Published in final edited form as:

J Magn Reson Imaging. 2014 July ; 40(1): 181–191. doi:10.1002/jmri.24365.

Fast Image Reconstruction with L2-Regularization

Berkin Bilgic, PhD^{1,2,*}, Itthi Chatnuntawech, BS¹, Audrey P. Fan, SM¹, Kavin Setsompop, PhD^{2,3}, Stephen F. Cauley, PhD², Lawrence L. Wald, PhD^{2,3,4}, and Elfar Adalsteinsson, PhD^{1,4}

¹Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

²A. A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Charlestown, Massachusetts, USA

³Harvard Medical School, Boston, Massachusetts, USA

⁴Harvard-MIT Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Abstract

Purpose—We introduce L2-regularized reconstruction algorithms with closed-form solutions that achieve dramatic computational speed-up relative to state of the art L1- and L2-based iterative algorithms while maintaining similar image quality for various applications in MRI reconstruction.

Materials and Methods—We compare fast L2-based methods to state of the art algorithms employing iterative L1- and L2-regularization in numerical phantom and in vivo data in three applications; 1) Fast Quantitative Susceptibility Mapping (QSM), 2) Lipid artifact suppression in Magnetic Resonance Spectroscopic Imaging (MRSI), and 3) Diffusion Spectrum Imaging (DSI). In all cases, proposed L2-based methods are compared with the state of the art algorithms, and two to three orders of magnitude speed up is demonstrated with similar reconstruction quality.

Results—The closed-form solution developed for regularized QSM allows processing of a 3D volume under 5 seconds, the proposed lipid suppression algorithm takes under 1 second to reconstruct single-slice MRSI data, while the PCA based DSI algorithm estimates diffusion propagators from undersampled q-space for a single slice under 30 seconds, all running in Matlab using a standard workstation.

Conclusion—For the applications considered herein, closed-form L2-regularization can be a faster alternative to its iterative counterpart or L1-based iterative algorithms, without compromising image quality.

Keywords

regularization; susceptibility mapping; diffusion imaging; spectroscopic imaging; lipid suppression

*Correspondence to: Berkin Bilgic, Massachusetts Institute of Technology, Room 36-776A, 77 Massachusetts Avenue, Cambridge, MA 02139, berkin@mit.edu, Fax: 617-324-3644, Phone: 617-866-8740.

Introduction

Regularized image reconstruction aims to improve image quality by imposing prior knowledge on the target signals. Regularization proves beneficial in a wide range of applications including parallel imaging (1,2), compressed sensing (3), denoising (4) and solution of inverse problems in general. Given a linear system $\mathbf{A} \cdot \mathbf{x} = \mathbf{b}$, where $\mathbf{A}$ is the observation model, $\mathbf{x}$ is the unknown signal and $\mathbf{b}$ are the acquired data, the most commonly encountered regularizers employ ℓ_2 or ℓ_1 penalties either on the reconstructed signal itself, or on its representation with respect to a transform $\mathbf{C}$ by solving,

$$\ell_2\text{-regularization: } \min \|\mathbf{A} \cdot \mathbf{x} - \mathbf{b}\|_2^2 + \alpha \cdot \|\mathbf{C} \cdot \mathbf{x}\|_2^2 \quad [1]$$

$$\ell_1\text{-regularization: } \min \|\mathbf{A} \cdot \mathbf{x} - \mathbf{b}\|_2^2 + \beta \cdot \|\mathbf{C} \cdot \mathbf{x}\|_1 \quad [2]$$

While ℓ_2 (Tikhonov)-regularized reconstruction in Eq.1 admits a closed-form solution $(\mathbf{A}^H\mathbf{A} + \alpha \cdot \mathbf{C}^H\mathbf{C})^{-1}\mathbf{A}^H\mathbf{b}$ when the inverse exists, established methods often operate iteratively, either because the system is too large to invert explicitly, or because $\mathbf{A} \cdot \mathbf{x}$ can be computed efficiently (e.g. Fast Fourier Transform) without having to store the matrix $\mathbf{A}$. On the other hand, ℓ_1 -penalized reconstruction in Eq.2 does not have a closed-form solution, and most compressed sensing algorithms operate iteratively by alternating between a soft thresholding step and ensuring consistency of the system $\mathbf{A} \cdot \mathbf{x} = \mathbf{b}$, e.g. (2,5,6).

This work presents ℓ_2 -based methods with closed-form solutions which can be computed efficiently, while retaining similar image quality as the iterative state of the art algorithms. We demonstrate the generality of the approach by assessing its performance and speed compared to more standard methods in three different applications. The main contributions include,

- i. **Derivation of a closed-form solution to regularized Quantitative Susceptibility Mapping (QSM):** The proposed method requires only two FFTs, and is three orders of magnitude faster than the existing iterative solvers (7,8), while giving the exact minimizer of the optimization problem.
- ii. **Fast lipid suppression for Magnetic Resonance Spectroscopic Imaging (MRSI):** By requiring the lipid and metabolite spectra to be approximately orthogonal, effective lipid artifact reduction is demonstrated. The present ℓ_2 -based method is three orders of magnitude faster than its previously proposed, iterative ℓ_1 -based counterpart (9,10), with similar lipid reduction performance. It can also be synergistically combined with the dual-density lipid suppression method (11,12) to yield artifact free spectra in less than a second per slice.
- iii. **Fast Diffusion Spectrum Imaging (DSI) reconstruction from undersampled q-space:** Applying Principal Component Analysis (PCA) on a training dataset from one subject captures the structure of diffusion propagators with a low-dimensional representation. Propagators of another subject are then estimated in this PCA space from undersampled q-space by solving a simple least-squares problem. The

proposed PCA-based method is two orders of magnitude faster with similar image quality compared to a previously proposed technique that applies compressed sensing with a dictionary trained for sparse representation. Dictionary training and reconstruction from undersampled q-space are also much simpler, requiring only linear algebra operations.

