Prospective Motion Correction for Functional MRI Using Sparsity and Kalman Filtering

Daniel S. Weller[a], Douglas C. Noll[b], and Jeffrey A. Fessler[a,b]

[a]EECS Department, University of Michigan, Ann Arbor, MI 48109, USA

[b]BME Department, University of Michigan, Ann Arbor, MI 48109, USA

ABSTRACT

We propose a novel algorithm to adaptively correct head motion during functional magnetic resonance imaging scans. Our method combines a Kalman-filter-like motion tracker and a registration cost function based on a sparse residual image model. Using simulated data, we compare a time series correlation analysis of our prospectively corrected reconstruction against the same analysis using post-scan motion correction provided by standard software. Our experiments demonstrate our prospective correction method is capable of mitigating motion effects and improving the sensitivity and specificity of the correlation analysis, without relying on costly external tracking hardware or separate navigational data that would take extra time to acquire during each time frame.

Keywords: functional magnetic resonance imaging, sparsity, motion correction, Kalman filtering

1. INTRODUCTION

Head motion is an important source of artifacts affecting functional magnetic resonance imaging (fMRI).[1] Misaligned time series confound voxel-by-voxel statistical analysis of fMRI data, as task-related activations may move among different voxels over time. Conventional processing of fMRI data includes retrospective (post-scan) motion correction[2] and inclusion of motion estimates as nuisance regressors in the general linear model[3] used to identify task-correlated regions. In this work, we investigate using prospective (within-scan) adjustments of the scan prescription to correct for estimated motion, improving the quality of fMRI data analysis. In particular, we propose a novel combination of Kalman filter-like motion tracking and sparse residual-based motion estimation and describe an implementation suitable for real time prospective correction.

A comparison[4] of retrospective post-processing software for fMRI highlights some features of commonly-used registration algorithms. Motion correction methods generally involve spatial interpolation that cause smoothing and can distort the temporal statistics of the acquired time series. Motion correlated with the stimulus introduces additional problems when motion estimates are included as nuisance regressors, potentially confusing activations for motion artifacts. Prospective correction adjusts the sampling pattern rather than the acquired data, so the data remains unaffected by registration-induced smoothing. If done on a slice-by-slice basis, prospective motion correction also can avoid spin saturation effects, which can be difficult to correct post-scan.[5] However, prospective motion correction must be done in real time, which imposes practical limits on estimation accuracy, since registration methods can be time-consuming.

Numerous prospective motion correction approaches have been proposed previously for functional MRI.[6] Some prospective correction methods rely on external tracking hardware for estimating motion.[7] Methods like Prospective Acquisition CorrEction (PACE)[8] use the images themselves for prospective correction with a linearized motion model. PROspective MOtion correction (PROMO)[9] uses an extended Kalman filter to track head motion. PROMO also uses special navigators interleaved with the fMRI readouts for motion estimation, since the fMRI data are sensitive to other time-varying phenomena like activations and physiological signals. However, these navigators cost additional time between frames, reducing temporal resolution.

Further author information: (Send correspondence to DSW.)

DSW: E-mail: wellerds@umich.edu, Telephone: 1 734 615 5735

To improve upon existing prospective correction techniques and mitigate the need for navigator scans, we leverage the sparse residual image concept used for retrospective motion correction in dynamic MRI in methods like k-t FOCal Underdetermined System Solver (FOCUSS)[10] and motion-adaptive spatio-temporal regularization (MASTeR).[11] By using a sparsifying penalty on the residual between successive time frames, we suppress residual motion while preserving time-varying spatially sparse image content like functional activations. Since the registration is performed against the previous frame, it helps that activations are also sparse in the time-gradient domain. We expect the combination of sparse residual modeling and Kalman filter prediction to yield robust motion estimates suited to prospective motion correction.

In Section 2, we first construct our cost functions for image reconstruction and motion estimation. We propose a method based on the alternating direction method of multipliers (ADMM) to solve this joint optimization problem efficiently. In Section 3, we describe our simulation of a functional MRI time series, including activations and rigid frame-by-frame motion. We provide experimental results in Section 4 that support using our method for prospective motion correction, and we discuss these results and future directions in Section 5.

