

A Variational Approach to MR Bias Correction

Ayres Fan

Abstract

We propose a novel bias correction method for magnetic resonance (MR) imaging that uses complementary body coil and surface coil images. The former are spatially homogeneous but have low signal intensity; the latter provide excellent signal response but have large bias fields. We present a variational framework where we optimize an energy functional to estimate the bias field and the underlying image using both observed images. The framework is applicable to both 2D and 3D images. The energy functional contains smoothness-enforcing regularization for both the image and the bias field. We present extensions of our basic framework to a variety of imaging protocols. We solve the optimization problem using a computationally efficient numerical algorithm based on coordinate descent, preconditioned conjugate gradient, half-quadratic regularization, and multigrid techniques. We show qualitative and quantitative results demonstrating the effectiveness of the proposed method in producing debiased and denoised MR images.

I. INTRODUCTION

IN magnetic resonance (MR) image acquisition, there is a fundamental trade-off between noise and spatially-homogeneous signal response. An uncorrupted image (which we refer to as the true image or the intrinsic image) would depend solely on the underlying tissue properties (ρ , T_1 , T_2) and the imaging parameters (T_E , T_R). Intensity inhomogeneities can occur for a number of reasons, but we will only focus on one of the largest artifacts: the inhomogeneity caused by the spatially-dependent response of the receiving coil. Receiving with a body coil results in low signal-to-noise ratio (SNR) but good spatial homogeneity. Surface coils have strong signal response near the coil, but the intensity rapidly diminishes with distance [1]. This variable response allows better visualization of the region of interest (ROI) but results in a systematic intensity inhomogeneity known as the *bias field*. Surface coils are widely used in practice, and most users find the superior SNR outweighs the negative effects of the bias field. The intensity distortions caused by the bias field can significantly impair both visual inspection and image processing tasks. Separating the bias field from the true underlying image is an underconstrained and ill-posed problem—there are half the number of observations as there are free variables.

In order to build tractable bias correction techniques, simplifying assumptions of the true physical process are imposed. Most existing bias correction methods assume that the bias field is multiplicative, slowly varying, and tissue independent. Additionally, many techniques ignore the noise and apply a log transform to make the bias field additive. The earliest bias correction techniques relied on phantoms [2] or homomorphic unsharp filtering [3], [4], but both methods have severe limitations. For the former,

This work was supported by ONR N00014-00-10089, NIH 5 P41 RR13218-01, NIH 5 P01 CA67165-03, NIH 5 R01 AG19513-03, and a NDSEG graduate fellowship. Work was performed jointly with William M. Wells III, John W. Fisher III, Müjdat Çetin, Steven Haker, Robert Mulkern, Clare Tempany, and Alan S. Willsky

it is difficult to register the phantom image with subsequent patient scans (especially in multiple coil imaging setups). For the latter, homomorphic filtering can ameliorate the worst effects of the bias field, but generally is not very accurate.

Sled *et al.* [5] sharpen the histogram of the observed image using deconvolution and use the resulting *a priori* density to do Bayes least-squares estimation of the true image. Many recent techniques simplify the problem by using parametric representations for the bias field. Dawant *et al.* [6] fit thin-plate splines to the bias field using a least-squares penalty. Likar *et al.* [7] compute a parameterized bias field estimate that minimizes the entropy of the reconstructed image. Some methods couple segmentation and bias correction. Wells *et al.* [8] use the expectation-maximization (EM) algorithm [9] to estimate the bias field on the expectation step and perform a statistical segmentation of the brain on the maximization step. Many people have improved on this framework including Guillemaud and Brady [10] who adopt a better model for the tissue intensities; and Zhang *et al.* [11] who use a Markov random field to model the bias field.

Some people acquire extra information during the imaging process to aid in removing the bias field. Singh and NessAiver [12] use a tube filled with oil to mark the coil location and compute the sensitivity profile using the Biot-Savart Law. Moyher *et al.* [13] employ a similar technique. A few methods capture a body coil image to help correct the surface coil image. Brey and Narayana [14] low-pass filter the two observation images to minimize the noise and estimate the bias field as the ratio of the two images. Lai and Fang [15] divide the surface coil image by the body coil image and select a sparse set of reliable control points. They then estimate the bias field by fitting a membrane model to the control points. Pruessmann *et al.* [16] take a similar approach by fitting local polynomials at every point in the image.

Our method is related to the imaging framework proposed by Brey and Narayana. While there is additional time and expense associated with acquiring the body coil scan, very reliable bias correction results may be obtained¹. Our algorithm exploits the homogeneity of the body coil and the high SNR of the surface coil to create a composite image that has higher SNR than either observation image and a minimal bias field. We construct a general variational framework which can be adapted to a number of different imaging setups. We introduce a computationally efficient approach to solve the variational problem, and we demonstrate our algorithm on a variety of MR imaging applications.

In Sec. II, we describe our imaging model and construct an energy functional which we wish to minimize in order to produce estimates of the true image and the bias field. We generalize the energy functional to handle the cases with multiple surface coils and multiple pulse sequences. In Sec. III, we detail our algorithm to compute the solution to our variational problem. In Sec. IV we apply our algorithm to synthetic brain images and real prostate, heart, and brain images. We summarize our results in Sec. V.

¹Note that we acquire the body coil and surface coil images sequentially. It is technically feasible to receive from both coils simultaneously, but the mutual induction between the coils would cause the bias field from the surface coil to enter into the body coil image.

II. PROBLEM FORMULATION

A. Observation Model

We formulate our observation model in a discrete manner. The body coil and surface coil observation image pixels are placed into column vectors \mathbf{y}_B and \mathbf{y}_S respectively and are described by the following pair of equations:

$$\mathbf{y}_B = \mathbf{f}^* + \mathbf{n}_B \quad (1)$$

$$\mathbf{y}_S = \mathbf{b}^* \circ \mathbf{f}^* + \mathbf{n}_S . \quad (2)$$

In the above equation, \circ represents the Hadamard product [17] (or Schür product or entrywise product). We force both observations to have the same intrinsic image \mathbf{f}^* , and we assume that the surface coil has a bias field \mathbf{b}^* and the body coil has a constant gain field. Without loss of generality, we set the gain of the body coil to be unity². Each element of the noise vectors \mathbf{n}_B and \mathbf{n}_S is assumed to be independent and identically distributed (IID). This is justified by the thermal nature of the noise. Thus the surface coil image has a higher SNR in the ROI due to the fact that \mathbf{b}^* tends to be significantly larger than 1 there. An implicit assumption in (1) and (2) is that \mathbf{y}_B and \mathbf{y}_S are properly registered. This will generally be true if the images are captured in rapid succession.

We introduce two diagonal matrices \mathbf{B}^* and \mathbf{F}^* which have \mathbf{b}^* and \mathbf{f}^* respectively as their diagonal entries. We can then rewrite (2) as

$$\mathbf{y}_S = \mathbf{B}^* \mathbf{f}^* + \mathbf{n}_S = \mathbf{F}^* \mathbf{b}^* + \mathbf{n}_S . \quad (3)$$

The noise in magnitude MR images is accurately modeled by a Rician distribution [18]. Rician random variables are generated by taking the norm of a Gaussian random vector with arbitrary mean [19]. As the SNR increases, the Rician probability density function (PDF) approaches the Gaussian PDF. The Rician PDF is unwieldy to work with, so we treat the noise as Gaussian and zero-mean in our algorithm. Rician noise has a positive mean, so this assumption results in a biased estimator. In most applications, the SNR in tissue regions is high enough so that our Gaussian noise assumption is reasonable, and only a moderate upward bias is imparted.

B. Variational Formulation

In this section, we formulate a variational problem with a statistical interpretation which results in an energy functional that we seek to minimize. Unlike some other bias correction schemes, we do not take the log transform of our observations, but instead pose our energy functional directly in the original

²We can only specify \mathbf{f}^* and \mathbf{b}^* up to a multiplicative constant. In most medical image processing applications, only relative intensity values are important— \mathbf{f}^* and $k\mathbf{f}^*$ are equivalent. Note that if the body coil image actually has a spatially-varying bias field \mathbf{b}_B rather than a constant gain, then our algorithm actually estimates $\mathbf{b}_B \circ \mathbf{f}^*$ rather than just \mathbf{f}^* .

multiplicative form. This leads to a cleaner formalism but imposes the need to do nonlinear estimation. We define our energy functional as

$$E(\mathbf{f}, \mathbf{b}) = \lambda_B \|\mathbf{y}_B - \mathbf{f}\|^2 + \lambda_S \|\mathbf{y}_S - \mathbf{b} \circ \mathbf{f}\|^2 + \alpha \|\mathbf{L}\mathbf{b}\|^2 + \gamma \|\mathbf{D}\mathbf{f}\|_p^p \quad (4)$$

and choose our optimal estimates $\hat{\mathbf{f}}$ and $\hat{\mathbf{b}}$ as the vectors that minimize $E(\mathbf{f}, \mathbf{b})$:

$$\hat{\mathbf{f}}, \hat{\mathbf{b}} = \arg \min_{\mathbf{f}, \mathbf{b}} E(\mathbf{f}, \mathbf{b}) . \quad (5)$$

λ_B , λ_S , γ , and α are positive weights. $\|\cdot\|_p$ represents the ℓ_p norm, and $\|\cdot\|$ represents the ℓ_2 norm. We design \mathbf{L} and \mathbf{D} to approximate derivative operators (generally either gradient or Laplacian operators) as finite differences. More generally, \mathbf{L} and \mathbf{D} can implement arbitrary linear high-pass convolutional kernel. This ensures that high-frequency components of $\hat{\mathbf{f}}$ and $\hat{\mathbf{b}}$ are penalized. The ℓ_2 norms for our data fidelity terms (the first two terms) in (4) correspond to a Gaussian noise assumption if the problem is formulated as a maximum *a posteriori* (MAP) estimation problem. From this perspective, we see that the scalar weights λ_B and λ_S should be proportional to the inverse noise variances for each observation image.

