FULL PAPER

Magnetic Resonance in Medicine 00:00–00 (2013)

Fast Quantitative Susceptibility Mapping with L1-Regularization and Automatic Parameter Selection

Berkin Bilgic,[1]* Audrey P. Fan,[1,2] Jonathan R. Polimeni,[1,3] Stephen F. Cauley,[1] Marta Bianciardi,[1,3] Elfar Adalsteinsson,[1,2,4] Lawrence L. Wald,[1,3,4] and Kawin Setsompop[1,3]

Purpose: To enable fast reconstruction of quantitative susceptibility maps with total variation penalty and automatic regularization parameter selection.

Methods: ℓ_1-Regularized susceptibility mapping is accelerated by variable splitting, which allows closed-form evaluation of each iteration of the algorithm by soft thresholding and fast Fourier transforms. This fast algorithm also renders automatic regularization parameter estimation practical. A weighting mask derived from the magnitude signal can be incorporated to allow edge-aware regularization.

Results: Compared with the nonlinear conjugate gradient (CG) solver, the proposed method is 20 times faster. A complete pipeline including Laplacian phase unwrapping, background phase removal with SHARP filtering, and ℓ_1-regularized dipole inversion at 0.6 mm isotropic resolution is completed in 1.2 min using MATLAB on a standard workstation compared with 22 min using the CG solver. This fast reconstruction allows estimation of regularization parameters with the L-curve method in 13 min, which would have taken 4 h with the CG algorithm. The proposed method also permits magnitude-weighted regularization, which prevents smoothing across edges identified on the magnitude signal. This more complicated optimization problem is solved 5 times faster than the nonlinear CG approach. Utility of the proposed method is also demonstrated in functional blood oxygen level–dependent susceptibility mapping, where processing of the massive time series dataset would otherwise be prohibitive with the CG solver.

Conclusion: Online reconstruction of regularized susceptibility maps may become feasible with the proposed dipole inversion. **Magn Reson Med 000:000–000, 2013. © 2013 Wiley Periodicals, Inc.**

Key words: Quantitative susceptibility mapping; Regularization; Total variation; L-curve

INTRODUCTION

Quantitative susceptibility mapping (QSM) aims to estimate the underlying magnetic susceptibility of the tissues that give rise to subtle changes in the magnetic field. Mapping this property allows quantification of tissue iron concentration (1,2) and vessel oxygen saturation (3–6) and has found recent use in the investigation of neurodegenerative diseases such as multiple sclerosis (7) and characterization of brain lesions (8).

Tissue susceptibility χ is related to the measured field map ϕ via the relation $\mathbf{DF}\chi = \mathbf{F}\phi$, where $\mathbf{F}$ is the discrete Fourier transform operator and $\mathbf{D} = 1/3 - k_z^2/k^2$ is the susceptibility kernel in k-space (9). This kernel effectively undersamples the frequency content of the field map on the conical surface $3k_z^2 = k^2$, which makes the inversion of the relation ill-conditioned. Existing approaches improve conditioning either by imposing prior knowledge on the reconstructed susceptibility map (10–12) or by collecting additional data by changing the orientation of the sample with respect to the main field (13,14). Compared with the regularized reconstruction techniques that employ a signal prior, multi-orientation methods yield susceptibility maps with superior quality, though at the cost of increased scan time and patient discomfort (15).

Single-orientation, regularized QSM methods often impose sparsity or smoothness constraints on the spatial gradients of the reconstructed susceptibility map, which can be formulated by penalizing the ℓ_1- or ℓ_2-norm of the gradients in three dimensions (10–12,16). The two types of regularized reconstruction methods can be expressed as an unconstrained convex optimization problem, minimizing either

$$\ell_1\text{-regularized}: \frac{1}{2}||\mathbf{F}^{-1}\mathbf{DF}\chi\text{-}\phi||_2^2 + \alpha \cdot ||\mathbf{WG}\chi||_1 \qquad [1]$$

or

$$\ell_2\text{-regularized}: ||\ \mathbf{F}^{-1}\mathbf{DF}\chi\text{-}\phi||_2^2 + \beta \cdot ||\mathbf{WG}\chi||_2^2 \qquad [2]$$

where α and β are regularization parameters and $\mathbf{G} = [\mathbf{G_x};\mathbf{G_y};\mathbf{G_z}]$ is the gradient operator in three spatial dimensions. The diagonal weighting term $\mathbf{W}$ can either be taken to be the identity matrix $\mathbf{I}$, or a binary mask $\mathbf{W} = [\mathbf{W_x};\mathbf{W_y};\mathbf{W_z}]$ derived from the magnitude signal can be employed to prevent penalizing certain gradient features (10–12). These two optimization problems are conventionally solved with the iterative nonlinear conjugate

Additional Supporting Information may be found in the online version of this article.

[1]Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, Massachusetts, USA.

[2]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Massachusetts, USA.

[3]Department of Radiology, Harvard Medical School, Boston, Massachusetts, USA.

[4]Harvard-MIT Health Sciences and Technology, Cambridge, Massachusetts, USA

*Correspondence to: Berkin Bilgic, Building 75, Room 2.102, 13th Street, Charlestown, MA, 02129. E-mail: berkin@nmr.mgh.harvard.edu

Received 4 June 2013; revised 11 October 2013; accepted 14 October 2013

DOI 10.1002/mrm.25029

Published online 00 Month 2013 in Wiley Online Library (wileyonlinelibrary.com).

gradient (CG) algorithm. Reconstruction times reported in the literature range between 20 min (11,12) and 2–3 h (17), which may be a limiting factor for the online reconstruction of susceptibility maps required for routine use or in clinical settings. Moreover, the parameters α and β need to be determined for optimal regularization, which is usually addressed by performing multiple reconstructions while sweeping a range of parameters to trace the L-curve (18) or finding the operating point that satisfies the discrepancy principle (19). As such, the identification of a suitable amount of regularization would further increase the computation time.

We recently introduced a closed-form solution to the ℓ_2-regularized QSM problem in Equation [2] without magnitude weighting ($\mathbf{W} = \mathbf{I}$) (20). This method requires only two fast Fourier transform (FFT) operations and takes <1 s to compute for a high-resolution phase data. In contrast, a closed-form solution to the ℓ_1-constrained problem in Equation [1] does not exist, which forces the existing algorithms to operate iteratively. Previously, ℓ_1-regularized reconstruction was shown to be superior to ℓ_2-penalty both in image quality and quantification accuracy (12). The ℓ_1-regularized results presented herein also show substantially reduced reconstruction error on numerical simulation and better estimation of undersampled content near the magic angle on in vivo data. In this article, we propose a fast ℓ_1-regularized QSM algorithm that works iteratively; however, each iteration is computed efficiently in closed-form. We also extend these fast ℓ_1- and ℓ_2-regularized solvers to incorporate magnitude prior.

This study employed an efficient variable-splitting algorithm (21) to solve Equation [1] without magnitude weighting and yielded reconstruction time 20 times faster than the nonlinear CG solver (12,16). By introducing an auxiliary variable that replaced the image gradient, each iteration of the proposed algorithm was computed in closed-form, requiring only FFTs and soft thresholding operations. With the proposed formulation, reconstruction for high-resolution in vivo susceptibility mapping at 0.6 mm isotropic voxel size took 1 min (using MATLAB on a standard workstation). Combined with state-of-the-art phase unwrapping and background phase removal methods (14,22), this comprised a fast reconstruction pipeline that might facilitate clinical application of QSM. In the Discussion section, we outline the differences between the presented methods and a similar variable-splitting algorithm for QSM previously proposed in the elegant contribution by Chen and Calhoun (23).

With the addition of magnitude prior, solution of both ℓ_1- and ℓ_2-regularized problems became more involved, since the related linear systems that need to be inverted were no longer diagonal. By employing a preconditioner that facilitated the inversion of these systems, we obtained a rapid iterative solver that led to substantial computational savings. Relative to the nonlinear CG method, we demonstrate a magnitude-weighted ℓ_1- and ℓ_2-regularization that is 5 times and 15 times faster, respectively.

The specific contributions of this study are as follows:

1. Fast susceptibility mapping with ℓ_1- and ℓ_2-regularization. The proposed ℓ_1-regularized algorithm is 20 times faster than the conventional CG solver. We also introduce ℓ_1- and ℓ_2-based solvers with magnitude prior that allow edge-aware regularization, while achieving 5× and 15× speed-up over the CG approach.
2. Automatic regularization parameter selection. Using the L-curve heuristic, the parameter that yields the maximum curvature on the curve is selected as the optimal operating point. Sweeping the L-curve becomes practical with the proposed fast reconstruction, whereas parameter selection with CG solver would take several hours of computation.
3. Functional blood oxygen level–dependent (BOLD) susceptibility mapping (6,24,25). combination of fast phase unwrapping, background phase removal, regularization parameter estimation, and susceptibility inversion makes reconstruction of large time series datasets feasible, thus making investigation of functional BOLD QSM practical.
4. Reproducible research. The MATLAB code for the complete reconstruction pipeline (Laplacian unwrapping, SHARP filtering, ℓ_2- and ℓ_1- regularized fast susceptibility mapping with magnitude weighting and parameter estimation) is included as supplementary material and made available online at: http://web.mit.edu/berkin/www/Fast_QSM_Magnitude_Toolbox.zip. A copy of the source code is also deposited into a git archive at https://sourceforge.net/projects/fastqsm/

THEORY

Proposed ℓ_2-Regularized QSM Algorithms

The global optimizer of the convex optimization problem in Equation [2] can be obtained by taking the gradient and setting it to zero, i.e.,

$$(\mathbf{F}^{-1}\mathbf{D}^2\mathbf{F} + \beta \cdot \mathbf{G}^T\mathbf{W}^2\mathbf{G})\chi = \mathbf{F}^{-1}\mathbf{D}\mathbf{F}\phi \qquad [3]$$

The solution of this linear system is facilitated by evaluating the spatial gradient via multiplication in k-space rather than using convolution in image space. This can be accomplished by expressing the image gradient along the x-axis as $\mathbf{G}_x = \mathbf{F}^{-1}\mathbf{E}_x\mathbf{F}$, where $\mathbf{E}_x$ is a diagonal matrix that provides the k-space representation of the image-domain differencing operator $\delta_x - \delta_{x-1}$. $\mathbf{E}_x$ can be constructed by starting from the vector $\mathbf{v}_x$ with entries $\mathbf{v}_x(i) = 1 - exp(-2\pi\sqrt{-1} \cdot i/N_x)$, where $i = 0,\ldots,N_x - 1$ and N_x is the matrix size along x. When $\mathbf{v}_x$ is replicated along k_y and k_z dimensions in k-space to generate an image volume of size $N_x \times N_y \times N_z$, the matrix $\mathbf{E}_x$ can be obtained by assigning this volume to diagonal entries of $\mathbf{E}_x$. The operators $\mathbf{E}_y$ and $\mathbf{E}_z$ are similarly defined. Equation [3] then becomes

$$\mathbf{F}^{-1}(\mathbf{D}^2 + \beta \cdot \mathbf{E}^H\mathbf{F}\mathbf{W}^2\mathbf{F}^{-1}\mathbf{E})\mathbf{F}\chi = \mathbf{F}^{-1}\mathbf{D}\mathbf{F}\phi \qquad [4]$$

where $\mathbf{E} = [\mathbf{E}_x;\mathbf{E}_y;\mathbf{E}_z]$. We recently proposed (20) a fast solution to Equation [4] when no magnitude weighting is used ($\mathbf{W} = \mathbf{I}$). In this special case, further simplification of Equation [4] yields

$$\chi = \mathbf{F}^{-1} \cdot (\mathbf{D}^2 + \beta \cdot (\mathbf{E}_x^2 + \mathbf{E}_y^2 + \mathbf{E}_z^2))^{-1} \cdot \mathbf{DF}\phi \qquad [5]$$

Note that the matrix inversion now involves only diagonal matrices, hence it is straightforward to evaluate. The total cost of ℓ_2-regularized QSM is then two FFT operations and summation and multiplication of diagonal matrices. These usually take <1 s for three-dimensional high-resolution image volume.

