MIT Open Access Articles

Sparsity-Promoting Calibration for GRAPPA Accelerated Parallel MRI Reconstruction

Sparsity-Promoting Calibration for GRAPPA Accelerated Parallel MRI Reconstruction

Daniel S. Weller [Member, IEEE],
University of Michigan, 1301 Beal Avenue,Room 4125, Ann Arbor, MI, 48109 USA, phone: +1.734.615.5735

Jonathan R. Polimeni [Member, IEEE],
Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, and Harvard Medical School, Boston, MA

Leo Grady [Member, IEEE],
HeartFlow, Redwood City, CA

Lawrence L. Wald,
Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, and Harvard Medical School, Boston, MA

Elfar Adalsteinsson, and
Massachusetts Institute of Technology, Cambridge, MA

Vivek K Goyal [Senior Member, IEEE]
Massachusetts Institute of Technology, Cambridge, MA

Daniel S. Weller: dweller@alum.mit.edu

Abstract

The amount of calibration data needed to produce images of adequate quality can prevent auto-calibrating parallel imaging reconstruction methods like Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) from achieving a high total acceleration factor. To improve the quality of calibration when the number of auto-calibration signal (ACS) lines is restricted, we propose a sparsity-promoting regularized calibration method that finds a GRAPPA kernel consistent with the ACS fit equations that yields jointly sparse reconstructed coil channel images. Several experiments evaluate the performance of the proposed method relative to un-regularized and existing regularized calibration methods for both low-quality and underdetermined fits from the ACS lines. These experiments demonstrate that the proposed method, like other regularization methods, is capable of mitigating noise amplification, and in addition, the proposed method is particularly effective at minimizing coherent aliasing artifacts caused by poor kernel calibration in real data. Using the proposed method, we can increase the total achievable acceleration while reducing degradation of the reconstructed image better than existing regularized calibration methods.

I. Introduction

Parallel imaging with multi-channel receive array coils and auto-calibrating reconstruction methods like Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) [1]

enable the recovery of high-quality images from undersampled k-space data, accelerating magnetic resonance imaging (MRI) acquisitions. Undersampling by skipping phase-encode lines in Cartesian MRI reduces the field of view (FOV) of the associated image, aliasing objects larger than the reduced FOV. Parallel imaging uses the inhomogeneous receive field sensitivities of the array coil as an additional source of spatial encoding to resolve the aliasing caused by this undersampling. GRAPPA consists of both a kernel calibration step and a reconstruction step; for details of both steps, see [1]. The quality of the calibration, a least-squares fit from a densely-spaced block of auto-calibration signal (ACS) data, directly influences the ability of the reconstruction step to properly resolve coherent aliasing; an ill-conditioned fit also amplifies noise in the reconstructed images [2]. At higher accelerations, larger kernels are needed, and a greater quantity of ACS data must be collected to properly calibrate these kernels [3], [4]. At the high levels of undersampling motivating this work, the ACS size necessary for high-quality conventional calibration severely constrains the total acceleration [5]. In this paper, we apply sparsity to improve the calibration quality with limited ACS data, enabling greater acceleration than with existing methods.

Tikhonov regularization [6] improves SENSE [7], [8] and can be applied to GRAPPA kernel calibration as well. Truncating the singular value decomposition (SVD) of the ACS source matrix can improve the conditioning of the least-squares calibration, helping robustness to noise [4]. A nonlinear technique [9] enforces the frequency-shift interpretation of the GRAPPA operator [10]. The proposed approach does not model the kernel directly, but it relies on the sparsity of the reconstructed images to impose indirectly a model on the kernel used to reconstruct that image. The sparsity of a variety of MRI images was demonstrated previously in the context of reconstruction using compressed sensing [11]. The underlying assumption of this work to be validated is that coil images reconstructed using GRAPPA with the correct kernel inherit the sparsity of the object being imaged.

Statistically speaking, GRAPPA is a low-complexity compromise of SENSE. SENSE roughly inverts multiplicative sensitivities in the image domain, which corresponds to deconvolution in the frequency domain. GRAPPA also performs convolution in the frequency domain, but with much smaller convolution kernels. Since SENSE is optimal in the mean squared error (MSE) sense, GRAPPA calibration could be interpreted as finding the spectrum of the best low-resolution approximation to the MSE-optimal reconstruction. However, such optimality requires a high quality calibration, like SENSE requires high quality coil sensitivities. By adding a sparse reconstruction condition to the GRAPPA kernel calibration, we are effectively using prior information that the optimal kernel produces transform sparse images, hence overcoming the statistical inadequacies of a low quality calibration based only on the k-space measurements, which may be too noisy or too few in number to yield useful kernels by themselves.

Sparsity-promoting regularization has other uses in parallel MRI reconstruction. Implementations combining sparsity with SENSE [12], [13] provide effective regularization for challenging problems like cardiac perfusion imaging. L_1-SPIRiT [14], [15] regularizes data-preserving SPIRiT reconstruction using a transform-domain joint sparsity-promoting $\ell_{1,2}$ norm, just like this work. However, this is the first incorporation of sparse modeling in the calibration stage of GRAPPA-like parallel imaging. L_1-SPIRiT applies sparsity after the SPIRiT kernels have already been calibrated, which means that inconsistencies in the SPIRiT kernel calibration may remain and propagate to the L_1-SPIRiT reconstruction, unlike in our method. Also, SPIRiT reconstruction applies the kernels to all data (known and unknown), yielding an iterative method without a cost-effective direct solution. In our work, we make use of the fast GRAPPA reconstruction enabled by uniform undersampling to transform our kernels into images whose transform sparsity can be evaluated. These key differences separate our work from previous endeavors to improve parallel imaging using

sparsity. Other denoising work such as expanding the GRAPPA kernel to use quadratic weighting terms [16] also bears mention, and the same approach to improving calibration can be applied to such extensions of GRAPPA as well.

We begin by introducing our notation for GRAPPA calibration and reconstruction. Our algorithm for using joint sparsity of the reconstructed channel images to regularize the GRAPPA kernel calibration, originally proposed in [17], is explained, and the consequences of uniform and nonuniform undersampling are discussed. We validate our assumption about the sparsity of GRAPPA reconstruction on real data. We present performance comparisons for both simulated and real brain data, including portraying the trade-off between reconstructed image quality and total acceleration studied in [18]. New in this paper, we study the effects of choosing an appropriate tuning parameter more thoroughly and include comparisons of g-factor maps and reconstructed image error autocorrelations.

II. Theory

We first provide our notation for the GRAPPA parallel imaging acquisition, calibration, and reconstruction process. A more thorough discussion of GRAPPA is provided in [1]. Existing methods for calibration, including Tikhonov, truncated-SVD, and nonlinear GRAPPA regularization, and the proposed sparsity-promoting regularization are briefly described. The proposed method relies on uniform undersampling, and extensions to nonuniform undersampling are explored.

A. Acquisition and GRAPPA Reconstruction

Let us consider the 2-D imaging case; 3-D reconstruction involves performing an inverse discrete Fourier transform (DFT) along the fully-sampled frequency encode direction and slice-by-slice reconstructions of the resulting 2-D hybrid k-space (other generalizations are possible; see [19]). The full-FOV true 2-D image $m(x, y)$ is sampled in the k-space domain by a P-channel receive array coil, with sensitivity profiles $s_1(x, y),…, s_P(x, y)$. Uniform Cartesian sampling of k-space yields noisy observations $d_p[k_x, k_y]$ of the samples $y_p[k_x, k_y]$ of k-space. Our samples are perturbed by additive complex Gaussian noise $n_p[k_x, k_y]$, uncorrelated across k-space frequencies. For a given frequency, the measurements from different coil channels have covariance Λ (assume this covariance does not vary across k-space). For accelerated parallel imaging, we undersample in one or both phase encode directions (by factors of R_x and R_y, respectively). This acquisition model is depicted in Fig. 1.

