{-1} w_2. \quad (30)$$

The variables Δu can then be recovered from $\Delta\chi$ by

$$\Delta u = \Sigma_{11} w_2 - \Sigma_{11}^{-1} \Sigma_{12} \Delta\chi. \quad (31)$$

These $\Delta\chi$ and Δu vectors are then used as an update to the initial χ and u vectors.

The log barrier algorithm employed is summarized in Algorithm (1). The input parameters to the main function `L1sol` are an initial feasible point $(\chi^{(0)}, u^{(0)})$, an initial τ , the step size by which τ is decreased in each iteration μ , an error bound ϵ on the optimality of the solution, the regularization parameter ν , line search parameters $\alpha \in (0, 0.5)$ and $\beta \in (0, 1)$, and the number of log barrier iterations N_{barrier} to run. Any initial guess suffices for $\chi^{(0)}$. Here, 100 iterations of the conjugate gradients algorithm with no regularization are used to calculate an initial estimate of $\chi^{(0)}$. The vector $u^{(0)}$ is chosen so that

Algorithm 1

ℓ_1 regularized solver

```

1:  function L1sol( $(\chi^{(0)}, u^{(0)})$ ,  $\tau^{(0)}$ ,  $\mu$ ,  $\epsilon$ ,  $\nu$ ,  $\alpha$ ,  $\beta$ ,  $N_{\text{barrier}}$ )
2:    for  $i \leftarrow 1, N_{\text{barrier}}$  do
3:       $(\chi^{(i)}, u^{(i)}) \leftarrow \text{NEWT}((\chi^{(i-1)}, u^{(i-1)}), \epsilon, \nu, \alpha, \beta)$ 
4:       $\tau^{(i)} \leftarrow \mu \tau^{(i-1)}$ 
5:    end for
6:    return  $(\chi^{(N_{\text{barrier}})}, u^{(N_{\text{barrier}})})$ 
7:  end function

8:  function NEWT( $(\chi, u)$ ,  $\epsilon$ ,  $\nu$ ,  $\alpha$ ,  $\beta$ )
9:    repeat
10:      $\Sigma_{11} \leftarrow \text{diag}(\chi - u)^{-2} + \text{diag}(-\chi - u)^{-2}$ 
11:      $\Sigma_{12} \leftarrow -\text{diag}(\chi - u)^{-2} + \text{diag}(-\chi - u)^{-2}$ 

```

```

12:    $w_1 \leftarrow 2\tau C^T W^T r - (\chi - u)^{-1} + (-\chi - u)^{-1}$ 
13:    $w_2 \leftarrow \tau v + (\chi - u)^{-1} + (-\chi - u)^{-1}$ 
14:    $A \leftarrow 2\tau C^T W^T W C + \Sigma_{11} - \Sigma_{12}^2 \Sigma_{11}^{-1}$ 
15:    $b \leftarrow w_1 - \Sigma_{12} \Sigma_{11}^{-1} w_2$ 
16:    $\Delta \chi \leftarrow \text{CONJUGATEGRADIENTS}(A, b)$ 
17:    $\Delta u \leftarrow \Sigma_{11} w_2 - \Sigma_{11}^{-1} \Sigma_{12} \Delta \chi$ 
18:    $(\chi, u) \leftarrow \text{LINESEARCH}((\chi, u), (\Delta \chi, \Delta u), \alpha, \beta)$ 
19:   until  $-\nabla G(\chi, u)^T [\Delta \chi^T \Delta u^T]^T < 2\epsilon$ 
20:   return  $(\chi, u)$ 
21: end function

22: function LINESEARCH $((\chi, u), (\Delta \chi, \Delta u), \alpha, \beta)$ 
23:    $t \leftarrow 1$ 
24:    $\hat{\chi} \leftarrow \chi + t\Delta \chi$ 
25:    $\hat{u} \leftarrow u + t\Delta u$ 
26:   while  $\max_i(|\hat{\chi}_i - \hat{u}_i|) > 0$  do
27:      $t \leftarrow \beta t$ 
28:      $\hat{\chi} \leftarrow \chi + t\Delta \chi$ 
29:      $\hat{u} \leftarrow u + t\Delta u$ 
30:   end while
31:   while  $G(\chi, u) + \alpha t \nabla G(\chi, u)^T [\Delta \chi^T \Delta u^T]^T > G(\hat{\chi}, \hat{u})$  do
32:      $t \leftarrow \beta t$ 
33:      $\hat{\chi} \leftarrow \chi + t\Delta \chi$ 
34:      $\hat{u} \leftarrow u + t\Delta u$ 
35:   end while
36:   return  $(\hat{\chi}, \hat{u})$ 
37: end function