With the inclusion of magnitude weighting **W**, the system that needs to be solved becomes

$$(\mathbf{D}^2 + \beta \cdot \mathbf{E}^H \mathbf{F} \mathbf{W}^2 \mathbf{F}^{-1} \mathbf{E}) \mathbf{F}\chi = \mathbf{DF}\phi \qquad [6]$$

Because this system is no longer diagonal, its inversion is more involved. The linear CG algorithm is a popular technique used for solution of such symmetric, positive, definite systems. The convergence speed of this method is determined by the condition number of the matrix $\mathbf{A} = (\mathbf{D^2} + \boldsymbol{\beta} \cdot \mathbf{E}^{\boldsymbol{H}} \mathbf{F} \mathbf{W}^{\mathbf{2}} \mathbf{F}^{-\mathbf{1}} \mathbf{E})$. The conditioning can be improved if the matrix **A** can be brought closer to being identity **I**. We propose use of the closed-form solution in Equation [5] as preconditioner and solve the modified system

$$(\mathbf{D}^2 + \beta \cdot \mathbf{E}^2)^{-1} \cdot \{\mathbf{A} \cdot \mathbf{F}\chi\text{-}\mathbf{DF}\phi\} = 0 \qquad [7]$$

Because $(\mathbf{D^2} + \boldsymbol{\beta} \cdot \mathbf{E^2})$ is diagonal, it is straightforward to compute its inverse and convenient to use it as a preconditioner. When the gradients of the magnitude image are thresholded so that the strongest 30% voxels within the brain mask are considered to be edges as suggested by Liu et al. (26), the resulting **W** matrix is equal to the identity **I** except for ~5% of its entries. This makes the approximation $(\mathbf{D^2} + \boldsymbol{\beta} \cdot \mathbf{E^2})^{-1} \approx \mathbf{A}^{-1}$ valid and renders the preconditioner useful. The closed-form solution in Equation [5] can be used as an initial guess to further accelerate convergence.

Herein, the native MATLAB function pcg.m is used for solving Equation [7]. The major cost of each iteration is 6 FFTs required for the operator **A**.

ℓ_1-Regularized QSM with Nonlinear CG Solver

Nonlinear CG is a popular method that is especially suitable for solution of large-scale optimization problems. It requires evaluation of only the first-order derivatives and has better convergence properties than simple gradient descent techniques. The algorithm can be summarized as follows:

$$\begin{aligned}
\textit{precomputation}: &\quad \mathbf{DF}\chi,\ \mathbf{DF}(\Delta\chi),\ \mathbf{WG}\chi,\ \mathbf{WG}(\Delta\chi) \\
\textit{line search for stepsize } \tau: &\quad \operatorname{argmin}_\tau\ 0.5 \cdot ||\mathbf{DF}(\chi + \tau \cdot \Delta\chi)\text{-}\mathbf{F}\phi||_2^2 + \alpha ||\mathbf{WG}(\chi + \tau \cdot \Delta\chi)||_1 \\
\chi\ \textit{update}: &\quad \chi = \chi + \tau \cdot \Delta\chi \\
\textit{gradient update}: &\quad g_{\text{new}} = \mathbf{F}^{-1}\mathbf{D}(\mathbf{DF}\chi\text{-}\mathbf{F}\phi) + \alpha \cdot \mathbf{G}^T\mathbf{W}^T \operatorname{sign}(\mathbf{WG}\chi) \\
\textit{parameter update}: &\quad b = ||g_{\text{new}}||_2 / ||g_{\text{prev}}||_2 \\
&\quad g_{\text{prev}} = g_{\text{new}} \\
\Delta\chi\ \textit{update}: &\quad \Delta\chi = b \cdot \Delta\chi\text{-}g_{\text{new}}
\end{aligned}$$

Note that the data consistency term for the CG algorithm is taken to be $||\mathbf{DF}\chi\text{-}\mathbf{F}\phi||_2^2$ instead of $||\mathbf{F}^{-1}\mathbf{DF}\chi\text{-}\phi||_2^2$ as in Equation [1]. Due to Parseval's relation, the two terms are equivalent; however, using this alternative formulation avoids four redundant FFT operations per iteration. The algorithm terminates when a user-specified convergence criterion is met. Per iteration, it requires four FFT operations and four spatial gradient evaluations (two for *precomputation* and two for *gradient update*). Because each gradient computation involves finite differences in three dimensions, the total cost of the algorithm is 12 differencing operations and four FFT operations per iteration.

Proposed ℓ_1-Regularized QSM Algorithms

The proposed approach relies on a fast algorithm that was initially proposed for reconstruction of undersampled k-space data in the context of compressed sensing (21). Here, this method is extended to susceptibility mapping using the following formulation:

$$min_{\chi,y} \frac{1}{2} ||\mathbf{F}^{-1}\mathbf{DF}\chi\text{-}\phi||_2^2 + \lambda \cdot ||y||_1 \quad \textit{such that} \quad \mathbf{WG}\chi = y \qquad [8]$$

This constrained objective is first relaxed to an unconstrained optimization problem, then solved with the two-phase "split-Bregman" iteration:

at iteration t,

$$(\chi_{t+1}, y_{t+1}) = argmin_{\chi,y}\ \frac{1}{2}||\mathbf{F}^{-1}\mathbf{DF}\chi - \phi||_2^2 + \lambda||y||_1 + \frac{\mu}{2}||\mathbf{y} - \mathbf{WG}\chi - \eta_t||_2^2 \qquad [9]$$

$$\eta_{t+1} = \eta_t + \mathbf{WG}\chi_{t+1} - y_{t+1} \qquad [10]$$

Here, $\boldsymbol{y}$ is an auxiliary variable that replaces the magnitude-weighted gradient of the susceptibility map $\boldsymbol{\chi}$. This idea of variable splitting will allow closed-form optimization via the soft thresholding operator. The variable $\boldsymbol{\eta}$ adds the mismatch in $\mathbf{WG}\boldsymbol{\chi} = \boldsymbol{y}$ back to the

unconstrained problem to enforce this constraint. Equation [9] can be solved efficiently by iteratively minimizing with respect to χ and y separately:

$$\chi_{t+1} = argmin_{\chi} \ \frac{1}{2}||\mathbf{F}^{-1}\mathbf{DF}\chi - \phi||_2^2 + \frac{\mu}{2}||\mathbf{y}_t - \mathbf{WG}\chi - \eta_t||_2^2 \quad [11]$$

$$y_{t+1} = argmin_{y} \quad \lambda \cdot ||y||_1 + \frac{\mu}{2}||y - \mathbf{WG}\chi_{t+1} - \eta_t||_2^2 \quad [12]$$

The optimality condition for Equation [11] can be found by taking the gradient and setting it to zero:

$$(\mathbf{F}^{-1}\mathbf{D}^2\mathbf{F} + \mu \cdot \mathbf{G}^T\mathbf{W}^2\mathbf{G})\chi_{t+1} = \mathbf{F}^{-1}\mathbf{DF}\phi + \mu \cdot \mathbf{G}^T\mathbf{W}^T(y_t - \eta_t) \quad [13]$$

For the special case where magnitude weighting is not utilized ($\mathbf{W = I}$), the solution to Equation [13] can be evaluated in closed-form to obtain the following update rule for χ:

$$\begin{aligned}\mathbf{F}\chi_{t+1} =& (\mathbf{D}^2 + \mu \cdot (\mathbf{E}_x^2 + \mathbf{E}_y^2 + \mathbf{E}_z^2))^{-1} \\ & \cdot \left(\mathbf{DF}\phi + \mu \cdot \mathbf{E}^H\mathbf{F}(y_t - \eta_t)\right)\end{aligned} \quad [14]$$

Here, $\mathbf{E}^H = [\mathbf{E}_x^H, \ \mathbf{E}_y^H, \ \mathbf{E}_z^H]$ and the variables y_t and η_t have components in three dimensions, $y_t = [y_{t,x};y_{t,y};y_{t,z}]$ and $\eta_t = [\eta_{t,x};\eta_{t,y};\eta_{t,z}]$. The matrix inversion involves only diagonal matrices and needs to be computed once. The cost of evaluating Equation [14] is three FFT operations (one for each spatial axis) and element-wise multiplications. The update for the susceptibility map is performed in k-space, thereby avoiding one redundant inverse FFT operation per iteration.

In the presence of magnitude weighting, Equation [13] no longer admits a closed-form solution. However, further simplification leads to a system similar to Equation [6] encountered while solving the magnitude-weighted ℓ_2-based problem,

$$(\mathbf{D}^2 + \mu \cdot \mathbf{E}^H\mathbf{FW}^2\mathbf{F}^{-1}\mathbf{E})\mathbf{F}\chi_{t+1} = \mathbf{DF}\phi + \mu \cdot \mathbf{E}^H\mathbf{FW}^T(y_t - \eta_t) \quad [15]$$

This equation can be efficiently solved with the linear conjugate gradient method by again employing the preconditioner ($\mathbf{D^2+\mu \cdot E^2}$). Convergence behavior can be substantially enhanced if the solution from the previous iteration χ_t is used as an initial guess. As the variable-splitting algorithm iterates, the difference between the successive solutions χ_t and χ_{t+1} decrease progressively. Using such a suitable initial guess, preconditioned CG can solve Equation [15] within 1% tolerance in a couple of steps.

Regardless of magnitude weighting, the solution to the ℓ_1-regularized least squares problem in Equation [12] is given by the following element-wise soft thresholding operation:

$$y_{t+1} = max\left(|\mathbf{WG}\chi_{t+1} + \eta_t| - \frac{\lambda}{\mu}, 0\right) \cdot sign(\mathbf{WG}\chi_{t+1} + \eta_t) \quad [16]$$

This operator is applied to all three components of $\mathbf{y}_{t+1}$. To compute the gradient in x direction $\mathbf{G}_x\chi_{t+1}$, the k-space representation of the current susceptibility estimate $\mathbf{F}\chi_{t+1}$ is utilized due to $\mathbf{G}_x\chi_{t+1} = \mathbf{F}^{-1}\mathbf{E}_x(\mathbf{F}\chi_{t+1})$. This way, the update rule in Equation [16] requires three inverse FFT operations (one for each spatial axis) and simple thresholding operations.

Combining the solutions to Equations [10] and [11] with the update rule in Equation [16] yields the complete algorithm for the proposed ℓ_1-regularized reconstruction:

$$initialization: y_{0,x} = y_{0,y} = y_{0,z} = \eta_{0,x} = \eta_{0,y} = \eta_{0,z} = 0$$

$$\mathbf{F}\chi \ \ update: \mathbf{Solve} \quad (\mathbf{D}^2 + \mu \cdot \mathbf{E}^H\mathbf{FW}^2\mathbf{F}^{-1}\mathbf{E})\mathbf{F}\chi_{t+1} = \mathbf{DF}\phi + \mu \cdot \mathbf{E}^H\mathbf{FW}^T(y_t - \eta_t)$$

$$gradient \ \ update: \begin{cases} \mathbf{W}_x\mathbf{G}_x\chi_{t+1} = \mathbf{W}_x\mathbf{F}^{-1}\mathbf{E}_x(\mathbf{F}\chi_{t+1}) \\ \mathbf{W}_y\mathbf{G}_y\chi_{t+1} = \mathbf{W}_y\mathbf{F}^{-1}\mathbf{E}_y(\mathbf{F}\chi_{t+1}) \\ \mathbf{W}_z\mathbf{G}_z\chi_{t+1} = \mathbf{W}_z\mathbf{F}^{-1}\mathbf{E}_z(\mathbf{F}\chi_{t+1}) \end{cases}$$

$$soft \ \ thresholding: \begin{cases} y_{t+1,x} = max(|\mathbf{W}_x\mathbf{G}_x\chi_{t+1} + \eta_{t,x}| - \lambda/\mu, 0) \cdot sign(\mathbf{W}_x\mathbf{G}_x\chi_{t+1} + \eta_{t,x}) \\ y_{t+1,y} = max(|\mathbf{W}_y\mathbf{G}_y\chi_{t+1} + \eta_{t,y}| - \lambda/\mu, 0) \cdot sign(\mathbf{W}_y\mathbf{G}_y\chi_{t+1} + \eta_{t,y}) \\ y_{t+1,z} = max(|\mathbf{W}_z\mathbf{G}_z\chi_{t+1} + \eta_{t,z}| - \lambda/\mu, 0) \cdot sign(\mathbf{W}_z\mathbf{G}_z\chi_{t+1} + \eta_{t,z}) \end{cases}$$

$$residual \ \ update: \begin{cases} \eta_{t+1,x} = \eta_{t,x} + \mathbf{W}_x\mathbf{G}_x\chi_{t+1} - y_{t+1,x} \\ \eta_{t+1,y} = \eta_{t,y} + \mathbf{W}_y\mathbf{G}_y\chi_{t+1} - y_{t+1,y} \\ \eta_{t+1,z} = \eta_{t,z} + \mathbf{W}_z\mathbf{G}_z\chi_{t+1} - y_{t+1,z} \end{cases}$$

Per iteration, the proposed method without magnitude weighting requires six FFT operations (three for $\mathbf{F}\chi$ *update* due to Equation [14] and another three for *gradient update*) and simple thresholding operations. Compared with the nonlinear CG method, there is no need to compute the 12 differencing operations in image space; however, there are two additional FFT operations per iteration. On the other hand, one iteration of the proposed method with magnitude prior requires $(6 + 6s)$ FFTs, where s is the number of steps it takes to solve Equation [15].

