Published in final edited form as:

IEEE Trans Comput Imaging. 2017 March ; 3(1): 11–21. doi:10.1109/TCI.2016.2626999.

Efficient, Convergent SENSE MRI Reconstruction for Nonperiodic Boundary Conditions via Tridiagonal Solvers

Mai Le [Student Member, IEEE] and Jeffrey A. Fessler [Fellow, IEEE]

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, 48109 USA

Abstract

Undersampling is an effective method for reducing scan acquisition time for MRI. Strategies for accelerated MRI such as parallel MRI and Compressed Sensing MRI present challenging image reconstruction problems with non-differentiable cost functions and computationally demanding operations. Variable splitting (VS) can simplify implementation of difficult image reconstruction problems, such as the combination of parallel MRI and Compressed Sensing, CS-SENSE-MRI. Combined with augmented Lagrangian (AL) and alternating minimization strategies, variable splitting can yield iterative minimization algorithms with simpler auxiliary variable updates. However, arbitrary variable splitting schemes are not guaranteed to converge. Many variable splitting strategies are combined with periodic boundary conditions. The resultant circulant Hessians enable $\mathcal{O}(n \log n)$ computation but may compromise image accuracy at the spatial boundaries. We propose two methods for CS-SENSE-MRI that use regularization with non-periodic boundary conditions to prevent wrap-around artifacts. Each algorithm computes one of the resulting variable updates efficiently in $\mathcal{O}(n)$ time using a parallelizable tridiagonal solver. AL-tridiag is a VS method designed to enable efficient computation for non-periodic boundary conditions. Another proposed algorithm, ADMM-tridiag, uses a similar VS scheme but also ensures convergence to a minimizer of the proposed cost function using the Alternating Direction Method of Multipliers (ADMM). AL-tridiag and ADMM-tridiag show speeds competitive with previous VS CS-SENSE-MRI reconstruction algorithm AL-P2. We also apply the tridiagonal VS approach to a simple image inpainting problem.

Index Terms

image reconstruction; augmented Lagrangian (AL); alternating direction method of multipliers (ADMM); non-periodic boundaries; variable splitting; tridiagonal solvers; parallel magnetic resonance imaging (MRI)

I. Introduction

Accelerated scan time benefits many applications of Magnetic Resonance Imaging (MRI), reducing cost and motion blurring. A popular strategy for reducing scan time is to acquire fewer k-space samples. To compensate for the reduced sampling, parallel MRI and Compressed Sensing (CS) are often used. SENSitivity Encoded (SENSE) MRI is a popular

framework for parallel MRI [1]. By simultaneously acquiring data from multiple receive coils with spatially varying sensitivities, more data can be collected without additional scan time. SENSE can be combined with CS-inspired techniques to further reduce acquisition time [2]. Irregular undersampling patterns enable higher acceleration but require iterative model-based image reconstruction.

CS MRI reconstructs images assuming image sparsity in some transform domain. To balance adherence of the estimated image to the noisy, undersampled data with the prior assumption of sparsity, we seek to minimize a cost function that describes both. These methods employ regularization with ℓ_1 norms that promote sparsity but also present computational challenges.

Variable splitting (VS) [3] [4] [5] is a versatile optimization approach for these cost functions. VS decouples a costly nonlinear optimization into simpler problems via the augmented Lagrangian (AL) framework. VS converts the original cost function into a constrained cost function involving additional auxiliary variables updated with alternating minimization. The AL-P2 algorithm proposed in [4] demonstrated that VS combined with AL can yield significant speed gains over conjugate gradients (CG) and monotone fast iterative shrinkage-thresholding algorithm (MFISTA) [6] for CS-SENSE-MRI. However, AL-P2 lacks a convergence guarantee and uses periodic boundary conditions when applying sparsifying transforms. This assumption leads to a circulant Hessian and an $\mathcal{O}(n \log n)$ FFT-based solution for one of the inner problems, where n is the number of pixels.

The use of periodic boundary conditions is not reasonable for reconstructing 2D axial slices of the brain, when it is surrounded by air at all boundaries. However, when reconstructing an entire 3D volume, it could be undesirable to impose periodic boundary conditions across the top and bottom slices. Non-periodic boundary conditions are also useful in dynamic imaging with acyclic temporal behavior. For example, Dynamic Contrast Enhanced (DCE) images differ significantly in the final frame compared to the initial frame.

Here we present two related variable splitting methods for CS-SENSE-MRI: ADMM-tridiag and AL-tridiag. ADMM-tridiag leverages [7] to ensure convergence. Both proposed algorithms use a regularizer with non-periodic boundary conditions to more accurately reflect the reality of the unknown image. Both algorithms incorporate parallelizable tridiagonal solvers that efficiently handle the non-periodic boundary conditions with $\mathcal{O}(n)$ operations. We present numerical results with real *in vivo* data to demonstrate the efficacy of the proposed methods.

We also apply the principles to describe novel VS algorithms for the simpler “special case” of image inpainting using regularization based on a combination of wavelets and anisotropic total variation (TV) with non-periodic boundary conditions.

II. Problem Formulation

The analysis formulation often used for SENSE MRI reconstruction estimates the unknown image, $\hat{x}$, by seeking the minimizer of a cost function consisting of a datafit term plus a regularizer. Regularization is particularly important for undersampled problems.

Many regularizers have been used for SENSE MRI reconstruction. For CS-SENSE-MRI, sparsifying transforms such as wavelets and first-order finite differences are often used with an ℓ_1 norm to promote sparsity [2]. This paper focuses on Cartesian SENSE reconstruction with regularizers having non-periodic boundary conditions.

Let N_c denote the number of sensitivity coils, N_s the number of samples received from each coil, and $N_r = N_x N_y$ the number of pixels in the latent image x . We formulate regularized SENSE reconstruction as the following optimization problem:

$$\hat{x} = \arg\min_x \frac{1}{2} \|y - \mathbf{F}\mathbf{S}x\|_2^2 + \lambda \|\mathbf{C}_H x\|_1 + \lambda \|\mathbf{C}_V x\|_1 \quad (1)$$

where $y \in \mathbb{C}^{N_c N_s}$ is the undersampled k-space data from all coils, $\mathbf{F} \in \mathbb{C}^{N_c N_s \times N_c N_r}$ is a block diagonal matrix consisting of undersampled DFT matrices, $\mathbf{S} \in \mathbb{C}^{N_c N_r \times N_r}$ is a stack of diagonal matrices containing the sensitivity maps, $\mathbf{C}_H \in \mathbb{R}^{N_r \times N_r}$ and $\mathbf{C}_V \in \mathbb{R}^{N_r \times N_r}$ denote finite differences in the horizontal and vertical directions (equivalent to anisotropic TV), and $\hat{x}$ is the reconstructed image. The regularization parameter $\lambda > 0$ balances adherence to noisy data and prior assumptions that inform choice of regularizers. Section III-D extends (1) to include discrete wavelet transforms. Although we focus on 2D imaging for notational simplicity, the methods generalize readily to 3D problems by adding another term to (1).

We choose $\mathbf{C}_H$ and $\mathbf{C}_V$ to have non-periodic boundary conditions. This is in contrast to many proposed algorithms that use finite differences with periodic boundary conditions, such as AL-P2 [4], RecPF [8], and recMRI [9]. Periodic boundary conditions have been used for computational convenience, despite being physically unnatural. Non-periodic boundary conditions are preferable in most applications [10], [11], [12]. Differences across the boundaries of medical images do not provide useful information for reconstruction. Furthermore, penalizing differences across spatial boundaries may degrade image quality when the region of interest extends to the boundary. For example, a coronal abdominal image depicts different anatomy at the top and bottom boundaries. Thus this paper focuses on developing methods that are fast yet suitable for non-periodic boundary conditions.