GRAPPA reconstructs missing (target) k-space points using a weighted linear combination of neighboring acquired (source) points from all the coils. We use a $B_x \times B_y$-point kernel g_{p,q,r_x,r_y} for each pair of input $p \in \{1, …, P\}$ and output $q \in \{1,…,P\}$ coil channels and each 2-D target offset $r_x \in \{0,…, R_x-1\}$ and $r_y \in \{0,…, R_y-1\}$. To compute the vector of kernel weights $[\mathbf{G}]_{q,r_x,r_y} = \text{vec}([g1_{,q,r_x,r_y}, …, gP_{,q,r_x,r_y}])$ ($\text{vec}(\cdot)$ reshapes a matrix into a column vector) for a given offset r_x, r_y and output coil q, we form a column vector of known target points $\left[\mathbf{D}_t^{ACS}\right]_{q,r_x,r_y}$ and a matrix of source points $\mathbf{D}_s^{ACS}$ from the ACS data. The resulting linear system of N_{Fit} fit equations $\left[\mathbf{D}_t^{ACS}\right]_{q,r_x,r_y} = \mathbf{D}_s^{ACS}[\mathbf{G}]_{q,r_x,r_y}$ has $B_x B_y P$ unknowns. If the system has full column rank, the least-squares solution is $[\mathbf{G}]_{q,r_x,r_y} = \left(\mathbf{D}_s^{ACS}\right)^{\dagger}\left[\mathbf{D}_t^{ACS}\right]_{q,r_x,r_y}$ where $(A)^{\dagger}$ is the left pseudo-inverse $(\mathbf{A}^H \mathbf{A})^{-1} \mathbf{A}^H$, and $\mathbf{A}^H$ is the conjugate transpose of complex matrix $\mathbf{A}$. Collecting the column vectors into a matrix of target points $\mathbf{D}_t^{ACS}$ and kernel weights $\mathbf{G}$ yields the linear system $\mathbf{D}_t^{ACS} = \mathbf{D}_s^{ACS}\mathbf{G}$ which can be solved for all the kernel weights needed to fill in the un-acquired k-space.

How many fit equations are there? In the 1-D case, a set of N_{ACS} points yields $N_{Fit} = N_{ACS} - (B - 1)R$ fit equations for $B > 1$, and $N_{Fit} = N_{ACS} - R + 1$ fits for $B = 1$. For higher dimensions, N_{Fit} is the product of the number of fits for each dimension considered individually.

The GRAPPA reconstruction $\mathbf{Y} = \text{GRAPPA}(\mathbf{D}, \mathbf{G})$ forms the full multi-coil k-space $\mathbf{Y}$ from the undersampled data $\mathbf{D}$ and the kernel $\mathbf{G}$. More specifically, $\mathbf{Y} = \bar{\mathbf{K}}^T \mathbf{D}_{conv} \mathbf{G} + \mathbf{K}^T \mathbf{D}$, where $\bar{\mathbf{K}}^T$ and $\mathbf{K}^T$ embed the interpolated and known k-space, respectively, into the full k-space, and $\mathbf{D}_{conv}$ is a block convolution matrix for multiplication of the acquired data and the kernel weights from all the coils.

B. Regularization of the Kernel Calibration

The least-squares problem formulation of conventional GRAPPA kernel calibration can be regularized using a penalty function $\mathcal{R}(\mathbf{G})$:

$$\widehat{\mathbf{G}} = \arg_{\mathbf{G}} \min \frac{1}{2} \| \mathbf{D}_s^{ACS} \mathbf{G} - \mathbf{D}_t^{ACS} \|_F^2 + \mathrm{R}(\mathbf{G}). \quad (1)$$

The Frobenius norm $\|\mathbf{A}\|_F$ is the ℓ_2 norm of the vector formed by stacking the columns of $\mathbf{A}$.

As mentioned in the discussion of previous work, several choices of regularization have been investigated in the literature. Four methods compared in this work are basic Tikhonov regularization, truncated-SVD-based regularization, nonlinear GRAPPA-operator-based regularization, and our novel sparsity-promoting regularization. Energy-minimizing basic Tikhonov regularization has the form $\mathscr{R}(\mathbf{G}) = \frac{1}{2}\|\alpha\mathbf{G}\|_F^2$, ($\alpha > 0$). The Tikhonov-regularized solution has the closed form

$$\widehat{\mathbf{G}} = ((\mathbf{D}_s^{ACS})^H (\mathbf{D}_s^{ACS}) + \alpha^2 \mathbf{I})^{-1} (\mathbf{D}_s^{ACS})^H (\mathbf{D}_t^{ACS}). \quad (2)$$

Note that the regularized solution is guaranteed to exist, even when the un-regularized least-squares problem is underdetermined. Tikhonov regularization mitigates noise amplification, since an energy-minimizing model reduces the gain of the kernel. The truncated-SVD approach [4] zeros out the singular values of the source matrix $\mathbf{D}_s^{ACS}$ less than a threshold τ relative to the maximum singular value (one can also retain a fixed number of singular values), reducing the gain of the kernel. The nonlinear GRAPPA-operator-based regularization method [9] penalizes the difference between the GRAPPA kernel applied R times, which maps acquired data to the same data shifted one block over, and a permutation matrix representing that shift. This method is effective at improving image quality, including reducing aliasing, with fewer ACS lines at low accelerations [9].

We propose leveraging the transform-domain sparsity of MRI images to improve calibration. We consider the shared sparsity of the true full-FOV channel images and use an $\ell_{1,2}$ hybrid norm to promote joint sparsity of a suitable transform of the coil channel images: for a length-N, P-channel joint sparse representation $\mathbf{W}$, $\|\mathbf{W}\|_{1,2} = \sum_{n=1}^{N} \|W_{n,1,...,}W_{n,P}\|_2$. As we are regularizing the kernel, we first "transform" the kernel into full-FOV reconstructed multi-channel k-space by convolving the kernel set with the acquired undersampled data, and the $\ell_{1,2}$ norm is evaluated on the sparse transform Ψ of the inverse DFT $\mathcal{F}^{-1}$ of the full reconstructed k-space. We can use any suitable linear transform or trained dictionary to sparsify the reconstruction coil images; in this paper, we investigate total variation and the

four-level bi-orthogonal '9-7' discrete wavelet transform (DWT). The optimization problem becomes

$$\widehat{\mathbf{G}}=\arg_{\mathbf{G}} \min \frac{1}{2}\|\mathbf{D}_s^{\text{ACS}}\mathbf{G} - \mathbf{D_t}^{\text{ACS}}\|_F^2 + \lambda\|\psi \text{F}^{-1}\text{GRAPPA}(\mathbf{D},\mathbf{G})\|_{1,2}, \quad (3)$$

where the tuning parameter $\lambda > 0$ controls the level of regularization.

Although more efficient computational methods based on message passing algorithms for compressed sensing [20] may be applicable for solving Eq. (3), we use half-quadratic minimization [21], [22], which iteratively approximates the $\ell_{1,2}$ norm with a weighted least-squares term and updates the weights after each iteration. Each iteration has the form

$$\widehat{\mathbf{G}}^{(i+1)}=\arg_{\mathbf{G}} \min \frac{1}{2}\|\mathbf{D}_s^{\text{ACS}}\mathbf{G} - \mathbf{D_t}^{\text{ACS}}\|_F^2 + \frac{\lambda}{2}\|(\Delta^{(i)})^{1/2}\psi \text{F}^{-1}\text{ GRAPPA}(\mathbf{D},\mathbf{G})\|_F^2. \quad (4)$$

Algorithm 1 GRAPPA with sparsity-promoting kernel calibration.

Require: D, $\mathbf{D}_s^{\text{ACS}}$, $\mathbf{D}_t^{\text{ACS}}$ $\hat{G}^{(0)}$, λ, ε, I, tol

1	Set $f^{(0)}$ to objective in Eq. (3) evaluated for G = $\hat{G}^{(0)}$.
2	**for** $i = 1 : I$ **do**
3	$\mathbf{W}^{(i-1)} \leftarrow \Psi \mathcal{F}^{-1} \text{GRAPPA}(\mathbf{D}, \hat{\mathbf{G}}^{(i-1)})$
4	$\Delta_{n,n}^{(i-1)} \leftarrow \|[W_{n,1}^{(i-1)}, \ldots, W_{n,P}^{(i-1)}, \varepsilon]\|_2^{-1}$ for $n=0, \ldots, N-1$.
5	Run LSMR to solve Eq. (4) for $\hat{\mathbf{G}}^{(i)}$ using $\mathbf{\Delta}^{(i-1)}$.
6	Compute $f^{(i)}$ using $\hat{\mathbf{G}}^{(i)}$ in the objective in Eq. (3).
7	**if** $f^{(i-1)} - f^{(i)} \leq$ tol. $f^{(i-1)}$ **then**
8	**break**
9	**end if**
10	**end for**
11	$\hat{\Upsilon} \leftarrow \text{GRAPPA}(\mathbf{D}, \hat{\mathbf{G}}^{(i)})$.
12	**return** $\hat{\Upsilon}$.