Moreover, the proposed method introduces a second regularization parameter μ that must be determined. Goldstein and Osher (21) showed that this parameter does not change the solution to the optimization problem, but affects the convergence speed. Empirically, a parameter that leads to fast convergence can be identified on one subject, which might then generalize to different datasets. Alternatively, an automatic way to select this parameter is proposed in the Methods section, which was observed to yield fast convergence speed.

METHODS

The performance of the proposed method was validated on a numerical phantom with known susceptibility. In addition, two in vivo datasets were reconstructed: 1) high-resolution whole brain three-dimensional (3D) gradient echo (GRE) acquired at 3T with 0.6 mm isotropic resolution and 2) whole brain gradient echo two-dimensional (2D) echo-planar at 7T with 1.5 mm isotropic resolution acquired as a time series. Subjects were scanned under the approval of the local Institutional Review Board.

All computations were performed in MATLAB using a workstation with 32 processors (AMD Opteron 6282 SE) and 128 GB memory.

Numerical Susceptibility Phantom

A three-compartment phantom of size $246 \times 246 \times 162$ with known susceptibility was generated. Within each compartment, susceptibility was fixed and constant at $\chi_{gray}=-0.023$, $\chi_{white}=0.027$ and $\chi_{csf}=-0.018$ ppm (27). From this distribution χ, the field map ϕ was generated using forward dipole model $\phi = \mathbf{F}^{-1}\mathbf{D}\mathbf{F}\chi$, to which random Gaussian i.i.d. noise with a peak signal-to-noise ratio (PSNR) of 100 was added (Fig. 1a), where PSNR was defined to be the ratio of the maximum value in the field map to the noise standard deviation. Starting from this noisy field map, susceptibility maps were reconstructed using three different methods:

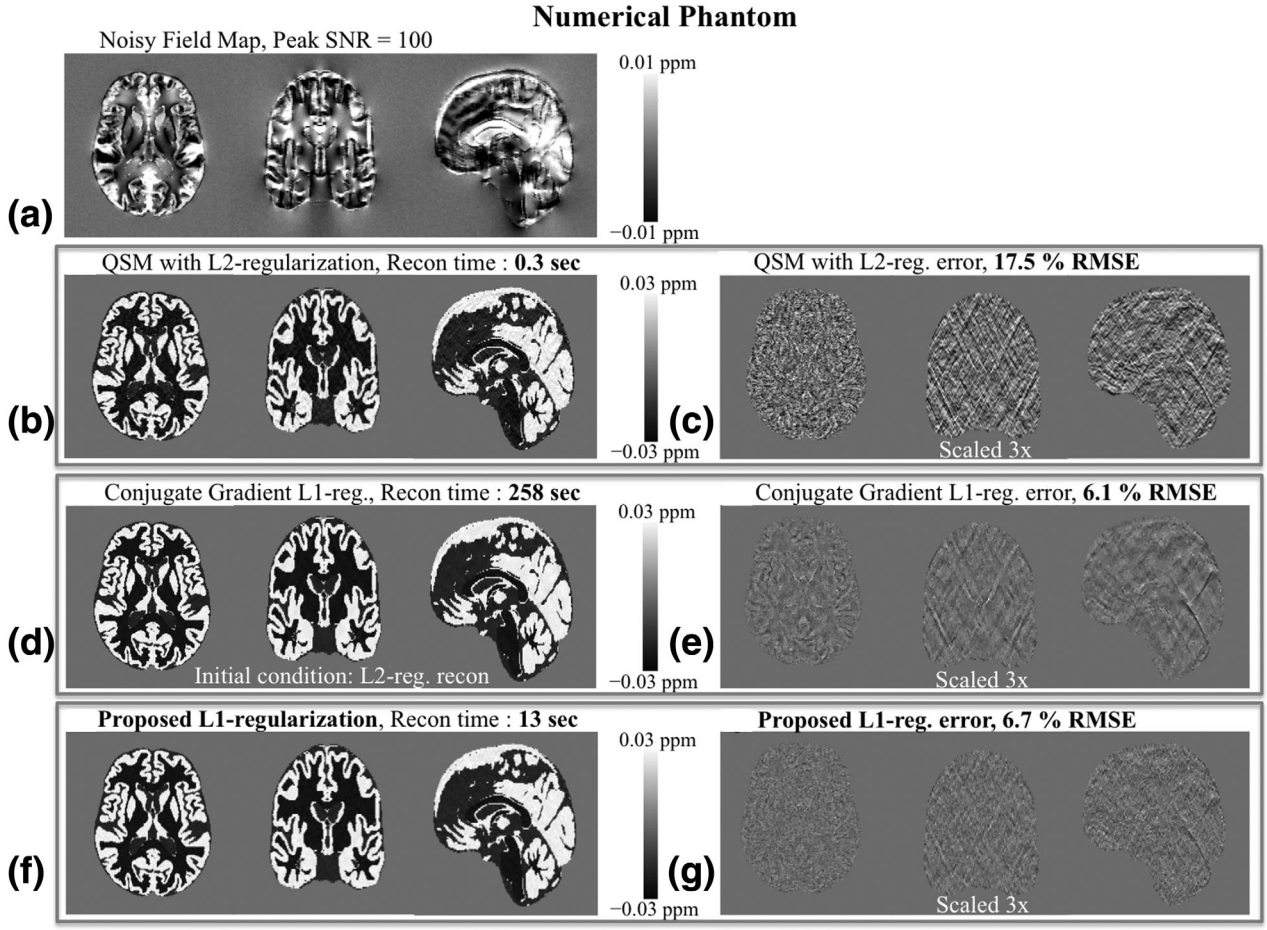

FIG. 1. QSM with numerical phantom. RMSE values were computed relative to the known, ground-truth susceptibility map. A field map was simulated using the true susceptibility, to which noise with PSNR = 100 was added to obtain the noisy field map (**a**). QSM with closed-form L2-regularization (**b**) took 0.3 s, and the reconstruction error (**c**) was 17.5%. L1-regularized CG reconstruction took 258 s (**d**), and the error (**e**) was 6.1%. The proposed L1-constrained QSM was completed in 13 s (**f**), with an RMSE of 6.7% (**g**).

1. ℓ_2-regularized QSM with closed-form solution,
2. ℓ_1 -regularization with nonlinear CG solver, and
3. Proposed ℓ_1-regularization with variable-splitting.

The regularization parameters (β for ℓ_2-regularization, α for CG and λ for the proposed method) were determined by parameter sweeping and the values that minimized the normalized root mean square error (RMSE) relative to the true χ were selected to be the optimal setting.

At the first iteration of the proposed algorithm, the initial estimate of the susceptibility map is found by using the initial condition $\mathbf{y_0 = \eta_0 = 0}$ in Equation [14]. This yields

$$\mathbf{F}\chi_1 = (\mathbf{D}^2 + \mu \cdot (\mathbf{E}_x^2 + \mathbf{E}_y^2 + \mathbf{E}_z^2))^{-1} \cdot \mathbf{DF}\phi \qquad [17]$$

Note that this is exactly the same expression as Equation [5], which gave the optimizer for the ℓ_2-regularized QSM formulation. Hence, it is seen that the first iteration of the proposed ℓ_1-constrained algorithm is actually the solution to the original ℓ_2-regularized problem in Equation [2]. Based on this, we propose to set μ to the optimal value of β found for the closed-form ℓ_2-reconstruction. The algorithm terminates when the change in k-space of the susceptibility falls below 1%.

To hasten convergence, the nonlinear CG algorithm used the closed-form ℓ_2-reconstruction as an initial guess, and the termination criterion was again <1% change in susceptibility.

Goldstein and Osher (21) showed that the parameter μ controlling the gradient consistency term $\mu \cdot ||\mathbf{y}\text{-}\mathbf{G}\chi\text{-}\eta_t||_2^2$ in the variable-splitting algorithm does not affect the solution but the speed of convergence. To empirically test this property, the numerical phantom was also reconstructed with various μ settings while keeping the ℓ_1-parameter λ fixed at 10^{-5}. To enforce convergence, the algorithm was run for 300 iterations.

High-Resolution In Vivo 3D GRE

A 26-year-old, female, healthy volunteer (having given informed consent) was scanned using a dual-echo gradient echo sequence with full flow compensation along all three axes (28) at a Siemens 3T system equipped with 32 receive channels. Imaging parameters were as follows: 0.6 mm isotropic resolution, repetition time (TR) = 26 ms, echo times (TE) = 8.1 and 20.3 ms, matrix size = 384 × 336 × 224, GRAPPA acceleration factor = 2, phase partial Fourier = 75%, and acquisition time = 15 min 42 sec. Additional GRE data were collected with the same spatial coverage but at a lower resolution of 1.8 × 1.8 × 2.4 mm^3 at five different TEs (6–10 ms with 1-ms intervals). The radiofrequency phase offset of each channel at TE = 0 was estimated using a linear fit to the phase images corresponding to the five acquired echoes. The estimated offset per channel was then removed from the 0.6 mm high-resolution phase data (29). Finally, magnitude-weighted, phase-valued coil combination (30) was performed and is given by

$$\psi_{combo} = \measuredangle\left(\sum\nolimits_i m_i \cdot \exp(\sqrt{-1} \cdot \psi_i)\right) \qquad [18]$$

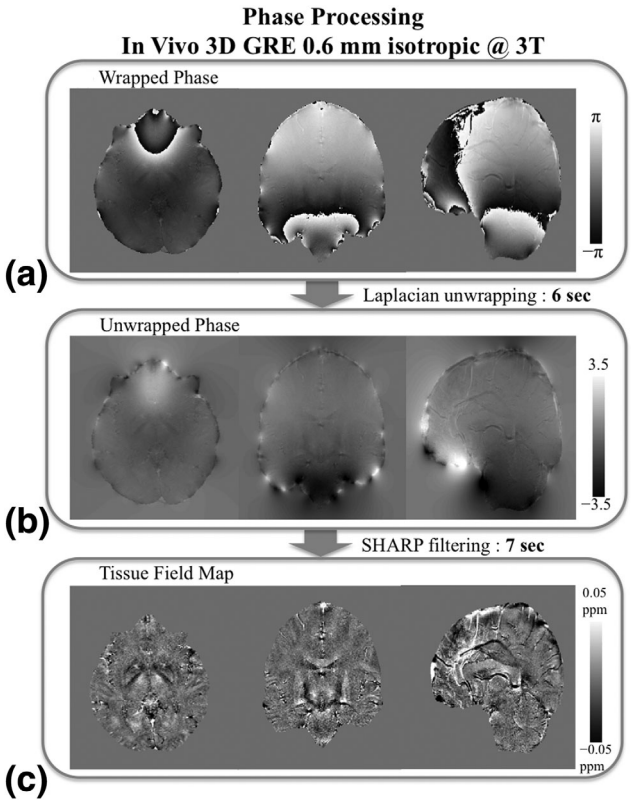

FIG. 2. Phase processing steps for in vivo 3D GRE data at 0.6 mm resolution. Starting from the coil-combined, wrapped phase in (**a**), unwrapped phase data are obtained with Laplacian unwrapping (**b**) in 6 s. Further processing with SHARP filtering yielded the tissue field map (**c**) in 7 s.

Combined phase image at TE = 8.1 ms was masked using a binary-valued region of interest generated with FSL-BET (31) (Fig. 2a). This wrapped-phase image was unwrapped with the fast Laplacian unwrapping algorithm (22) (Fig. 2b) and the background contributions due to air-tissue interfaces were removed with SHARP filtering (14) (Fig. 2c). The truncation level for SHARP was set to 0.05 as in Schweser et al. (14) and following Sun and Wilman (32), the kernel size was set to 9 × 9 × 9 pixels so that the filter length was about 5 mm isotropic.