III. Variable Splitting Methods

The non-differentiable ℓ_1 norms make (1) a challenging optimization problem. Variable splitting methods like [13], [14], [4] are useful for such problems. Here, we reformulate (1) in an equivalent constrained form using the following novel variable splitting scheme:

$$\begin{aligned} \hat{u} &= \arg\min_{\underline{u}} f(\underline{u}) \\ f(\underline{u}) &= \frac{1}{2} \|y - \mathbf{F}u_2\|_2^2 + \lambda \|u_0\|_1 + \lambda \|u_1\|_1 \end{aligned} \quad (2)$$

$$\begin{aligned} \text{s. t. } u_0 &= \mathbf{C}_H x, & u_1 &= \mathbf{C}_V u_3, \\ u_2 &= \frac{1}{2} \mathbf{S} u_3 + \frac{1}{2} \mathbf{S} x, & u_3 &= x. \end{aligned} \quad (3)$$

For convenience, we group x and the auxiliary variables $u_0, \dots, u_3$ into one column vector: $\underline{u} \triangleq (u_0, u_1, u_2, u_3, x)$.

This variable splitting scheme intentionally separates the horizontal and vertical finite differences operators, applying them to different auxiliary variables. This separates the tridiagonal structures in the Hessians resulting from AL, permitting decoupled, computationally efficient variable updates further detailed in Section III-C. If the finite difference matrices were combined in the same auxiliary variable as in AL-P2 [4], the resulting Hessians would have a block-tridiagonal with tridiagonal blocks (BTTB) structure that cannot take advantage of an $\mathcal{O}(N_r)$ tridiagonal solver.

A. Direct AL Approach: AL-tridiag

Here we detail the algorithm resulting from directly applying AL with alternating minimization to (2). The resulting algorithm, AL-tridiag, does not satisfy the sufficient conditions for convergence in [7], so currently it lacks convergence guarantees. However, it has worked well in all of our experiments, so it is possible that future generalized convergence proofs could be applicable.

We rewrite constrained cost function (2) with a matrix constraint as follows:

$$\min_{\underline{u}} f(\underline{u}) \text{ s. t. } \mathbf{P}\underline{u} = \mathbf{0} \quad \mathbf{P} \triangleq \begin{bmatrix} -\mathbf{I}_{N_r} & 0 & 0 & 0 & \mathbf{C}_H \\ 0 & -\mathbf{I}_{N_r} & 0 & \mathbf{C}_V & 0 \\ 0 & 0 & -\mathbf{I}_{N_r N_c} & \frac{1}{2} \mathbf{S} & \frac{1}{2} \mathbf{S} \\ 0 & 0 & 0 & -\mathbf{I}_{N_r} & \mathbf{I}_{N_r} \end{bmatrix}. \quad (4)$$

The constraint $\mathbf{P}\underline{u} = \mathbf{0}$ enforces (3). The augmented Lagrangian corresponding to (4) is:

$$\mathcal{L}(\underline{u}, \underline{\eta}; \mathbf{M}) = f(\underline{u}) + \frac{1}{2} \|\mathbf{P}\underline{u} - \underline{\eta}\|_{\mathbf{M}}^2. \quad (5)$$

This formulation introduces dual variables, stacked in a vector: $\underline{\eta} \triangleq (\eta_0, \dots, \eta_3) \in \mathbb{C}^{(3+N_c)N_r}$. Matrix $\mathbf{M}$ is a positive definite diagonal matrix, consisting of user-selected AL penalty parameters. The diagonal block corresponding to the i th segment of $\underline{u}$ is denoted $\mathbf{M}_i$. The choice of $\mathbf{M}$ does not affect the final solution of (5), but it can affect the convergence rate of the resulting algorithm. For many multiplier methods, using positive penalty parameters guarantees convergence to the solution of the original problem that does not involve the penalty parameters [15]. Section III-G discusses heuristics for selecting $\mathbf{M}$.

Ideally, an AL method would update block variables $\underline{u}$ and $\underline{\eta}$ at iteration $n+1$ as follows:

$$\underline{u}^{(n+1)} = \underset{\underline{u}}{\operatorname{argmin}} \mathcal{L}(\underline{u}, \underline{\eta}^{(n)}) \quad (6)$$

$$\underline{\eta}^{(n+1)} = \underline{\eta}^{(n)} - \mathbf{P} \underline{u}^{(n+1)}. \quad (7)$$

Our proposed AL-tridiag algorithm uses alternating minimization across $u_0, \dots, u_3, x$ to descend the AL term in (6). Section III-C and the supplement describe the $\underline{u}$ variable updates in more detail. Due to the variable splitting design of (3), each variable update has a direct, closed-form solution with an efficient implementation, e.g., by FFTs or a parallelizable tridiagonal solver.

B. ADMM Equivalence for ADMM-tridiag

To design a minimization algorithm with convergence guarantees, we reformulate (2) as an instance of the generalized Alternating Direction Method of Multipliers (ADMM) [7] [16]. This formulation allows us to invoke the convergence proof in [7] for our second proposed algorithm, ADMM-tridiag.

We first express the matrix constraint $\mathbf{P}$ in (4) as a product of two matrices, $\mathbf{P} = \mathbf{B}\mathbf{A}$, and incorporate the left matrix $\mathbf{B}$ into the following convex cost function:

$$\min_{\underline{u}, \underline{v}} f(\underline{u}) + g(\underline{v}) \text{ s. t. } \mathbf{A}\underline{u} = \underline{v} \quad (8)$$

$$g(\underline{v}) = \begin{cases} 0, & \mathbf{B}\underline{v} = 0 \\ +\infty, & \mathbf{B}\underline{v} \neq 0. \end{cases} \quad (9)$$

For (8) to satisfy the sufficient conditions for convergence of ADMM in [7], $\mathbf{A}$ must have full rank. For our proposed algorithm, ADMM-tridiag, we design $\mathbf{B}$ and $\mathbf{A}$ as follows:

$$\mathbf{B} \triangleq \begin{bmatrix} \mathbf{I}_{N_r} & 0 & 0 & 0 & 0 \\ 0 & \mathbf{I}_{N_r} & 0 & 0 & 0 \\ 0 & 0 & \mathbf{I}_{N_r N_c} & 0 & 0 \\ 0 & 0 & 0 & \mathbf{I}_{N_r} & \mathbf{I}_{N_r} \end{bmatrix}$$

$$\mathbf{A} \triangleq \begin{bmatrix} -\mathbf{I}_{N_r} & 0 & 0 & 0 & \mathbf{C}_H \\ 0 & -\mathbf{I}_{N_r} & 0 & \mathbf{C}_V & 0 \\ 0 & 0 & -\mathbf{I}_{N_r N_c} & \frac{1}{2}\mathbf{S} & \frac{1}{2}\mathbf{S} \\ 0 & 0 & 0 & -\mathbf{I}_{N_r} & 0 \\ 0 & 0 & 0 & 0 & \mathbf{I}_{N_r} \end{bmatrix}. \quad (10)$$

For convenience, we describe $\underline{v}$ in terms of its block elements: $\underline{v} \triangleq (v_0, \dots, v_4)$. The full rank of $\mathbf{A}$, combined with the following alternating minimization framework, satisfies the convergence conditions for ADMM [7]. Thus, (8)–(10) describe an instance of ADMM and guarantees convergence to a minimizer of (8). To handle the constraints of (8), ADMM-tridiag uses the following augmented Lagrangian, similar to Section III-A:

$$\mathcal{L}(\underline{u}, \underline{v}, \eta; \mathbf{M}) = f(\underline{u}) + g(\underline{v}) + \frac{1}{2} \|\mathbf{A}\underline{u} - \underline{v} - \underline{\eta}\|_{\mathbf{M}}^2. \quad (11)$$

The dual variables $\underline{\eta} \triangleq (\eta_0, \dots, \eta_4) \in \mathbb{C}^{N_r(4+N_c)}$ have an additional block-element compared to the AL case. For convenience, we reuse the notation for the dual variables for both algorithms. The matrix $\mathbf{M}$ for ADMM-tridiag is also larger than in the AL case, consisting of blocks $\{\mathbf{M}_j\}_{j=0,\dots,4}$.