The diagonal weight matrix $\Delta_{n,n}^{(i)}=\|[W_{n,1}^{(i)}, \ldots, W_{n,P}^{(i)}, \in]\|_2^{-1}$, the sparse transform $\mathbf{W}^{(i)} = \Psi \mathcal{F}^{-1}$ GRAPPA($\mathbf{D}$, $\hat{\mathbf{G}}^{(i)}$), and $\varepsilon > 0$ is small to smooth the ℓ_1 norm. The resulting least-squares problems are solved iteratively using the LSMR [23] solver. The LSMR method is a descent method for minimizing $\|\text{A}x - \mathbf{b}\|_2^2$ that requires performing matrix-vector multiplication involving $\mathbf{A}$ and $\mathbf{A}^H$ efficiently. For Eq. (4), $\mathbf{x}$ corresponds to vec($\mathbf{G}$), and $\mathbf{A}x - $b requires computing

$$\mathbf{R}_1=\mathbf{D}_s^{\text{ACS}}\mathbf{G} - \mathbf{D_t}^{\text{ACS}}, \text{and} \quad (5)$$

$$\mathbf{R}_2= \sqrt{\lambda}(\Delta^{(i)})^{1/2}\psi \mathcal{F}^{-1} \text{GRAPPA}(\mathbf{D},\mathbf{G}). \quad (6)$$

Multiplying the adjoint $\mathbf{A}^H$ by the residual uses

$$(\mathbf{D}_s^{\mathrm{ACS}})^H \mathbf{R}_1 + \mathrm{GRAPPA} * (\mathbf{D}, \sqrt{\lambda} \mathrm{F} \psi^H (\Delta^{(i)})^{1/2} \mathbf{R}_2). \quad (7)$$

Since the GRAPPA reconstruction step consists of convolution of the data with the different kernels, the adjoint of the GRAPPA reconstruction operation GRAPPA* ($\mathbf{D}$, $\mathbf{Y}$) consists of convolution of $\mathbf{Y}$ with the conjugated and reversed acquired k-space data. Both these operations can be implemented efficiently using the fast Fourier transform (FFT). The proposed method is summarized in Algorithm 1. The reconstructed full multi-channel k-space $\hat{Y}$ is then post-processed to combine the channels into a single image $\hat{m}[x, y]$ for evaluation.

C. Uniform vs. Nonuniform Undersampling

This implementation suffices when the spacing between acquired k-space samples is uniform, since we need to calibrate only one set of kernels, and the reconstruction step can be very efficiently implemented using the FFT. Since the GRAPPA reconstruction is performed within the inner iterations of sparsity-promoting calibration, an efficient implementation is essential for practical use.

When k-space is undersampled nonuniformly, GRAPPA can be adjusted in a couple ways. One approach uses different kernels for each target point or block of target points. While the missing k-space still is interpolated via direct computation, the need for many sets of kernels greatly expands the number of variables that need to be optimized during the calibration phase. Another approach [24] reconstructs the full k-space by finding values for the missing data consistent with the acquired data using a single kernel set relating source and target points throughout k-space. The implementation of this method solves a large-scale (regularized) least-squares problem, which no longer has an efficient direct implementation.

Tiled small random patterns [25] would reduce the number of kernels needed for direct nonuniform GRAPPA while preserving significant randomness in the undersampling pattern. In this case, the GRAPPA reconstruction step remains relatively fast, so the proposed method can be used without significant modification. When greater randomness in the sampling pattern is desired, a joint calibration and reconstruction method can be performed, alternating between updating the calibrated kernel according to the sparsity of the current estimate of full k-space and updating the full k-space according to iterative GRAPPA with the current kernel.

III. Methods

Both simulated and real data (see Fig. 2) are used to analyze the proposed method. We simulated a noise-free T_1-weighted normal brain with 2.0 mm isotropic resolution using the Brain-Web database (http://www.bic.mni.mcgill.ca/brainweb/) [26], and we multiplied a 128×128 voxel axial slice of this volume (the ground truth) with coil sensitivities simulated for a 16-channel circular array receive coil. This synthetic data set (with 20 dB SNR noise added) is phase-encoded in the AP-direction and frequency-encoded in the RL-direction. We employed total variation to sparsify the simulated data due to its piece-wise smooth nature.

We also acquired a three-dimensional brain image of a consented subject using a conventional T_1-weighted MPRAGE (TR = 2.53 s, TI = 1.1 s, TE = 3.45 ms, bandwidth = 190 Hz/pixel) pulse sequence with 256 × 256 × 176 mm FOV and 1.0 mm isotropic spatial resolution with a Siemens 32-channel receive array coil in a Siemens Tim Trio (3 T) MRI scanner. This acquisition completed in about eight minutes without acceleration. We

processed the raw data later using MATLAB. First, we normalized the data post-scan by dividing by the sum-of-squares of smoothed separately measured coil sensitivities (these sensitivities are then discarded) to undo the attenuation in the center of the head from a sum-of-squares combination of inhomogeneous coil sensitivities. We extracted an axial slice from the volume; since the head-to-foot direction was frequency-encoded, the axial slice plane contained both phase encode directions. The slice was cropped in the image domain before undersampling. By undersampling both phase encode directions, we imitated slice-by-slice processing of a 3-D data set that is acquired with 2-D undersampling and transformed to a hybrid image/k-space domain by taking the inverse-DFT along the fully-sampled frequency-encoded z-direction. We combined the cropped and normalized non-undersampled k-space using a sum-of-squares method [27] to form a magnitude image, retained as a reference for image quality comparisons. Without highly accurate sensitivities, the complex-valued combined image phase may be inaccurate. Fortunately, the magnitude is sufficient for evaluating these T_1-weighted brain images, although phase information is important for other applications. We used the bi-orthogonal '9-7' DWT to sparsify reconstructions with this data set.

In this paper, we measure total acceleration R by dividing the total number of points in the full-FOV k-space by the total number of acquired points:

$$R = (R_x R_y) \frac{XY}{XY + N_{\mathrm{ACS}}^2 (R_x R_y - 1)}, \quad (8)$$

for a square $N_{\mathrm{ACS}} \times N_{\mathrm{ACS}}$ block of ACS lines and $R_x \times R_y$ nominal undersampling.

To combine the reconstructed multi-channel k-space into a single image, we applied un-accelerated SENSE using low-resolution coil sensitivity profiles estimated from the ACS lines. These ACS lines were apodized with a Blackman window to mitigate ringing artifacts. When comparing reconstructions of the same image using different numbers of ACS lines (as in Fig. 5 and 6), we used the same set of sensitivities for all coil combinations, as to not affect the comparison. When evaluating the quality of reconstructed images, visual comparisons using both magnitude and difference images are standard. From these images, residual aliasing artifacts should be readily identifiable. However, visual comparison is not conducive towards establishing quantitative trends in image quality. For such trends, the peak-signal-to-noise ratio (PSNR) is used; for magnitude images, it is defined in dB units as

$$\mathrm{PSNR} = \mathbf{10} \log_{\mathbf{10}} \left(\frac{\max_{\mathbf{x,y}} |\mathbf{m}[\mathbf{x,y}]|^2}{\frac{1}{XY} \sum_{\mathbf{x,y}} \||\mathbf{m}[\mathbf{x,y}]| - |\widehat{\mathbf{m}}[\mathbf{x,y}]|\|^2} \right), \quad (9)$$

where $|m[x, y]|$ is the magnitude combined "ground-truth" image, and $|\hat{m}[x,y]|$ is the magnitude combined image from the GRAPPA-reconstructed k-space. PSNR has several limitations: it does not correlate strongly with visual image quality, and it is not a robust statistical measure of spatially varying noise. PSNR also does not distinguish between amplified noise and aliasing errors. We evaluated noise amplification by forming Monte-Carlo estimates of the g-factors, the spatially varying quantity of noise amplification beyond the factor of $\sqrt{R}$ from undersampling, using the pseudo multiple replica method [28]. Analytical g-factors can also be computed [29], although such an analysis ignores the effects of including ACS data in the reconstruction. To quantify aliasing, we compared the sample autocorrelations of the complex difference images. Since spatially-varying noise is not strongly correlated, it will be spread throughout the autocorrelation, while coherent aliasing

should be strongly correlated, causing peaks at offsets of $\pm X/R_x$ and $\pm Y/R_y$ and multiples thereof.

As mentioned in the introduction, we assume in developing our method that a GRAPPA-reconstructed image inherits the transform-domain sparsity of the object, even after under-sampling, interpolation, and the resulting noise amplification. To test this hypothesis, we applied the sparsity-promoting GRAPPA calibration with a 36×36 ACS block to the 4×4-undersampled real data set. After applying the resulting kernel to the undersampled data, with and without including the ACS in the reconstruction, we transformed the noisy reconstructed images into the DWT sparse transform domain. We compared the coil-combined magnitudes of the transform coefficients against the same magnitudes for the fully-sampled reference data by plotting their empirical cumulative distribution functions (CDFs) on a log scale.