The regularization parameter β for closed-form ℓ_2-regularization without magnitude prior was selected using the L-curve heuristic (18) (Fig. 3). Since the L-curve traced by the data consistency and regularization terms as β varied did not have a clear elbow point, the point with the largest curvature was selected as the optimal operating point. The curvature κ was computed as in (18),

$$\kappa = 2\frac{\ddot{\rho}\dot{\omega} - \ddot{\omega}\dot{\rho}}{(\dot{\rho}^2 + \dot{\omega}^2)^{1.5}} \qquad [19]$$

Here, $\rho = \log(||\mathbf{F}^{-1}\mathbf{DF}\chi - \phi||_2^2)$ and $\omega = \log(||\mathbf{G}\chi||_2^2)$ and $\dot{\omega}$, $\dot{\rho}$, $\ddot{\omega}$ and $\ddot{\rho}$ represent first and second derivatives of ω and ρ with respect to β. Since the L-curve is only sampled at discrete points, we fit cubic splines to ρ and ω to express them as smooth functions of β, which can then be differentiated.

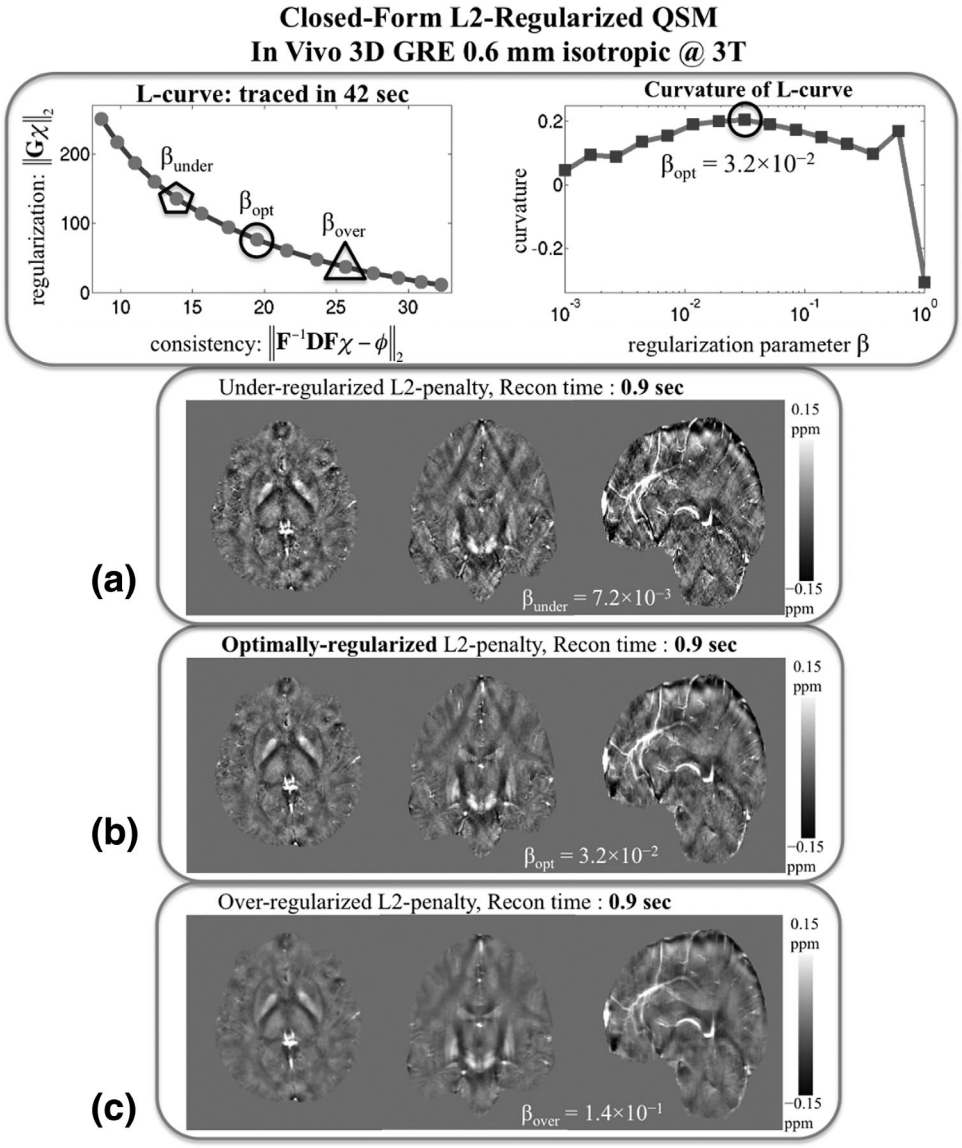

FIG. 3. Closed-form L2-constrained reconstruction for 3D GRE. Upper panel: the L-curve is traced in 42 s, and the parameter value that maximized the curvature was $\beta = 3.2 \cdot 10^{-2}$, corresponding to the optimal level of regularization. **a–c:** Underregularized (a), optimally regularized (b), and overregularized (c) susceptibility maps. Each reconstruction took 0.9 s of computation time.

The proposed magnitude-weighted ℓ_2-based reconstruction with the preconditioned conjugate gradient solver used the same β parameter. The termination criterion for this iterative approach was to attain 0.1% tolerance, where tolerance is defined as $||\mathbf{A}x - b||_2/||b||_2$ regarding the solution of the system $\mathbf{Ax = b}$. The magnitude prior was derived by thresholding the spatial gradients of the magnitude signal so that 30% voxels in the brain mask with the strongest gradients were considered to be edges (26).

Regarding the proposed ℓ_1-based method without magnitude weighting, the smoothing parameter λ was similarly determined by maximizing the curvature of the L-curve (Fig. 4). Each point on the L-curve was reconstructed using 10 iterations of the proposed algorithm. Based on the observation that the initial iteration yields the ℓ_2-regularized reconstruction, the consistency parameter μ was set to the optimal β determined for the closed-form solution. The algorithm terminated when there was <1% change in the k-space of the susceptibility map.

The dependence of convergence speed of the proposed ℓ_1-based method on gradient consistency μ parameter was also investigated. Reconstructions with parameters that are 10 times larger and 10 times smaller than the optimal choice of μ were performed with the same 1% change convergence criterion.

Reconstruction with ℓ_1-regularization using magnitude prior employed the same λ parameter identified from the L-curve analysis. The tolerance of the preconditioned conjugate gradient used at each iteration was set to 1%, and the algorithm terminated when the updates between iterations was below 1%.

The smoothing parameter α for the ℓ_1-regularized nonlinear CG method was selected so that the data consistency of the final susceptibility map, $||\mathbf{F}^{-1}\mathbf{DF}\chi - \phi||_2^2$, matched the consistency of the proposed reconstruction (Fig. 5).

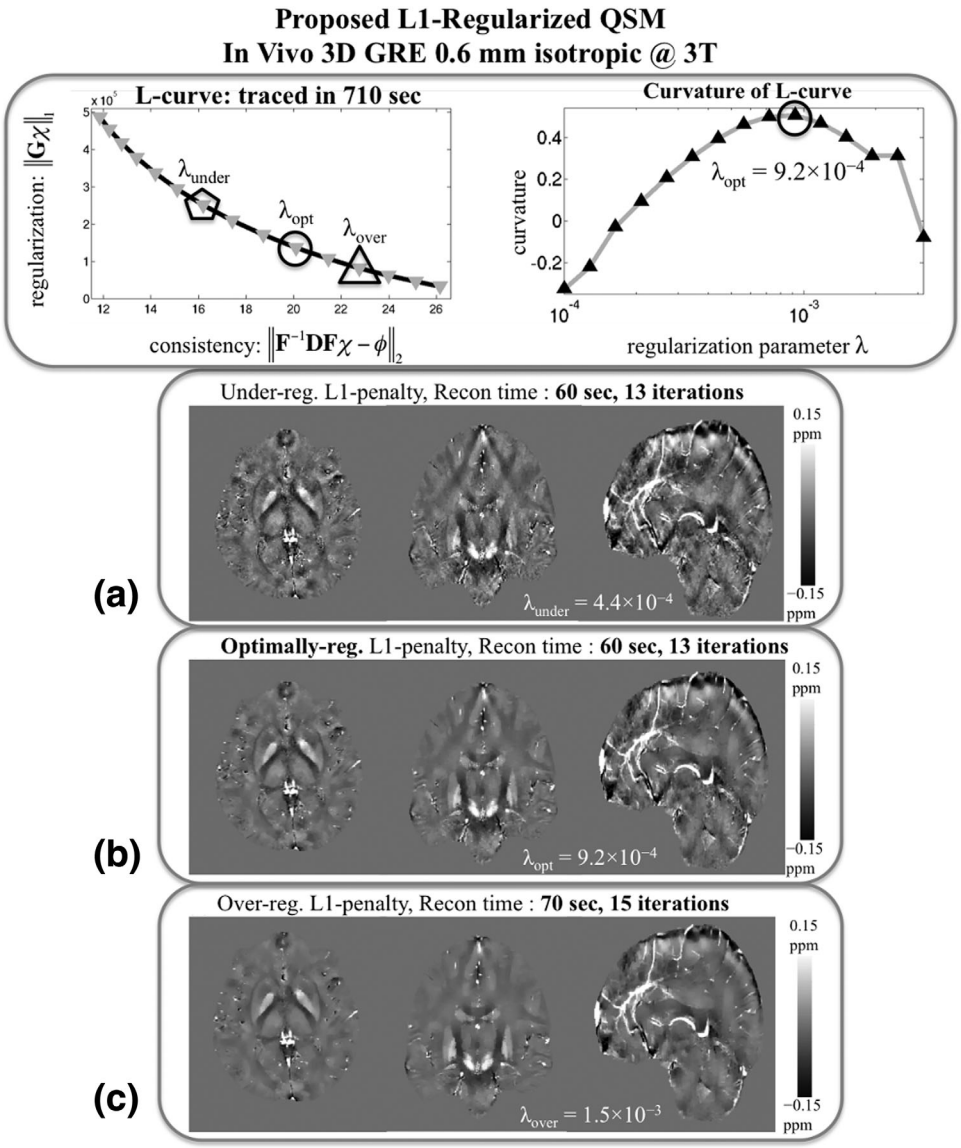

FIG. 4. Proposed L1-constrained reconstruction for 3D GRE. Upper panel: the L-curve is traced in 710 seconds, and the parameter value that maximized the curvature was $\lambda = 9.2 \cdot 10^{-4}$, corresponding to the optimal level of regularization. **a–c:** Underregularized (a), optimally regularized (b), and overregularized (c) susceptibility maps. Underregularized and optimally regularized reconstructions took 60 s and 13 iterations to converge, whereas optimization took 70 s and 15 iterations for the overregularized case.

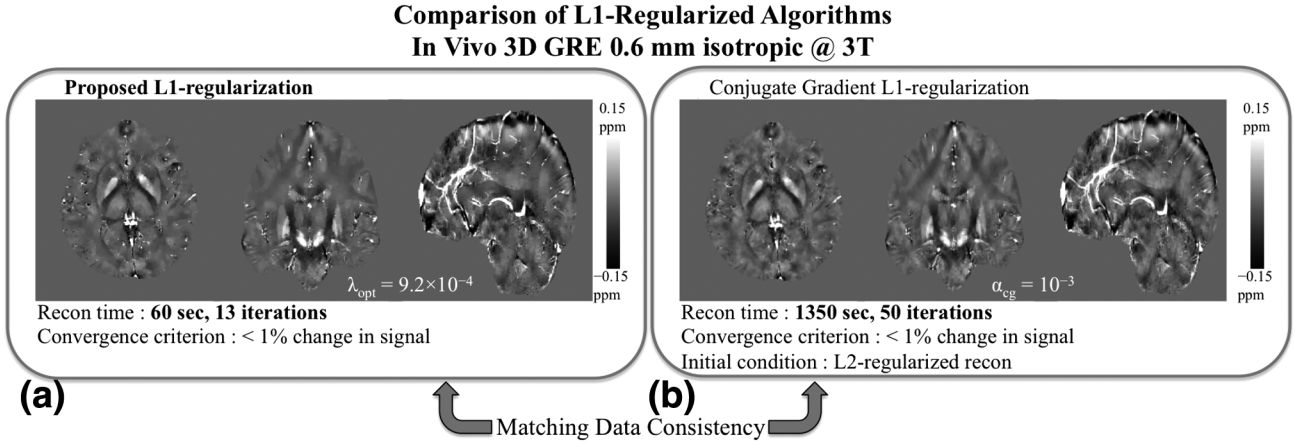

FIG. 5. Comparison of L1-regularized dipole inversion methods for in vivo 3D GRE. The proposed algorithm (**a**) converged in 60 s and 13 iterations, whereas it took 1350 s and 50 iterations for the CG algorithm to finish (**b**).