Following [15] [7], ADMM-tridiag alternates between updating $\underline{u}$, $\underline{v}$, and the dual variables $\underline{\eta}$. Examining the $\underline{v}$ update and recalling (9), we see that the role of $\underline{v}$ simplifies greatly by considering the feasible set of $\underline{v}$, $\Omega \triangleq \{\underline{v} \in \mathbb{C}^{(4+N_c)N_r} : \mathbf{B}\underline{v} = 0\}$, resulting in the following alternating updates:

$$\underline{u}^{(n+1)} \underset{\varepsilon_n}{\approx} \underset{\underline{u}}{\operatorname{argmin}} f(\underline{u}) + \frac{1}{2} \|\mathbf{A}\underline{u} - \underline{v}^{(n)} - \underline{\eta}^{(n)}\|_{\mathbf{M}}^2 \quad (12)$$

$$\underline{v}^{(n+1)} \underset{\zeta_n}{\approx} \underset{\underline{v} \in \Omega}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{A}\underline{u}^{(n+1)} - \underline{v} - \underline{\eta}^{(n)}\|_{\mathbf{M}}^2 \quad (13)$$

$$\underline{\eta}^{(n+1)} = \underline{\eta}^{(n)} - (\mathbf{A}\underline{u}^{(n+1)} - \underline{v}^{(n+1)}). \quad (14)$$

Here, we allow for some inexactness in the updates of $\underline{u}$ and $\underline{v}$ at each iteration, $\{\varepsilon_n\}$ and $\{\zeta_n\}$, respectively. ADMM convergence holds if the inexactness sequences are summable [3, Theorem 8] [7, Theorem 2.1].

The update for $\underline{v}$ is simple and exact to machine precision. Due to the simple structure of $\mathbf{B}$, we have $v_0, v_1, v_2 = 0$ and $v_4 = -v_3$. The entire update for $\underline{v}$ reduces to a quadratic minimization problem for the auxiliary variable v_3 of size N_r with simple, closed-form solution:

$$v_3^{(n+1)} = (\mathbf{M}_3 + \mathbf{M}_4)^{-1} \left[\mathbf{M}_3 \left(-u_3^{(n+1)} - n_3^{(n)} \right) + \mathbf{M}_4 \left(-x^{(n+1)} + \eta_4^{(n)} \right) \right]. \quad (15)$$

For the joint minimization of $\underline{u}$ in (12), ADMM-tridiag employs alternating minimization, as detailed in Section III-C. If one round of alternating minimization sufficiently approximates

the joint minimizer of (12), the iterates of ADMM-tridiag, $\{x^{(n)}\}$, converge to a minimizer of (1), per [7]. This alternating minimization approach is common in other ADMM methods, such as [10] [12].

Compared to AL-tridiag in Section III-A, the direct AL approach to (3), ADMM-tridiag involves one additional variable split. An alternative way to understand the effect of (10) is to describe this ADMM algorithm as the result of applying AL to the following variable splitting scheme for (1):

$$\min_{\underline{u}} \frac{1}{2} \|y - \mathbf{F}u_2\|^2 + \lambda \|u_0\|_1 + \lambda \|u_1\|_1 \quad (16)$$

$$\begin{aligned} \text{s. t. } u_0 &= \mathbf{C}_H x, & u_2 &= \frac{1}{2} \mathbf{S}u_3 + \frac{1}{2} \mathbf{S}x, \\ u_1 &= \mathbf{C}_V u_3, & u_3 &= -v_3, & v_3 &= -x \end{aligned} \quad (17)$$

This formulation indirectly enforces $u_3 = x$ through additional auxiliary variable v_3 . However, written in this form it is less clear that the full-rank condition of [7] is satisfied, whereas that is clear in (8)–(10).

C. Variable Updates with Parallelizable Tridiagonal Solvers

This section describes the block-variable updates for alternating minimization of (12). These variable updates are very similar for both proposed methods, AL-tridiag and ADMM-tridiag, differing only in tuning parameter indices and use of $\underline{v}$ and $\underline{u}$ on the right-hand side of these updates. For brevity, we give only the alternating minimization updates for ADMM-tridiag in (12). The variable updates for AL-tridiag in (6) are very similar as shown in Algorithm 2 of the Supplement¹.

First, we consider the special case where some blocks of $\mathbf{M}$ are constant diagonal matrices. We leave $\mathbf{M}_3$ and $\mathbf{M}_4$ as general positive definite diagonal matrices for reasons explained in Section III-G. Letting $\mathbf{M}_i \triangleq \mu_i \mathbf{I}$, $\mu_i > 0$ for $i = 0, 1, 2$ and leveraging the constraint that $v_4 = -v_3$, (12) expands to:

$$\begin{aligned} \underline{u}^{(n+1)} = \operatorname{argmin}_{x, u_0, \dots, u_3} & \frac{1}{2} \|y - \mathbf{F}u_2\|^2 + \lambda \|u_0\|_1 + \lambda \|u_1\|_1 \\ & + \frac{\mu_0}{2} \|-u_0 + \mathbf{C}_H x - \eta_0^{(n)}\|^2 \\ & + \frac{\mu_1}{2} \|-u_1 + \mathbf{C}_V u_3 - \eta_1^{(n)}\|^2 \\ & + \frac{\mu_2}{2} \|-u_2 + \frac{1}{2} \mathbf{S}u_3 + \frac{1}{2} \mathbf{S}x - \eta_2^{(n)}\|^2 \\ & + \frac{1}{2} \|-u_3 - v_3^{(n)} - \eta_3^{(n)}\|_{\mathbf{M}_3}^2 \\ & + \frac{1}{2} \|x + v_3^{(n)} - \eta_4^{(n)}\|_{\mathbf{M}_4}^2. \end{aligned} \quad (18)$$

¹Supplementary material available in the supplementary files/multimedia tab.

ADMM-tridiag uses alternating minimization to update the blocks of $\underline{u}$. The resulting variable updates are:

$$u_0^{(n+1)} = \text{soft} \left(\mathbf{C}_H x^{(n)} - \eta_0^{(n)}, \frac{\lambda}{\mu_0} \right) \quad (19)$$

$$u_1^{(n+1)} = \text{soft} \left(\mathbf{C}_V u_3^{(n)} - \eta_1^{(n)}, \frac{\lambda}{\mu_1} \right) \quad (20)$$

$$u_2^{(n+1)} = \mathbf{H}_2^{-1} \left(\mathbf{F}' y + \mu_2 \left(\frac{1}{2} \mathbf{S} u_3^{(n)} + \frac{1}{2} \mathbf{S} x^{(n)} - \eta_2^{(n)} \right) \right) \quad (21)$$

$$u_3^{(n+1)} = \mathbf{H}_3^{-1} \left(\mu_1 \mathbf{C}_V' \left(u_1^{(n+1)} + \eta_1^{(n)} \right) + \frac{\mu_2}{2} \mathbf{S}' \left(u_2^{(n+1)} - \frac{1}{2} \mathbf{S} x^{(n)} + \eta_2^{(n)} \right) + \mathbf{M}_3 \left(-v_3^{(n)} - \eta_3^{(n)} \right) \right) \quad (22)$$

$$x^{(n+1)} = \mathbf{H}_x^{-1} \left(\mu_0 \mathbf{C}_H' \left(u_0^{(n+1)} + \eta_0^{(n)} \right) + \frac{\mu_2}{2} \mathbf{S}' \left(u_2^{(n+1)} - \frac{1}{2} \mathbf{S} u_3^{(n+1)} + \eta_2^{(n)} \right) + \mathbf{M}_4 \left(-v_3^{(n)} + \eta_4^{(n)} \right) \right). \quad (23)$$