We use Tikhonov-type regularization [20] to make our intrinsic image and bias field estimates conform to our prior knowledge of the signals. Specifically, we ensure that our bias field estimate is smooth and our intrinsic image estimate is piecewise constant. Without regularization on $\hat{\mathbf{f}}$ or $\hat{\mathbf{b}}$ (i.e., $\alpha = \gamma = 0$), the minimization is a well-posed problem (in the sense the solution is attainable and unique), but it produces a trivial result. In this case, the minimum of E is 0 and is achieved for $\hat{\mathbf{f}} = \mathbf{y}_B$ and $\hat{\mathbf{b}} = \mathbf{y}_S \oslash \mathbf{y}_B$ (where \oslash indicates element-by-element division). The ℓ_p regularization on $\hat{\mathbf{f}}$ is similar to putting an anisotropic edge-preserving filter into our method [21]. It is well known that ℓ_2 norms tend to overpenalize large derivative values associated with edges. Hence, using ℓ_2 regularization in image reconstruction tends to oversmooth edges, and ℓ_p norms with $p < 2$ are said to be edge preserving. The main advantage of fully integrating the denoising operation into our algorithm is that the amount of filtering applied varies depending on the SNR at each pixel. This will be discussed more fully in Sec. III-B.2.

C. Extension to Multiple Surface Coils

Multiple surface coil images can be simultaneously captured using carefully crafted coil arrays without requiring additional image acquisition time [22]. Multiple coils are used due to the typically sharp drop-off in sensitivity far away from surface coils. By distributing the coils spatially, better signal coverage can be achieved. One way to process multiple surface coil images is to combine them into one composite surface coil image using a method such as Roemer's sum-of-squares technique [22] and then use our formulation in (4). However, there are advantages to processing the surface coil measurements individually. We introduce a new measurement model where we receive one body coil image and K surface coil images:

$$\mathbf{y}_B = \mathbf{f}^* + \mathbf{n}_B \quad (6)$$

$$\mathbf{y}_{S,k} = \mathbf{b}_k^* \circ \mathbf{f}^* + \mathbf{n}_{S,k}, \quad k \in \{1, 2, \dots, K\} . \quad (7)$$

We can extend (4) to handle this more general case:

$$E(\mathbf{f}, \mathbf{b}_1, \dots, \mathbf{b}_K) = \lambda_B \|\mathbf{y}_B - \mathbf{f}\|^2 + \sum_{k=1}^K \lambda_{S,k} \|\mathbf{y}_{S,k} - \mathbf{b}_k \circ \mathbf{f}\|^2 + \sum_{k=1}^K \alpha_k \|\mathbf{L}_k \mathbf{b}_k\|^2 + \gamma \|\mathbf{D}\mathbf{f}\|_p^p . \quad (8)$$

We obtain superior results minimizing (8) because we can optimally combine the surface coil observations by waiting until we have each $\hat{\mathbf{b}}_k$. The sum-of-squares method implicitly uses $\mathbf{y}_{S,k}$ as an estimate of $\hat{\mathbf{b}}_k$ when combining the images. Additionally, processing the surface coil images individually allows us to choose α_k and \mathbf{L}_k to tune the regularization for each coil. With a composite surface coil image, if some of the surface coils have very different reception profiles, a single choice of α and \mathbf{L} will oversmooth or undersmooth in some regions of the image.

D. Extension to Multiple Pulse Sequences

Multiple scans of the same location using different pulse sequences (*e.g.*, T_1 -weighted and T_2 -weighted) are commonly acquired. One possibility would be to acquire a body coil image and a surface coil image for each pulse sequence and perform each bias correction independently. But the bias fields in all of the surface coil images should be nearly identical, so we can achieve satisfactory results using only one body coil image. Our measurement model for this case again involves one body coil image and K surface coil images, but this time each surface coil image has the same bias field but different intrinsic images:

$$\mathbf{y}_B = \mathbf{f}_1^* + \mathbf{n}_B \quad (9)$$

$$\mathbf{y}_{S,k} = \mathbf{b}^* \circ \mathbf{f}_k^* + \mathbf{n}_{S,k}, \quad k \in \{1, 2, \dots, K\} . \quad (10)$$

Without loss of generality, we have assigned \mathbf{f}_1^* to correspond to the intrinsic image in the body coil image. We can again generalize (4) to handle this case:

$$E(\mathbf{f}_1, \dots, \mathbf{f}_K, \mathbf{b}) = \lambda_B \|\mathbf{y}_B - \mathbf{f}_1\|^2 + \sum_{k=1}^K \lambda_{S,k} \|\mathbf{y}_{S,k} - \mathbf{b} \circ \mathbf{f}_k\|^2 + \alpha \|\mathbf{L}\mathbf{b}\|^2 + \sum_{k=1}^K \gamma_k \|\mathbf{D}_k \mathbf{f}_k\|_p^p . \quad (11)$$

E. General Imaging Setups

More complex permutations beyond the two extensions we have presented can also be handled in a straightforward manner. Assume we have N_o observation images, N_f intrinsic images, and N_b surface coils:

$$\mathbf{y}_i = \mathbf{b}_{m_i}^* \circ \mathbf{f}_{n_i}^* + \mathbf{n}_i, \quad i \in \{1, 2, \dots, N_o\} . \quad (12)$$

For a body coil image, we let $m_i = 0$ and define $\mathbf{b}_0^* = \mathbf{1}$ (this is the bias field of the body coil). Thus $m_i \in \{0, 1, \dots, N_b\}$ and $n_i \in \{1, 2, \dots, N_f\}$. For any given set of observation images, the choice of m_i and n_i is not unique because we may place the observation images in any order. For notational simplicity, define sets $\mathcal{F} = \{\mathbf{f}_1^*, \mathbf{f}_2^*, \dots, \mathbf{f}_{N_f}^*\}$ and $\mathcal{B} = \{\mathbf{b}_1^*, \mathbf{b}_2^*, \dots, \mathbf{b}_{N_b}^*\}$. The energy functional then

consists of a ℓ_2 error term for each observation image, a ℓ_2 regularizing term for each surface coil, and a ℓ_p regularizing term for each intrinsic image:

$$E(\mathcal{F}, \mathcal{B}) = \sum_{i=1}^{N_o} \lambda_i \|\mathbf{y}_i - \mathbf{b}_{m_i} \circ \mathbf{f}_{n_i}\|^2 + \sum_{m=1}^{N_b} \alpha_m \|\mathbf{L}_m \mathbf{b}_m\|^2 + \sum_{n=1}^{N_f} \gamma_n \|\mathbf{D}_n \mathbf{f}_n\|_p^p. \quad (13)$$

For a discussion of the conditions necessary for a well-posed problem, see the Appendix.

III. SOLUTION OF THE OPTIMIZATION PROBLEM

This section details the solution to the optimization problem defined in Sec. II. We will begin by describing the solution to (4) and then extend the result for (13). A closed-form solution for any form of our optimization problem does not exist, and gradient descent on the full energy functional is slow and cumbersome. Therefore, we minimize our energy functional using coordinate descent. This is an iterative technique that minimizes the energy by alternating between \mathbf{f} and \mathbf{b} minimizations. This results in estimates $\hat{\mathbf{f}}^{(i)}$ and $\hat{\mathbf{b}}^{(i)}$ at each iteration i . Coordinate descent is useful in problems where computing solutions over all of the variables is difficult, but computing solutions over a subset is relatively easy. A stationary point obtained through coordinate descent is also a stationary point of the overall minimization problem. In order for coordinate descent to terminate, the derivative for each coordinate must be zero. Thus the gradient of the complete energy functional is zero.

At each iteration, we refer to the computation of $\hat{\mathbf{f}}^{(i)}$ and $\hat{\mathbf{b}}^{(i)}$ as a f-step and a b-step respectively. Note that for the more general case (e.g., (8), (11), or (13)), each f- or b-step may involve updating multiple $\hat{\mathbf{f}}_k^{(i)}$ or $\hat{\mathbf{b}}_k^{(i)}$. It is easy to show that within each f- or b-step, the updates of each $\hat{\mathbf{f}}_k^{(i)}$ or $\hat{\mathbf{b}}_k^{(i)}$ can be done independently. Given all of the $\hat{\mathbf{b}}_k^{(i)}$, the $\hat{\mathbf{f}}_k^{(i)}$ are conditionally independent of each other and vice versa.