```

each component is slightly larger than the magnitude of the corresponding component of $\chi^{(0)}$. In this implementation, the value of $\tau^{(0)}$ is chosen so that the duality gap after the first barrier iteration $m/\tau^{(0)}$ is approximately equal to $\|\chi^{(0)}\|_1$, as in [46]. The parameters α and β are parameters for the backtracking line search described below. N_{barrier} is the number of log barrier iterations necessary to achieve the desired accuracy ϵ , and can be calculated from [27]

$$N_{\text{barrier}} = \lceil \frac{\log(2m) - \log(\epsilon) - \log(\tau^{(0)})}{\log(\mu)} \rceil. \quad (32)$$

The algorithm begins with an initial coarse approximation to the ideal barrier function using $\tau^{(0)}$. From the initial point, a Newton step (function NEWT) is calculated using the conjugate gradients solver. The Newton step indicates a search direction, but does not indicate the exact distance to step in that direction since the function is not quadratic. Computing the exact minimum in that direction is not computationally feasible, so an approximate backtracking line search is employed (function LINESEARCH) [27]. In

essence, this procedure decreases the distancing being traveled by a factor of β in each iteration, until the reduction in the objective function estimated from a linear fit at the current point is at most a factor of α less than that of the exact decrease in the function value at that point (Algorithm 1.31). Here the parameters $\alpha = 0.01$ and $\beta = 0.5$ were chosen as in [46]. Prior to the backtracking line search, a simple search is performed to ensure that the step remains within the feasible region of the problem (Algorithm 1.26). The appropriate length step is then taken, and Newton's method is then performed from the updated point until the Newton decrement is less than the desired accuracy (Algorithm 1.19). The Newton decrement

$$\left(\nabla G(\chi, u)^T \nabla^2 G(\chi, u)^{-1} \nabla G(\chi, u)\right)^{\frac{1}{2}} \quad (33)$$

provides an estimate of the duality gap derived from a local second order approximation of the function being minimized [27]. Once the Newton decrement is sufficiently small, the parameter τ is updated to better approximate the ideal barrier (Algorithm 1.4), and the process is repeated using the solution from the previous barrier iteration as the starting point for the current barrier iteration.

The log barrier is iteratively updated to trade off the accuracy of the approximation employed and the speed with which the Newton steps converge. If τ is chosen too large, the function being minimized is not accurately approximated by a quadratic, and a large number of iterations in the NEWT function are required, slowing the convergence the algorithm. If τ is chosen too small, the function minimized does not accurately approximate the ℓ_1 norm. Here, a coarse approximation (small τ) is used initially, and the solutions for coarse approximations are then used as starting points for finer approximations (larger τ). In this implementation, $\mu = 10$, so the parameter τ is increased by a factor of 10 with each iteration. Typically the number of outer iterations required for the log barrier method is 8 – 10 for a matrix size of $128 \times 128 \times 16$ and $\epsilon = 10^{-3}$.

REFERENCES

1. Chung HW, Hwang SN, Yeung HN, Wehrli FW. Mapping of the magnetic-field distribution in cancellous bone. *J Magn Reson B*. 1996; vol. 113(no. 2):172–176. [PubMed: 8948141]
2. Techawiboonwong A, Song HK, Wehrli FW. In vivo MRI of submillisecond T(2) species with two-dimensional and three-dimensional radial sequences and applications to the measurement of cortical bone water. *NMR Biomed*. 2008; vol. 21(no. 1):59–70. [PubMed: 17506113]
3. Hsiao JK, Tai MF, Chu HH, Chen ST, Li H, Lai DM, Hsieh ST, Wang JL, Liu HM. Magnetic nanoparticle labeling of mesenchymal stem cells without transfection agent: cellular behavior and capability of detection with clinical 1.5 T magnetic resonance at the single cell level. *Magn Reson Med*. 2007; vol. 58(no. 4):717–724. [PubMed: 17899592]
4. Frank JA, Miller BR, Arbab AS, Zywicke HA, Jordan EK, Lewis BK, Bryant JLH, Bulte JW. Clinically applicable labeling of mammalian and stem cells by combining superparamagnetic iron oxides and transfection agents. *Radiology*. 2003; vol. 228(no. 2):480–487. [PubMed: 12819345]
5. Shapiro EM, Skrtic S, Sharer K, Hill JM, Dunbar CE, Koretsky AP. MRI detection of single particles for cellular imaging. *Proc Natl Acad Sci*. 2004; vol. 101(no. 30):10901–10906. [PubMed: 15256592]
6. Kraitchman DL, Gilson WD, Lorenz CH. Stem cell therapy: MRI guidance and monitoring. *J Magn Reson Imaging*. 2008; vol. 27(no. 2):299–310. [PubMed: 18219684]
7. Weisskoff RM, Kiihne S. MRI susceptometry: image-based measurement of absolute susceptibility of MR contrast agents and human blood. *Magn Reson Med*. 1992; vol. 24(no. 2):375–383. [PubMed: 1569876]