The soft-thresholding operator performs element-wise shrinkage for the ℓ_1 norm using a given threshold τ : $\text{soft}(x, \tau) \triangleq \text{sign}(x) \max(|x| - \tau, 0)$. Thus, (19) and (20) provide simple, direct solutions for updating u_0 and u_1 . The bulk of the computation is “inverting” the following Hessians:

$$\mathbf{H}_2 \triangleq \mathbf{F}' \mathbf{F} + \mu_2 \mathbf{I} = \mathbf{Q} (\mathbf{\Lambda}_F + \mu_2 \mathbf{I}) \mathbf{Q}' \quad (24)$$

$$\mathbf{H}_3 \triangleq \mu_1 \mathbf{C}_V' \mathbf{C}_V + \frac{\mu_2}{4} \mathbf{S}' \mathbf{S} + \mathbf{M}_3 \quad (25)$$

$$\mathbf{H}_x \triangleq \mu_0 \mathbf{C}'_H \mathbf{C}_H + \frac{\mu_2}{4} \mathbf{S}' \mathbf{S} + \mathbf{M}_4. \quad (26)$$

Due to Cartesian undersampling, $\mathbf{H}_2$ is diagonalizable via N_c FFTs, each of which operate efficiently in $\mathcal{O}(N_r \log N_r)$ time. The multi-coil FFT operator is denoted $\mathbf{Q}$. The x update in (23) uses $\mathbf{H}_x$ in (26), a block diagonal matrix with tridiagonal blocks (BDTB) that can be “inverted” in $\mathcal{O}(N_r)$ time via Gaussian elimination. Because it is block diagonal, we parallelize this variable update over the N_y independent blocks. We reformulate the u_3 update as another instance of the same BDTB inverse problem through permutation and solve (22) using a tridiagonal solver parallelized over N_x blocks. Each variable update is exact and easy to implement.

We designed the proposed variable splitting in (2) to enable these efficient variable updates. The separation of horizontal and vertical finite differences into u_0 and u_1 allows for BDTB structures in $\mathbf{H}_3$ and $\mathbf{H}_x$. If the finite difference matrices were combined in the same auxiliary variable as in AL-P2 [4], the resulting Hessian would have a block-tridiagonal with tridiagonal blocks (BTTB) structure and the associated variable update would require a more computationally costly solution [17] [18].

Algorithm 1 summarizes the overall procedure for ADMM-tridiag. All of the variable updates are done in place, so the memory requirements for storing x , u_0 , ..., u_3 , v_3 , and η are $8N_r(4 + N_c)$ bytes for $N_r(4 + N_c)$ complex single-precision values. For AL-tridiag, the $\underline{u}$ updates are very similar.

Algorithm 1

ADMM-Tridiag

-
- 1: Initialize x to square root of sum-of-squares (SoS) of zero-filled iFFT images.
 - 2: Initialize $u_0 = \mathbf{C}_H x$, $u_1 = \mathbf{C}_V u_3$, $u_3 = x$, $u_2 = \frac{1}{2} \mathbf{S} u_3 + \frac{1}{2} \mathbf{S} x$, $v_3 = -x$, $\eta = \underline{0}$
 - 3: **for** n total iterations **do**
 - 4: Compute $u_0^{(n+1)}$ via soft-thresholding (19)
 - 5: Compute $u_1^{(n+1)}$ via soft-thresholding (20)
 - 6: Compute $u_2^{(n+1)}$ via FFTs (21)
 - 7: Compute $u_3^{(n+1)}$ via tridiagonal solver (22)
 - 8: Compute $x^{(n+1)}$ via tridiagonal solver (23)
 - 9: Compute $v_3^{(n+1)}$ via (15)
 - 10: Compute $\eta^{(n+1)}$ via (14)
 - 11: **end for**
-

D. Regularization with Finite Differences and Wavelets

The variable splitting scheme in (17) readily generalizes to combinations of finite difference and orthonormal wavelet regularization. Let $\mathbf{W}$ be an orthonormal wavelet transform (e.g., Haar wavelets). Then the following combined TV/wavelet sparsity cost-function can be manipulated to resemble (1):

$$\frac{1}{2} \|y - \mathbf{F}\mathbf{S}x\|_2^2 + \lambda_1 \|\mathbf{C}_H x\|_1 + \lambda_1 \|\mathbf{C}_V x\|_1 + \lambda_2 \|\mathbf{W}x\|_1 \quad (27)$$

$$= \frac{1}{2} \|y - \mathbf{F}\mathbf{S}x\|_2^2 + \lambda_1 \|\tilde{\mathbf{C}}_H x\|_1 + \lambda_1 \|\tilde{\mathbf{C}}_V x\|_1 \quad (28)$$

$$\tilde{\mathbf{C}}_H \triangleq \begin{bmatrix} \mathbf{C}_H \\ \frac{\alpha_w \lambda_2}{\lambda_1} \mathbf{W} \end{bmatrix}; \quad \tilde{\mathbf{C}}_V \triangleq \begin{bmatrix} \mathbf{C}_V \\ \frac{(1-\alpha_w)\lambda_2}{\lambda_1} \mathbf{W} \end{bmatrix}. \quad (29)$$

An additional spatial regularization parameter, λ_2 , controls the weight of the wavelet regularization. Due to the orthonormality of $\mathbf{W}$, we can use this regularizer for both AL-tridiag and ADMM-tridiag, with only minor changes to variable updates (19) – (23). For ADMM-tridiag, the only Hessians affected by introducing wavelets are:

$$\tilde{\mathbf{H}}_3 \triangleq \mu_1 \mathbf{C}'_V \mathbf{C}_V + \frac{\mu_2}{4} \mathbf{S}' \mathbf{S} + \tilde{\mathbf{M}}_3 \quad (30)$$

$$\tilde{\mathbf{H}}_x \triangleq \mu_0 \mathbf{C}'_H \mathbf{C}_H + \frac{\mu_2}{4} \mathbf{S}' \mathbf{S} + \tilde{\mathbf{M}}_4, \quad (31)$$

with positive definite matrices $\tilde{\mathbf{M}}_3 = \mathbf{M}_3 + \frac{(1-\alpha_w)^2 \lambda_2^2}{\lambda_1^2} \mathbf{I}$ and $\tilde{\mathbf{M}}_4 = \mathbf{M}_4 + \frac{\alpha_w^2 \lambda_2^2}{\lambda_1^2} \mathbf{I}$. Hessians $\tilde{\mathbf{H}}_3$ and $\tilde{\mathbf{H}}_x$ are still BDTB, and they can be “inverted” efficiently with a parallelizable tridiagonal solver routine. This wavelet-inclusive variation is featured in experimental results in Section IV-C.

E. Special Case: Image Inpainting

To highlight the value of non-periodic boundary conditions, we examine a specific application of the proposed variable splitting scheme in (2), namely inpainting. Image inpainting fills in image data that is lost or corrupted. Many image inpainting, deblurring, and denoising methods use finite difference regularizers like anisotropic total variation (TV) [19] [11] [20].

Let $\mathbf{D}$ be a binary, diagonal matrix whose nonzero entries denote the set of indices in the inpainting domain. Setting $\mathbf{FS} = \mathbf{D}$ in the CS-SENSE-MRI cost function (1) leads to the following simpler image inpainting problem:

$$\hat{x} = \underset{x}{\operatorname{argmin}} \frac{1}{2} \|y - \mathbf{D}x\|_2^2 + \lambda \|\tilde{\mathbf{C}}_H x\|_1 + \lambda \|\tilde{\mathbf{C}}_V x\|_1, \quad (32)$$

where $\tilde{\mathbf{C}}_H$ and $\tilde{\mathbf{C}}_V$ are defined in (29).

For inpainting, we simplify the VS scheme developed for SENSE MRI in (2) to:

$$\begin{aligned} \hat{u} &= \underset{u}{\operatorname{argmin}} f(u) \\ f(u) &= \frac{1}{2} \|y - \mathbf{D} \left(\frac{1}{2} u_2 + \frac{1}{2} x \right)\|_2^2 + \lambda \|u_0\|_1 + \lambda \|u_1\|_1 \\ \text{s. t. } & u_0 = \tilde{\mathbf{C}}_H x, u_1 = \tilde{\mathbf{C}}_V u_2, u_2 = x. \end{aligned} \quad (33)$$

The resulting VS algorithm is a simplification of AL-tridiag, which we denote AL-tridiag-inpaint. Due to the entirely diagonal system matrix, the variable updates consist only of shrinkage and tridiagonal solver updates, eliminating the need for any FFT-based updates. Similarly, we can generalize ADMM-tridiag to the inpainting problem by applying an extra variable splitting, resulting in an additional quadratic minimization problem in each iteration. Section IV-C illustrates the effect of non-periodic boundary conditions for noisy inpainting.