A. Bias Field Solution

For a given \mathbf{f} , (4) is quadratic in terms of \mathbf{b} . Thus setting the gradient of E with respect to \mathbf{b} equal to zero results in a simple linear equation:

$$(\lambda_S \mathbf{F}^2 + \alpha \mathbf{L}^T \mathbf{L}) \hat{\mathbf{b}}^{(i)} = \lambda_S \mathbf{F} \mathbf{y}_S \quad (14)$$

Although we could solve (14) by direct matrix inversion, we note that $(\lambda_S \mathbf{F}^2 + \alpha \mathbf{L}_b^T \mathbf{L}_b) \geq 0$, so the subproblem is convex. Hence we can use an iterative algorithm such as preconditioned conjugate gradient [24] to efficiently compute approximate solutions. For unconstrained quadratic optimization, line minimizations can be easily computed in closed form, and conjugate gradient has a superlinear convergence rate. Additionally, because we use coordinate descent, finding the exact solution to (14) may actually overshoot the minimum of the global energy functional. For our conjugate gradient implementation, we use as a preconditioner the tridiagonal matrix composed of the main diagonal and the adjacent subdiagonals of $(\lambda_S \mathbf{F}^2 + \alpha \mathbf{L}_b^T \mathbf{L}_b)$. We chose this preconditioner because it is extremely fast to construct and apply.

We find that this preconditioner doubles our convergence rate while increasing our time per conjugate gradient iteration by 5%. This technique is effective because a large percentage of the non-zero entries are located in our preconditioning matrix.

For the general form in (13), we obtain the following linear equation for each $\hat{\mathbf{b}}_k$:

$$\left(\sum_{\substack{i \text{ s.t.} \\ m_i=k}} \lambda_i \mathbf{F}_{n_i}^2 + \alpha \mathbf{L}_k^T \mathbf{L}_k \right) \hat{\mathbf{b}}_k^{(i)} = \sum_{\substack{i \text{ s.t.} \\ m_i=k}} \lambda_i \mathbf{F}_{n_i} \mathbf{y}_i . \quad (15)$$

This can be solved in an identical manner as described above.

B. Intrinsic Image Solution

1) *No Regularization on $\hat{\mathbf{f}}$* : To provide some insight into the nature of the solution, we begin by examining the minimization of (4) for a given \mathbf{b} and $\gamma = 0$. We take the gradient of E with respect to \mathbf{f} and set it equal to zero to obtain a pointwise solution at each pixel index n :

$$\hat{\mathbf{f}}^{(i)}[n] = \frac{\lambda_B \mathbf{y}_B[n] + \lambda_S \mathbf{b}[n] \mathbf{y}_S[n]}{\lambda_B + \lambda_S \mathbf{b}^2[n]} \quad (16)$$

Because λ_B and λ_S are related to the inverse noise variances, $\hat{\mathbf{f}}^{(i)}[n]$ is the noise-weighted convex combination of $\mathbf{y}_B[n]$ and $\mathbf{y}_S[n]/\mathbf{b}[n]$. The weighting factor varies in space depending on the strength of $\mathbf{b}[n]$, so some locations in $\hat{\mathbf{f}}$ may take most of their value from \mathbf{y}_S while others may take most of their value from \mathbf{y}_B . Note that we primarily think of \mathbf{b} as an intensity artifact, but it can also be thought of as the gain field of the surface coil.

In contrast to our method, Brey-Narayana [14] only use the data from \mathbf{y}_S to construct $\hat{\mathbf{f}}$. When the SNR for \mathbf{y}_S is much larger than the SNR for \mathbf{y}_B (i.e., $\mathbf{b}^2[n] \gg \lambda_B/\lambda_S$), $\hat{\mathbf{f}}^{(i)}[n] \approx \mathbf{y}_S[n]/\mathbf{b}[n]$ which corresponds to the Brey-Narayana method. But in regions where the surface coil response is weak, using both observation images in the reconstruction can be advantageous. Reconstructing from both observation images ensures that every point in $\hat{\mathbf{f}}$ will have SNR ranging from 0 to 3 dB better than the SNR from either of our observation images. The maximal 3 dB gain comes when the bias field is 1 which results in the body coil and surface coil images having the same SNR.

For (13) with a given set of $\{\mathbf{b}_m\}$, we have the following solution:

$$\hat{\mathbf{f}}_k^{(i)}[n] = \frac{\sum_{\substack{i \text{ s.t.} \\ n_i=k}} \lambda_i \mathbf{b}_{m_i}[n] \mathbf{y}_i[n]}{\sum_{\substack{i \text{ s.t.} \\ n_i=k}} \lambda_i \mathbf{b}_{m_i}^2[n]} \quad (17)$$

This is the same equation obtained by Roemer *et al.* [22] for the reconstruction that maximizes SNR when dealing with the magnitude MR images.

2) *Half-Quadratic Solution*: We now describe the f-step for $\gamma \neq 0$. When $p \neq 2$, the optimization problem for \mathbf{f} with a given \mathbf{b} is non-quadratic, and we obtain a nonlinear condition for the minimum. The ℓ_p norm for $p \leq 1$ is non-differentiable at zero, so we use a smoothed approximation:

$$\|\mathbf{x}\|_p^p \approx \sum_n (\mathbf{x}^2[n] + \xi)^{p/2} . \quad (18)$$

As $\xi \rightarrow 0$, the approximation approaches the unsmoothed norm.

Half-quadratic optimization is a fixed-point iterative scheme pioneered by Geman and Reynolds [21] that constructs a weighted- ℓ_2 approximation at each sub-iteration j . It has been demonstrated [25] that half-quadratic optimization provides superior convergence rates compared with gradient descent. Using half-quadratic optimization results in a linear condition on $\hat{\mathbf{f}}^{(i,j)}$:

$$\left(\lambda_B \mathbf{I} + \lambda_S \mathbf{B}^2 + \gamma \mathbf{D}^T \mathbf{W}^{(i,j)} \mathbf{D}\right) \hat{\mathbf{f}}^{(i,j)} = \lambda_B \mathbf{y}_B + \lambda_S \mathbf{B} \mathbf{y}_S \quad (19)$$

with the weighting matrix $\mathbf{W}^{(i,j)}$ being diagonal with the following entries:

$$\mathbf{W}^{(i,j)}[n, n] = \frac{p}{2} \left(((\mathbf{D} \hat{\mathbf{f}}^{(i,j-1)})[n])^2 + \xi \right)^{p/2-1} . \quad (20)$$

This preserves edges by weighting the ℓ_2 norm less in regions with large derivatives. Equation (19) is a positive definite linear system which we can again solve using preconditioned conjugate gradient.

One of the key features of (19) is that the effective amount of regularization is spatially varying—less smoothing is performed in regions where \mathbf{B} is large. This is superior to applying an anisotropic post-processing filter to the output of our algorithm. Depending on the regularization strength, post-processing will either oversmooth in high SNR regions or undersmooth in low SNR regions.

For the general form in (13), we obtain the following linear equation:

$$\left(\sum_{\substack{i \text{ s.t.} \\ n_i=k}} \lambda_i \mathbf{B}_{m_i}^2 + \gamma_k \mathbf{D}_k^T \mathbf{W}_k^{(i,j)} \mathbf{D}_k \right) \hat{\mathbf{f}}_k^{(i,j)} = \sum_{\substack{i \text{ s.t.} \\ n_i=k}} \lambda_i \mathbf{B}_{m_i} \mathbf{y}_i . \quad (21)$$

$\mathbf{W}_k^{(i,j)}$ is defined in an identical manner as in (20).

C. Convergence and Speed

The energy functional E in (4) is non-convex due to the cross-multiplication between \mathbf{b} and \mathbf{f} . Our algorithm possesses convergence qualities similar to the EM algorithm [9]. Each f- and b-step decreases the energy, so our algorithm will at least find a local minimum of E . In practice, we have found excellent convergence properties with the algorithm converging to identical reasonable solutions for random initializations. In fact, initialization with random noise converges to the correct result.

Multigrid techniques [26] can help avoid local minima and improve computation speed for large problems. We use a basic form of multigrid with a single coarse-to-fine sweep. We downsample our data to the coarsest level we wish to process. We then run our coordinate descent solver at this level and upsample the results to the next finest level. This cycle repeats until we have a solution at the original scale. The key advantage of multigrid is that the low-frequency components of the solution, which typically converge slowly at the original scale, can be more efficiently computed at the coarser scales. Curiously, we find that upsampling and downsampling using nearest neighbor interpolation produces superior results to more complex methods such as bilinear interpolation.