Accompanying Matlab code submitted as Supplementary Material reproduces the key figures presented in the manuscript, and will also be available online on the author's website.

The proposed and the state of the art algorithms are detailed in the Materials and Methods section along with data acquisition and reconstruction methods. Quantitative comparisons of computation time and image quality are presented under the Results header.

Materials and Methods

All computational results reported in this work were obtained in Matlab environment running on a workstation with 48 GB memory and 12 processors.

Regularized Quantitative Susceptibility Mapping with closed-form solution

QSM aims to map the tissue magnetic susceptibility χ based on the measured tissue phase ϕ . Susceptibility is highly correlated with tissue iron deposition especially in the deep gray matter structures (13), hence its estimation allows quantification of tissue iron concentration. As the innate paramagnetic effect of deoxyhemoglobin leads to a susceptibility shift between the veins and the surrounding tissue, susceptibility mapping also permits estimation of blood oxygenation level (14). However, the mapping requires the solution of the system

$$\mathbf{F}^H \mathbf{D} \mathbf{F} \chi = \phi \quad [3]$$

for the unknown χ , where $\mathbf{F}$ is the Fourier transform and $\mathbf{D}$ is a diagonal matrix with entries $1/3 - k_z^2/k^2$. As the kernel $\mathbf{D}$ equals to zero on the conical surface $2 \cdot k_z^2 = k_x^2 + k_y^2$ and effectively undersamples the frequency content of χ . As such, Eq.3 is an ill-posed problem and its solution is facilitated by additional information about the underlying susceptibility map. This information is either provided by acquiring additional observations where the object is tilted at various angles with respect to the main field (15), or by imposing a spatial prior about the susceptibility distribution via regularization (7). While the maps obtained from multi-orientation measurements were seen to have higher quality than the regularized single-orientation reconstructions (16), this benefit comes at the expense of substantially increased scan time. As such, regularized QSM remains an important tool that aims to solve,

$$\min \|\mathbf{F}^H \mathbf{D} \mathbf{F} \chi - \phi\|_2^2 + \lambda \cdot \mathcal{R}(\chi) \quad [4]$$

Regularizers employed in previous works are of the form $\mathcal{R}(\chi) = \|\mathbf{W} \mathbf{G} \chi\|_2^2$ or $\|\mathbf{W} \mathbf{G} \chi\|_1$, where $\mathbf{W}$ is either the identity $\mathbf{I}$ or a diagonal weighting matrix derived from the magnitude signal (7) and $\mathbf{G} = [\mathbf{G}_x; \mathbf{G}_y; \mathbf{G}_z]$ is the gradient operator in three dimensions. Average susceptibility values in iron rich gray matter structures have been computed with ℓ_2 - and ℓ_1 -

based reconstructions in the literature, and were reported to yield similar correlations with tissue iron concentration (8). However, a closed-form solution to Eq.4 exists in the case of ℓ_2 -regularization:

$$\hat{\chi} = (\mathbf{F}^H \mathbf{D}^H \mathbf{D} \mathbf{F} + \lambda \cdot \mathbf{G}^H \mathbf{G})^{-1} \mathbf{F}^H \mathbf{D}^H \mathbf{F} \phi, \text{ with } \mathbf{W} = \mathbf{I} \quad [5]$$

Since the matrix inversion involved in Eq.5 is computationally prohibitive, existing methods operate iteratively to minimize the objective in Eq.4. The proposed closed-form solution relies on computing the image gradients in k-space rather than the image space, thus making direct inversion of the system possible. The gradient along the x-axis can be expressed as

$\mathbf{G}\mathbf{x} = \mathbf{F}^H \mathbf{E}_x \mathbf{F}$, where $\mathbf{E}_x$ is a diagonal matrix with entries $E_x(i, i) = 1 - e^{(-2\pi \sqrt{-1} k_x(i, i)/N_x)}$, which is the k-space representation of the difference operator $\delta_x - \delta_{x-1}$. Here, k_x is the k-space index and N_x is the matrix size along x, and $\mathbf{G}_y$ and $\mathbf{G}_z$ are similarly defined. With this formulation, the term $\mathbf{G}^H \mathbf{G}$ becomes $\mathbf{F}^H (\mathbf{E}_x^2 + \mathbf{E}_y^2 + \mathbf{E}_z^2) \mathbf{F}$, and after simplifications the closed-form solution can be expressed as

$$\hat{\chi} = \mathbf{F}^H \mathbf{D} [\mathbf{D}^2 + \lambda \cdot (\mathbf{E}_x^2 + \mathbf{E}_y^2 + \mathbf{E}_z^2)]^{-1} \mathbf{F} \phi \quad [6]$$

The matrix inversion now involves only diagonal matrices, hence it is straight forward to evaluate. The total cost of susceptibility mapping is two FFTs and multiplication of diagonal matrices.

QSM Reconstruction Experiments

Two reconstruction methods that minimize the objective function in Eq.4 are considered:

- i. Nonlinear conjugate gradient (CG) algorithm (7,8) using 100 iterations, and
- ii. Proposed closed-form solution.

Experiments were performed on a numerical phantom and on in vivo data.

- i. **Numerical phantom:** consists of 3-compartments (gray and white matter, CSF) with a matrix size of $240 \times 240 \times 154$. Within each compartment, χ is constant and equal to $\chi_{gray} = -0.023$, $\chi_{white} = 0.027$, $\chi_{CSF} = -0.018$ ppm (17). The field map ϕ was computed from this ground truth χ map using forward dipole model $\phi = \mathbf{F}^H \mathbf{D} \mathbf{F} \chi$, and Gaussian noise with peak-SNR = 100 was added, so that the normalized root-mean-square-error (RMSE) of the noisy field map was 5.9 % relative to the noise free phase (Fig.1, first row). λ was chosen to minimize the RMSE in the χ maps reconstructed with the closed-form method, and was found to be $\lambda = 2 \cdot 10^{-4}$. The same λ was used for both the closed-form and iterative CG reconstructions.
- ii. **In vivo data:** 3D SPGR (spoiled gradient echo) data were acquired on a healthy subject at 1.5T with resolution $0.94 \times 0.94 \times 2.5 \text{ mm}^3$, matrix size $256 \times 256 \times 62$ and TR/TE = 58ms/40ms. Background phase was removed using dipole fitting (18) to yield the tissue phase (Fig.2, first row). $\lambda = 1.5 \cdot 10^{-2}$ was chosen based on the L-curve heuristic (19). Data were zero-padded to twice the size to avoid aliasing with circular convolution.