F. Comparison with AL-P2

We compare ADMM-tridiag with AL-P2 [4], a fast VS scheme designed for CS-SENSE-MRI. The original AL-P2 version in [4] used periodic boundary conditions, whereas here we modify it for the non-periodic conditions of (1) and call the modified algorithm AL-P2-NC. The suffix “NC” refers to the non-circulant Hessian we describe in this section. To define AL-P2-NC, we stack the finite difference matrices into a tall matrix, $\mathbf{R} \triangleq [\mathbf{C}_H; \mathbf{C}_V]$. Applying the AL-P2 variable splitting scheme to (1) yields the following constrained cost function:

$$\min_{x,u,v,z} \|y - \mathbf{F}u\|_2^2 + \lambda \|v\|_1 \quad (34)$$

$$\text{s. t. } u = \mathbf{S}x, v = \mathbf{R}u, z = x. \quad (35)$$

Applying the augmented Lagrangian to the constrained cost function (34) and using alternating minimization results in variable updates like in [4]:

$$x^{(n+1)} = \mathbf{H}_x^{-1} \left(\mathbf{F}' y + \mu_u \left(\mathbf{S} x^{(n)} + \eta_u^{(n)} \right) \right) \quad (36)$$

$$u^{(n+1)} = \mathbf{H}_u^{-1} \left(\mu_u \mathbf{S}' \left(u^{(n)} - \eta_u^{(n)} \right) + \mu_z \left(z^{(n)} - \eta_z^{(n)} \right) \right) \quad (37)$$

$$v^{(n+1)} = \text{soft} \left(\mathbf{R} z^{(n)} + \eta_v^{(n)}, \frac{\lambda}{\mu_u} \right) \quad (38)$$

$$z^{(n+1)} = \mathbf{H}_z^{-1} \left(\mu_v \mathbf{R}' \left(v^{(n+1)} - \eta_v^{(n)} \right) + \mu_z \left(x^{(n+1)} + \eta_z^{(n)} \right) \right). \quad (39)$$

The scalar AL penalty parameters, $\mu_u, \mu_v, \mu_z > 0$, do not affect the final solution but do affect convergence rate. The Hessians for x and u are simple to invert and are the same as in ADMM-tridiag:

$$\mathbf{H}_x = \mu_u \mathbf{S}' \mathbf{S} + \mu_z \mathbf{I} \quad (40)$$

$$\mathbf{H}_u = \mathbf{F}' \mathbf{F} + \mu_u \mathbf{I} = \mathbf{Q} (\Lambda_F + \mu_u \mathbf{I}) \mathbf{Q}'. \quad (41)$$

The Hessian $\mathbf{H}_z$ for the z update is BTTB (non-circulant) as follows:

$$\mathbf{H}_z = \mu_v \mathbf{R}' \mathbf{R} + \mu_z \mathbf{I}, \quad (42)$$

for which there is no $\mathcal{O}(N_r)$ solver. To “invert” $\mathbf{H}_z$ for (39) we applied one iteration of preconditioned gradient descent with a circulant preconditioner. The resulting computation per iteration is essentially identical to that of the original AL-P2 with periodic boundary conditions in [4].

G. Parameter Selection

For all of the experiments shown in Section IV and in the supplement, we manually chose the spatial regularization parameter, λ , so that the converged image $x^{(\infty)}$ resembled the true image (for inpainting and CS-SENSE-MRI simulations in the supplement) or the fully-sampled body coil image (for *in vivo* data).

As with other AL algorithms, the tuning parameters $\mathbf{M}_0, \dots, \mathbf{M}_4 \geq 0$ do not affect the solution $\hat{x}$, but they can greatly affect the convergence rate. To facilitate comparison with

AL-P2-NC, we chose the AL tuning parameters of AL-P2-NC based on the guidelines provided in [4, Eqns. (40)–(41)] as follows:

$$\begin{aligned}\kappa(\mathbf{H}_u) &= 24; \quad \kappa(\mathbf{H}_z) = 12; \\ \kappa(\mathbf{H}_x) &= 0.9\kappa(\mathbf{S}'\mathbf{S}).\end{aligned}\quad (43)$$

For consistency, we selected the AL tuning parameters of AL-tridiag and ADMM-tridiag with a similar heuristic strategy. For the Hessian of the multi-coil FFT step (21) we selected its associated parameter, μ_2 , such that $\kappa(\mathbf{H}_2) = 24$, exactly as in AL-P2. Our proposed variable splitting (2) results in Hessians in which the regularizer is combined with the sensitivity encoding, so the remaining AL-P2 tuning rules are inapplicable. Instead, we selected the remaining tuning parameters (μ_0 , μ_1 , $\mathbf{M}_3$, and $\mathbf{M}_4$) to enforce $\kappa(\mathbf{H}_3) = \kappa(\mathbf{H}_x) = 12$. Because our remaining Hessians $\mathbf{H}_3$ and $\mathbf{H}_x$ partially consist of $\mathbf{C}'_v \mathbf{C}_v$ and $\mathbf{C}'_h \mathbf{C}_h$, respectively, we choose to enforce the condition number of 12 used for $\mathbf{H}_z$ of AL-P2, which is characterized by the periodic boundary finite differences. We choose to apply this heuristic over the alternative guideline based on $\kappa(\mathbf{S}'\mathbf{S})$ because our choices of $\mathbf{M}_3$ and $\mathbf{M}_4$ in (49)–(50) below make $\mathbf{H}_3$ and $\mathbf{H}_x$ approximately circulant, but far from diagonal.

First we chose the parameters that interact with the thresholding steps, μ_0 and μ_1 , based on the maximum value of the initial image, $x_{\max}$, and spatial regularization parameter, λ :

$$\mu_0 = \mu_1 = \frac{\lambda}{0.02 x_{\max}}. \quad (44)$$

This sets the threshold of the shrinkage step in (19) and (20) at 2% of the maximum initial image value. This threshold worked well for the noise level of the following simulated and *in vivo* experiments. Recalling the BDTB structures of $\mathbf{H}_3$ (25) and $\mathbf{H}_x$ (26), we designed $\mathbf{M}_3$ and $\mathbf{M}_4$ to enforce the following conditions for scalar $c_3, c_4 > 0$:

$$c_3 \mathbf{I} = \frac{\mu_2}{4} \mathbf{S}'\mathbf{S} + \mathbf{M}_3 \quad (45)$$

$$c_4 \mathbf{I} = \frac{\mu_2}{4} \mathbf{S}'\mathbf{S} + \mathbf{M}_4. \quad (46)$$

The constant diagonal term results from allowing a spatially varying $\mathbf{M}_3$ and $\mathbf{M}_4$. Therefore, $\mathbf{M}_3$ is higher in regions where the sum-of-squares (SoS) of the sensitivity maps is low and vice versa. Intuitively this results in stronger enforcement of the $u_3 = -v_3$ and $x = -v_3$ constraints in spatial regions where the $u_2 = \frac{1}{2}\mathbf{S}u_3 + \frac{1}{2}\mathbf{S}x$ constraint provides less information. We use [21] for an analytical solution for maximum and minimum eigenvalues of $\mathbf{C}'_h \mathbf{C}_h$ and

$\mathbf{C}_V' \mathbf{C}_V$. Due to (45)–(46), $\mathbf{H}_3$ and $\mathbf{H}_x$ are approximately circulant, and eigenvalue analysis of $\mathbf{H}_3$ and $\mathbf{H}_x$ becomes simple:

$$\kappa(\mathbf{H}_3) = \frac{\mu_1 \lambda_{\max}(\mathbf{C}_V' \mathbf{C}_V) + c_3}{\mu_1 \lambda_{\min}(\mathbf{C}_V' \mathbf{C}_V) + c_3} \quad (47)$$

$$\kappa(\mathbf{H}_x) = \frac{\mu_0 \lambda_{\max}(\mathbf{C}_H' \mathbf{C}_H) + c_4}{\mu_0 \lambda_{\min}(\mathbf{C}_H' \mathbf{C}_H) + c_4}. \quad (48)$$