8. Beuf O, Briguët A, Lissac M, Davis R. Magnetic resonance imaging for the determination of magnetic susceptibility of materials. *J Magn Reson B*. 1996; vol. 112(no. 2):111–118. [PubMed: 8812894]
9. Wang ZJ, Li S, Haselgrove JC. Magnetic resonance imaging measurement of volume magnetic susceptibility using a boundary condition. *J Magn Reson*. 1999; vol. 140(no. 2):477–481. [PubMed: 10497053]
10. Cheng YN, Hsieh C, Neelavalli J, Li Q, Dawood MS, Haacke EM. A complex sum method of quantifying susceptibilities in cylindrical objects: the first step toward quantitative diagnosis of small objects in MRI. *Magn Reson Imaging*. 2007; vol. 25:1171–1180. [PubMed: 17905248]
11. Li L. Magnetic susceptibility quantification for arbitrarily shaped objects in inhomogeneous fields. *Magn Reson Med*. 2001; vol. 46(no. 5):907–916. [PubMed: 11675642]
12. Li L, Wang ZJ. Magnetic susceptibility quantitation with MRI by solving boundary value problems. *Med Phys*. 2003; vol. 30(no. 3):449–453. [PubMed: 12674246]
13. Li L, Leigh JS. Quantifying arbitrary magnetic susceptibility distributions with MR. *Magn Reson Med*. 2004; vol. 51(no. 5):1077–1082. [PubMed: 15122694]
14. Liu, T.; Spincemille, P.; de Rochefort, L.; Kressler, B.; Wang, Y. *Proc Intl Soc Magn Reson Med*. Toronto: 2008 May. Multiple orientation acquisition to invert dipole field for quantitative susceptibility mapping; p. 643
15. Morgan, J.; Irarrazaval, P. *Proc Intl Soc Magn Reson Med*. Berlin: 2007 May. Efficient solving for arbitrary susceptibility distributions using residual difference fields; p. 35
16. Salomir R, de Senneville BD, Moonen CTW. A fast calculation method for magnetic field inhomogeneity due to an arbitrary distribution of bulk susceptibility. *Concepts Magn Reson B*. 2003; vol. 19B(no. 1):26–34.
17. Haacke, EM.; Brown, RW.; Thompson, MR.; Venkatesan, R. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. New York: Wiley-Liss; 1999.
18. Marques JP, Bowtell R. Application of a Fourier-based method for rapid calculation of field inhomogeneity due to spatial variation of magnetic susceptibility. *Concepts Magn Reson B*. 2005; vol. 25B(no. 1):65–78.
19. Haacke EM, Cheng NY, House MJ, Liu Q, Neelavalli J, Ogg RJ, Khan A, Ayaz M, Kirsch W, Obenaus A. Imaging iron stores in the brain using magnetic resonance imaging. *Magn Reson Imaging*. 2005; vol. 23(no. 1):1–25. [PubMed: 15733784]
20. Hanke M, Hansen P. Regularization methods for large-scale problems. *Surv Math Ind*. 1993; vol. 3:253–315.
21. Moon, TK.; Stirling, WC. *Mathematical Methods and Algorithms for Signal Processing*. Upper Saddle River, NJ: Prentice Hall; 2000.
22. Hanke M, Nagy J. Inverse Toeplitz preconditioners for ill-posed problems. *Linear Algebra Appl*. 1998; vol. 284:137–156.
23. Taylor HL, Banks SC, McCoy JF. Deconvolution with the ℓ_1 norm. *Geophysics*. 1979; vol. 44(no. 1):39–52.
24. Tropp JA. Just relax: convex programming methods for identifying sparse signals in noise. *IEEE Trans Inf Theory*. 2006; vol. 52(no. 3):1030–1051.
25. Candès EJ, Romberg J. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans Inf Theory*. 2006; vol. 52(no. 2):489–509.
26. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med*. 2007; vol. 58(no. 6):1182–1195. [PubMed: 17969013]
27. Boyd, S.; Vandenberghe, L. *Convex Optimization*. Cambridge: Cambridge University Press; 2004.
28. Conturo TE, Smith GD. Signal-to-noise in phase angle reconstruction: dynamic range extension using phase reference offsets. *Magn Reson Med*. 1990; vol. 15(no. 3):420–437. [PubMed: 2233221]
29. Glover GH, Schneider E. Three-point Dixon technique for true water/fat decomposition with B0 inhomogeneity correction. *Magn Reson Med*. 1991; vol. 18(no. 2):371–383. [PubMed: 2046518]