D. Three Dimensions

Note that nowhere have we made any assumption about the dimensionality of our data. Three dimensions (3D) is just as easy as two dimensions (2D). The only thing that needs to change is the structure of the D and L matrices which now must implement 3D kernels. 3D processing allows us to couple the regularization of both $\hat{\mathbf{f}}$ and $\hat{\mathbf{b}}$ in the z-direction which can help us in regions where the data does not guide the solution very well (e.g., in air-filled regions). Because we constrain $\hat{\mathbf{f}}$ to be close to \mathbf{y}_B and $\hat{\mathbf{b}}$ to be close to $\mathbf{y}_S \circledast \hat{\mathbf{f}}$, estimates for $\hat{\mathbf{f}}$ and $\hat{\mathbf{b}}$ obtained via 2D or 3D processing do not differ tremendously. Additionally, in MR imaging, the inter-slice length is much larger than the intra-slice dimensions, so this also limits the gain from 3D processing. Still, 3D processing does eliminate some discontinuities in the bias field in the z-direction. The main difficulty with 3D is convergence speed. Our algorithm does not converge in linear time, so processing a volume in 3D is slower than processing all of the individual slices in 2D.

In many MR imaging protocols, different planes are used to image the same location (e.g., axial, sagittal, and coronal). Generally the aspect ratio (the ratio between the slice thickness and pixel dimension) is much larger than one (for SNR purposes), so acquiring image sequences in orthogonal planes can provide additional information. Now, assume that we capture a body coil volume \mathbf{y}_B for the axial slices, and surface coil volumes for axial and sagittal slices ($\mathbf{y}_{S,a}$ and $\mathbf{y}_{S,s}$ respectively). Then we can estimate the bias field $\hat{\mathbf{b}}_a$ for the axial volume using our standard algorithm. Note that the bias fields for both surface coil volumes are the exact same, just sampled on different grids. So we can estimate the bias field $\hat{\mathbf{b}}_s$ for the sagittal volume by simply resampling $\hat{\mathbf{b}}_a$ on the appropriate grid. We find trilinear interpolation works sufficiently well, but alternatives can be found in [27].

Now we have a bias field estimate and a surface coil image for the sagittal volume, and we simply need to perform one f-step to obtain an estimate of the true sagittal image $\hat{\mathbf{f}}_s$:

$$\hat{\mathbf{f}}_s = \arg \min_{\mathbf{f}} \left[\lambda_{S,s} \|\mathbf{y}_{S,s} - \hat{\mathbf{b}}_s \circ \mathbf{f}\|^2 + \gamma \|\mathbf{D}\mathbf{f}\|_p^p \right]. \quad (22)$$

E. Parameter Selection and Initialization

There are a number of parameters that need to be set in our energy functional: $\{\lambda_i\}$, $\{\alpha_m\}$, $\{\gamma_n\}$, and p (for the most general form). We generally use $p = 1$ because it is the smallest value of p that allows the f-step to remain convex. Ideally, we would specify $\{\alpha_m\}$ and $\{\gamma_n\}$ based on training data (e.g., phantom scans of the surface coil profiles, long acquisition-time body coil images). In practice, we choose the regularization parameters based on subjective visual assessment of the results. Because we use an iterative solver, we must specify initial values for both $\hat{\mathbf{f}}^{(i)}$ and $\hat{\mathbf{b}}^{(i)}$. The convergence speed of our solver can be greatly impacted by these choices. We use the bias correction method of Brey-Narayana to produce simple and effective initializations.

We stated that the λ_i should be related to the inverse noise variances of the corresponding observation images. We can estimate the noise variances directly from the images using the method from Nowak

[18]. The true signal should be uniformly zero in air-filled regions, so the observed signal should then be Rayleigh distributed. Let the noise in both the real and imaginary portions of the complex signal be Gaussian distributed with variance σ_i^2 , and let n index into a air-filled region of \mathbf{y}_i . Then the second moment of $\mathbf{y}_i[n]$ is

$$\mathbb{E}[\mathbf{y}_i^2[n]] = \mathbb{E}[n_{\text{real}}^2 + n_{\text{imag}}^2] = 2\sigma_i^2 . \quad (23)$$

We can approximate the expected value by taking the sample average over a large air-filled region to estimate $2\sigma_i^2$. Note that the bias field has no effect in air-filled regions, so we can perform this technique for all observation images regardless of coil configuration.

When using a multigrid solver, we choose $\{\lambda_i\}$, $\{\alpha_m\}$, and $\{\gamma_n\}$ at the original scale. We must also choose these parameters at each scale s so that the solutions at the coarser and original scales are similar. The λ_i 's should scale by 4^s (or 8^s in 3D) due to noise reduction from spatial averaging. For wavelet-based reconstruction, others have found that multiplicative scaling of the regularization parameters is effective [28], [29]. Hence we multiply α_m and γ_n at each scale s by experimentally determined positive scalars ζ_m^s and η_n^s respectively:

$$\alpha_m^{[s]} = (\zeta_m)^s \alpha \quad (24)$$

$$\gamma_n^{[s]} = (\eta_n)^s \gamma \quad (25)$$

Note the $[s]$ on the left-hand side of the equations indexes scale while the s on the right-hand side indicates exponentiation.

IV. RESULTS

In this section, we demonstrate results on real and synthetic data, along with comparisons with Brey-Narayana bias correction³. All real data in this section were captured on General Electric Signa 1.5-T machines. We computed results on a Pentium 4 1.8 GHz workstation using our 2D multigrid solver and stopped the algorithm when the energy changed by less than 0.01%. For all results, we use Laplacian regularization on $\hat{\mathbf{b}}$ and gradient regularization on $\hat{\mathbf{f}}$. The numerical values of α presented in this section provide an indication of relative smoothness across examples because the bias field is a ratio between the body coil image and a surface coil image. Thus the bias field is unchanged if \mathbf{y}_B and \mathbf{y}_S are both scaled equally. On the other hand, the numerical values of γ are not very informative because of scaling variations in the examples.

³Even though Brey-Narayana is an older technique, we have found that its performance is superior to that of newer body coil/surface coil correction methods such as Lai-Fang. Additionally, we do not compare with other bias correction techniques that only have access to the surface coil image. The goal of our algorithm and Brey-Narayana is to produce an intrinsic image estimate that is as close to the true intrinsic image as possible, and we can ensure this closeness because we have access to the body coil image which is a noisy, bias-free version of the true image. Bias correction techniques that do not have access to the body coil image may in fact reliably increase the intensity homogeneity within regions, but the corrected image may not correspond well in a mean-squared error sense with the true image.

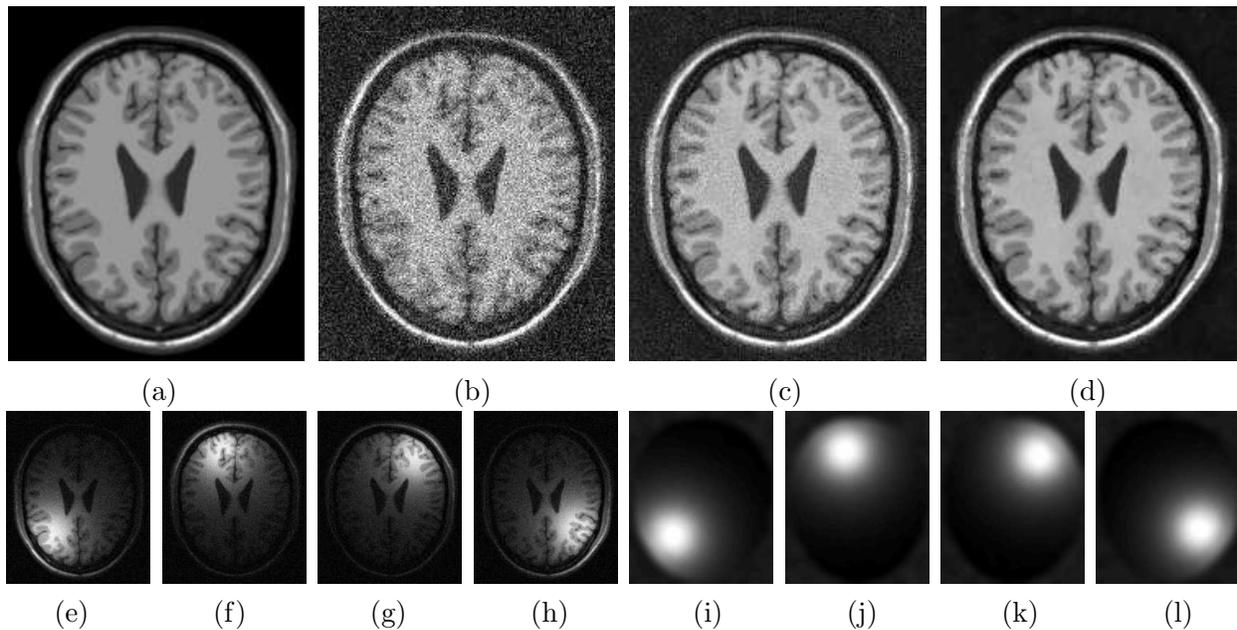


Fig. 1. Synthetic axial T_1 -weighted brain images. (a) True image (f^*). (b) Body coil image (y_B). Estimated intrinsic image (\hat{f}) computed with (c) Brey-Narayana and (d) proposed method using $\gamma = 0.014$. (e)–(h) Surface coil images ($y_{S,1}$ – $y_{S,4}$). (i)–(l) Estimated bias fields (\hat{b}_1 – \hat{b}_4). $\alpha_k = 2000$.