Let $\hat{c}_3$ and $\hat{c}_x$ be the respective solutions for (47) and (48) for $\kappa(\mathbf{H}_3) = \kappa(\mathbf{H}_x) = 12$. Then the values for $\mathbf{M}_3$ and $\mathbf{M}_4$ are as follows:

$$\mathbf{M}_3 \triangleq \max \left(\hat{c}_3 \mathbf{I} - \frac{\mu_2}{4} \mathbf{S}' \mathbf{S}, 10^{-3} \right) \quad (49)$$

$$\mathbf{M}_4 \triangleq \max \left(\hat{c}_4 \mathbf{I} - \frac{\mu_2}{4} \mathbf{S}' \mathbf{S}, 10^{-3} \right). \quad (50)$$

This choice of $\mathbf{M}_3$ and $\mathbf{M}_4$ is informed by the thresholding levels through (44) but also enforces the positive-definite condition for $\mathbf{M}_3$ and $\mathbf{M}_4$. We selected AL-tridiag parameters using the same procedure.

IV. Results

A. in vivo Experiment Setup

Following [4], we used a 3D in-vivo volunteer data set acquired from a GE 3T scanner ($T_R = 25$ ms, $T_E = 5.172$ ms, voxel size = $1 \times 1.35 \times 1$ mm³) with an 8-channel head coil. A corresponding body coil dataset was also acquired for sensitivity map estimation and image quality comparison. The fully-sampled data was 256×144 with 128 samples in the read-out direction along z . We performed the proposed reconstruction algorithms for two 2D axial slices from retrospectively undersampled data. To promote FFT efficiency, we resampled the data to correspond to an image size of 256×128 . Fig. 1 shows the Poisson-disk-based undersampling pattern (reduction factor 6) in the k_x - k_y phase-encode plane that included the central 16×16 phase-encodes, pictured in Figure 1. This sampling corresponds to one slice of a 3D acquisition with frequency encoding in k_z .

We used the central 16×16 phase-encodes to generate low resolution images that were then used with the body coil image to estimate smooth sensitivity maps [12], shown in Figure 9 of the Supplement. The sensitivity values in the air regions were truncated in magnitude to

control the maximum value in the sum-of-squares of the sensitivity maps and aid in tuning AL penalty parameters.

Computation was done on a Genuine Intel Xeon CPU E5-2680 with a 2.8 GHz 20 core machine with hyper-threading. The operating system was 64-bit Red Hat 6.7 running gcc version 4.4.7. All algorithms were implemented in Matlab version 8.6 using the image reconstruction toolbox [22], and all algorithms operated on single precision data. We performed parallelization of variable updates in (22) and (23) with a Pthreaded MEX function. The Pthreaded MEX function performed blockwise Gaussian elimination in parallel across each of the tridiagonal blocks of $\mathbf{H}_3$ and $\mathbf{H}_x$. We allocated 20 Pthreads for these operations.

The initial estimate for axial slice 38 was the square-root of the sum-of-squares of the zero-filled iFFT coil images, shown in Figure 2. Figure 2 also shows the qualitative similarity between the separately acquired body coil image and the MFISTA solution of (1).

We repeated the experiment with axial slice 90 from the same *in vivo* dataset and using the same sampling pattern. The sensitivity maps estimated for this axial slice are shown in Figure 9 of the supplement. Figure 3 shows the initial sum-of-squares estimate, the separately acquired body coil image, and MFISTA solution, and the difference between the body coil and MFISTA solution. For both slice 38 and 90, the converged MFISTA solution shows lower noise than the body coil image.

For both slice 38 and 90, both proposed algorithms reached the solution $x^{(\infty)}$ (to within machine precision), shown in Figure 2c and 3c. For brevity, Figures 2c and 3c do not include $x^{(\infty)}$ for AL-tridiag or ADMM-tridiag, because they are visually indistinguishable from the MFISTA solution. The solutions $x^{(\infty)}$ were also visually similar to the fully sampled SENSE reconstruction without regularization, $\hat{x}_{\text{SENSE}}$. Image comparisons are presented in the supplement.

B. Computation Speed Results for in vivo MRI data

We quantified the convergence rate of these algorithms using the normalized root mean squared distance (NRMSD) between a given iterate $x^{(n)}$ and the converged solution, $x^{(\infty)}$, in decibels:

$$\text{NRMSD} \triangleq 20 \log_{10} \left(\frac{\|x^{(n)} - x^{(\infty)}\|_2}{\|x^{(\infty)}\|_2} \right) \quad (51)$$

To generate the solution, $x^{(\infty)}$, we ran MFISTA for 50000 outer iterations with 5 inner NCG iterations. We also calculated the normalized root mean squared error (NRMSE) between a given iterate $x^{(n)}$ and the fully-sampled SENSE reconstruction, $\hat{x}_{\text{SENSE}}$, computed without any regularization.

For computation speed, we measured the wall time of each algorithm. For the AL and ADMM algorithms, we omitted time spent tuning AL penalty parameters and compiling

Pthreaded MEX functions. The MFISTA method requires precomputation of the maximum eigenvalue of $\mathbf{S}'\mathbf{F}'\mathbf{F}\mathbf{S}$ via power iteration, which took approximately 4.3 seconds for *in vivo* experiments, whereas all the VS methods avoid this overhead. Computation time excludes time spent computing this maximum eigenvalue. For all algorithms, we omitted the time spent computing the initial sum-of-squares estimate.

As demonstrated in Figure 4, MFISTA is costlier per iteration than the proposed methods and AL-P2-NC. AL-tridiag converges slightly faster than ADMM-tridiag due to having fewer auxiliary and dual variables to update. AL-P2-NC converged the fastest for this slice.

Axial slice 90 presented a change in relative computation speed toward $x^{(\infty)}$: AL-tridiag and ADMM-tridiag converge faster than AL-P2-NC down to -65 dB NRMSD and up to 500 iterations. In the simulation results shown in the supplement, ADMM-tridiag also converged faster than AL-P2-NC in the early iterations. Overall, the speed of ADMM-tridiag is generally comparable to that of AL-P2-NC. For all AL/ADMM methods, the convergence rate depends on parameter selection; the heuristics used for parameter design in [4] may perform better under some conditions than others. One possible reason for the difference in relative convergence speeds in experiments for slice 38 and slice 90 is the smaller anatomical support in slice 90. Due to the head coil geometry, the smaller head circumference at slice 90 results in a lower signals from the surface coils, which may present a more difficult reconstruction problem.

For these axial slices, we also examined the NRMSE between iterates $x^{(n)}$ and the fully sampled SENSE reconstruction $\hat{x}_{\text{SENSE}}$ without regularization. AL-tridiag and ADMM-tridiag reach the minimum NRMSE after similar amounts of computation as AL-P2-NC, approximately 4 seconds and 60 iterations. By this metric, MFISTA performs slightly worse as a function of computation time and iterations, and would be far worse when one accounts for the overhead of running the power iteration to find the maximum eigenvalue of $\mathbf{S}'\mathbf{F}'\mathbf{F}\mathbf{S}$.

C. Image Inpainting

This section illustrates the benefits of non-periodic boundary conditions and the proposed variable splitting scheme for an inpainting problem. Unlike medical images that often have air at one or more boundaries, natural scenes typically contain useful, distinct information at the boundaries.

To test the effect of AL-tridiag-inpaint, we took a 432×540 digital photograph using a Samsung SM-G930V camera, randomly discarded 75% of the pixels, and added white Gaussian noise corresponding to 20 dB SNR to the remaining pixels. We used 2D nearest neighbor interpolation to initialize the inpainting estimate, pictured in Fig. 6.