30. Chen Q, Schneider E, Aghazadeh B, Weinhaus MS, Humm J, Ballon D. An automated iterative algorithm for water and fat decomposition in three-point Dixon magnetic resonance imaging. *Med Phys.* 1999; vol. 26(no. 11):2341–2347. [PubMed: 10587215]
31. Szumowski J, Coshov WR, Li F, Quinn SF. Phase unwrapping in the three-point Dixon method for fat suppression MR imaging. *Radiology.* 1994; vol. 192(no. 2):555–561. [PubMed: 8029431]
32. Moon-Ho Song S, Napel S, Pelc N, Glover G. Phase unwrapping of MR phase images using Poisson equation. *IEEE Trans Image Proc.* 1995; vol. 4(no. 5):667–676.
33. Chavez S, Xiang QS, An L. Understanding phase maps in MRI: a new outline phase unwrapping method. *IEEE Trans Med Imaging.* 2002; vol. 21(no. 8):966–977. [PubMed: 12472269]
34. Bagher-Ebadian H, Jiang Q, Ewing JR. A modified Fourier-based phase unwrapping algorithm with an application to MRI venography. *J Magn Reson Imaging.* 2008; vol. 27(no. 3):649–652. [PubMed: 18183579]
35. Reeder SB, Wen Z, Yu H, Pineda AR, Gold GE, Markl M, Pelc NJ. Multicoil Dixon chemical species separation with an iterative least-squares estimation method. *Magn Reson Med.* 2004; vol. 51(no. 1):35–45. [PubMed: 14705043]
36. Barmet, C.; De Zanche, N.; Sasportas, LS.; Pruessmann, KP. *Proc Intl Soc Magn Reson Med.* Berlin: 2007 May. A model-free method for high-precision MR susceptometry; p. 36
37. Hansen, PC. The L-curve and its use in the numerical treatment of inverse problems. In: Johnston, P., editor. *Computational Inverse Problems in Electrocardiology*, ser. *Advances in Computational Bioengineering*. Vol. vol. 4. WIT Press; 2000.
38. Jesmanowicz, A.; Roopchansingh, V.; Cox, RW.; Starewicz, P.; Punchard, WFB.; Hyde, JS. *Proc Intl Soc Magn Reson Med.* Glasgow: 2001 May. Local ferrosims using office copier toner; p. 617
39. Hornak JP. Labels printed with magnetic toner. *Magn Reson Imaging.* 2007; vol. 25(no. 10):1459–1460. [PubMed: 17881177]
40. Lide, DR. *CRC Handbook of Chemistry and Physics*, 88th Edition (Internet Version 2008). Boca Raton, FL: CRC Press/Taylor and Francis; 2008. Magnetic susceptibility of the elements and inorganic compounds.
41. O’Handley, RC. *Modern Magnetic Materials: Principles and Applications*. New York: Wiley; 2000.
42. Bertero, M.; Boccacci, P. *Introduction to Inverse Problems in Imaging*. Bristol, UK: Institute of Physics Publishing; 1998.
43. Rubinstein RY. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability.* 1999; vol. 1:127–190.
44. Trzasko, J.; Manduca, A.; Borisch, E. *IEEE/SP 14th Workshop on Statistical Signal Processing.* Madison, WI: 2007. Sparse mri reconstruction via multiscale l0-continuation; p. 176-180.
45. Li M, Yang H, Kudo H. An accurate iterative reconstruction algorithm for sparse objects: application to 3D blood vessel reconstruction from a limited number of projections. *Phys Med Biol.* 2002; vol. 47(no. 15):2599–2609. [PubMed: 12200927]
46. Candès, E.; Romberg, J. ℓ_1 -MAGIC : Recovery of Sparse Signals via Convex Programming. Caltech, CA: 2005 Oct. [Online]. Available <http://www.acm.caltech.edu/l1magic/downloads/l1magic.pdf>