The Tikhonov, truncated SVD, and proposed sparsity-promoting regularization methods all depend on carefully selecting the tuning parameter (denoted a, τ, and λ, respectively) optimal for kernel calibration. To avoid reliance on heuristic methods particular to a given regularizer, we chose the optimal tuning parameter for each method via two-stage coarse-then-fine parameter sweeps. We then portrayed the effects of varying the tuning parameters on the visual quality of the reconstructed images. We expect that the optimal choice of tuning parameter varies with the image and the sampling pattern used, including the acceleration factor and number of ACS lines. We performed this experiment on the simulated data set with an ACS block of reasonable size, where sweeping the parameters yielded visible variations in image quality.

For our comparisons of GRAPPA kernel calibration regularization techniques, we began with the phantom data. First, we varied the number of ACS lines used for calibration. We calculated the PSNR values for reconstructions using un-regularized, Tikhonov-regularized, truncated-SVD-based, and sparsity-promoting calibration. We also considered L_1-SPIRiT reconstructions for the same range of ACS lines, using a 7×7-point kernel (Tikhonov-regularized by default). Unfortunately, nonlinear GRAPPA was too computationally intensive to generate such a trade-off curve. The trade-off curves correspond to $R_y = 3$ nominal undersampling with the number of ACS lines varying from 10 to 30 full lines. We repeated this experiment on our real data set for the same methods, using 4×4 nominal undersampling and varying the number of ACS lines from 20×20 to 64×64. To find the best achievable image quality for any desired total acceleration, we performed this experiment on the real data for a couple different levels of nominal undersampling and constructed trade-off curves for each regularizer that represent the maximum PSNRs of all the trade-off curves for that technique.

Next, we performed visual comparisons for both simulated and real data sets. We compared all the methods for the simulated data with $R_y = 3$ 1-D nominal undersampling and 20 and 10 ACS lines. We also compared all the methods, minus nonlinear GRAPPA, for the real data with 4×4 2-D nominal undersampling and 36×36 and 24×24 ACS lines. We chose the larger ACS blocks to yield low-quality calibrations without any regularization, and the smaller blocks to produce underdetermined calibrations. Bydder's nonlinear GRAPPA [9] only supports 1-D undersampling of k-space at present, so we did not apply it to the 2-D undersampled real data.

We generated G-factor maps for the kernels calibrated from low-quality (but not underdetermined) sets of ACS lines in the previous experiments, except nonlinear GRAPPA. G-factors generated from the pseudo multiple replica method for each kernel confirm that all the regularization methods help defeat noise amplification. As is done in [29], we ignored

the effect of noise amplification on the calibration of the kernel. Our implementation of the pseudo multiple replica method repeated the GRAPPA reconstruction 400 times for each calibrated kernel, adding complex Gaussian noise with covariance Λ. We also included G-factor maps for L_1-SPIRiT for comparison.

Finally, we computed autocorrelation maps of the difference images for the underdetermined reconstructions to measure residual aliasing. Significant offset peaks signify noticeable residual aliasing in the reconstructed image. Since we expect the residual aliasing to vary with the tuning parameter, we computed curves for each choice across our parameter sweep and show the best curve for each method.

IV. Results

The curves in Fig. 3 demonstrate very similar behavior for the fully-sampled reference data and the GRAPPA reconstructed images. The empirical CDF curves are nearly identical when the ACS data was included in the reconstruction. A compressible signal would be expected to have sparse transform coefficients that decay exponentially, which coincides with $1 - CDF$ being linear on a log-log scale. As is apparent from the curves, the GRAPPA reconstruction including the 36×36 block of ACS data in the output has nearly identical sparsity to the original data, while the data with no ACS included in the output are similar, but with a slight bias towards larger-magnitude coefficients. One would expect a reconstruction calibrated with fewer ACS lines to yield a curve between these two, so the assumption appears validated, at least for this real data set. One must remain concerned about even greater noise amplification, but considering the noise level present in the real data at this acceleration, one probably would not be interested in even greater undersampling.

In Fig. 4, we display the reconstructed simulated images after varying the tuning parameters for Tikhonov, truncated SVD, and sparsity-promoting regularized calibration. For overdetermined calibration (20 ACS lines with $R_y = 3$ 1-D nominal acceleration), very low values of the tuning parameters for any of these three regularizers reduce that regularizer's effect on the calibration. All three methods also introduce significant residual aliasing for larger values of the parameters, suggesting that overemphasizing the regularization can undermine even overdetermined calibrations.

In the next series of simulations on the phantom data, we varied the number of ACS lines from 10 samples up to 30 samples, holding fixed the 1-D nominal undersampling at $R_y = 3$. After computing the GRAPPA reconstructions with un-regularized, Tikhonov-regularized, TSVD-based, and sparsity-promoting kernel calibrations, we plot their PSNRs versus the total acceleration corresponding to the nominal undersampling and number of ACS lines (see (8)). We added the L_1-SPIRiT reconstructions for the same range of ACS lines for reference. The PSNRs (not shown) for the GRAPPA calibration methods are nearly identical, trending from between 27-28 dB with 10 ACS lines up to 36 dB with 30 lines. L_1-SPIRiT fares slightly better, increasing in the same pattern from 30 dB to nearly 37 dB over the same range.

In Fig. 5, a similar series of simulations were run on the real data, varying the number of ACS lines from 20×20 samples up to 64×64 samples, and holding fixed the nominal undersampling at $R_x = R_y = 4$. As the number of ACS lines increased, the PSNR also increased, as we observe for the simulated data. The difference between un-regularized and regularized methods is more significant here, although PSNR improvement still does not compare to the level achieved by L_1-SPIRiT. As the number of ACS lines decreases, the PSNR decreases at a far more diminished rate for the regularized calibration techniques.

Because un-regularized calibration cannot be performed when the number of ACS lines yields too few fit equations, un-regularized calibration cannot be performed for total accelerations beyond $R = 10.5$ in this plot. At the highest accelerations, sparsity-promoting regularization produces images with better PSNR than Tikhonov regularization. The trade-off shift due to sparsity-promoting regularization suggests that the PSNR for the real data is no worse than for other regularized GRAPPA calibration methods and is significantly improved over un-regularized GRAPPA kernel calibration.

The plot in Fig. 5 only considers the performance achievable with the nominal undersampling held fixed. In Fig. 6, we computed trend lines for different levels of nominal undersampling and generated curves by taking the greatest PSNR for each level of total acceleration, merging trends for 4×3 and 4×4 nominal undersampling to produce a staircase effect. Lower nominal undersampling with fewer ACS lines tends to produce higher PSNR.

The previous experiments establish trends for different calibration and reconstruction methods using PSNR. However, as PSNR is not sensitive to the type of coherent aliasing artifacts we expect with extreme undersampling, we should inspect the reconstructed images visually. Visual comparisons allow us to readily evaluate the significance of coherent aliasing artifacts and noise amplification on overall image quality. First, we explored regularizing overdetermined and underdetermined ACS fits on simulated data. We used a 4×3-block GRAPPA kernel with $R_y = 3$ nominal undersampling and 20 and 10 ACS lines for the overdetermined and underdetermined fits, respectively. We also used L_1-SPIRiT with a 7×7-point kernel for comparison. Reconstructions and difference images using Tikhonov-regularized, truncated-SVD, nonlinear, and sparsity-promoting GRAPPA are shown, along with L_1-SPIRiT, for both overdetermined and underdetermined cases in Fig. 7. In the overdetermined case, the images appear very similar. In the underdetermined case, none of the methods completely remove the aliasing, but the nonlinear GRAPPA method does slightly better than the rest at mitigating the aliasing within the simulated brain.

The visual comparisons using real data more effectively demonstrate the advantages of sparsity-promoting calibration on image quality. Using a 4×4-block GRAPPA kernel and 5×5-point SPIRiT kernel with $R_y = R_z = 4$ 2-D nominal undersampling, the ACS block measured 36×36 samples in the overdetermined case and 24×24 samples in the underdetermined case. All the methods shown for the simulated data with the exception of nonlinear GRAPPA, which does not support 2-D acceleration, were repeated for the real data set. The reconstructed images and difference images shown in Figure 8 show effective denoising and diminished aliasing for the sparsity-promoting method. In particular, the aliased structure is significant in the other methods in the underdetermined case. The L_1-SPIRiT method is capable of even greater denoising, most likely accounting for the PSNR difference, but the result displays aliasing artifacts in both overdetermined and underdetermined cases. We would expect this aliasing to be mitigated if L_1-SPIRiT were applied to randomly undersampled k-space, rather than the uniform Cartesian undersampling employed here.