Fast Lipid Suppression for MR Spectroscopic Imaging

Estimation and visualization of biochemical metabolites in the brain, especially the NAA (N-acetyl Aspartate) peak at 2 ppm is made difficult by the ringing artifacts caused by the subcutaneous lipid signals around the skull. These artifacts are caused by side lobes of the point spread function arising from the limited spatial resolution of spectroscopic imaging techniques. Because the lipid signals have much higher signal amplitude than the cortical metabolites, especially at short echo times, the ringing artifacts severely contaminate the brain spectra and impede the detection of metabolite signals. In addition to important contributions on the excitation side (e.g. outer volume suppression (20), inversion recovery (21)), several post-processing methods have been proposed to mitigate lipid artifacts. These algorithms include data extrapolation (22), dual-density reconstruction (11,12), and lipid-basis penalty (9,10). In particular, lipid-basis method operates iteratively, and can be synergistically combined with dual-density sampling to yield effective lipid suppression (9). In more detail,

- i. **Dual-density reconstruction:** requires acquisition of high-resolution spectra csi_{high} to generate a lipid image, and low-resolution data csi_{low} with adequate SNR for metabolite signal quantification. With the help of a binary mask M_{lipid} that selects the lipid ring, a high-resolution lipid image is generated as:

$$lipid = M_{lipid} \cdot csi_{high} \quad [7]$$

This masked lipid image is then combined with the low-resolution spectra csi_{low} in k-space via

$$dual = F^{-1}(F_{high} \cdot lipid + F_{low} \cdot csi_{low}) \quad [8]$$

Here, F_{low} samples the low-resolution k-space, while F_{high} selects the peripheral k-space. The dual-density image $dual$ is then generated by combining the low spatial frequency content in the metabolite image csi_{low} and the high frequency content of the $lipid$ image.

- ii. **Lipid-basis penalty:** relies on the approximation that lipid and metabolite spectra are orthogonal to each other. This prior is enforced via the following optimization problem,

$$\min \|x - dual\|_2^2 + \alpha \cdot \sum_{i \in M_{brain}} \|L^H x_i\|_1 \quad [9]$$

where x_i is the spectrum in the i^{th} voxel, α is a regularization parameter, M_{brain} is the binary brain mask, and L is the lipid-basis matrix. Spectra from $dual$ inside the lipid mask are used to generate L , so that each column of L is a lipid spectrum sampled from the dual-density image. In essence, Eq.9 minimizes the sum of inner products between lipid and target metabolite spectra and is solved iteratively by gradient descent methods (10).

- iii. **Proposed lipid-basis reconstruction with ℓ_2 -regularization:** Instead of summing the absolute value of inner products, a simplified closed-form solution can be obtained by considering the *square* of inner product terms:

$$\min \|x - \mathbf{dual}\|_2^2 + \beta \cdot \sum_{i \in \mathbf{M}_{brain}} \|\mathbf{L}^H x_i\|_2^2 \quad [10]$$

This can be further simplified by noticing that the optimization problem is independent across voxels,

$$\begin{aligned} \text{if } i \in \mathbf{M}_{brain}, \quad & \min \|x_i - \mathbf{dual}_i\|_2^2 + \beta \cdot \|\mathbf{L}^H x_i\|_2^2, \\ \text{other wise,} \quad & x_i = \mathbf{dual}_i \end{aligned} \quad [11]$$

The solution can be evaluated in closed-form,

$$\begin{aligned} \text{if } i \in \mathbf{M}_{brain}, \quad & x_i = (\mathbf{I} + \beta \cdot \mathbf{L}\mathbf{L}^H)^{-1} \mathbf{dual}_i \\ \text{other wise,} \quad & x_i = \mathbf{dual}_i \end{aligned} \quad [12]$$

The reconstruction matrix $(\mathbf{I} + \beta \cdot \mathbf{L}\mathbf{L}^H)^{-1}$ needs to be computed only once, and the matrix inversion is of modest size (e.g. 512×512 for a 512-point frequency axis). Per voxel, the computational cost is a matrix-vector multiplication.

Lipid Suppression Experiments

Single slice constant density spiral MRSI data were acquired in vivo at 3T using 32-channel receive array with a voxel size of 0.16 mL (FOV = 24 cm, slice thickness = 1 cm, echo time = 50 ms, repetition time = 2 s, number of averages = 20, acquisition time = 33 min). Chemical shift selective suppression (CHESS (23)) pulse was applied for water suppression, and PRESS-box (point resolved spectroscopy (24)) excited the entire field-of-view; however, no lipid suppression was applied during acquisition. This spiral acquisition was coil combined after gridding onto a Cartesian grid, and all subsequent processing was applied on this grid. Lipid suppression was performed with three different methods:

- i. **Dual-density reconstruction:** To emulate dual-density sampling strategy, two datasets were derived from the 20-average, 0.16 mL resolution image. The high-resolution $\mathbf{csi}_{high}$ has 2 averages of 0.16 mL spectra and corresponds to a 3.3 min acquisition. The low-resolution $\mathbf{csi}_{low}$ has 20 averages of 0.56 mL data with a corresponding acquisition time of 10 min. The artifact reduced image $\mathbf{dual}$ was obtained by the combination of the two images due to Eq.8 (Fig.3). The lipid and brain masks were generated manually with the aid of a thresholded lipid image derived from non-lipid-suppressed dataset.
- ii. **Lipid-basis reconstruction with ℓ_1 -penalty:** Further lipid suppression was applied to $\mathbf{dual}$ due to Eq.9, and optimization was performed with the conjugate gradient algorithm. The regularization parameter α was chosen with the L-curve heuristic (19) and was found to be $\alpha = 10^{-3}$.
- iii. **Proposed lipid-basis reconstruction with ℓ_2 -penalty:** Further suppression was applied to $\mathbf{dual}$ due to closed-form expression in Eq.12. The regularization

parameter β was tuned so that the data consistency $\|x - dual\|_2^2$ term was the same for the ℓ_1 - and ℓ_2 -regularized reconstructions and was found to be $\beta = 0.65$. Setting the data consistency levels to be the same aimed to allow a fair comparison of lipid suppression performances of the two methods.

Fast undersampled Diffusion Spectrum Imaging (DSI) reconstruction

DSI is a diffusion imaging technique that involves sampling of full q-space and yields a complete description of water diffusion in terms of diffusion propagators that represent the local probability of water motion across space. The relation between the q-space samples and the diffusion propagator inside a voxel is given by $\mathbf{F} \cdot \mathbf{p} = \mathbf{q}$, where $\mathbf{F}$ represents the Fourier transform, while $\mathbf{p}$ and $\mathbf{q}$ are the probability-space and q-space descriptions of diffusion. However, the full-sampling requirement of DSI is prohibitive for clinical applications, as data acquisition takes ~1 hour. Existing compressed sensing algorithms aim to recover diffusion propagators from undersampled q-space. These include,

- i. **Wavelet & TV regularization (25):** places a sparsity prior on the diffusion propagators with respect to wavelet and total variation (TV) transforms and iteratively solves,

$$\min \|\mathbf{F}_\Omega \cdot \mathbf{p} - \mathbf{q}\|_2^2 + \alpha \cdot \|\Psi \cdot \mathbf{p}\|_1 + \beta \cdot \mathbf{TV}(\mathbf{p}) \quad [13]$$

where $\mathbf{F}_\Omega$ is the undersampled Fourier transform, Ψ is a wavelet operator, and $\mathbf{TV}(\cdot)$ is the total variation penalty.

- ii. **Dictionary-FOCUSS (26):** Starting from a training dataset of propagators $\mathbf{P}$, the K-SVD algorithm (26)
- iii. is employed to generate a dictionary $\mathbf{D}$ for sparse representation of these propagators. Compressed sensing reconstruction with respect to this dictionary is carried out via the FOCUSS (27) algorithm,

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{x}\|_1 \text{ such that } \mathbf{F}_\Omega \mathbf{D} \cdot \mathbf{x} = \mathbf{q} \quad [14]$$

Here, $\mathbf{x}$ are the dictionary transform coefficients and the reconstructed diffusion propagator $\mathbf{p}$ is obtained by the mapping $\mathbf{p} \hat{=} \mathbf{D} \cdot \mathbf{x}$.

Both of these sparsity-based methods operate iteratively and are computationally expensive. The proposed PCA-based algorithm simplifies both the training and reconstruction steps, and the solution can be computed in closed-form:

PCA-based fast DSI reconstruction—Again starting with a training set $\mathbf{P}$ consisting of L propagators, PCA finds a low-dimensional representation that captures most of the variation in the dataset. First, the mean signal $\mathbf{p}_{mean}$ is subtracted from each training propagator $\mathbf{p}_i$, and the covariance matrix $\mathbf{Z}\mathbf{Z}^H$ composed of these mean-subtracted descriptors is orthogonalized,

$$\mathbf{z}_i = \mathbf{p}_i - \frac{1}{L} \cdot \sum_{j=1}^L \mathbf{p}_j = \mathbf{p}_i - \mathbf{p}_{mean} \quad [15]$$

$$\mathbf{Z}\mathbf{Z}^H = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^H, \text{ where } \mathbf{z}_i \text{ is the } i^{th} \text{ column of } \mathbf{Z} \quad [16]$$

This yields an orthonormal matrix $\mathbf{Q}$ that contains the diffusion eigenvectors as its columns. Selecting the submatrix $\mathbf{Q}_T$ that contains the eigenvectors corresponding to the T largest eigenvalues in $\mathbf{\Lambda}$, PCA coefficients of a target propagator $\mathbf{p}$ can be obtained using,

$$\mathbf{pca} = \mathbf{Q}_T^H (\mathbf{p} - \mathbf{p}_{mean}) \quad [17]$$

The location of $\mathbf{pca}$ in the probability-space $\mathbf{p}_T$ can be recovered via,

$$\mathbf{p}_T = \mathbf{Q}_T \mathbf{Q}_T^H (\mathbf{p} - \mathbf{p}_{mean}) + \mathbf{p}_{mean} \quad [18]$$

Next, a low-dimensional fit to the undersampled q-space can be computed in the least-squares sense, $\min \|\mathbf{F}_\Omega \mathbf{p}_T - \mathbf{q}\|_2^2$, which can equivalently be expressed in PCA space,

$$\min \|\mathbf{F}_\Omega \mathbf{Q}_T \mathbf{pca} - (\mathbf{q} - \mathbf{F}_\Omega \mathbf{p}_{mean})\|_2^2 \quad [19]$$

A closed-form solution to Eq.19 can be obtained by using the pseudo inverse of $\mathbf{F}_\Omega \mathbf{Q}_T$,

$$\mathbf{pca} = \text{pinv}(\mathbf{F}_\Omega \mathbf{Q}_T) (\mathbf{q} - \mathbf{F}_\Omega \mathbf{p}_{mean}) \quad [20]$$

The result in the diffusion probability space is finally found by the mapping

$\tilde{\mathbf{p}} = \mathbf{Q}_T \mathbf{pca} + \mathbf{p}_{mean}$. The reconstruction matrix $\text{pinv}(\mathbf{F}_\Omega \mathbf{Q}_T)$ needs to be computed only once, and the computational cost is one matrix-vector multiplication per voxel.