This strategy aimed to obtain comparable amounts of regularization with the two ℓ_1-penalized algorithms. To facilitate convergence, closed-form ℓ_2-reconstruction was used as initial guess to start the CG iterations, which terminated when there was <1% change in the susceptibility map.

In Vivo 2D Echo-Planar Imaging for Functional BOLD QSM

To emulate functional BOLD QSM data acquisition, a 35-year-old healthy male volunteer (having given informed consent) was scanned with a 2D echo-planar imaging (EPI) trajectory using a Siemens 7T system with a 32-channel receive coil array (33) and a birdcage transmit coil at 1.5 mm isotropic resolution. The imaging parameters were as follows: TR = 3660 ms, TE = 21 ms, matrix size = 128 × 128 × 76, total number of time points = 30, GRAPPA acceleration factor = 3 (reconstruction performed offline in MATLAB), slice acquisition order: ascending, slice orientation: axial. Following Hammond et al. (30), the relative phase offset of each coil image was estimated from the center 3 × 3 × 3 pixels and removed from each coil. Magnitude-weighted phase-valued coil combination was performed using Equation [18]. A brain mask was generated using FSL BET (31) based on the first frame in the time series. After masking, coil-combined wrapped-phase image of each frame was processed with Laplacian unwrapping (22) and SHARP filtering (14) (Fig. 7) with truncation level = 0.05 and kernel size = 3 × 3 × 3, so that the resulting filter size was about 5 mm (32).

The resulting tissue phase series was inverted using closed-form ℓ_2-regularized QSM and the proposed ℓ_1-based reconstruction without magnitude prior. Smoothing parameters β and α were selected with L-curve analysis using the mean tissue phase averaged over the 30 frames as the input. The gradient consistency parameter μ for the ℓ_1-based method was again set to the β value determined by L-curve. The proposed iterative algorithm terminated when there was <1% signal change.

To quantify the stability of phase and susceptibility time series, time-SNR (t-SNR) and standard deviation maps were computed for the raw unwrapped phase, tissue phase after background removal, and ℓ_1- and ℓ_2-regularized susceptibility maps. The "noise" term involved in t-SNR computation was estimated by subtracting the time-average from each time point in the series. Taking the standard deviation of noise over time yielded the standard deviation maps. The "signal" term in the t-SNR metric was obtained by averaging the data over time, then taking the absolute value of this result. The ratio of the signal and standard deviation components yielded the t-SNR maps.

RESULTS

Numerical Susceptibility Phantom

Starting from the noisy field map in Figure 1a, the closed-form ℓ_2-reconstruction was completed in 0.3 seconds with 17.5% RMSE (Fig. 1b and 1c). The nonlinear CG method converged in 50 iterations and 258 s with 6.1% error (Fig. 1d and 1e). The proposed variable-splitting reconstruction was completed in 10 iterations and 13 s with 6.7% RMSE (Fig. 1f and 1g). The parameter setting that was used in these experiments was $\beta = 2.2 \cdot 10^{-4}$ (ℓ_2-recon), $\alpha = 1.5 \cdot 10^{-5}$ (CG), $\lambda = 10^{-5}$ (proposed), and these values were chosen to minimize the reconstruction error of each method.

When the proposed algorithm was run for 20 iterations (10 extra iterations), the error decreased to 6.1%, and the reconstruction time increased to 25 s. To test the effect of the parameter μ in the final susceptibility solution, a wide range of μ values were swept with 300 iterations. The value of $\mu = 2.2 \cdot 10^{-4}$ used in Fig. 1f yielded 5.95% RMSE. The same error of 5.95% was obtained at $\mu = 2.2 \cdot 10^{-3}$ and $2.2 \cdot 10^{-2}$, while this was 6.02% for $2.2 \cdot 10^{-5}$.

High-Resolution *In Vivo* 3D GRE

The proposed fast phase processing pipeline allowed unwrapping and background removal to be completed in 13 s for the high-resolution anatomical dataset (Fig. 2). Tracing the L-curve with the ℓ_2-based method took 42 seconds, where 15 reconstructions were made for logarithmically spaced β values between 10^{-3} and 1 (Fig. 3 top panel). The setting that maximized the curvature of the L-curve was found to be $\beta = 3.2 \cdot 10^{-2}$, which was taken to the optimal smoothing parameter value (Fig. 3b). Two additional reconstructions with under- and overregularization are depicted in Fig. 3a and 3c for comparison. The computation time for each ℓ_2-based susceptibility map was 0.9 s.

The smoothing parameter value corresponding to the largest curvature on the L-curve for proposed ℓ_1-regularization was found to be $\lambda = 9.2 \cdot 10^{-4}$. The total reconstruction time for 15 reconstructions with logarithmically spaced λ values between 10^{-4} and $10^{-2.5}$ was 710 s (Fig. 4, top panel). Underregularized, optimally regularized, and overregularized susceptibility maps are plotted in Figure 4a–4c for comparison. Processing time for the optimally smoothed reconstruction was 60 s, and the convergence criterion was met in 13 iterations.

The convergence speed of the ℓ_1-based method was dependent on the selection of μ parameter. The result reported in Figure 4b was obtained with the heuristically selected value of $\mu = 3.2 \cdot 10^{-1}$. If the value of μ were instead chosen to be 10 times larger ($3.2 \cdot 10^{-1}$), the 1% convergence criterion would be met in 20 iterations and 91 s. Using a parameter value 10 times smaller than the optimal ($3.2 \cdot 10^{-3}$), it would take 25 iterations and 114 seconds to converge. These are both slower than the optimal setting that required 13 iterations and 60 s.

The nonlinear CG smoothing parameter value that yielded matching data consistency with the proposed method was found to be $\alpha = 10^{-3}$. Computation took 1350 s for the nonlinear CG method, corresponding to 50 iterations (Fig. 5b). If the regularization parameter α had been estimated with the L-curve technique, the parameter sweep would have taken ~4 h with the CG algorithm. Regularized susceptibility maps with and without magnitude prior are presented in Figure 6. To facilitate comparison, maximum intensity projections are taken over 3-mm-thick slabs. Compared with the closed-form solution

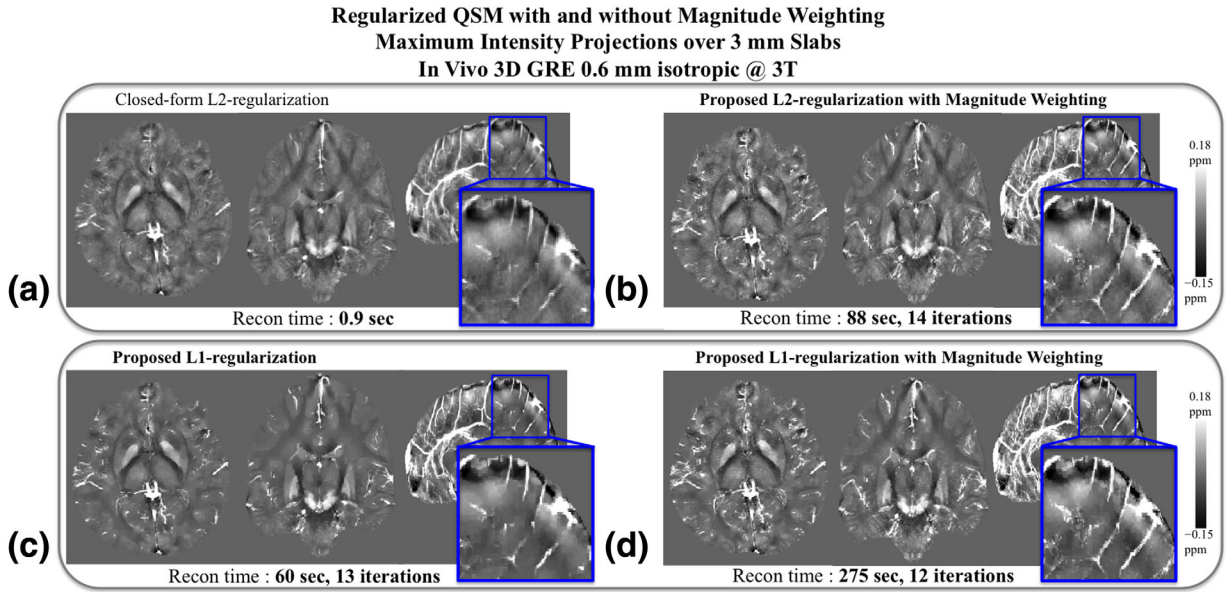

FIG. 6. L2- and L1-regularized QSM with and without magnitude prior. Compared with closed-form reconstruction (**a**) that is completed in 0.9 s, magnitude-weighted L2-regularization (**b**) requires 88 s of processing while increasing conspicuity of high-frequency structures like vessels. Relative to the proposed L1-based method (**c**), inclusion of magnitude prior (**d**) is computationally more demanding, requiring 275 s of reconstruction.

in Figure 6a that was computed in 0.9 s, ℓ_2-based reconstruction with magnitude prior in Figure 6b took 88 s and 14 iterations. While the proposed ℓ_1-based method without magnitude weighting (Fig. 6c) converged in 60 s, the processing time increased to 275 s when magnitude prior was included (Fig. 6d). To demonstrate the effect of magnitude weighting, sagittal planes are further zoomed in.

Maximum intensity projections of tissue phase and ℓ_2- and ℓ_1-based susceptibility maps with and without

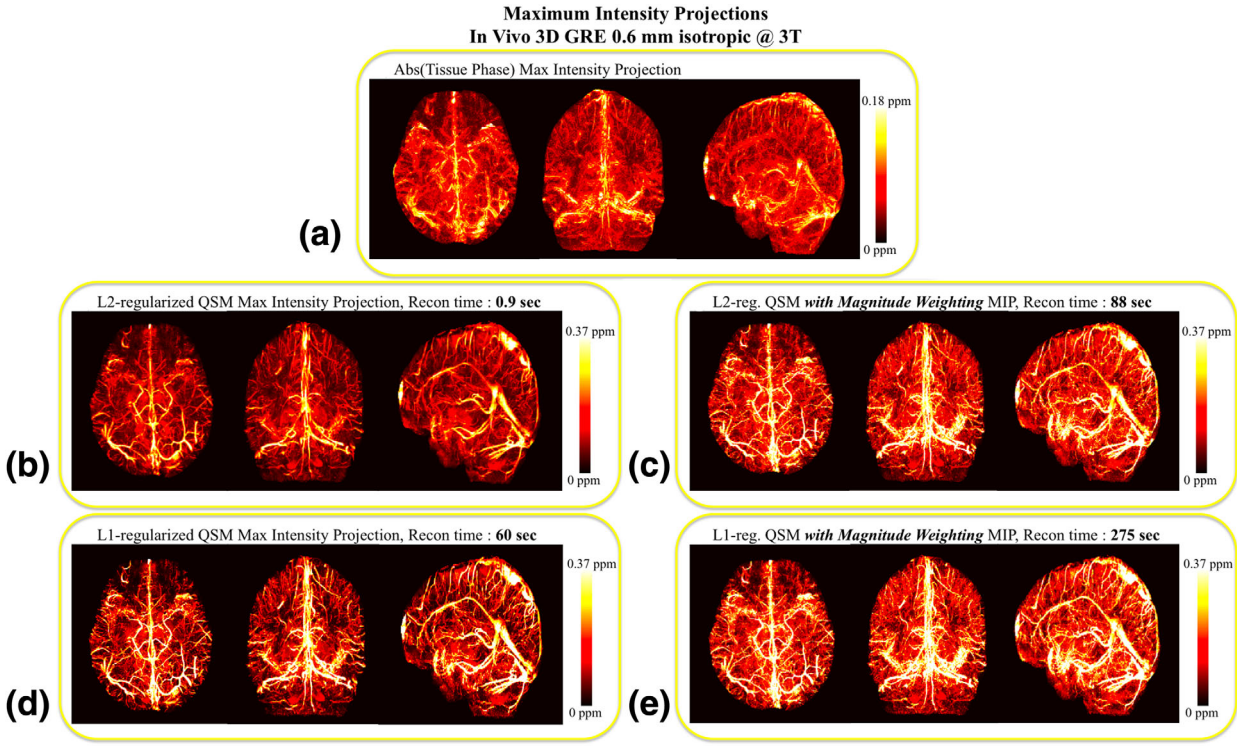

FIG. 7. Maximum intensity projections (MIPs) of in vivo 3D GRE dataset. **a:** Tissue phase MIP. **b:** Closed-form L2-based susceptibility map. **c:** Projection for the proposed L2-regularized QSM with magnitude prior. **d, e:** L1-based reconstruction without (d) and with (e) magnitude weighting. Note the increase in the vessel susceptibility values estimated with the methods that use magnitude prior.