Next, we quantified the improvement in noise amplification for these calibration methods using simulated g-factors. Spatial g-factor maps for the simulated and real combined images are shown in Fig. 9 and Fig. 10 for the overdetermined calibrated kernels used in Fig. 7 and Fig. 8. We masked the g-factor maps to ignore pixels outside the brain, as what happens to noise outside the brain is less important. We report the mean and maximum observed g-factors using this mask. We observe little difference between Tikhonov-regularized, truncated SVD-based, and sparsity-promoting GRAPPA calibration methods for either data set, so our method does not greatly sacrifice noise reduction by improving aliasing

suppression. L_1-SPIRiT provides improved g-factors over all the GRAPPA-based methods, at the expense of the residual aliasing observed in the reconstructions of uniformly undersampled real data.

Lastly, we measured the residual aliasing in the reconstructed images from each regularized GRAPPA kernel used in Fig. 7 (simulated brain) and Fig. 8 (real data) for the underdetermined calibration regime. We depict one side of the center slice in the y-direction of the two-dimensional normalized autocorrelation map for the 1-D undersampled simulated data, and we show center slices in the x- and y-directions for the 2-D undersampled real data. Since only relative magnitudes are important, we normalized the autocorrelation curves so that the largest value across all the curves is **1.0**. The autocorrelation magnitudes in Fig. 11 and Fig. 12 portray significant residual aliasing in the reconstruction using Tikhonov-regularized or truncated SVD-based kernels, while the nonlinear GRAPPA kernel has a slightly lower peak for the simulated data, and the sparsity-promoting kernel has much less of a peak for the real data. While tuning the regularization parameters may be thought to reduce this aliasing, these are the minimum peaks observed varying a or τ. Together with the visual comparisons, the reduced autocorrelation peaks justify using sparsity-promoting regularization to calibrate GRAPPA, even though the acquired data are uniformly spaced. While coherent aliasing is visible in the image reconstructed using Li-SPIRiT, the autocorrelation plot for Li-SPIRiT does not appear to have coherent peaks in the same places.

V. Discussion

The success of GRAPPA accelerated parallel imaging depends on the quality of the calibration. At high levels of acceleration, acquiring sufficiently many ACS lines to reconstruct a high quality image may be impractical, and regularization can help improve image quality without sacrificing for more ACS lines. The sparsity-promoting regularization method described in this paper produces a GRAPPA kernel that mitigates both noise amplification and residual coherent aliasing in the reconstructed image, even when the number of ACS lines would provide only enough fit equations to form an underdetermined least-squares problem without regularization. The proposed method enables calibration using this under-determined system of least-squares fit equations because the regularized calibration takes advantage of the joint transform-domain sparsity expected of the images reconstructed using the calibrated kernels. The plots of PSNR/total acceleration trade-off curves for the different calibration methods demonstrate the significant impact of regularized calibration on the performance of accelerated parallel imaging using GRAPPA, especially when calibrating using the proposed method. The effects of regularization on noise amplification are confirmed using simulated g-factor maps, and the residual aliasing is measured using the difference image autocorrelation maps. Altogether, these experiments describe the benefits of a novel approach to calibrating the GRAPPA kernels for superior image reconstruction quality using fewer ACS lines. By successfully reducing the number of ACS lines required for a high-quality reconstruction, this method will enable more effective calibration in studies that can only accommodate the acquisition of limited ACS data.

Compared to incorporating a sparsity-promoting regularizer into the reconstruction of another parallel imaging method like SENSE or SPIRiT, promoting sparsity at the GRAPPA calibration step may appear suboptimal. Indeed, methods like L_1-SPIRiT should have an advantage in denoising the result due to being able to modify the un-acquired data more freely. Regularizing the kernel calibration constrains the sparse reconstructed coil images to reside in a subspace spanned by the acquired (aliased) image data from all the coils, which is more restrictive than finding the sparsest of any image consistent with the acquired data. If we assume that a GRAPPA kernel does exist that would yield the full-FOV k-space from

these acquired data, this subspace should contain the sparse reconstructed coil images of interest. If the acquired data are particularly noisy, or not enough coils are used, this assumption may be violated, so this calibration method, like GRAPPA itself, does have limitations. We studied the sparsity-preserving assumption for GRAPPA applied to real data, and our results appear to validate this assumption for at least the level of undersampling addressed in this paper. To overcome this suboptimality, either more coils should be used, or the reconstructed k-space should be denoised using a method like DESIGN denoising [30].

Other limitations of this sparsity-promoting calibration method remain. First, we have no analytical formula for the optimal choice of λ. In this work, a coarse-then-fine parameter sweep was used to identify a useful value, but such sweeps can be time-consuming. Fortunately, iterative reweighting methods like the half-quadratic minimization algorithm used to implement the optimization are compatible with a warm-start approach; we can slowly scale up the value of λ, and successive optimization problems take less time to solve. We aim to investigate automatic parameter selection methods in future work. Another challenge is the efficient or real-time implementation of this method. Typical usage requires a few hundred inner iterations of the LSMR least-squares solver to converge to the final solution, and each iteration can take several seconds because of the need to perform a GRAPPA reconstruction in each step. Since the computationally intensive operations are highly parallelizable for implementation using GPUs, combining parallel computing with alternative implementations of the cost function in Eq. (3) such as variable-splitting-based alternating minimization methods could accelerate the method significantly, for practical use.

As discussed earlier, the efficient implementation of this method is limited to uniform Cartesian sampling patterns by the existence of efficient GRAPPA reconstruction methods and the number of kernels that need to be calibrated. More development is necessary to generalize this sparsity-promoting calibration method to nonuniform or non-Cartesian sampling patterns or indirect reconstruction methods like SPIRiT in a practical manner.

Acknowledgments

References

1. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magn Reson Med. Jun.2002 47, no. 6:1202–10. [PubMed: 12111967]

2. Ding, Y.; Xue, H.; Chang, TC.; Guetter, C.; Simonetti, O. A quantitative study of Sodickson's paradox. Proc ISMRM 20th Scientific Meeting; May 2012; p. 3352

3. Nana R, Zhao T, Heberlein K, LaConte SM, Hu X. Cross-validation-based kernel support selection for improved GRAPPA reconstruction. Magn Reson Med. Apr.2008 59, no. 4:819–25. [PubMed: 18383295]

4. Samsonov AA. On optimality of parallel MRI reconstruction in k-space. Magn Reson Med. Jan. 2008 59, no. 1:156–64. [PubMed: 18058935]

5. Bauer S, Markl M, Honal M, Jung BA. The effect of reconstruction and acquisition parameters for GRAPPA-based parallel imaging on the image quality. Magn Reson Med. Aug.2011 66, no. 2:402–9. [PubMed: 21381100]

6. Tikhonov, AN.; Arsenin, VIA. Solutions of Ill-Posed Problems. New York; Winston; distributed solely by Halsted Press: 1977.

7. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: Sensitivity encoding for fast MRI. Magn Reson Med. Nov.1999 42, no. 5:952–62. [PubMed: 10542355]

8. Lin FH, Kwong KK, Belliveau JW, Wald LL. Parallel imaging reconstruction using automatic regularization. Magn Reson Med. Mar.2004 51, no. 3:559–67. [PubMed: 15004798]

9. Bydder M, Jung Y. A nonlinear regularization strategy for GRAPPA calibration. Magn Reson Imaging. Jan.2009 27, no. 1:137–41. [PubMed: 18583080]

10. Griswold MA, Blaimer M, Breuer F, Heidemann RM, Mueller M, Jakob PM. Parallel magnetic resonance imaging using the GRAPPA operator formalism. Magn Reson Med. Dec.2005 54, no. 6:1553–6. [PubMed: 16254956]

11. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magn Reson Med. Dec.2007 58, no. 6:1182–95. [PubMed: 17969013]

12. Liang D, Liu B, Wang J, Ying L. Accelerating SENSE using compressed sensing. Magn Reson Med. Dec.2009 62, no. 6:1574–84. [PubMed: 19785017]

13. Otazo R, Kim D, Axel L, Sodickson DK. Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI. Magn Reson Med. 2010; 64, no. 3:767–76. [PubMed: 20535813]

14. Lustig M, Pauly JM. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. Magn Reson Med. Aug.2010 64, no. 2:457–71. [PubMed: 20665790]

15. Murphy M, Alley M, Demmel J, Keutzer K, Vasanawala S, Lustig M. Fast ℓ_1-SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime. Jun.2012 31, no. 6:1250–62.

16. Chang Y, Liang D, Ying L. Nonlinear GRAPPA: A kernel approach to parallel MRI reconstruction. Magn Reson Med. Sep.2012 68:730–40. [PubMed: 22161975]

17. Weller DS, Polimeni JR, Grady L, Wald LL, Adalsteinsson E, Goyal VK. Regularizing GRAPPA using simultaneous sparsity to recover de-noised images. Proc SPIE Wavelets and Sparsity XIV. Aug.2011 8138:81381M–1–9.