DSI Reconstruction Experiments

Diffusion imaging data were acquired at a 3T system equipped with the Connectome gradients ($G_{\max} = 300$ mT/m and Slew = 200 T/m/s) from two healthy subjects using a 64-channel receive array (28). Echo-planar imaging parameters were: 2.3 mm isotropic voxel size, field of view = 220×220×130, matrix size = 96×96×57, and $b_{\max} = 8000$ s/mm² (acquired using $G_{\max} = 200$ mT/m). 515 directions full sphere q-space sampling with interspersed $b = 0$ images every 20 pulse repetition times (for motion correction using FLIRT (29)) was applied using in-plane acceleration = 2, pulse repetition time/echo time = 5.4 s/60 ms, for a total imaging time of 50 minutes. Eddy current related distortions were corrected using the reversed polarity method (30). Variable-density undersampling using a power-law density function (3) with $R = 3$ acceleration was applied in q-space. Dictionary-FOCUSS and PCA methods used training diffusion propagators obtained from a single slice of the training subject that is different from the test subject on which reconstructions were applied.

- i. **Wavelet & TV regularization:** The objective function in Eq.13 is minimized iteratively with the sparse MRI toolbox (3) using 100 conjugate gradient iterations,

which were seen to be sufficient for convergence. In addition to the total variation constraint, two different wavelet transforms (Ψ) were considered: Haar wavelets (using Matlab's native wavelet decomposition) and CDF 9–7 wavelets (using Matlab code available at <http://www.getreuer.info/home/waveletcdf97>) which were reported to yield better sparse approximations for simulated diffusion propagators than the Haar transform (31). In the experiments herein, CDF 9–7 implementation was seen to be computationally more efficient than the Haar transform (23 sec/voxel for Haar & TV and 0.8 sec/voxel for CDF 9–7 & TV, Fig.5). For both types of wavelet bases, a single level of decomposition yielded the lowest reconstruction error (results not shown). The regularization parameters α and β were determined by parameter sweeping. In the case of CDF 9–7 & TV method, whole slice was reconstructed with different parameters, and the setting that minimized the average reconstructed error ($\alpha = 10^{-3}$ and $\beta = 3 \cdot 10^{-4}$) was chosen to be the optimal one. For the Haar & TV method, only 100 voxels randomly selected within the slice were reconstructed for each parameter setting due to the high computational cost. The optimal selection was found to be $\alpha = 3 \cdot 10^{-4}$ and $\beta = 10^{-4}$.

- ii. **Dictionary-FOCUSS:** Compared to the implementation in (26) that uses 10 outer and 50 inner loops for the FOCUSS algorithm, 5 outer and 30 inner iterations were seen to be sufficient for convergence. This way, reconstruction time was reduced from 12 sec/voxel to 2.2 sec/voxel without affecting the reconstruction quality. As in (26), an overcomplete dictionary with 3191 columns was trained with K-SVD and used in reconstruction.
- iii. **PCA reconstruction:** The number of eigenvalues to keep (T) was determined using the training dataset, which was assumed to be fully-sampled. The value of T that minimized the reconstruction error in the training dataset for the same undersampling mask was chosen to be the optimal parameter ($T=50$). PCA reconstruction time was about 30 sec *per slice*.

Results

Regularized Quantitative Susceptibility Imaging

Reconstruction of the numerical phantom was completed in 3.3 sec with the proposed closed-form solution, while the conjugate gradient algorithm with 100 iterations took 65 min of computation (Fig.1). The normalized RMSE relative to the true susceptibility map was 17.4 % with the proposed method and 18.0 % with the conjugate gradient algorithm.

Processing the in vivo dataset took 1.3 sec with the closed-form method and 29 min with the iterative solver (Fig.2). Bottom row of Fig.2 depicts the 250-fold magnified difference between the closed-form and conjugate gradient reconstructions.

Lipid suppression for MR Spectroscopic Imaging

Lipid maps obtained by summing the absolute value of the reconstructed spectra over the lipid resonance frequency range are shown in top row of Fig.3, where images without suppression, dual-density reconstruction, lipid-basis penalty with ℓ_1 regularization, and the

proposed ℓ_2 -based method are compared. The reconstruction times were 0.1 sec for dual-density, 7 min for lipid-basis with ℓ_1 -penalty, and 0.2 sec for the proposed method. Regarding the lipid signal profile inside the brain (bottom half of Fig.3), dual-density reconstruction obtained 6.59 dB (2.1 times) average reduction in artifacts compared to the non-lipid-suppressed image. ℓ_1 -based lipid-basis yielded 19.56 dB (9.5 times) reduction of lipid signal compared to using no suppression and 12.98 dB (4.5 times) reduction relative to dual-density technique inside the brain. Similarly, ℓ_2 -based lipid-basis obtained 19.53 dB (9.5 times) and 12.95 dB (4.4 times) reduction in lipid power relative to using no suppression and dual-density sampling.

Example spectra marked on the anatomical image are depicted in Fig.4, where the dual-density and the proposed ℓ_2 -based method are compared in the top row, while the bottom row overlays the ℓ_2 - and ℓ_1 -regularized reconstructions.

Undersampled Diffusion Spectrum Imaging

Normalized RMSE values in the reconstructed probability density functions (propagators) at each voxel are presented in Fig.5 for the algorithms under consideration. Using CDF 9–7 wavelets with total variation yielded 15.9 % error on average, and the computation time was 35 min for this slice. Haar wavelets & TV had 15.6 % average RMSE and required 950 min of processing. Dictionary-FOCUSS method obtained 7.6 % average RMSE in 90 min, and the proposed PCA-based reconstruction yielded 7.8 % error in 0.4 min of computation time.