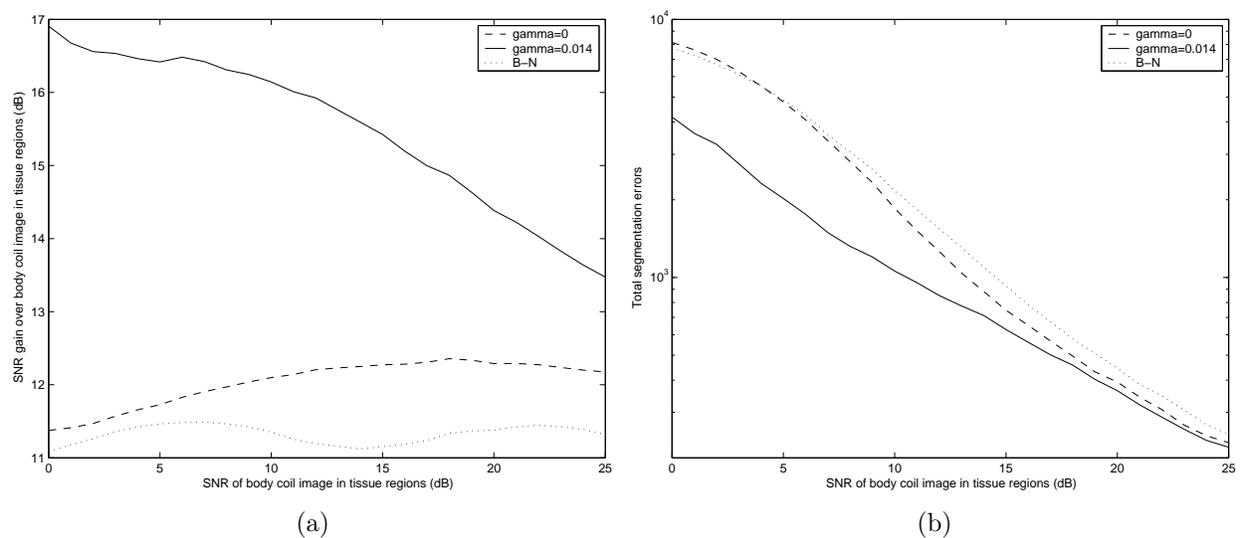


Fig. 2. Performance provided by Brey-Narayana and proposed correction method with varying SNR levels. (a) SNR gain over the body coil image (y_B). (b) Total GM and WM segmentation errors. Averaged over 10 Monte Carlo trials.

	\mathbf{y}_B	B-N	$\hat{\mathbf{f}}, \gamma = 0$	$\hat{\mathbf{f}}, \text{post-filtered}$	$\hat{\mathbf{f}}, \gamma = 0.014$
MSE (tissue)	196,542	20,820	20,428	15,853	10,901
MAE (tissue)	353.91	113.59	112.64	98.41	81.99
GM errors	64.4%	14.6%	14.3%	12.4%	9.9%
WM errors	24.7%	3.6%	3.3%	2.8%	2.4%

TABLE I

QUANTITATIVE COMPARISONS USING THE MNI BRAIN PHANTOM. CORRECTED IMAGES ARE GENERATED USING BREY-NARAYANA AND THE PROPOSED METHOD WITH $\gamma = 0$ (WITH AND WITHOUT POST-PROCESSING USING ANISOTROPIC FILTERING) AND $\gamma = 0.014$. ALL METHODS HAD PARAMETERS ADJUSTED TO PRODUCE OPTIMAL RESULTS. THE FIRST TWO LINES ARE THE MEAN SQUARED ERROR AND MEAN ABSOLUTE ERROR (BASED ON THE TRUE IMAGE \mathbf{f}^*) COMPUTED ONLY IN TISSUE REGIONS. THE LAST TWO LINES ARE THE PERCENTAGE OF MISCLASSIFIED POINTS IN GM AND WM REGIONS. RESULTS AVERAGED OVER 20 RANDOM TRIALS.

We begin with synthetic results using the Montreal Neurological Institute (MNI) [30], [31] BrainWeb simulator. We used the T_1 -weighted images with 1 mm slice thickness and constructed synthetic bias fields that simulate a four-coil phased array. We then added Rician noise to obtain our body coil and surface coil images. For tissue regions of \mathbf{y}_B , we set the SNR at 13 dB. Rician noise at this level imparts an upward bias, increasing the mean signal level by 2-3%. We computed bias field and intrinsic image estimates by minimizing (8), our multiple surface coil framework.

We present the observed and corrected images in Fig. 1. The bias field estimates are largely independent of the tissue, and our method produces a $\hat{\mathbf{f}}$ with noticeably superior noise properties than Brey-Narayana. These visual impressions are confirmed with our quantitative results in Table I with mean squared error in tissue regions 48% lower than Brey-Narayana. Segmentation accuracy is another way to quantify the quality of the bias correction. We generated gray matter (GM) and white matter (WM) segmentation results using a simple thresholding scheme with manual skull peeling. Ground truth was determined using \mathbf{f}^* . Compared with Brey-Narayana, we reduce overall segmentation errors by 33%.

Most of the gain from our method comes from the integrated denoising rather than from superior bias field estimates. But the integrated denoising provides superior results to simple post-processing using an anisotropic edge-preserving filter. To obtain the post-processing results, we first computed an intrinsic image estimate $\hat{\mathbf{f}}_0$ with $\gamma = 0$. We then chose $\hat{\mathbf{f}}$ as

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} \left[\|\hat{\mathbf{f}}_0 - \mathbf{f}\|^2 + \zeta \|\mathbf{D}\mathbf{f}\|_p^p \right]. \quad (26)$$

Table I shows that post-processing is able to improve upon the $\gamma = 0$ results, but is inferior to our integrated denoising method. This is because integrating the denoising allows us to spatially vary the amount of regularization based on the local signal level.

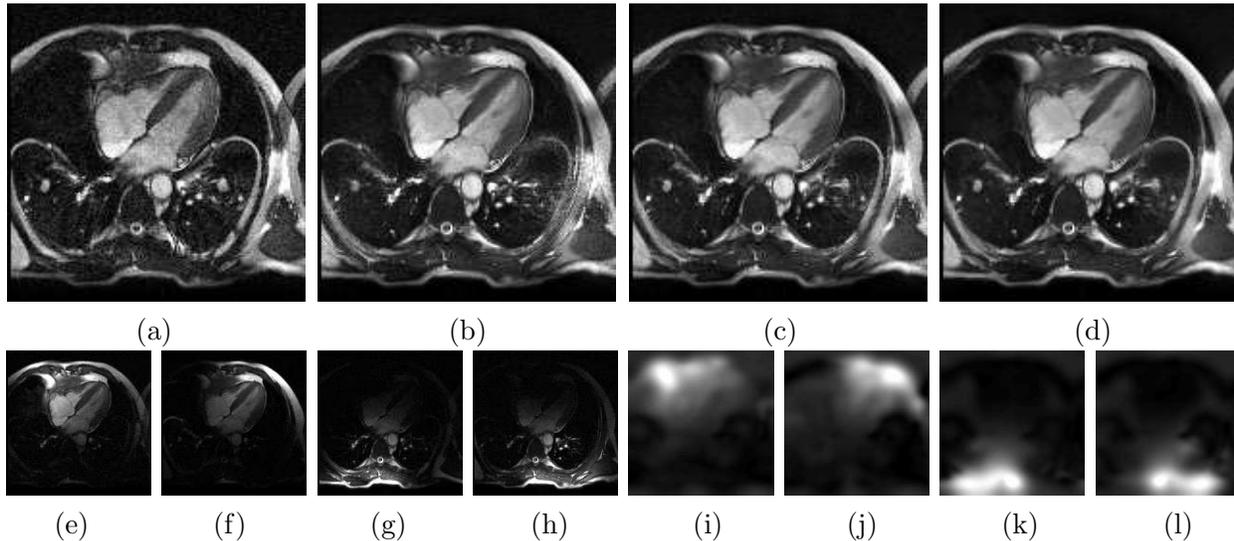


Fig. 3. Gated cardiac MR images. (a) Body coil image (\mathbf{y}_B). Estimates of the intrinsic image ($\hat{\mathbf{f}}$) using (b) Brey-Narayana and proposed method with (c) $\gamma = 0$ and (d) $\gamma = 1800$. (e)–(h) Surface coil images ($\mathbf{y}_{S,1}$ – $\mathbf{y}_{S,4}$). (i)–(l) Estimated bias fields ($\hat{\mathbf{b}}_1$ – $\hat{\mathbf{b}}_4$). $\alpha_k = 3000$.