Fig. 1.

Demonstration of ill conditioned susceptibility estimation without regularization. (a) True susceptibility distribution with value 1 ppm foreground and 0 ppm background. (b) Field shift map. (c) Field shift map with zero mean additive white Gaussian noise, $\sigma^2 = 2.5 \times 10^{-5}$ ppm. (d) Estimated susceptibility distribution by inversion. (e) Residual in estimated field map relative to noisy true map. (f) Field map calculated from the estimated susceptibility distribution. (a) and (d) share the same window level and (b), (c), (e), and (f) share the same window level.

Fig. 2.

Plots and images for the numerical phantom. L-curves for the ℓ_2 (a) and ℓ_1 (b) regularized inversions. Plot of slope vs. regularization parameter for ℓ_2 (c) and ℓ_1 (d) regularized inversions. (e) True image, field map, and reconstructed images for a number of regularization parameters. The best regularization parameter (chosen to match the residual to the noise variance) is indicated with a black box.

Fig. 3.

(a) Curves showing the integral of susceptibility vs. dot size for the ℓ_1 and ℓ_2 regularized inversion techniques. (b) Field shift map. Maximum intensity projection image generated using the ℓ_2 (c) and ℓ_1 (d) regularization techniques. Maximum intensity projection is used for visualization because the printed dots do not fall exactly on a plane.

Fig. 4.

Results for the Gd doped tube phantom with 5 echos. Fit parameters are ℓ_2 : slope = 0.92, intercept = 0.54; ℓ_1 : slope = 1.44, intercept = 0.19. (a) Plots of estimated susceptibility vs. true susceptibility. (b) First echo, illustrating phantom configuration. (c) Field inhomogeneity map. (d) ℓ_2 regularized inversion result. (e) ℓ_1 regularized inversion result.

(a)

(b)

(c)

(d)

Fig. 5.

Results for the Gd doped tube phantom with 3 echos. Fit parameters are ℓ_2 : slope = 0.65, intercept = 0.56; ℓ_1 : slope = 1.43, intercept = -0.64. (b) ℓ_2 regularized inversion result. (c) ℓ_1 regularized inversion result.

Fig. 6. Gd tube phantom with regularization parameters tuned to achieve slopes of 1 between estimated and true susceptibilities. (a) ℓ_2 regularization with 5 echos. (b) ℓ_1 regularization with 5 echos. (c) ℓ_2 regularization with 3 echos. (d) ℓ_1 regularization with 3 echos.

Fig. 7. Susceptibility estimation in an excised mouse brain. (a) Intensity image. (b) Field shift map. Estimated susceptibility distribution using ℓ_2 (c) and ℓ_1 (d) regularized inversions. The black boxes on the intensity image and field shift map indicate the regions over which susceptibilities were estimated.

TABLE I

Slopes and residuals for the Gd doped tubes phantom estimated using regularization parameters determined from noise in the frequency map (var) or by matching the slope to approximately one (slope).

norm	method	# echoes	α^2 or ν	slope	residual
ℓ_2	var	5	$10^{-3.4}$	0.92	8.16
ℓ_2	slope	5	$10^{-3.8}$	1.02	6.06
ℓ_2	var	3	$10^{-2.8}$	0.65	15.50
ℓ_2	slope	3	$10^{-4.0}$	1.03	5.59
ℓ_1	var	5	$10^{-2.8}$	1.44	7.80
ℓ_1	slope	5	$10^{-1.4}$	1.08	22.77
ℓ_1	var	3	$10^{-2.2}$	1.43	16.44
ℓ_1	slope	3	$10^{-1.4}$	1.07	26.68