Orientation distribution function (ODF) glyphs inside the region of interest marked on the average fractional anisotropy image are depicted in Fig.6. To facilitate comparison between the ODFs produced by different methods, three voxels are further magnified.

Discussion

Quantitative Susceptibility Mapping

A closed-form expression for ℓ_2 -regularized susceptibility kernel inversion was derived and shown to be three orders of magnitude faster than the iterative conjugate gradient solution. This allowed computation of 3-d susceptibility maps to be performed in a few seconds. The proposed method finds the exact minimizer of the optimization problem, while the iterative algorithms aim to converge to this solution (Figs. 1 and 2), which explains the reduced RMSE in Fig.1. This closed-form solution may facilitate clinical application of regularized QSM. However, it is noted that obtaining the tissue field map from the unwrapped phase images require a pre-processing step that removes the contribution of phase background effects. The background component arises from air-tissue and bone-tissue interfaces, and can be eliminated to a great extent using the dipole fitting algorithm (32). This method, however, works iteratively and processing a 3-d dataset takes at least half an hour. Recently proposed fast background removal methods may greatly mitigate the computational burden of this pre-processing step (33,34).

Recent articles investigate employing ℓ_1 -regularized image gradients (i.e. total variation) for susceptibility inversion (8,35). For estimation of average χ values inside iron rich gray matter structures, both choices of regularization style (ℓ_1 and ℓ_2) were reported to generate

similar results (8). However, the ℓ_2 -regularization over the gradients assumes smoothly varying signal characteristics, therefore using ℓ_1 penalty for a piece-wise constant solution might be more suitable for susceptibility quantification in narrow structures such as vessels. On the other hand, ℓ_1 -based methods are constrained to work iteratively, and hence can be time consuming, thereby limiting their practicality.

Lipid Suppression for MRSI

The proposed ℓ_2 -based lipid suppression algorithm attains three orders of magnitude speed up relative to its previously proposed ℓ_1 -regularized counterpart (10) with slightly improved artifact reduction performance (Fig.4). Since reconstruction of a single slice MRSI data is completed under a second, clinical application of the method to 3-d spectroscopic imaging would be feasible.

Another artifact reduction method, dual-density sampling (12), was seamlessly merged in the proposed reconstruction pipeline. This method provides a partially lipid-suppressed starting point for the lipid-basis method, and enhances the end result. Since the dual-density acquisition may be difficult to realize on the scanner, it can be by-passed or a different initial starting point might be considered. For instance, the data extrapolation method (36) can be synergistically combined with the proposed lipid-basis reconstruction. Another alternative could be variable density spiral acquisition, which was shown to have more benign point spread function (37).

The fast lipid suppression method demonstrated effective artifact reduction, at a relatively short echo time of 50 ms, near the skull where the contamination is the largest (Fig. 4). A drawback of the lipid-basis and data extrapolation post-processing algorithms is that they require a binary mask that marks the location of the lipid layer. Herein, thresholding-guided manual segmentation was employed. More sophisticated methods, such as acquiring an accompanying structural image or a non-water-suppressed reference data, could allow automatic generation of a lipid mask.

The regularization parameter α in the ℓ_1 -based lipid-basis reconstruction was selected using the L-curve method. This technique involves reconstruction of the dataset several times and finding the optimal balance between the data consistency and regularization terms. For the proposed method, the regularization parameter β was chosen so that ℓ_1 - and ℓ_2 -based algorithms yielded the same data consistency level. While this was done to allow a fair comparison of the methods, the parameter β will still need to be determined for different in vivo settings. As the ℓ_2 -based algorithm takes less than a second of computation time, it will be feasible to trace the L-curve and find a suitable parameter. If an estimate of the noise variance is available, an alternative would be to set β so that the data consistency term is at noise level (38). Based on the lipid signal profiles in Fig.3, ℓ_1 - and ℓ_2 -based suppression yield comparable artifact reduction inside the brain (19.56 dB for ℓ_1 and 19.53 dB reduction for ℓ_2 method), but the proposed ℓ_2 -based technique remains computationally much faster.

Undersampled Diffusion Spectrum Imaging

The proposed PCA-based method demonstrated two orders of magnitude speed-up relative to Dictionary-FOCUSS algorithm, while obtaining similar image quality. Both of these

methods involve extracting priors from a training dataset, and thus achieve much reduced reconstruction error compared to using fixed transforms such as wavelets and total variation. This may indicate that the key to obtaining high fidelity diffusion propagators from undersampled q-space is encouraging the reconstructed signals to lie in the column space of a trained dictionary, rather than the selection of regularization style (ℓ_1 or ℓ_2) with respect to this dictionary.

As the proposed method is simple and fast in terms of training and reconstruction, it might allow online reconstruction of 3-dimensional DSI datasets. Since each voxel is processed independently, another potential source of performance gain is parallel processing.

For results reported herein, training data were obtained from a subject different from the test subject. Even though this might indicate that the method generalizes across subjects, further validation across healthy versus patient and young versus elderly populations is necessary. The number of eigenvectors used in the PCA representations was also determined based on the training subject and was found to be $T = 50$. If this was determined based on the test subject, $T = 49$ would be obtained. Both settings lead to the same RMSE, indicating that parameter selection on the training dataset might be feasible.

Regarding the ODF visualization in Fig.6, it can be seen that glyphs obtained with Dictionary-FOCUSS and PCA have higher fidelity than the Wavelet & TV method.

In conclusion, ℓ_2 -based reconstruction is shown to work just as well as ℓ_1 -regularization, given that the imaging application is suitable. Moreover, the fact that ℓ_2 -regularization admits a closed-form solution can be exploited to achieve substantial computational savings. In the contexts of QSM, MRSI, and DSI, two to three orders of magnitude speed up was demonstrated relative to the state of the art algorithms. In the spirit of reproducible research, Matlab code and example datasets for the proposed methods will be offered online.