In Fig. 2, we show how the different bias correction schemes function as the SNR is varied. Our method with $\gamma = 0$ consistently outperforms Brey-Narayana due to better bias field estimates. In high SNR regions, all methods provide similar results. As the SNR is decreased, our method with regularization on $\hat{\mathbf{f}}$ builds up a significant advantage over the other methods. At 0 dB SNR, Brey-Narayana and our method with $\gamma = 0$ produce segmentation error rates of 51% (which is approximately equivalent to random guessing), while using regularization on $\hat{\mathbf{f}}$ reduces the error to 27%.

Next, we apply our algorithm to one time step from a gated cardiac MR sequence in Fig. 3. The heart is the object in the upper-middle portion of the image. For the surface coil images, a four-element phased array was used. The images have a field of view (FOV) of $32 \text{ cm} \times 32 \text{ cm}$, resolution of 160×192 , and slice thickness of 8 mm. To obtain our results, we applied our multiple surface coil correction framework and minimized (8). The main differences between the Brey-Narayana estimate in Fig. 3(b) and our result in Fig. 3(c) are in regions where none of the surface coils have good response such as the middle and the right-hand side of the image. Our method results in higher SNR because we use the body coil information while Brey-Narayana does not. Fig. 3(d) (using $\hat{\mathbf{f}}$ regularization) is moderately better than Fig. 3(c), but the high SNR in \mathbf{y}_B and $\mathbf{y}_{S,k}$ limit the benefits of filtering.

In Fig. 4, we display the results of our algorithm on a real prostate image. The rectum is the black circular object, and the prostate is directly above it. The surface coils used were an endorectal coil along with a four element pelvic phased-array coil. We captured T_2 -weighted images using the body coil and surface coils and T_1 -weighted images using just the surface coils. The FOV is $12 \text{ cm} \times 12 \text{ cm}$, resolution

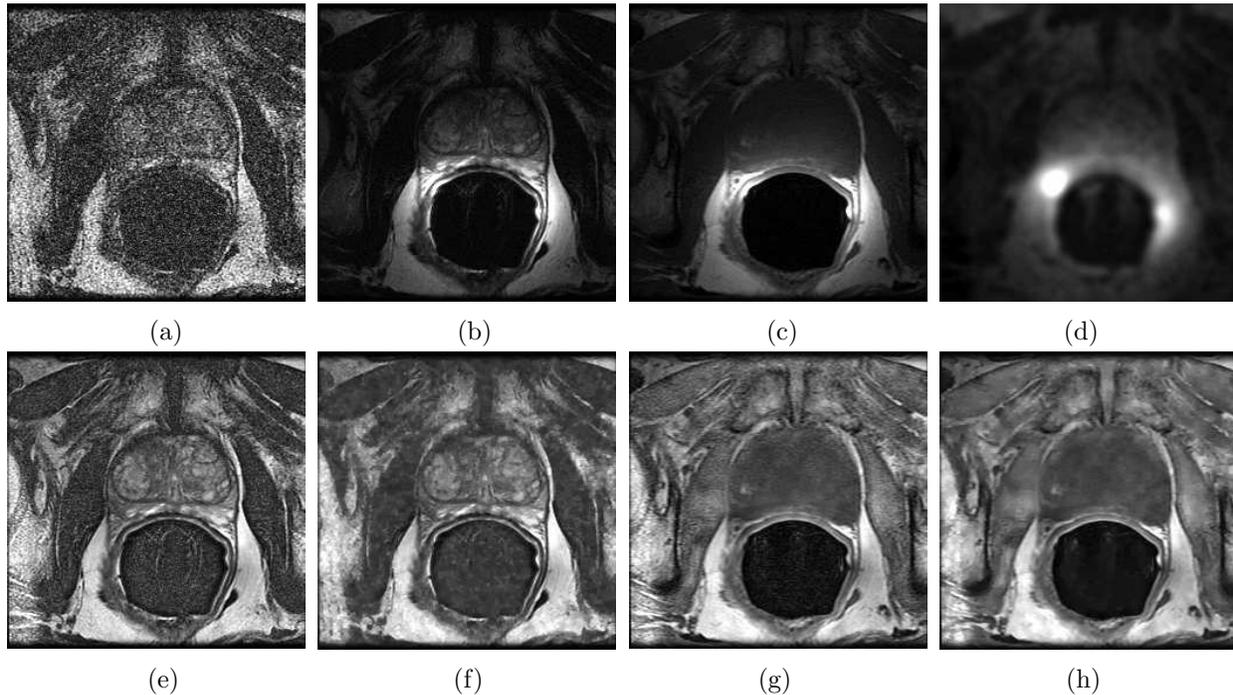


Fig. 4. Axial mid-gland prostate images. T_2 -weighted (a) body coil (\mathbf{y}_B) and (b) composite surface coil ($\mathbf{y}_{S,1}$) images. (c) Composite surface coil T_1 -weighted image ($\mathbf{y}_{S,2}$). (d) Estimated bias field ($\hat{\mathbf{b}}$). T_2 -weighted intrinsic image estimates ($\hat{\mathbf{f}}_1$) using (e) Brey-Narayana and (f) proposed method with $\gamma_1 = 0.018$. T_1 -weighted intrinsic image estimates ($\hat{\mathbf{f}}_2$) using (g) Brey-Narayana and (h) proposed method with $\gamma_2 = 0.010$. $\alpha = 125$.

is 256×256 , and slice thickness is 3 mm. Estimates were computed by minimizing (11) using composite surface coil images because individual surface coil data were not available to us.

The prostate is the most challenging example we consider here. The FOV is small so \mathbf{y}_B has very low SNR (about 7 dB in the prostate). To compensate, the endorectal coil produces a strong local response profile which results in a severe intensity artifact. Because the reception profile of the endorectal coil is much less homogeneous than that of the pelvic phased-array coil, the prostate would probably benefit significantly from processing each coil separately. Fig. 4(d) shows that using a composite surface coil image causes $\hat{\mathbf{b}}$ to be under-regularized in regions where the endorectal coil does not dominate. Figs. 4(e)–(h) (when viewed under sufficiently high resolution) demonstrate that our method preserves edges while resulting in lower noise than Brey-Narayana. Fig. 4(h) shows that even without a body coil image, we can obtain reasonable intrinsic image estimates for the T_1 -weighted sequence simply by minimizing (11). This largely confirms our assumption that \mathbf{f}^* and \mathbf{b}^* are independent. If there was dependence, we would not be able to correct the T_1 -weighted image using a bias field largely estimated from the T_2 -weighted images.

Fig. 5 shows an example of using the axial-plane bias field estimates obtained in Fig. 4 to process

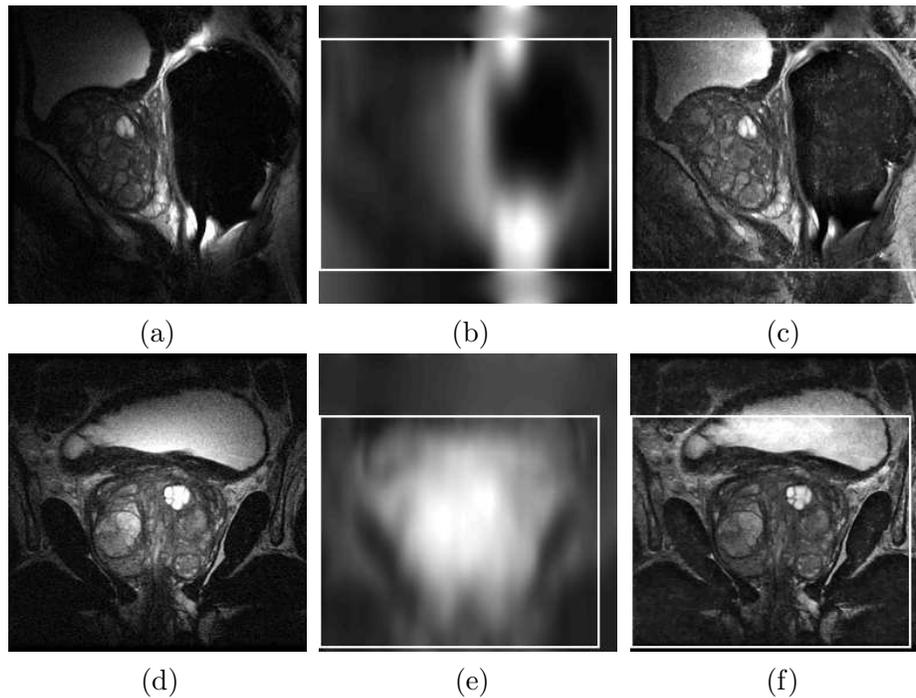


Fig. 5. Sagittal and coronal T_2 -weighted prostate images corrected using bias field estimates obtained from axial body coil and surface coil images. Spatial regions where there is overlap with the axial volume are enclosed within the white boxes. The top row (a-c) contains sagittal images, the bottom row (d-f) contains coronal images. The left column (a,d) contains composite surface coil images, middle column (b,e) contains interpolated bias field estimate, and the right column (c,f) contains intrinsic image estimates.

sagittal and coronal volumes of the same patient. Because the bias field is largely a low-frequency phenomenon, the lack of resolution in the z -axis for the axial images does not severely impact the quality of the interpolated bias fields. One of the issues with this type of correction is that the volumes of space imaged for each plane do not perfectly coincide. Hence there are some points in each sagittal and coronal image where we do not have any body coil information. This is generally acceptable because these locations should be on the boundary of the image and should not involve the main object of interest. We fill in the missing data in the interpolated bias fields using harmonic interpolation.