To demonstrate the ease with which the proposed formulation accommodates orthonormal wavelets, we estimated the inpainted image using AL-tridiag-inpaint with the modified regularization operators in (29). We selected regularization parameters λ_1 and λ_2 for good image reconstruction quality, and set $\alpha_w = 1$ to limit additional memory usage. We also applied the AL-P2 variable splitting scheme to the inpainting problem, using finite-differences with periodic boundary conditions and Haar wavelets. We show the inpainting

images estimated using non-periodic boundary conditions in Fig. 6. All images are displayed on the same grayscale axis as the original image, unless otherwise noted.

Figure 7 shows the error between the inpainting estimates using non-periodic vs. periodic boundary conditions. The use of periodic boundary conditions results in higher error near the boundaries of the image.

We conducted the inpainting computational speed experiments on the machine described in Section IV-A, and we compute wall time using the same rules as in Section III-G, excluding time spent tuning AL parameters. We measure NRMSD to the MFISTA solution, $x^{(\infty)}$, as a function of wall time. As in the CS-SENSE-MRI experiments, we compare AL-tridiag-inpaint to a variant of AL-P2 to better understand the effect of the proposed variable splitting scheme. We apply the AL-P2 variable splitting scheme to the inpainting problem with non-periodic boundary conditions (32), and we call the resulting AL algorithm AL-P2-NC-inpaint. Similar to AL-P2-NC, it requires an inner iterative variable update due to the non-circulant Hessian. We solve this inner step using one iteration of preconditioned gradient descent with a circulant preconditioner.

Supposing that boundary artifacts are a secondary concern to computational speed, we also compare the speed of AL-tridiag-inpaint to AL-P2-inpaint. AL-P2-inpaint is distinct from AL-P2-NC-inpaint due to its cost function, which uses regularizers with periodic boundary conditions. AL-P2-inpaint is not handicapped by an inner iterative update, because the circulant Hessian can be diagonalized efficiently via FFTs. Though AL-tridiag-inpaint must complete two tridiagonal solver variable updates for each of AL-P2-inpaint's FFT-based variable updates, the $\mathcal{O}(n)$ runtime of the tridiagonal solver and the parallelized implementation result in comparably fast iterations for AL-tridiag-inpaint. The average computation time for each iteration was 0.0543 seconds for AL-P2-inpaint, 0.1193 seconds for AL-P2-NC-inpaint, 0.0583 seconds for AL-tridiag-inpaint, and 0.0745 seconds for ADMM-tridiag-inpaint.

Figure 8 also demonstrates convergence benefits of the proposed variable splitting scheme. Both AL-tridiag-inpaint and ADMM-tridiag-inpaint were able to reach the same $x^{(\infty)}$ as MFISTA, unlike AL-P2-NC.

V. Discussion

The ADMM-tridiag algorithm provides a simple way to ensure convergence for variable splitting methods. By examining the constraint matrix and designing $\mathbf{B}$, we show equivalence between the variable splitting scheme in (8) and ADMM. The additional variable split and variable update led to parallelizability of two of the resulting variable updates. For applications as sensitive as medical diagnosis, an algorithm with convergence guarantees may be preferable to those having unknown convergence properties.

Unlike AL-P2 [4], the proposed algorithm, ADMM-tridiag, has a convergence guarantee and addresses non-periodic boundary conditions, while demonstrating comparable computational speed. Using heuristic parameter tuning based on condition numbers of

variable update Hessians, we demonstrated that the speeds of AL-P2-NC and ADMM-tridiag are similar but can vary depending on experimental conditions.

The proposed variable splitting scheme can be useful for a variety of image processing problems, because many natural scenes have non-zero values at the boundaries that do not relate periodically to opposite boundaries. As shown in Section III-E, the proposed method AL-tridiag is readily adapted to image denoising and inpainting problems. Image deblurring is also a good candidate for this proposed variable splitting. Separation of horizontal and vertical differences has also been explored in image segmentation [23]. For MRI reconstruction, the benefit of non-periodic boundary conditions is evident for anatomy that is not entirely surrounded by air. Additional simulation experiments in the supplement illustrate the effect of periodic boundary conditions in reconstructing a sagittal slice of the brain.

To fully benefit from the $\max(N_x, N_y)$ parallelizable tridiagonal updates of AL-tridiag and ADMM-tridiag, one should use a highly parallel computing platform, such as a GPU.

The algorithms proposed in this work have several limitations. Though the proposed algorithms are designed to facilitate fast computation, the convergence speed is highly dependent on good penalty parameter choice. Though we present some useful heuristics for choosing the AL convergence parameters, the optimal procedure for designing these parameters is a difficult analysis problem and still unknown. (For some simpler ADMM methods, optimal parameter tuning has been analyzed [24].) Moreover, the complexity of convergence parameter design increases with the number of variable splits, and this work is built around an additional separation of horizontal and vertical differences into distinct auxiliary variables. The increased number of tuning parameters introduces another degree of freedom. Using non-scalar penalty parameter matrices $\mathbf{M}_3$ and $\mathbf{M}_4$ further increases the degrees of freedom compared to the simple scalar choice used in most AL methods.

Though the proposed variable splitting scheme can be easily extended to 3D reconstruction problems, this would require introducing two additional variable splits, separating each of the three finite difference directions and introducing a second auxiliary variable proxy for x . Though the corresponding variable updates can be quickly computed with shrinkage and the parallelizable tridiagonal solver, this 3D variable splitting scheme could further complicate penalty parameter analysis.

Finally, the variable splitting scheme at the center of the proposed algorithms is applicable only for regularization with first-order finite differences. Though Section III-D shows that the formulation also accommodates orthonormal wavelet penalties combined with finite differences, this variable splitting scheme yields no benefit for other sparsity transforms, e.g., non-orthonormal wavelets and learned dictionaries.

VI. Summary and Conclusion

This work proposes a variable splitting algorithm for SENSE MRI reconstruction, ADMM-tridiag. The proposed method offers convergence guarantees and efficient variable updates for non-periodic boundary conditions. ADMM-tridiag efficiently handles the non-periodic

boundary conditions by separating the finite differences in the horizontal and vertical directions to create easily solvable and parallelizable tridiagonal problems. The method for inducing ADMM equivalence requires only one additional variable split and variable update.

We also presented a simpler variation of this algorithm: AL-tridiag. AL-tridiag was derived from the same variable splitting scheme as ADMM-tridiag, but has a simpler update procedure, resulting in a slight speed increase albeit without any convergence guarantees. We showed a simple relationship between AL-tridiag and ADMM-tridiag and compared their convergence speeds to that of AL-P2-NC and MFISTA. Convergence speed was evaluated in terms of distance to the solution of the proposed cost function (1) as well as to the fully sampled SENSE reconstruction. For retrospectively undersampled *in vivo* data, the proposed algorithms demonstrated comparable convergence speed and produced reconstructed images with good image quality. AL-tridiag was also applied to a noisy image inpainting problem, demonstrating faster convergence speed than AL-P2-NC, and improved image fidelity at the boundaries than AL-P2 with periodic boundary conditions.

All of the proposed algorithms require selection of penalty parameters. We use heuristics determined in [4] to select these AL penalty parameters, although we exploited a more general version of tuning parameters to enable computation of condition numbers for tridiagonal Hessians. Using methods that adapt the parameters as a function of iteration [25] might simplify and accelerate AL methods.

The supplement also describes a fully parallelizable variant of ADMM-tridiag inspired by [16] that converged slower than the proposed methods, as well as additional simulation results and image quality comparisons.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors would like to thank Sathish Ramani for helpful discussions regarding variable splitting methods for SENSE MRI reconstruction.

Supported in part by NIH grant P01 CA 059827.

References

1. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Mag Res Med*. Nov; 1999 42(5):952–62.
2. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Mag Res Med*. Dec; 2007 58(6):1182–95.
3. Eckstein J, Bertsekas DP. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*. Apr; 1992 55(1–3):293–318.
4. Ramani S, Fessler JA. Parallel MR image reconstruction using augmented Lagrangian methods. *IEEE Trans Med Imag*. Mar; 2011 30(3):694–706.
5. Chen Y, Hager W, Huang F, Phan D, Ye X, Yin W. Fast algorithms for image reconstruction with application to partially parallel MR imaging. *SIAM J Imaging Sci*. 2012; 5(1):90–118.