18. Weller, DS.; Polimeni, JR.; Grady, L.; Wald, LL.; Adalsteinsson, E.; Goyal, VK. Greater acceleration through sparsity-promoting GRAPPA kernel calibration. Proc. ISMRM 20th Scientific Meeting; May 2012; p. 3354

19. Brau AC, Beatty PJ, Skare S, Bammer R. Comparison of reconstruction accuracy and efficiency among autocalibrating data-driven parallel imaging methods. Magn Reson Med. Feb.2008 59, no. 2:382–95. [PubMed: 18228603]

20. Rangan S. Generalized approximate message passing for estimation with random linear mixing. Proc IEEE Int Symp Inform Theory. Aug.2011 :2168–72.

21. Geman D, Reynolds G. Constrained restoration and the recovery of discontinuities. Mar.1992 14, no. 3:367–83.

22. Geman D, Yang C. Nonlinear image recovery with half-quadratic regularization. Jul.1995 4, no. 7:932–46.

23. Fong DCL, Saunders MA. LSMR: An iterative algorithm for sparse least-squares problems. SIAM J Sci Comput. 2011; 33, no. 5:2950–71.

24. Lustig, M.; Pauly, JM. Iterative GRAPPA: a general solution for the GRAPPA reconstruction from arbitrary k-space sampling. Proc. ISMRM 15th Scientific Meeting; May 2007; p. 333

25. Lai, P.; Zhang, T.; Lustig, M.; Vasanawala, SS.; Brau, AC. Improving compressed sensing parallel imaging using autocalibrating parallel imaging initialization with variable density tiled random k-space sampling; Proc. ISMRM 19th Scientific Meeting; May 2011; p. 68

26. Kwan RKS, Evans AC, Pike GB. MRI simulation-based evaluation of image-processing and classification methods. Nov.1999 18, no. 11:1085–97.

27. Roemer PB, Edelstein WA, Hayes CE, Souza SP, Mueller OM. The NMR phased array. Magn Reson Med. Nov.1990 16, no. 2:192–225. [PubMed: 2266841]

28. Robson PM, Grant AK, Madhuranthakam AJ, Lattanzi R, Sodickson DK, McKenzie CA. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-

based parallel imaging reconstructions. Magn Reson Med. Oct.2008 60, no. 4:895–907. [PubMed: 18816810]

29. Breuer FA, Kannengiesser SAR, Blaimer M, Seiberlich N, Jakob PM, Griswold MA. General formulation for quantitative G-factor calculation in GRAPPA reconstructions. Magn Reson Med. Sep.2009 62, no. 3:739–46. [PubMed: 19585608]

30. Weller DS, Polimeni JR, Grady L, Wald LL, Adalsteinsson E, Goyal VK. Denoising sparse images from GRAPPA using the nullspace method. Magn Reson Med. Oct.2012 68, no. 4:1176–89. [PubMed: 22213069]
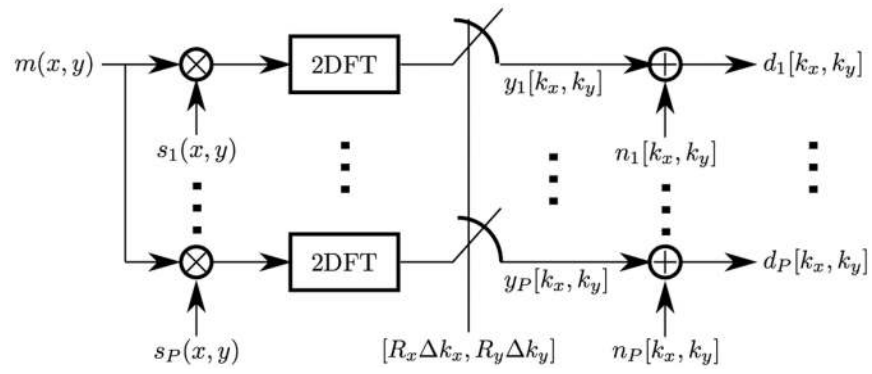
Fig. 1.
Abstract model relating the original image $m(x,y)$ to the acquired data $d_p[k_x, k_y]$ for $p = 1,$ …, P.
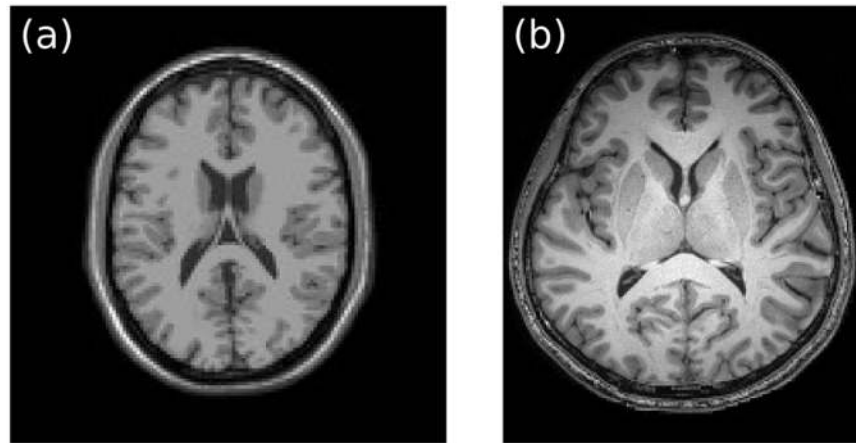
Fig. 2.
Magnitude combined image of reference T_1-weighted (a) simulated BrainWeb phantom and (b) real data.
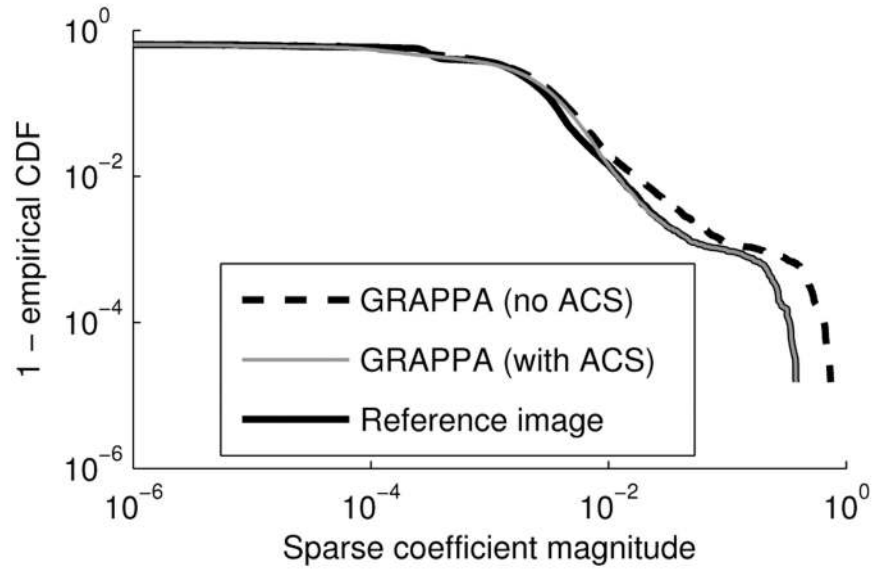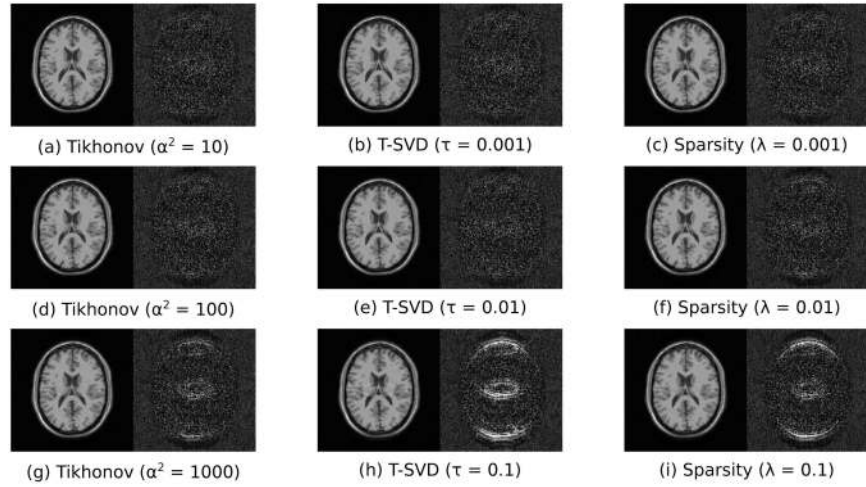
Fig. 3.
The curves representing 1— CDF for the fully-sampled reference and GRAPPA reconstructions' sparse transform coefficient magnitudes are plotted on a log-log scale. At worst, such a curve representing sparse coefficients would be linear on this scale. The curve for GRAPPA with ACS data included in the reconstruction nearly matches the fully-sampled data curve exactly, while the reconstruction without calibration data appears to favor sparse coefficients slightly larger in magnitude.