The sagittal image exhibits a fairly large bias field, but the coronal image is not affected as much. The intersection between an axial image and a coronal image is a horizontal line in the axial image, and the intersection between a sagittal image and a coronal image is a vertical line in the sagittal image. The distance between points in the coronal image and the endorectal coil tends to be fairly constant throughout the plane. Because the main source of the bias field in these images is caused by the endorectal coil, the bias field tends to be more homogeneous. Hence, correcting the sagittal images is much more beneficial than correcting the coronal images for this particular example. Of course, this kind of characterization

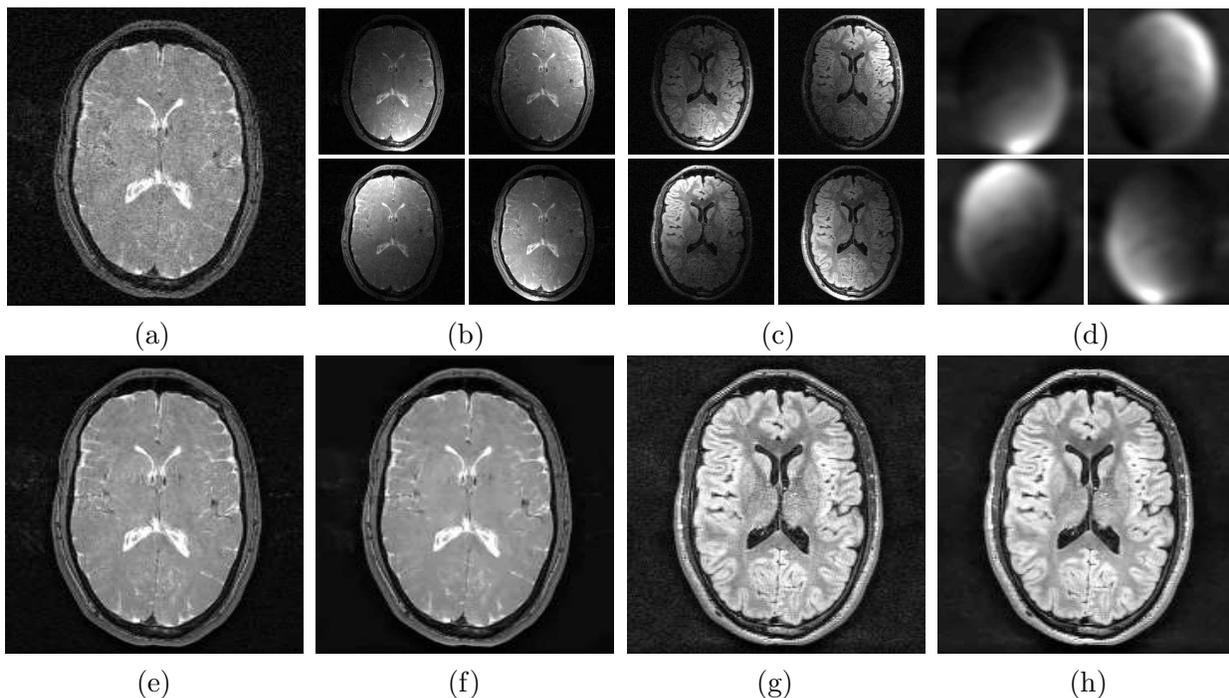


Fig. 6. Axial brain images. (a) GRE body coil image (\mathbf{y}_B). (b) GRE surface coil images ($\mathbf{y}_{S,1}-\mathbf{y}_{S,4}$). (c) FLAIR surface coil images ($\mathbf{y}_{S,5}-\mathbf{y}_{S,8}$). (d) Estimated bias fields ($\hat{\mathbf{b}}_1-\hat{\mathbf{b}}_4$). Estimated GRE intrinsic images ($\hat{\mathbf{f}}_1$) using (e) Brey-Narayana and (f) proposed method with $\gamma_1 = 1000$. Estimated FLAIR intrinsic images ($\hat{\mathbf{f}}_2$) using (g) Brey-Narayana and (h) proposed method with $\gamma_2 = 1200$. $\alpha_k = 1000$.

will vary from application to application due to coil geometry.

We show axial brain images in Fig. 6. The surface coils are a four-element phased array. We captured gradient-recalled echo (GRE) images using both the body coil and surface coils and fluid-attenuated FLAIR images using the surface coils. FOV is $24 \text{ cm} \times 24 \text{ cm}$, resolution is 192×256 , and slice thickness is 3 mm. We minimize (13) using 8 surface coil images, 1 body coil image, 2 intrinsic images, and 4 bias fields. All of the surface coils have weak signal strength in the middle of the brain, so our final estimate of the FLAIR image in Fig. 6(h) is still noisy in the middle, even with the ℓ_p reconstruction. This artifact is not present in our GRE estimate in Fig. 6(f) because the body coil image ensures a minimum SNR level.

V. CONCLUSION

The main contribution of this paper is a fully-automatic non-parametric approach to MR bias correction. We presented a unified approach that simultaneously debiases and denoises the MR images. We constructed an energy functional that ensured our estimates were consistent both with the observed data and with prior knowledge about the signals. We produced our corrected images by iteratively minimizing

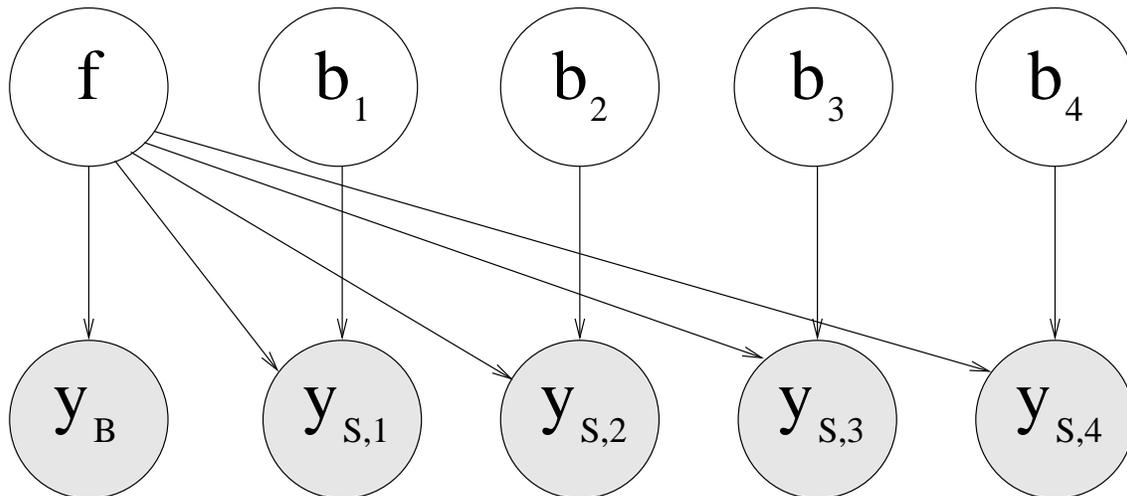


Fig. 7. Bayesian network illustrating the graph structure of an acquisition scenario with one intrinsic image and four surface coils resulting in one body coil image and four surface coil images. White nodes indicate unknown variables, and shaded nodes indicate observations.

the energy functional using coordinate descent, conjugate gradient, half-quadratic regularization, and multigrid. In practice, we have found our algorithm to have nice convergence properties.

We presented results on a number of examples. We found that the ℓ_p regularization helped reduce the noise on relatively high SNR images such as cardiac and brain images, but the largest gains came on lower SNR examples such as the prostate. We demonstrated the advantage quantitatively using the MNI brain phantom. We feel that the quality of the reconstructed true images justifies the extra time and expense associated with capturing the additional body coil scan, especially for certain applications such as the prostate which have severe intensity inhomogeneities.

APPENDIX

Obviously with our most general framework described in Sec. II-E, not every choice of parameters will produce a well-posed minimization problem. As a trivial example, look at the case with one surface coil and one pulse sequence. This is the standard bias correction framework with twice as many unknowns as observations. We can describe the relationships between our observations and our unknowns using graphical models [23]. Graphical models specify the dependencies of random variables by representing the random variables as nodes and dependencies as edges connecting two nodes. We use Bayesian networks (i.e., directed graphs) which explicitly contain causality information. In Fig. 7, we show one possible representation for the example covered in Sec. II-C of one intrinsic image and four surface coils. Here the direction of the arrows shows that the observation images are causally dependent on the intrinsic image and bias fields. This perspective makes sense from a physical interpretation.