6. Beck A, Teboulle M. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Trans Im Proc.* Nov; 2009 18(11):2419–34.
7. Eckstein J. Parallel alternating direction multiplier decomposition of convex programs. *J Optim Theory Appl.* Jan; 1994 80(1):39–62.
8. Yang J, Zhang Y, Yin W. A fast alternating direction method for TVL1-L2 signal reconstruction from partial Fourier data. *IEEE J Sel Top Sig Proc.* Apr; 2010 4(2):288–97.
9. Chen Y, Ye X, Huang F. A novel method and fast algorithm for MR image reconstruction with significantly undersampled data. *Inverse Problems and Imaging.* 2010; 4(2):223–240.
10. Matakos A, Ramani S, Fessler JA. Accelerated edge-preserving image restoration without boundary artifacts. *IEEE Trans Im Proc.* May; 2013 22(5):2019–29.
11. Almeida MSC, Figueiredo MAT. Deconvolving images with unknown boundaries using the alternating direction method of multipliers. *IEEE Trans Im Proc.* Aug; 2013 22(8):3074–86.
12. Allison MJ, Ramani S, Fessler JA. Accelerated regularized estimation of MR coil sensitivities using augmented Lagrangian methods. *IEEE Trans Med Imag.* Mar; 2013 32(3):556–64.
13. Afonso MV, Bioucas-Dias José M, Mário A, Figueiredo T. Fast image recovery using variable splitting and constrained optimization. *IEEE Trans Im Proc.* Sep; 2010 19(9):2345–56.
14. Ye X, Chen Y, Huang F. Computational acceleration for MR image reconstruction in partially parallel imaging. *IEEE Trans Med Imag.* May; 2011 30(5):1055–63.
15. Bertsekas DP. Multiplier methods: A survey. *Automatica.* Mar; 1976 12(2):133–45.
16. Ramani, S., Fessler, JA. Regularized parallel MRI reconstruction using an alternating direction method of multipliers. *Proc. IEEE Intl. Symp. Biomed. Imag;* 2011; p. 385-8.
17. Fischer B, Modersitzki J. Fast inversion of matrices arising in image processing. *Numer Algorithms.* 1999; 22(1):1–11.
18. Davis, PJ. Circulant matrices. New York: Wiley; 1983.
19. Chan RH, Yang J, Yuan X. Alternating direction method for image inpainting in wavelet domains. *SIAM Journal on Imaging Sciences.* 2011; 4(3):807–826.
20. Dahl J, Hansen PC, Jensen SH, Jensen TL. Algorithms and software for total variation image reconstruction via first-order methods. *Numer Algorithms.* Jan; 2010 51(1):67–92.
21. Strang G. The discrete cosine transform. *SIAM Review.* 1999; 41(1):135–47.
22. Fessler, JA. Matlab tomography toolbox. 2004. available from <http://web.eecs.umich.edu/~fessler/code/>
23. Storath M, Weinmann A. Fast partitioning of vector-valued images. *SIAM J Imaging Sci.* 2014; 7(3):1826–52.
24. Ghadimi, E., Teixeira, A., Shames, I., Johansson, M. Optimal parameter selection for the alternating direction method of multipliers (ADMM): quadratic problems. 2013, optimization Online. [Online]. Available: http://www.optimization-online.org/DB_HTML/2013/06/3917.html
25. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found & Trends in Machine Learning.* 2010; 3(1):1–122.

Biographies

Mai Le received the B. Sc. in Electrical Engineering from Stanford University in 2012 and the M.S. also in Electrical Engineering from the University of Michigan in 2013. Currently she is a Ph. D. candidate in Electrical Engineering: Systems at the University of Michigan. Her current research interests include model based, iterative reconstruction methods for MRI.

Jeffrey A. Fessler is the William L. Root Professor of EECS at the University of Michigan. He received the BSEE degree from Purdue University in 1985, the MSEE degree from Stanford University in 1986, and the M.S. degree in Statistics from Stanford University in 1989. From 1985 to 1988 he was a National Science Foundation Graduate Fellow at Stanford, where he earned a Ph.D. in electrical engineering in 1990. He has worked at the University of Michigan since then. From 1991 to 1992 he was a Department of Energy Alexander Hollaender Post-Doctoral Fellow in the Division of Nuclear Medicine. From 1993 to 1995 he was an Assistant Professor in Nuclear Medicine and the Bioengineering Program. He is now a Professor in the Departments of Electrical Engineering and Computer Science, Radiology, and Biomedical Engineering. He became a Fellow of the IEEE in 2006, for contributions to the theory and practice of image reconstruction. He received the Francois Erbsmann award for his IPMI93 presentation, and the Edward Hoffman Medical Imaging Scientist Award in 2013. He has served as an associate editor for IEEE Transactions on Medical Imaging, the IEEE Signal Processing Letters, and the IEEE Transactions on Image Processing, and is currently serving as an associate editor for the IEEE Transactions on Computational Imaging. He has chaired the IEEE T-MI Steering Committee and the ISBI Steering Committee. He was co-chair of the 1997 SPIE conference on Image Reconstruction and Restoration, technical program co-chair of the 2002 IEEE International Symposium on Biomedical Imaging (ISBI), and general chair of ISBI 2007. His research interests are in statistical aspects of imaging problems, and he has supervised doctoral research in PET, SPECT, X-ray CT, MRI, and optical imaging problems.

Fig. 1. Retrospective Poisson-disk-based undersampling pattern used for *in vivo* experiments, with reduction factor of 6 and fully sampled central 16×16 phase-encodes.

Fig. 2.

Axial slice 38 of volunteer data (a) square root of sum-of-squares of the zero-filled iFFT coil images, used as an initial estimate for all algorithms; (b) separately acquired body coil image; (c) $x^{(\infty)}$ calculated by MFISTA; (d) difference between body coil and $x^{(\infty)}$.

Fig. 3.

Axial slice 90 of volunteer data (a) square root of sum-of-squares of the zero-filled iFFT coil images, used as an initial estimate for all algorithms; (b) separately acquired body coil image; (c) $x^{(\infty)}$ calculated by MFISTA; (d) difference between body coil and $x^{(\infty)}$.

Fig. 4.

Axial slice 38: (a, b) NRMSE comparison of AL-tridiag, ADMM-tridiag, MFISTA, and AL-P2-NC to $x^{(\infty)}$; (c, d) NRMSE comparison of AL-tridiag, ADMM-tridiag, MFISTA, and AL-P2-NC to fully sampled SENSE reconstruction.

Fig. 5.

Axial slice 90: (a, b) NRMSD comparison of AL-tridiag, ADMM-tridiag, MFISTA, and AL-P2-NC to $x^{(\infty)}$; (c, d) NRMSE comparison of AL-tridiag, ADMM-tridiag, MFISTA, and AL-P2-NC to fully sampled SENSE reconstruction.

Fig. 6.

Inpainting images (a) true image; (b) nearest neighbor interpolation of noisy, partial data (SNR = 20 dB, 75% discarded); (c) inpainting estimate using finite differences with non-periodic boundary conditions and Haar wavelet regularization; (d) inpainting estimate using finite differences with periodic boundary conditions and Haar wavelet regularization.

Fig. 7.

Absolute difference image between true image and inpainting reconstruction ($\times 10$) (a) using finite differences with non-periodic boundary conditions and Haar wavelets (NRMSE = 0.153); (b) using finite differences with periodic boundary conditions and Haar wavelets (NRMSE = 0.155); (c) corner detail of (a); (d) corner detail of (b).

Fig. 8. NRMSD to $x^{(\infty)}$ as a function of elapsed computational time for the proposed inpainting algorithms, AL-tridiag-inpaint and ADMM-tridiag-inpaint, compared with existing methods AL-P2-NC, and MFISTA.