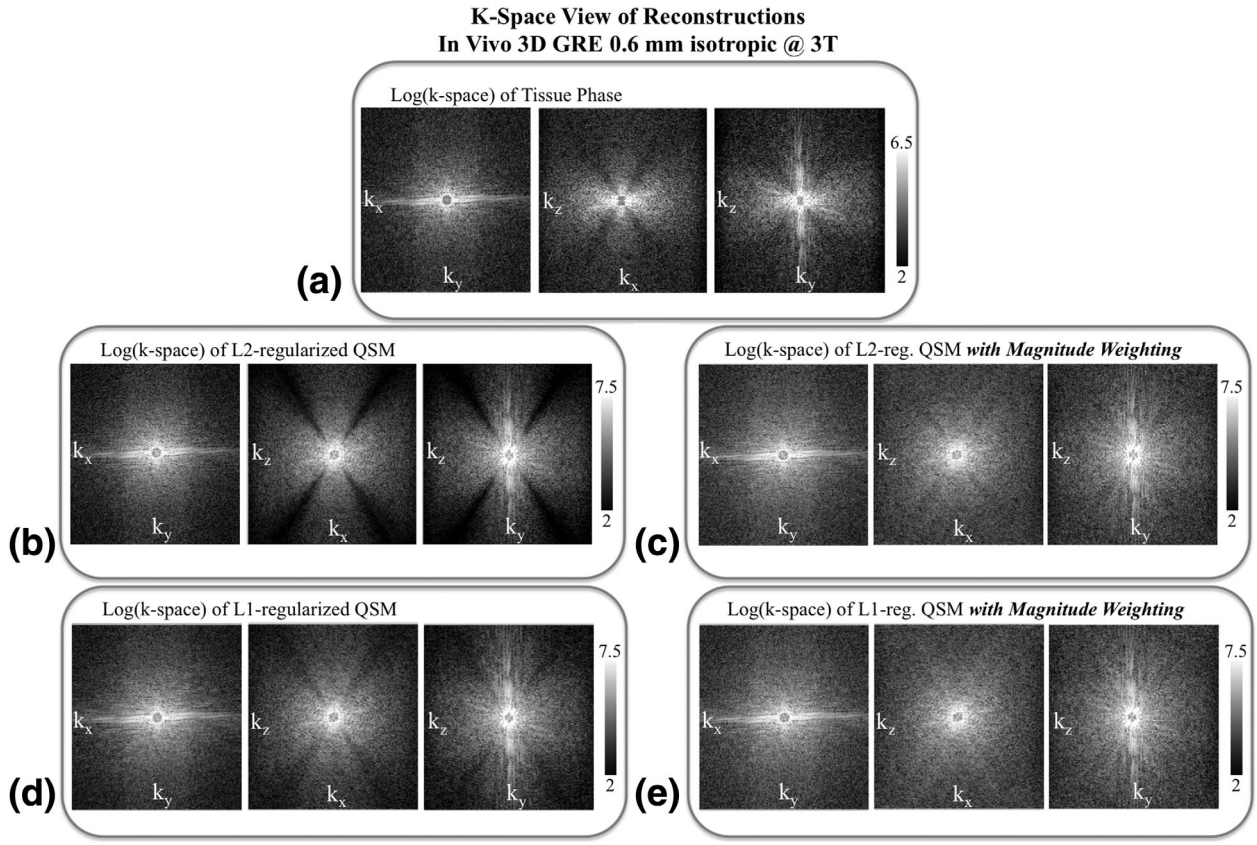

FIG. 8. k-Space views for tissue phase (**a**), closed-form QSM (**b**), proposed L2-regularization with magnitude weighting (**c**), L1-based reconstruction without (**d**), and with magnitude prior (**e**). Note the increase in the k-space content near the magic angle for the methods that use magnitude prior.

magnitude prior through axial, coronal, and sagittal planes are depicted in Figure 7. The corresponding k-space views are plotted in Figure 8.

In Vivo 2D EPI for Functional QSM

Obtaining the tissue phase for the 30 time frames using Laplacian unwrapping and SHARP filtering took 9 s in total (0.3 s per frame; Fig. 9). L-curve parameter value estimation was performed once on the average tissue phase, and the resulting parameter values were applied to the 30 frames. ℓ_2-Based parameter estimation took 2.7 s, whereas tracing the L-curve took 44 s with the proposed ℓ_1-based method. The ranges of the parameter sweeps were again $\beta \in \{10^{-3},1\}$ and $\lambda \in \{10^{-4},10^{-2.5}\}$ and the optimal parameter values were found to be $\beta = 3.2 \cdot 10^{-2}$ and $\lambda = 9.2 \cdot 10^{-4}$.

Closed-form ℓ_2-reconstruction completed in a total of 2.1 s (0.07 s per frame, Fig. 10a), whereas the proposed method completed in 192 s (6.4 s per frame, Fig. 10b).

Time-series standard deviation maps for phase and susceptibility signals are depicted in Figure 10c–10f. The standard deviation and t-SNR values averaged over the brain mask were $\sigma = 6 \cdot 10^{-3}$ and t-SNR = 19.1 for the raw phase (Fig. 10c), $\sigma = 3.3 \cdot 10^{-3}$ and t-SNR = 3.8 for the tissue phase (Fig. 10d), $\sigma = 4.4 \cdot 10^{-3}$ and t-SNR = 5.4 for ℓ_2-regularized susceptibility maps (Fig. 10e), and $\sigma = 4.7 \cdot 10^{-3}$ and t-SNR = 4.3 for ℓ_1-based susceptibility reconstruction (Fig. 10f).

Regularization parameters for the QSM algorithms employed in the reconstruction of the three datasets are summarized in Table 1.

DISCUSSION

This study introduces a fast total variation regularized QSM algorithm that achieves processing 20 times faster than a conventional nonlinear CG solver. Compared with the previously reported computation times [20 min to several hours (12,17)], the proposed algorithm drastically reduces the processing time down to 1 min for a whole brain high-resolution dataset. This efficient formulation of ℓ_1-constrained reconstruction might therefore facilitate the clinical use of regularized susceptibility mapping. The increased computational efficiency stems from two sources: To reach the same convergence criterion, the proposed method requires $\sim 4\times$ fewer iterations. At each iteration, this algorithm uses 6 FFT operations, whereas the nonlinear CG solver needs 4 FFTs and 12 spatial differencing operations. The combined effect of the smaller number of iterations at reduced computational load leads to the observed speed-up. The convergence criterion employed for the nonlinear CG algorithm was <1% change in the image-space representation of the susceptibility map, while this criterion was 1% change in the k-space representation of the

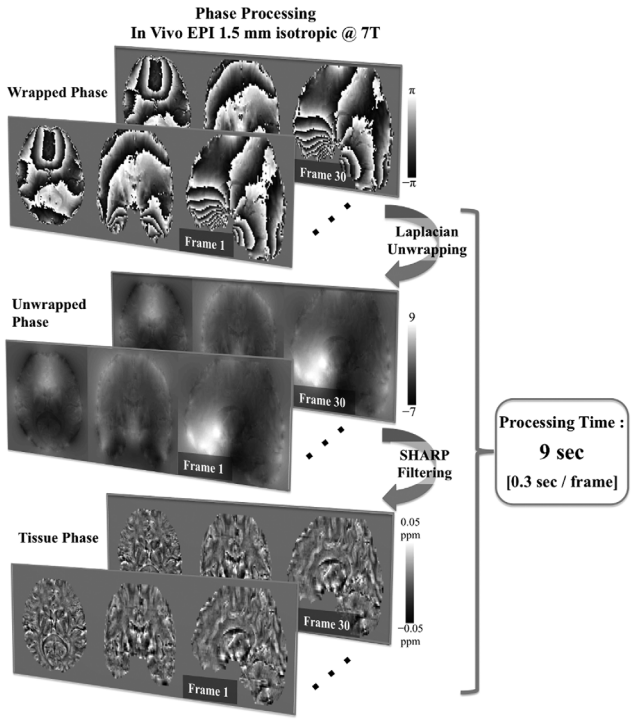

FIG. 9. Phase processing steps for in vivo EPI at 1.5 mm isotropic resolution acquired as a time series (frames 1–30). Starting from the coil-combined wrapped phase, application of Laplacian unwrapping and SHARP filtering took 9 s for the 30 frames, corresponding to a processing time of 0.3 s per frame.

susceptibility reconstruction for the split-Bregman algorithm. More explicitly, we compute $||\chi_{new}\text{-}\chi_{prev}||_2 \,/||\, \chi_{new}||_2$ for nonlinear CG, and $||\mathbf{F}\cdot\chi_{new}\text{-}\mathbf{F}\cdot\chi_{prev}||_2 \,/||\, \mathbf{F}\cdot\chi_{new}||_2$ for split-Bregman reconstruction. While the two stopping criteria are seemingly different, they are in fact equivalent due to Parseval's theorem, as $||\mathbf{F}\cdot\chi||_2 = ||\chi||_2$.

Combined with the previously proposed Laplacian phase unwrapping and SHARP filtering steps (14,22), the proposed algorithm comprises a fast processing pipeline that produces susceptibility maps from the wrapped phase data. As each dipole inversion can be completed rapidly, automatic estimation of the regularization parameter also becomes practical. The proposed method without magnitude prior is capable of tracing the L-curve faster than the time it would take the CG solver to complete a single reconstruction. For in vivo whole brain susceptibility mapping at 0.6 mm isotropic resolution, the complete pipeline of phase processing, parameter estimation, and regularized QSM takes under 14 min. Because the optimal smoothing parameters would be expected to remain within a certain range across different datasets, it would be possible to further reduce the parameter estimation time by considering a smaller interval of values. For studies that involve scanning multiple subjects with similar imaging parameters, it might also be possible to determine the regularization amount on one subject and apply this to the rest of the subjects in the study.

Variable-splitting methods are well studied and popular in signal processing community. These are also known as *split-Bregman* iterations (21) and have been deployed successfully in compressed sensing applications for MRI (34). A variable-splitting method for susceptibility mapping similar to the proposed algorithm was proposed in (23), which aims to solve

$$min_{\chi,y,v}\ \frac{1}{2}||v-\mathbf{F}^{-1}\mathbf{D}\mathbf{F}\chi-z_t||_2^2+\lambda\cdot||y||_1+\frac{\mu}{2}||\mathbf{y}-\mathbf{G}\chi-\eta_t||_2^2+\frac{\gamma}{2}||v\text{-}\phi||_2^2 \qquad [20]$$

Compared with our formulation in Equation [9], this introduces a new variable v and a third regularization parameter γ. Because the tissue phase ϕ is the observed data, it should not be necessary to further introduce an auxiliary variable v to penalize these observations. By omitting the term $\frac{\gamma}{2}||v\text{-}\phi||_2^2$, we avoid selection of the third parameter γ and simplify the reconstruction. Further, we extend the variable-splitting algorithm to admit magnitude weighting on the spatial gradients. Originally proposed in the influential MEDI papers (10–12), magnitude-weighted regularization aims to avoid smoothing across strong edges identified on the magnitude signal. Because the linear system involved in the reconstruction is no longer diagonal, we propose a preconditioned conjugate gradient solver for rapid matrix inversion. Despite the increased complexity of the optimization problem, the resulting ℓ_1-regularized algorithm with magnitude prior is still 5 times faster than the nonlinear CG technique.