Acknowledgments

Grant Support: NIH R01 EB007942; NIBIB K99EB012107; NIBIB R01EB006847; K99/R00 EB008129; NCRR P41RR14075 and the NIH Blueprint for Neuroscience Research U01MH093765 The Human Connectome project; Siemens Healthcare; Siemens-MIT Alliance; MITCIMIT Medical Engineering Fellowship

References

1. Lin FH, Kwong KK, Belliveau JW, Wald LL. Parallel imaging reconstruction using automatic regularization. *Magnet Reson Med*. 2004; 51(3):559–567.
2. Lustig M, Pauly JM. SPIRiT: Iterative Self-consistent Parallel Imaging Reconstruction From Arbitrary k-Space. *Magnet Reson Med*. 2010; 64(2):457–471.
3. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnet Reson Med*. 2007; 58(6):1182–1195.
4. Haldar JP, Wedeen VJ, Nezamzadeh M, Dai GP, Weiner MW, Schuff N, Liang ZP. Improved diffusion imaging through SNR-enhancing joint reconstruction. *Magnet Reson Med*. 2013; 69(1): 277–289.
5. Akcakaya M, Basha TA, Goddu B, Goepfert LA, Kissinger KV, Tarokh V, Manning WJ, Nezafat R. Low-dimensional-structure self-learning and thresholding: regularization beyond compressed sensing for MRI reconstruction. *Magn Reson Med*. 2011; 66(3):756–767. [PubMed: 21465542]

6. Otazo R, Li F, Chandarana H, Block T, Axel L, Sodickson DK. Combination of compressed sensing and parallel imaging for highly-accelerated dynamic MRI. 2012 May 2–5.:980–983. 2012.
7. de Rochefort L, Liu T, Kressler B, Liu J, Spincemaille P, Lebon V, Wu JL, Wang Y. Quantitative Susceptibility Map Reconstruction from MR Phase Data Using Bayesian Regularization: Validation and Application to Brain Imaging. *Magnet Reson Med*. 2010; 63(1):194–206.
8. Bilgic B, Pfefferbaum A, Rohlfing T, Sullivan EV, Adalsteinsson E. MRI estimates of brain iron concentration in normal aging using quantitative susceptibility mapping. *NeuroImage*. 2012; 59(3): 2625–2635. [PubMed: 21925274]
9. Bilgic B, Gagoski B, Kok T, Adalsteinsson E. Lipid suppression in CSI with spatial priors and highly undersampled peripheral k-space. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*. 2012
10. Lee, J.; Adalsteinsson, E. Iterative CSI Reconstruction with High-Resolution Spatial Priors for Improved Lipid Suppression; International Society for Magnetic Resonance in Medicine 19th Scientific Meeting; 2010. p. 965
11. Metzger G, Sarkar S, Zhang X, Heberlein K, Patel M, Hu X. A hybrid technique for spectroscopic imaging with reduced truncation artifact. *Magn Reson Imaging*. 1999; 17(3):435–443. [PubMed: 10195587]
12. Sarkar S, Heberlein K, Hu X. Truncation artifact reduction in spectroscopic imaging using a dual-density spiral k-space trajectory. *Magn Reson Imaging*. 2002; 20(10):743–757. [PubMed: 12591570]
13. Langkammer C, Schweser F, Krebs N, Deistung A, Goessler W, Scheurer E, Sommer K, Reishofer G, Yen K, Fazekas F, Ropele S, Reichenbach JR. Quantitative susceptibility mapping (QSM) as a means to measure brain iron? A post mortem validation study. *NeuroImage*. 2012; 62(3):1593–1599. [PubMed: 22634862]
14. Fan AP, Benner T, Bolar DS, Rosen BR, Adalsteinsson E. Phase-based regional oxygen metabolism (PROM) using MRI. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*. 2012; 67(3): 669–678.
15. 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. *Magnet Reson Med*. 2009; 61(1):196–204.
16. Wharton S, Bowtell R. Whole-brain susceptibility mapping at high field: A comparison of multiple- and single-orientation methods. *NeuroImage*. 2010; 53(2):515–525. [PubMed: 20615474]
17. Duyn JH, van Gelderen P, Li TQ, de Zwart JA, Koretsky AP, Fukunaga M. High-field MRI of brain cortical substructure based on signal phase. *Proceedings of the National Academy of Sciences of the United States of America*. 2007; 104(28):11796–11801. [PubMed: 17586684]
18. Liu T, Khalidov I, de Rochefort L, Spincemaille P, Liu J, Tsiouris AJ, Wang Y. A novel background field removal method for MRI using projection onto dipole fields (PDF). *Nmr Biomed*. 2011; 24(9):1129–1136. [PubMed: 21387445]
19. Hansen PC. The L-Curve and its Use in the Numerical Treatment of Inverse Problems. *Computational inverse problems in electrocardiology*. 2000:119–142.
20. Shungu DC, Glickson JD. Sensitivity and Localization Enhancement in Multinuclear in-Vivo Nmr-Spectroscopy by Outer Volume Presaturation. *Magnet Reson Med*. 1993; 30(6):661–671.
21. Ebel A, Govindaraju V, Maudsley AA. Comparison of inversion recovery preparation schemes for lipid suppression in H-1 MRSI of human brain. *Magnet Reson Med*. 2003; 49(5):903–908.
22. Haupt CI, Schuff N, Weiner MW, Maudsley AA. Removal of lipid artifacts in H-1 spectroscopic imaging by data extrapolation. *Magnet Reson Med*. 1996; 35(5):678–687.
23. Haase A, Frahm J, Hanicke W, Matthaei D. 1H NMR chemical shift selective (CHESS) imaging. *Physics in medicine and biology*. 1985; 30(4):341–344. [PubMed: 4001160]
24. Bottomley PA. Spatial localization in NMR spectroscopy in vivo. *Annals of the New York Academy of Sciences*. 1987; 508:333–348. [PubMed: 3326459]