(a) Tikhonov ($\alpha^2 = 10$) (b) T-SVD ($\tau = 0.001$) (c) Sparsity ($\lambda = 0.001$)

(d) Tikhonov ($\alpha^2 = 100$) (e) T-SVD ($\tau = 0.01$) (f) Sparsity ($\lambda = 0.01$)

(g) Tikhonov ($\alpha^2 = 1000$) (h) T-SVD ($\tau = 0.1$) (i) Sparsity ($\lambda = 0.1$)

Fig. 4.
For simulated data, reconstructed and difference images for GRAPPA reconstructions with Tikhonov-regularized, T-SVD-regularized, and sparsity-promoting calibration are portrayed using a range of tuning parameter choices for the overdetermined calibration case (4×3-block GRAPPA kernel from 20 ACS lines with $R_y = 3$ 1-D nominal undersampling). The middle tuning parameters (d-f) have the best PSNR; the reconstructed images for smaller values of α, τ, or λ have slightly more noise, while reconstructed images for larger tuning parameters have residual aliasing visible in the difference images.
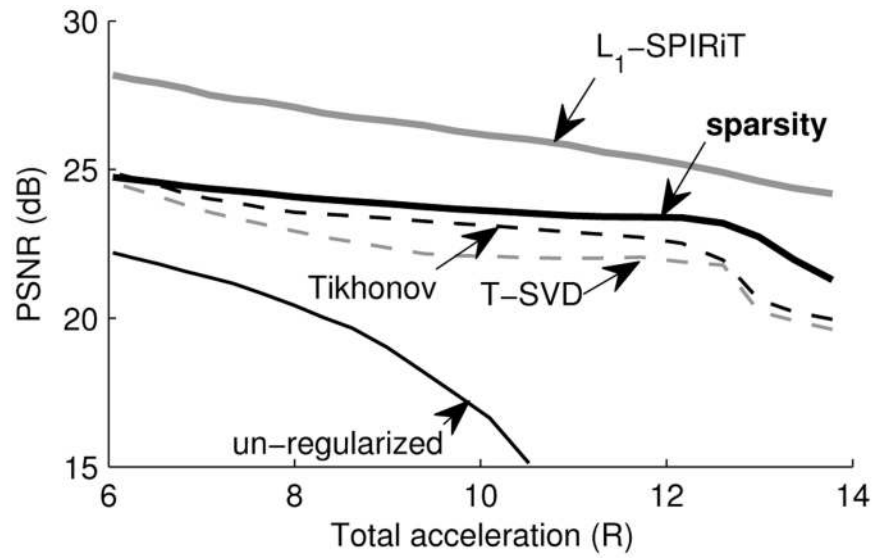
Fig. 5.
Regularized calibration improves the image quality (PSNR) of GRAPPA reconstructions of the real data over a broad range of total accelerations (corresponding to varying the number of ACS lines), with sparsity-promoting calibration yielding the best performance at the highest accelerations among the GRAPPA-based methods. However, the L_1-SPIRiT reconstruction has even higher PSNR. The nominal undersampling was held fixed ($R_x = R_y$ = 4 for the real data).
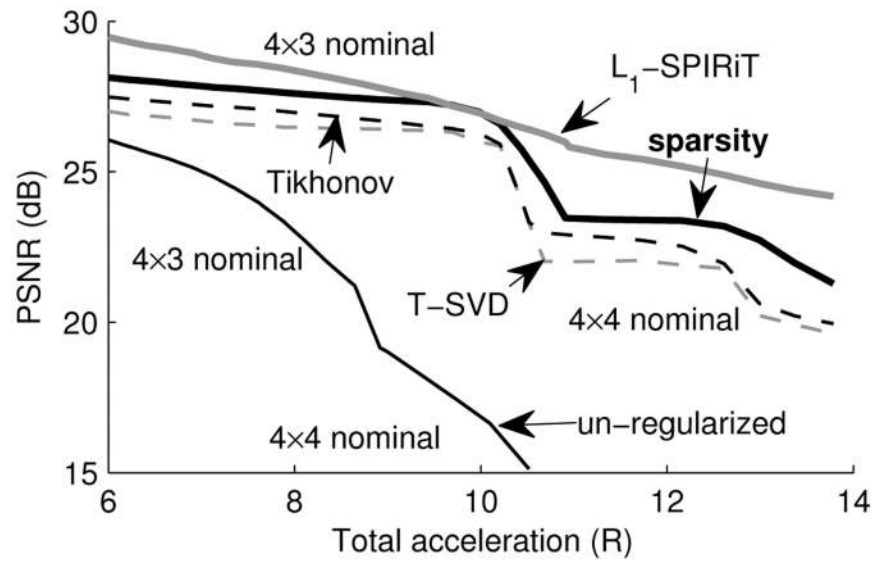
Fig. 6.
For real data, as the desired total acceleration increases (by reducing the number of ACS lines), the best PSNR is achieved using sparsity-promoting calibration with the lowest possible nominal undersampling. The greatest benefit of sparsity-promoting calibration appears at effective accelerations R between 12 and 14 with 4×4 nominal undersampling, although L_1-SPIRiT remains the clear winner in terms of overall PSNR.
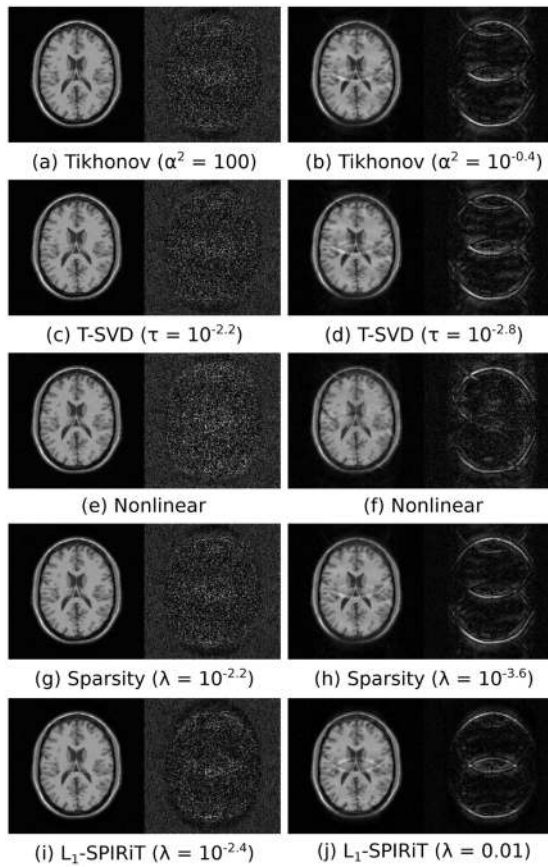
Fig. 7.
For the simulated data, all the regularized 4×3-block GRAPPA kernels for $R_y = 3$ 1-D nominal undersampling yield similar image quality, both in terms of noise reduction in the overdetermined case, and mitigating aliasing in the underdetermined case. The results for L_1-SPIRiT appear similar to the GRAPPA reconstructions, with less aliasing outside the brain.
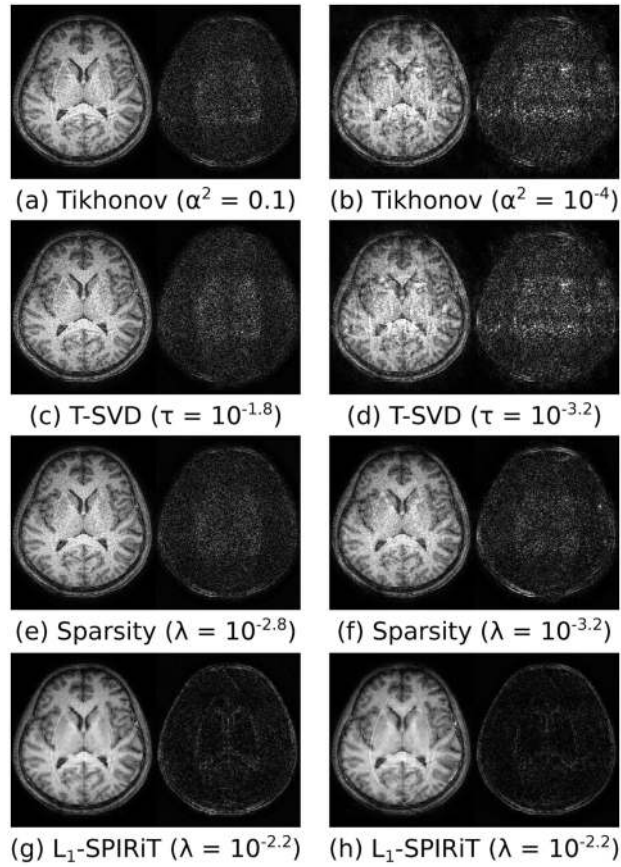
(a) Tikhonov ($\alpha^2 = 0.1$)   (b) Tikhonov ($\alpha^2 = 10^{-4}$)