Herein, ℓ_2-regularized QSM is also extended to admit prior information derived from the magnitude signal. While the resulting optimization problem is no longer solved in closed-form, the proposed preconditioned linear solver allows significant computational savings. To reach a tolerance level of 0.1%, preconditioned solver requires 14 iterations, which take 1.5 min for the high-resolution in vivo dataset. Without the preconditioner, the same tolerance level would have been reached with 30 iterations in more than twice the processing time. At 0.1% tolerance, the objective value of the ℓ_2-penalized minimization problem obtained with the proposed method was 381.9675. If the nonlinear CG algorithm were used, the objective value at the end of reconstruction would have been 381.9684, with a processing time of 22 min. Based on this, the proposed ℓ_2-based algorithm with magnitude prior is able to solve the optimization problem 15 times faster compared with nonlinear CG with similar accuracy.

Existing fast QSM algorithms include thresholded k-space division (TKD) (35,36) and the closed-form ℓ_2-constrained method that our group proposed recently (20). Although these techniques are computationally very efficient (requiring only two FFT operations), their regularization approaches might hamper the conspicuity of high spatial frequency components, such as vessels and small iron-rich structures like the red nucleus. This point can be appreciated on Figure 7b and 7d, where the vessels in the MIP image for ℓ_2-penalized reconstruction are visibly dimmer than their ℓ_1-penalized counterparts without magnitude prior.

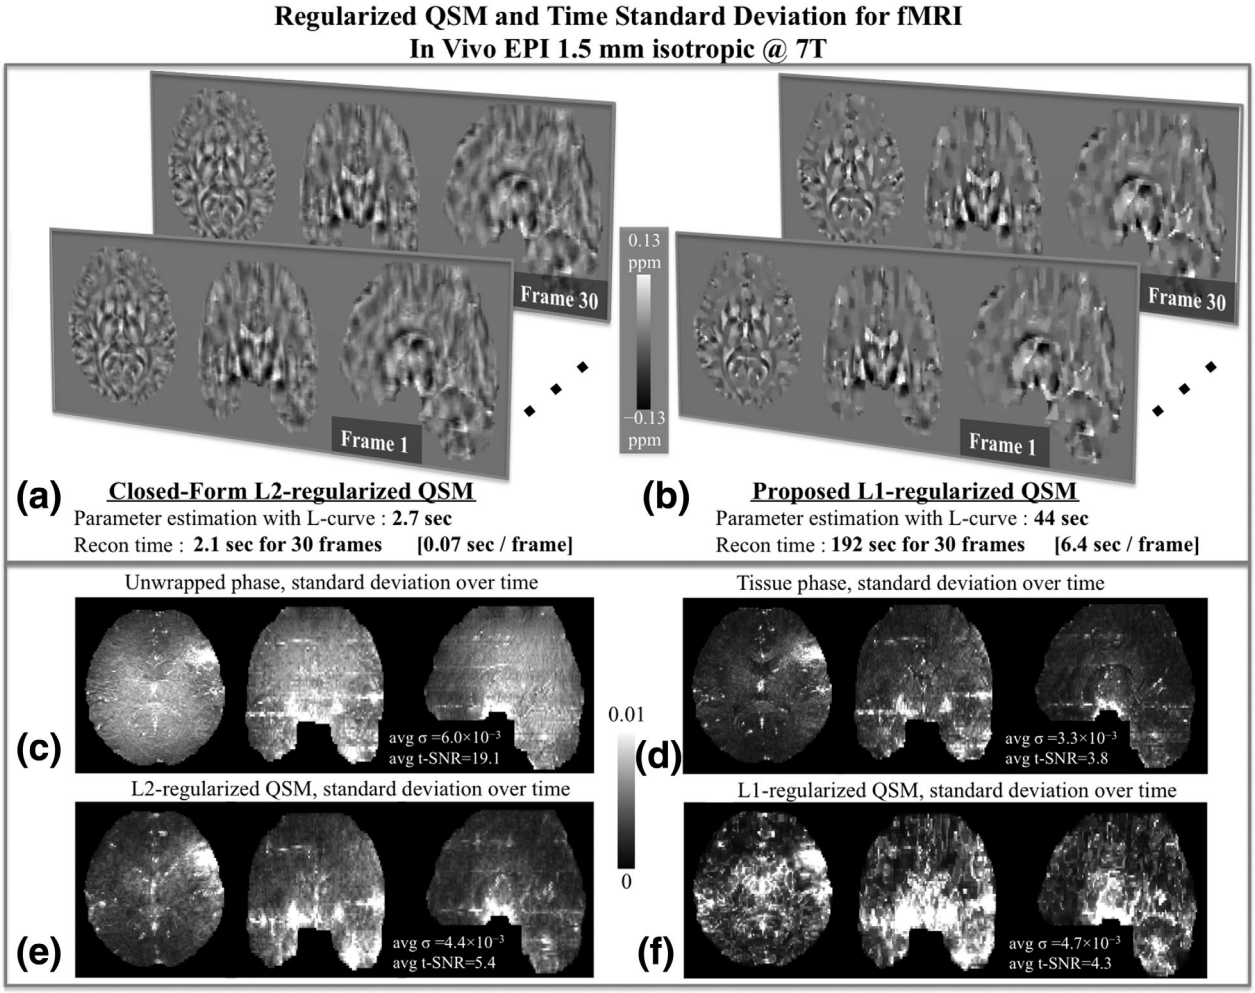

FIG. 10. Reconstruction of 30 frames of EPI data with the closed-form L2-regularized QSM (**a**) was completed in 2.1 s, corresponding to a speed of 0.07 s per frame. Using the proposed L1-based method (**b**), the reconstruction time was 192 s for the 30 frames, yielding a processing speed of 6.4 s per frame. L-curve parameter estimation took 2.7 s for L2-constrained and 44 s for L1-constrained reconstruction. **c–f:** Standard deviation maps of phase and susceptibility time series. The raw unwrapped phase (c) has a standard deviation of $6 \cdot 10^{-3}$ over time, and a t-SNR of 19.1 averaged inside the brain mask. For tissue phase (d), the values were $\sigma = 3.3 \cdot 10^{-3}$ and t-SNR $= 3.8$. L2-regularized susceptibility time series had $\sigma = 4.4 \cdot 10^{-3}$ and t-SNR $= 5.4$ (e), while L1-based reconstruction returned $\sigma = 4.7 \cdot 10^{-3}$ and t-SNR $= 4.3$ (f).

The difference stems from the fact that ℓ_2-regularization enforces a smooth susceptibility map, whereas ℓ_1-penalty imposes a piece-wise constant image model that preserves sharp edges. The k-space views in Figure 8b and 8d also reveal that ℓ_1-regularization is more successful at estimating the undersampled content in the vicinity of the magic angle relative to ℓ_2-based QSM when no magnitude weighting is employed. To quantify the improvement, k-space energy of the reconstructed maps at the frequencies where the susceptibility kernel dampens the field map [frequencies k that satisfy $\mathbf{D}(k_x, k_y, k_z) \leq 0.25$] was computed. In this case, ℓ_1-based maps had 11% higher energy than the ℓ_2-reconstruction. Considering only high frequency content at this threshold level ($\mathbf{D} \leq 0.25$ and $\sqrt{k_x^2 + k_y^2 + k_z^2} \geq 0.5 \cdot k_{max}$), ℓ_1-regularized map was found to have 69% higher k-space energy. However, the difference between the two regularization approaches would be smaller when average

Table 1
QSM Reconstruction Algorithms and Related Regularization Parameters

QSM Algorithm	Parameter	Numerical Phantom	In Vivo 3D GRE	In Vivo EPI
Closed-form ℓ_2-regularization	β (ℓ_2 penalty)	$2.2 \cdot 10^{-4}$	$3.2 \cdot 10^{-2}$	$3.2 \cdot 10^{-2}$
Nonlinear CG ℓ_1-regularization	α (ℓ_1 penalty)	$1.5 \cdot 10^{-5}$	$1.0 \cdot 10^{-3}$	—
Proposed ℓ_1-regularization	γ (ℓ_1 penalty)	$1.0 \cdot 10^{-5}$	$9.2 \cdot 10^{-4}$	$9.2 \cdot 10^{-4}$
	μ (gradient consistency)	$2.2 \cdot 10^{-4}$	$3.2 \cdot 10^{-2}$	$3.2 \cdot 10^{-2}$

susceptibility values are computed within anatomical boundaries of large iron-rich structures such as putamen and globus pallidus (16).

Magnitude-weighted regularization might mitigate this drawback of ℓ_2-based regularization through edge-aware smoothing. As seen in Figure 6a and 6b, the conspicuity of vessels increased substantially when magnitude prior was involved in ℓ_2-regularization. When Figure 6c and 6d are compared, the benefit of magnitude prior can also be observed in ℓ_1-regularization, albeit to a smaller degree. The effect of edge-aware dipole inversion can also be observed in the susceptibility MIPs (Fig. 7c and 7e) and in improved estimation of k-space content near the magic angle (Fig. 8c and 8e).

Increased computational efficiency would have a particular impact on processing large datasets. Recently introduced functional susceptibility mapping (6,24,25) involves reconstruction of a time series of three-dimensional phase images. The complete pipeline consisting of phase unwrapping, background removal, smoothing parameter value estimation, and dipole inversion for 30 time frames at 1.5 mm isotropic resolution required 4 min when state-of-the-art phase processing and the proposed algorithm were combined. The reconstruction time would otherwise exceed 1 h with the nonlinear CG algorithm. As such, the proposed algorithm is expected to facilitate the investigation of the relation between functional BOLD contrast and changes in the underlying susceptibility distribution. To mitigate the low-spatial frequency background variation in functional MRI (fMRI) phase images, a combination of phase images at different time points can be used (e.g., by subtraction of the first time point or of the average phase over time), and then χ can be computed (6). Whereas TKD and ℓ_2-regularization are linear in χ, ℓ_1-regularization is a nonlinear reconstruction technique. The resulting susceptibility maps will be independent of the order in which subtraction and dipole inversion are computed for TKD and ℓ_2-penalty; however, the ordering is important for ℓ_1-constrained inversion because of the thresholding step.

Phase processing and dipole inversion for EPI data acquired at 7T constitutes a challenging problem, particularly due to imperfections associated with echo-planar k-space trajectory. Although 1.5 mm isotropic voxel size (Figs. 9 and 10) is considered a high resolution for fMRI, it is considerably lower than the resolution of the 3D GRE data acquired at 0.6 mm isotropic voxel size (Figs. 2–8). As such, the EPI volumes do not have the same level of spatial detail as the GRE images; however, they are capable of representing variations in response to neural activity and physiological changes, because they constitute a time series dataset. Relative to the tissue phase images in Figure 9, loss of spatial resolution can be observed in the ℓ_2- and ℓ_1-constrained time series in Figure 10a and 10b. Although the smoothing parameters β and λ were selected with the L-curve method, the amount of regularization can also be tailored to a particular problem while using the L-curve parameter values as useful landmarks (e.g., a parameter 10% less than the L-curve selected value can be employed for reduced smoothing). Once the regularization parameter is selected on a single time point, the same value can be applied to the rest of the volumes in the time series. This point constitutes the major computational difference between the high-resolution 3D GRE dataset and the functional QSM experiment.

Based on confounding effects of instrumental and physiological origin in the EPI phase signal (6), we further acknowledge that extracting functional information is a difficult task that requires stability over the time points. To quantify the stability in the raw unwrapped phase, tissue phase, and ℓ_2- and ℓ_1-regularized time series, we report average time-SNR and maps of standard deviation over time in Figure 10c–10f. The streaking artifacts visible in these maps stem from imperfect estimation of the relative phase offset of each coil. This can be mitigated by estimating the offset from a dual-echo acquisition (29). As a result of deconvolution, ℓ_2- and ℓ_1-based susceptibility series exhibited a 33% and 42% increase in average standard deviation relative to the tissue phase. However, the t-SNR values in the susceptibility maps were larger by 42% and 13% when compared with the t-SNR value of the tissue phase, owing to the fact that the increase in the signal counterbalanced the increase in the noise. It can also be seen that the noise standard deviation for ℓ_1-based reconstruction had substantial spatial variation. Based on the nonlinear nature of regularization, smooth regions tended to remain below the ℓ_1-threshold, which led to small standard deviation over time. Further, a 5-fold reduction in t-SNR was observed between the raw unwrapped and the tissue phase. Because the tissue component is about an order of magnitude smaller than the background contribution, substantial reduction in the phase signal is expected when background phase is eliminated.