25. Menzel MI, Tan ET, Khare K, Sperl JI, King KF, Tao X, Hardy CJ, Marinelli L. Accelerated diffusion spectrum imaging in the human brain using compressed sensing. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*. 2011; 66(5):1226–1233.
26. Bilgic B, Setsompop K, Cohen-Adad J, Yendiki A, Wald LL, Adalsteinsson E. Accelerated diffusion spectrum imaging with compressed sensing using adaptive dictionaries. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*. 2012; 68(6):1747–1754.
27. Gorodnitsky IF, Rao BD. Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm. *Ieee T Signal Proces*. 1997; 45(3):600–616.
28. Keil B, Blau JN, Biber S, Hoeft P, Tountcheva V, Setsompop K, Triantafyllou C, Wald LL. A 64-channel 3T array coil for accelerated brain MRI. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*. 2012.
29. Jenkinson M, Bannister P, Brady M, Smith S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*. 2002; 17(2):825–841. [PubMed: 12377157]
30. Bodammer N, Kaufmann J, Kanowski M, Tempelmann C. Eddy current correction in diffusion-weighted imaging using pairs of images acquired with opposite diffusion gradient polarity. *Magnet Reson Med*. 2004; 51(1):188–193.
31. Merlet, SL.; Paquette, M.; Deriche, R.; Descoteaux, M. Ensemble Average Propagator Reconstruction via Compressed Sensing: Discrete or Continuous Bases?. 20th annual meeting of the International Society for Magnetic Resonance in Medicine; Melbourne, Australia. 2012. p. 2277
32. Liu T, Khalidov I, de Rochefort L, Spincemaille P, Liu J, Tsiouris AJ, Wang Y. A novel background field removal method for MRI using projection onto dipole fields (PDF). *NMR Biomed*. 2011; 24(9):1129–1136. [PubMed: 21387445]
33. Abuhashem, O.; Bilgic, B.; Adalsteinsson, E. GPU Accelerated Quantitative Susceptibility Mapping. 20th annual meeting of the International Society for Magnetic Resonance in Medicine; Melbourne, Australia. 2012. p. 3442
34. Schweser, F.; Deistung, A.; Sommer, K.; Reichenbach, JR. Superfast Dipole Inversion for real-time Quantitative Susceptibility Mapping. 20th annual meeting of the International Society for Magnetic Resonance in Medicine; Melbourne, Australia. 2012.
35. Liu J, Liu T, de Rochefort L, Ledoux J, Khalidov I, Chen W, Tsiouris AJ, Wisnieff C, Spincemaille P, Prince MR, Wang Y. Morphology enabled dipole inversion for quantitative susceptibility mapping using structural consistency between the magnitude image and the susceptibility map. *Neuroimage*. 2012; 59(3):2560–2568. [PubMed: 21925276]
36. Haupt CI, Schuff N, Weiner MW, Maudsley AA. Removal of lipid artifacts in 1H spectroscopic imaging by data extrapolation. *Magn Reson Med*. 1996; 35(5):678–687. [PubMed: 8722819]
37. Adalsteinsson E, Star-Lack J, Meyer CH, Spielman DM. Reduced spatial side lobes in chemical-shift imaging. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*. 1999; 42(2):314–323.
38. Hansen PC. Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion: Society for Industrial Mathematics. 1987

QSM: Numerical Phantom

Noisy phase ϕ

error due to noise: 5.9% RMSE

Closed-form QSM in 3.3 seconds

Closed-form QSM error relative to true χ

Iterative QSM: 65 minutes (not shown)

Fig.1.

QSM reconstruction of a numerical phantom. **Top row:** Phase map is simulated by the forward dipole model from the ground truth, piece-wise constant susceptibility phantom and Gaussian noise is added to this simulated phase. **Middle row:** Susceptibility map computed from the noisy phase using the proposed closed-form method with a processing time of 3.3 seconds. **Bottom row:** The difference between the ground truth susceptibility and the closed-form solution.

QSM: In Vivo

Tissue phase ϕ

Closed-form QSM in 1.3 seconds

Closed-form and Iterative QSM difference

Iterative QSM: 29 minutes (not shown)

Fig.2.

In vivo QSM at 1.5T. **Top row:** Tissue phase obtained after removal of background contributions to the field map with the dipole fitting algorithm. **Middle row:** Closed-form QSM solution obtained from the tissue phase in 1.3 seconds. **Bottom row:** 250-fold magnified difference between the proposed closed-form susceptibility and the iterative solution obtained with the conjugate gradient algorithm.

Lipid Maps at TE = 50ms

Fig.3.

Top row: Lipid maps in dB-scale obtained by summation over the lipid resonance frequencies without lipid suppression, with dual-density reconstruction, L1-based lipid-basis method and the proposed L2-based lipid-basis regularization. **Bottom row:** Overlay of signal profiles along the black horizontal line for the four reconstruction methods.

Fig.4.

Example cortical spectra inside the region of interest marked on the structural image. Top panel overlays the spectra from dual-density and the proposed L2-based lipid suppression method, while the bottom panel compares the results obtained with L1- and L2-based lipid-basis penalty reconstructions.

3-fold undersampled DSI: error maps in PDF space

Fig.5.

Pdf reconstruction error maps obtained with the four methods for 3-fold accelerated DSI. CDF 9–7 wavelets & TV regularization obtained 15.9% average RMSE in 35 minutes of computation time, Haar wavelets & TV reconstruction had 15.6% error in 950 minutes, and Dictionary-FOCUSS reconstruction yielded 7.6% error in 90 minutes for this slice. The proposed PCA-based reconstruction returned 7.8% average RMSE in 0.4 minutes of processing time.

Fig.6. Odf visualizations inside the region of interest marked on the Fractional Anisotropy (FA) map, comparing glyphs obtained from fully-sampled data, Haar & TV regularization, Dictionary-FOCUSS reconstruction and the proposed PCA method. Three voxels are further magnified to facilitate the comparison.