(c) T-SVD ($\tau = 10^{-1.8}$)   (d) T-SVD ($\tau = 10^{-3.2}$)

(e) Sparsity ($\lambda = 10^{-2.8}$)   (f) Sparsity ($\lambda = 10^{-3.2}$)

(g) L_1-SPIRiT ($\lambda = 10^{-2.2}$)   (h) L_1-SPIRiT ($\lambda = 10^{-2.2}$)

Fig. 8.
For the real data, all the sparsity-promoting 4×4 GRAPPA kernels for $R_y = R_z = 4$ 2-D nominal undersampling yield significantly improved image quality, in terms of mitigating aliasing in the underdetermined case. In the overdetermined case, all the methods are effective at denoising. In both cases, the L_1-SPIRiT method is even more effective at denoising, but the coherent aliasing that also results may be undesirable.
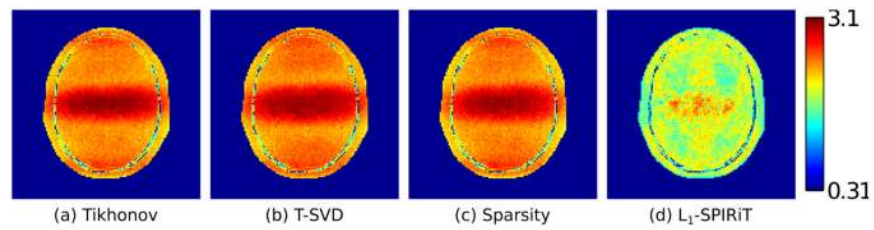
Fig. 9.
For simulated data, with 20 ACS lines, the g-factors for the 4×3-block GRAPPA kernel calibrated for $R_y = 3$ 1-D nominal undersampling of the simulated brain with (a) Tikhonov, (b) truncated SVD, and (c) sparsity-promoting regularization have minor noise amplification (mean = 1.8/1.9/1.8, respectively; max. = 3.4/3.3/3.3, respectively), while (d) L_1-SPIRiT greatly reduces noise amplification (mean = 1.1; max. = 2.2). The colors are logarithmically spaced to facilitate comparison.
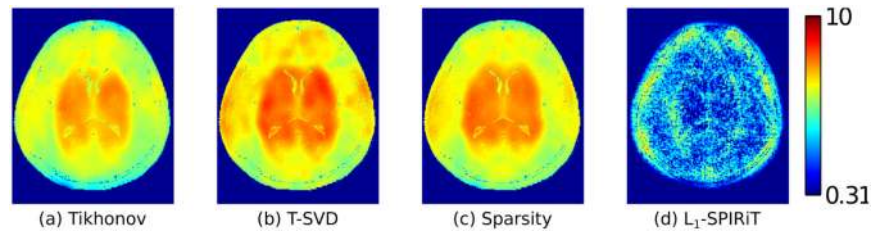
(a) Tikhonov   (b) T-SVD   (c) Sparsity   (d) L_1-SPIRiT

Fig. 10.
For real data, with 36×36 ACS lines, the g-factors for the 4×4 GRAPPA kernel calibrated for 4×4 nominal undersampling with (a) Tikhonov, (b) truncated SVD, and (c) sparsity-promoting regularization reduce noise amplification (mean = 2.5/3.1/2.8, respectively; max. = 4.5/5.9/5.0, respectively), while (d) L_1-SPIRiT greatly reduces noise amplification (mean = 0.86; max = 4.3). The colors are logarithmically spaced to facilitate comparison.
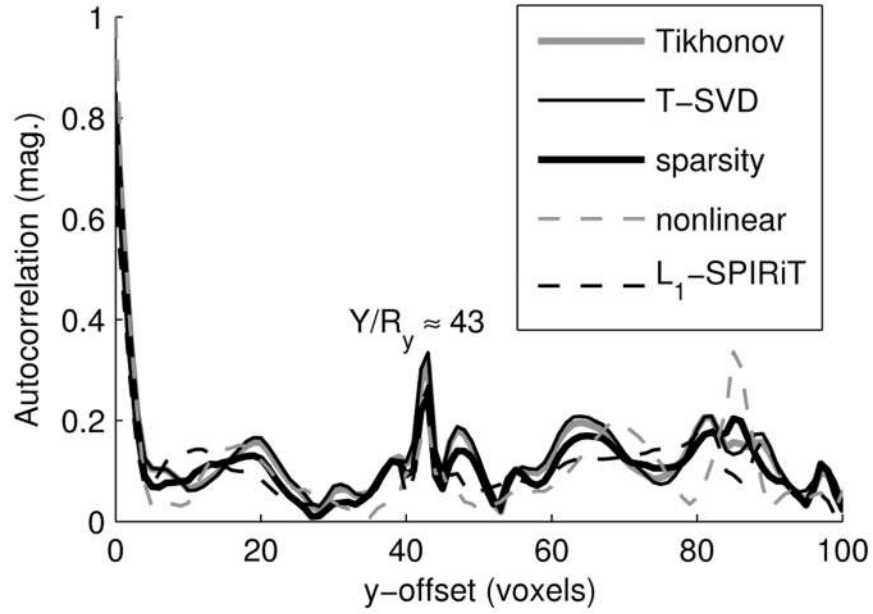

Fig. 11.
For simulated data, with 10 ACS lines, the Tikhonov-regularized and truncated SVD-based 4×3 GRAPPA kernels for $R_y = 3$ 1-D nominal undersampling yield images with noticeable aliasing artifacts that cause a significant peak at $Y/R_y \approx 43$ in the autocorrelation map in the y-direction, while L_1-SPIRiT and sparsity-promoting regularization have smaller peaks. Nonlinear GRAPPA also has a smaller peak, but another larger peak appears around $2Y/R_y$. The L_1-SPIRiT peak is just slightly larger than the sparsity-promoting peak in this example.
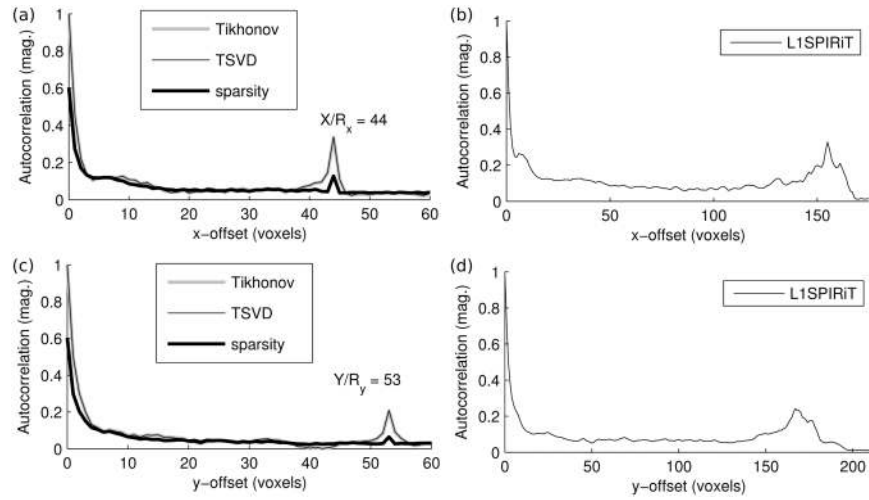
Fig. 12.
For real data, with 24 × 24 ACS lines, the Tikhonov-regularized and truncated SVD-based 4× 4 GRAPPA kernels for 4 × 4 nominal undersampling yield images with noticeable aliasing artifacts that cause a significant peak at $X/R_x = 44$ in the autocorrelation map in the (a) x-direction and at $Y/R_y = 53$ in the autocorrelation map in the (c) y-direction, while sparsity-promoting regularization has much smaller peaks. The Tikhonov and TSVD curves are on top of each other in the above figures. The L_1-SPIRiT curves in the (b) x-direction and (d) y-direction demonstrate different behavior, with additional autocorrelation energy clustered near the far edges of the curves.