In addition to the regularization parameter λ that adjusts the contribution of the signal prior to the reconstructed susceptibility map, the proposed variable-splitting formulation introduced a second parameter μ that weighted the gradient consistency due to $(\mu \cdot ||y\text{-}\mathbf{G}\chi\text{-}\eta||_2^2)$. While Goldstein and Osher (21) showed that μ does not affect the solution but the speed of convergence, a suitable parameter value still needs to be selected. For the in vivo setting, we addressed the parameter identification problem by setting μ to the optimal ℓ_2-parameter β that was determined with the L-curve method. This heuristic selection was seen to yield favorable convergence speed. Based on the numerical phantom experiments detailed in the Results section, the same reconstruction error was obtained when the value of μ varied within three orders of magnitude. This shows that the same susceptibility map is obtained regardless of the value of μ. Regarding the convergence speed of the heuristically selected μ parameter, the experiments performed on the in vivo dataset demonstrated that using parameters that were 10 times larger or 10 times smaller led to slower convergence. As such, the heuristically selected parameter had favorable convergence characteristics in practice.

Limitations

To compensate for the noise variation in the field map, a diagonal weighting $\mathbf{M}$ proportional to the image

magnitude can be included in the data consistency term (10,12,13). With this refinement, the ℓ_2-constrained problem becomes

$$min_\chi ||\mathbf{M}(\mathbf{F}^{-1}\mathbf{D}\mathbf{F}\chi - \phi)||_2^2 + \beta \cdot ||\mathbf{W}\mathbf{G}\chi||_2^2 \quad [21]$$

The optimizer of this expression is given by the solution of

$$(\mathbf{D}^T\mathbf{F}\mathbf{M}^2\mathbf{F}^{-1}\mathbf{D} + \beta \cdot \mathbf{E}^H\mathbf{F}\mathbf{W}^2\mathbf{F}^{-1}\mathbf{E})\mathbf{F}\chi = \mathbf{D}^T\mathbf{F}\mathbf{M}^2\mathbf{F}^{-1}\mathbf{D}\mathbf{F}\phi \quad [22]$$

This system needs to be solved iteratively, and the preconditioner ($\mathbf{D} + \beta \cdot \mathbf{E}^2$) is expected to be less effective compared with its use without noise weighting **M**. A similar system arises in the update formula for ℓ_1-based reconstruction,

$$\begin{aligned}&(\mathbf{D}^T\mathbf{F}\mathbf{M}^2\mathbf{F}^{-1}\mathbf{D} + \mu\mathbf{E}^T\mathbf{F}\mathbf{W}^2\mathbf{F}^{-1}\mathbf{E})\mathbf{F}\chi_{t+1} \\ &\quad = \mathbf{D}^T\mathbf{F}\mathbf{M}^2\mathbf{F}^{-1}\mathbf{D}\mathbf{F}\phi + \mu\mathbf{E}^T\mathbf{F}\mathbf{W}^T(y_t - \eta_t)\end{aligned} \quad [23]$$

Employing efficient matrix factorization algorithms could potentially facilitate these more challenging matrix inversion problems (37,38).

Extensions

L-curve parameter estimation entails reconstruction with varying levels of regularization. Because each reconstruction is independent of the others, they could be performed in parallel for increased time efficiency.

The phase processing pipeline employed in the current work could be further refined. A regularized version of the SHARP filter was proposed recently to enhance the quality of the background phase removal (32,39). This improvement would, however, come at the cost of additional processing time, as this regularized formulation is solved iteratively. A second refinement would be to use a spatially varying SHARP kernel size, which would yield higher quality tissue phase inside the brain and reduce the amount erosion that needs to be applied on the mask boundary (40). This improvement is included in the software package that accompanies this publication.

CONCLUSION

This study introduces a variable-splitting algorithm that reduces the processing time of ℓ_1-regularized QSM by 20 times relative to the conventional nonlinear CG solver. Such efficient optimization also renders regularization parameter estimation with the L-curve method practical. Combined with state-of-the-art phase unwrapping and background removal techniques, the proposed algorithm comprises a pipeline that might facilitate clinical use of susceptibility mapping. This method is also extended to admit prior information derived from the magnitude signal for edge-aware regularization. The developed fast dipole inversion methods are expected to facilitate the investigation of the relation between the BOLD signal and the underlying tissue susceptibility changes by reconstructing four-dimensional time-series datasets in feasible time.

REFERENCES

1. Langkammer C, Schweser F, Krebs N, et al. Quantitative susceptibility mapping (QSM) as a means to measure brain iron? A post mortem validation study. NeuroImage 2012;62:1593–1599.
2. Zheng W, Nichol H, Liu S, Cheng Y, Haacke E. Measuring iron in the brain using quantitative susceptibility mapping and X-ray fluorescence imaging. NeuroImage 2013;78:68–74.
3. Fan AP, Bilgic B, Benner T, Rosen BR, Adalsteinsson E. Regularized quantitative susceptibility mapping for phase-based regional oxygen metabolism (PROM) at 7T. In Proceedings of the 19th Annual Meeting of ISMRM, Montreal, Canada, 2011. p. 7942.
4. Fan AP, Benner T, Bolar DS, Rosen BR, Adalsteinsson E. Phase-based regional oxygen metabolism (PROM) using MRI. Magn Reson Med 2012;67:669–78.
5. Haacke E, Tang J, Neelavalli J, Cheng Y. Susceptibility mapping as a means to visualize veins and quantify oxygen saturation. J Magn Reson Imaging 2010;32:663–676.
6. Bianciardi M, van Gelderen P, Duyn JH. Investigation of BOLD fMRI resonance frequency shifts and quantitative susceptibility changes at 7 T. Hum Brain Mapp 2013. doi: 10.1002/hbm.22320.
7. Langkammer C, Liu T, Khalil M, Enzinger C, Jehna M, Fuchs S, Fazekas F, Wang Y, Ropele S. Quantitative susceptibility mapping in multiple sclerosis. Radiology 2013;267:551–559.
8. Schweser F, Deistung A, Lehr B, Reichenbach J. Differentiation between diamagnetic and paramagnetic cerebral lesions based on magnetic susceptibility mapping. Med Phys 2010;37:5165–178.
9. 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 Part B: Magn Reson Eng 2005;25B:65–78.
10. De Rochefort L, Liu T, Kressler B, Liu J, Spincemaille P, Lebon V, Wu J, Wang Y. Quantitative susceptibility map reconstruction from MR phase data using bayesian regularization: validation and application to brain imaging. Magn Reson Med 2010;63:194–206.
11. Liu T, Liu J, de Rochefort L, Spincemaille P, Khalidov I, Ledoux JR, Wang Y. Morphology enabled dipole inversion (MEDI) from a single-angle acquisition: comparison with COSMOS in human brain imaging. Magn Reson Med 2011;66:777–83.
12. Liu J, Liu T, de Rochefort L, et al. Morphology enabled dipole inversion for quantitative susceptibility mapping using structural consistency between the magnitude image and the susceptibility map. NeuroImage 2012;59:2560–2568.
13. Liu T, Spincemaille P, de Rochefort L, Kressler B, Wang Y. Calculation of susceptibility through multiple orientation sampling (COSMOS): a method for conditioning the inverse problem from measured magnetic field map to susceptibility source image in MRI. Magn Reson Med 2009;61:196–204.
14. Schweser F, Deistung A, Lehr BW, Reichenbach JR. Quantitative imaging of intrinsic magnetic tissue properties using MRI signal phase: an approach to in vivo brain iron metabolism? NeuroImage 2011;54:2789–2807.
15. Wharton S, Bowtell R. Whole-brain susceptibility mapping at high field: a comparison of multiple- and single-orientation methods. NeuroImage 2010;53:515–25.
16. Bilgic B, Pfefferbaum A, Rohlfing T, Sullivan E V, Adalsteinsson E. MRI estimates of brain iron concentration in normal aging using quantitative susceptibility mapping. NeuroImage 2012;59: 2625–35.
17. Wu B, Li W, Guidon A, Liu C. Whole brain susceptibility mapping using compressed sensing. Magn Reson Med 2012;67:137–147.
18. Hansen PC. The L-curve and its use in the numerical treatment of inverse problems. In Johnston P, editor. Computational inverse problems in electrocardiology. Southampton: WIT Press; 2000. p. 119–142.
19. Morozov V. On the solution of functional equations by the method of regularization. Soviet Math Dokl 1966;7:414–417.
20. Bilgic B, Chatnuntawech I, Fan AP, Adalsteinsson E. Regularized QSM in seconds. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. p. 168.
21. Goldstein T, Osher S. The split Bregman method for L1-regularized problems. SIAM J Imaging Sci 2009;2:323–343.
22. Li W, Wu B, Avram A V, Liu C. Magnetic susceptibility anisotropy of human brain in vivo and its molecular underpinnings. NeuroImage 2012;59:2088–2097.

23. Chen Z, Calhoun V. Computed inverse resonance imaging for magnetic susceptibility map reconstruction. J Comput Assist Tomogr 2012;36:265–274.
24. Balla DZ, Sanchez-Panchuelo RM, Wharton S, Hagberg GE, Scheffler K, Francis ST, Bowtell RW. Functional quantitative susceptibility mapping (fQSM). In Proceedings of the 20th Annual Meeting of ISMRM, Melbourne, Australia, 2012. p. 325.
25. Balla DZ, Sanchez-Panchuelo RM, Wharton S, Hagberg GE, Scheffler K, Francis ST, Bowtell RW. Experimental investigation of the relation between gradient echo BOLD fMRI contrast and underlying susceptibility changes at 7T. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. p. 300.
26. Liu T, Xu W, Spincemaille P, Avestimehr AS, Wang Y. Accuracy of the morphology enabled dipole inversion (MEDI) algorithm for quantitative susceptibility mapping in MRI. IEEE Trans Med Imaging 2012;31:816–824.
27. Duyn J, van Gelderen P, Tie-Qiang L, de Zwart JA, Koretsky AP, Fukunaga M. High-field MRI of brain cortical substructure based on signal phase. Proc Natl Acad Sci U S A 2007;104:11796–11801.
28. Deistung A, Dittrich E, Sedlacik J, Rauscher A, Reichenbach JR. ToF-SWI: simultaneous time of flight and fully flow compensated susceptibility weighted imaging. J Magnc Reson Imaging 2009;29:1478–1484.
29. Robinson S, Grabner G, Witoszynskyj S, Trattnig S. Combining phase images from multi-channel radiofrequency coils using 3D phase offset maps derived from a dual-echo scan. Magn Reson Med 2011;65: 1638–1648.
30. Hammond KE, Lupo JM, Xu D, Metcalf M, Kelley DAC, Pelletier D, Chang SM, Mukherjee P, Vigneron DB, Nelson SJ. Development of a robust method for generating 7.0 T multichannel phase images of the brain with application to normal volunteers and patients with neurological diseases. NeuroImage 2008;39:1682–1692.
31. Smith S. Fast robust automated brain extraction. Hum Brain Mapp 2002;17:143–155.
32. Sun H, Wilman AH. Susceptibility mapping using regularization enabled harmonic artifact removal. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. p. 169.
33. Keil B, Triantafyllou C, Hamm M, Wald LL. Design optimization of a 32-channel head coil at 7 T. In Proceedings of the 18th Annual Meeting of ISMRM, Stockholm, Sweden, 2010. p. 1493.
34. Montesinos P, Abascal JJ, Chamorro J, Chavarrias C, Benito M, Vaquero JJ, Desco M.High-resolution dynamic cardiac MRI on small animals using reconstruction based on Split Bregman methodology. Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) 2011:3462–3464.
35. Schweser F, Deistung A, Sommer K, Reichenbach JR. Toward online reconstruction of quantitative susceptibility maps: superfast dipole inversion. Magn Reson Med 2012;69:1581–1593.
36. Shmueli K, Zwart J De. Magnetic susceptibility mapping of brain tissue in vivo using MRI phase data. Magn Reson Med 2009;62:1510–1511.
37. Xia J, Chandrasekaran S, Gu M, Xiaoye S. Fast algorithms for hierarchically semiseparable matrices. Numerical Linear Algebra with Applications 2010;17:953–976.
38. Xia J, Chandrasekaran S, Gu M, Li X. Superfast multifrontal method for large structured linear systems of equations. SIAM Journal on Matrix Analysis and Applications 2009;31:1382–1411.
39. Sun H, Wilman A. Background field removal using spherical mean value filtering and Tikhonov regularization. Magn Reson Med 2013. doi: 10.1002/mrm.24765.
40. Wu B, Li W, Avram AV, Gho S-M, Liu C. Fast and tissue-optimized mapping of magnetic susceptibility and T2* with multi-echo and multi-shot spirals. NeuroImage 2012;59:297–305.