In order to have a well-posed optimization problem, we have three conditions:

- 1) There must be at least as many observation images as intrinsic images and bias fields to estimate (i.e., $N_o \geq N_f + N_b$).
- 2) There must be at least one body coil image.
- 3) The graph describing the problem must be connected.

A graph is connected if, regardless of the initial node, it is possible to reach any node in the graph solely by traversing along the edges of the graph. This condition is necessary to ensure that there is not an independent sub-problem contained in the overall problem that does not satisfy the first two conditions⁴. The first two conditions are necessary to ensure that an infinite subspace of the solution space are not all minimums of the energy functional. The second condition may not seem necessary, but can be illustrated by the case of two intrinsic images (\mathbf{f}_1 and \mathbf{f}_2) and two bias fields (\mathbf{b}_1 and \mathbf{b}_2) for each one, resulting in four observation images. Then condition one is satisfied, as is condition three. For simplicity, we consider the case where there is no additive noise (though this argument generalizes), and $\alpha = \gamma = 0$:

$$\mathbf{y}_1 = \mathbf{b}_1 \circ \mathbf{f}_1, \quad \mathbf{y}_2 = \mathbf{b}_1 \circ \mathbf{f}_2, \quad \mathbf{y}_3 = \mathbf{b}_2 \circ \mathbf{f}_1, \quad \mathbf{y}_4 = \mathbf{b}_2 \circ \mathbf{f}_2 .$$

The energy functional is

$$E(\mathcal{F}, \mathcal{B}) = \|\mathbf{y}_1 - \mathbf{b}_1 \circ \mathbf{f}_1\|^2 + \|\mathbf{y}_2 - \mathbf{b}_1 \circ \mathbf{f}_2\|^2 + \|\mathbf{y}_3 - \mathbf{b}_2 \circ \mathbf{f}_1\|^2 + \|\mathbf{y}_4 - \mathbf{b}_2 \circ \mathbf{f}_2\|^2 .$$

Then a minimum of the energy functional is the true intrinsic images and bias fields:

$$\hat{\mathbf{f}}_1 = \mathbf{f}_1, \quad \hat{\mathbf{f}}_2 = \mathbf{f}_2, \quad \hat{\mathbf{b}}_1 = \mathbf{b}_1, \quad \hat{\mathbf{b}}_2 = \mathbf{b}_2 .$$

But we can take any vector \mathbf{a} and have the following also produce the minimum value:

$$\hat{\mathbf{f}}_1 = \mathbf{a} \circ \mathbf{f}_1, \quad \hat{\mathbf{f}}_2 = \mathbf{a} \circ \mathbf{f}_2, \quad \hat{\mathbf{b}}_1 = \mathbf{b}_1 \oslash \mathbf{a}, \quad \hat{\mathbf{b}}_2 = \mathbf{b}_2 \oslash \mathbf{a} .$$

The modifying vector \mathbf{a} cancels out when computing the energy, and this solution also results in an energy of zero.

REFERENCES

- [1] L. Axel, "Surface coil magnetic resonance imaging," *J. of Comp. Asst. Tomography*, vol. 8, pp. 381–384, 1984.
- [2] L. Axel, J. Costantini, and J. Listerud, "Intensity correction in surface-coil MR imaging," *Am. J. Roentgenology*, vol. 148, pp. 418–420, 1987.
- [3] R. B. Lufkin, T. Sharpless, B. Flannigan, and W. Hanafee, "Dynamic-range compression in surface-coil MRI," *Am. J. Roentgenology*, vol. 147, pp. 379–382, 1986.
- [4] J. Haselgrove and M. Prammer, "An algorithm for compensation of surface-coil images for sensitivity of the surface coil," *Mag. Res. Imag.*, vol. 4, pp. 469–472, 1986.

⁴For example, we could have a body coil image and a surface coil image of the heart, two T_1 -weighted surface coil images of the brain, and two T_2 -weighted surface coil images of the brain. Then we have a body coil image (satisfying condition 2) and have 6 observation images, 3 surface coils, and 3 intrinsic images (satisfying condition 1). But the heart and brain bias correction problems are completely decoupled (which is indicated by the unconnected graph), and the brain bias correction problem is ill posed.

- [5] J. G. Sled, A. P. Zijdenbos, and A. C. Evans, "A nonparametric method for automatic correction of intensity nonuniformity in MRI data," *IEEE Trans. Med. Imag.*, vol. 17, pp. 87–97, 1998.
- [6] B. M. Dawant, A. P. Zijdenbos, and R. A. Margolin, "Correction of intensity variations in MR images for computer-aided tissue classification," *IEEE Trans. Med. Imag.*, vol. 12, pp. 770–781, 1993.
- [7] B. Likar, M. A. Viergever, and F. Pernus, "Retrospective correction of MR intensity inhomogeneity by information minimization," in *MICCAI 2000*, 2000, pp. 375–384.
- [8] W. M. Wells, W. E. L. Grimson, R. Kikinis, and F. Jolesz, "Adaptive segmentation of MRI data," *IEEE Trans. Med. Imag.*, vol. 15, no. 4, pp. 429–442, 1996.
- [9] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Stat. Soc.*, vol. 39, pp. 1–38, 1977.
- [10] R. Guillemaud and M. Brady, "Estimating the bias field of MR images," *IEEE Trans. Med. Imag.*, vol. 16, no. 3, pp. 238–251, 1997.
- [11] Y. Zhang, M. Brady, and S. Smith, "Segmentation of brain MR images through a hidden Markov random field model and the Expectation-Maximization algorithm," *IEEE Trans. Med. Imag.*, vol. 20, no. 1, pp. 45–57, 2001.
- [12] M. Singh and M. NessAiver, "Accurate intensity correction for endorectal surface coil MR imaging of the prostate," *IEEE Trans. Nucl. Sci.*, vol. 40, pp. 1169–1173, 1993.
- [13] S. E. Moyher, L. . Wald, S. J. Nelson, D. Hallam, W. P. Dillon, D. Norman, and D. B. Vigneron, "High resolution T2-weighted imaging of the human brain using surface coils and an analytical reception profile correction," *J. Mag. Res. Imag.*, vol. 7, no. 3, pp. 512–517, 1997.
- [14] W. W. Brey and P. A. Narayana, "Correction for intensity falloff in surface coil magnetic resonance imaging," *Med. Phys.*, vol. 15, no. 2, pp. 241–245, 1988.
- [15] S.-H. Lai and M. Fang, "Intensity inhomogeneity correction for surface-coil MR images by image fusion," in *Proceedings of the International Conference on Multisource-Multisensor Information Fusion*. CSREA Press, 1998, pp. 880–887.
- [16] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," *Mag. Res. Med.*, vol. 42, pp. 952–962, 1999.
- [17] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.
- [18] R. D. Nowak, "Wavelet-based Rician noise removal for magnetic resonance imaging," *IEEE Trans. Imag. Proc.*, vol. 8, no. 10, pp. 1408–1419, 1999.
- [19] S. Rice, "Mathematical analysis of random noise," *Bell Sys. Tech. J.*, vol. 23, pp. 282–332, July 1944.
- [20] A. Tikhonov and V. Arsenin, *Solution of Ill-Posed Problems*. Washington, D.C.: W.H. Winston, 1977.
- [21] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 14, no. 3, pp. 367–383, 1992.
- [22] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, "The NMR phased array," *Mag. Res. Med.*, vol. 16, pp. 192–225, 1990.
- [23] F. V. Jensen, *Bayesian Networks and Decision Graphs*. New York: Springer, 2001.
- [24] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, 1999.
- [25] C. R. Vogel and M. E. Oman, "Iterative methods for total variation denoising," *SIAM J. Sci. Comp.*, vol. 17, no. 1, pp. 227–238, 1996.
- [26] W. L. Briggs, *A Multigrid Tutorial*. Philadelphia: Society for Industrial and Applied Mathematics, 1987.
- [27] J. S. Lim, *Two-Dimensional Signal and Image Processing*. Prentice Hall, 1989.
- [28] M. Bhatia, W. C. Karl, and A. S. Willsky, "A wavelet-based method for multiscale tomographic reconstruction," *IEEE Trans. Med. Imag.*, vol. 15, no. 1, pp. 92–101, 1996.
- [29] M. Belge, M. E. Kilmer, and E. L. Miller, "Wavelet domain image restoration with adaptive edge-preserving regularization," *IEEE Trans. Imag. Proc.*, vol. 9, no. 4, pp. 597–608, 2000.
- [30] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans, "Design and construction of a realistic digital brain phantom," *IEEE Trans. Med. Imag.*, vol. 17, no. 3, pp. 463–468, June 1998.

- [31] R. K.-S. Kwan, A. C. Evans, and G. B. Pike, "MRI simulation-based evaluation of image-processing and classification methods," *IEEE Trans. Med. Imag.*, vol. 18, no. 11, pp. 1085–1097, Nov. 1999.