2. THEORY

To begin, consider frames $\boldsymbol{x}_0^t$, for $t = 0, \ldots, N_F - 1$, all in the scanner coordinate system. Due to head motion, the position of the head in these frames changes relative to its position in the initial frame, which we use as a common reference. Ignoring image content changes, we approximate this deformation from frame $\boldsymbol{x}_0^0$ to $\boldsymbol{x}_0^t$ as three-dimensional rigid-body motion described using the six-parameter vector $\boldsymbol{\alpha}^t$ (three translational parameters and three rotational parameters). From an estimate $\widehat{\boldsymbol{\alpha}}^t$ for this motion, we form the registered frame $\boldsymbol{x}_{\text{reg}}^t$, so the head ideally has the same position and orientation as $\boldsymbol{x}^0$ (we drop the subscript for the first frame). We use the b-spline interpolation[12, 13] code in the Image Reconstruction Toolbox available at `http://web.eecs.umich.edu/~fessler/code/index.html` to perform this transformation and let $\boldsymbol{T}(\boldsymbol{x}; \boldsymbol{\alpha})$ be the transformation of image $\boldsymbol{x}$ by motion $\boldsymbol{\alpha}$. For convenience, we introduce the notation $\boldsymbol{T}(\boldsymbol{x}; \boldsymbol{\alpha}_n, \ldots, \boldsymbol{\alpha}_1)$ to describe successive motion transformations, first by $\boldsymbol{\alpha}_1$, then by $\boldsymbol{\alpha}_2$, and last by $\boldsymbol{\alpha}_n$. By iterating motion transforms, we can register one frame to the next: $\boldsymbol{x}_0^t$ is registered to $\boldsymbol{T}(\boldsymbol{x}_0^{t-1}; \widehat{\boldsymbol{\alpha}}^t, (\widehat{\boldsymbol{\alpha}}^{t-1})^{(-1)})$, where $(\widehat{\boldsymbol{\alpha}}^{t-1})^{(-1)}$ is the inverse motion for frame $t - 1$.

With prospective correction, we measure our images not in the common coordinate system but in a measurement coordinate system corrected using our motion estimates for the previous frame. We define $\boldsymbol{x}_{\text{meas}}^t = \boldsymbol{T}(\boldsymbol{x}_0^t; (\widehat{\boldsymbol{\alpha}}^{t-1})^{(-1)})$ to be the image we actually sample in k-space. The assumption underlying using prospective correction is that the residual motion between frames is small compared to the overall motion relative to the initial frame. The data $\boldsymbol{y}^t$ acquired for frame t is corrupted by complex AWGN $\boldsymbol{\eta}^t$ with variance σ^2:

$$\boldsymbol{y}^t = \boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{x}_{\text{meas}}^t + \boldsymbol{\eta}^t, \tag{1}$$

where $\boldsymbol{\mathcal{F}}$ is the Discrete Fourier Transform (DFT), and $\boldsymbol{M}$ extracts the sampled k-space frequencies.

A sparse residual image model is used in retrospective methods like k-t FOCUSS[10] and MASTeR[11] because the difference between successive frames after correcting for motion can be modeled as sparse in problems like dynamic cardiac imaging. The cost function used for registration is $\|\boldsymbol{x} - \boldsymbol{T}(\boldsymbol{x}_{\text{reg}}^{t-1}; (\boldsymbol{\alpha}^{t-1})^{-1}, \boldsymbol{\alpha})\|_1$. This sparse residual model is a good fit for the functional imaging problem based on the spatial sparsity of the time-varying activations. The joint motion estimation/reconstruction problem is

$$\{\boldsymbol{x}_{\text{meas}}^t, \boldsymbol{\alpha}^t\} \in \arg\min_{\boldsymbol{x}, \boldsymbol{\alpha}} \frac{1}{2\sigma^2} \|\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{x} - \boldsymbol{y}^t\|^2 + \lambda \|\boldsymbol{x} - \boldsymbol{T}(\boldsymbol{x}_{\text{reg}}^{t-1}; (\boldsymbol{\alpha}^{t-1})^{(-1)}, \boldsymbol{\alpha})\|_1, \tag{2}$$

where the parameter λ controls the sparsity of the difference image. In the undersampled or noisy cases, one could add a suitable image regularization term to the objective function.[14]

Kalman filtering has been used to track motion in prospectively corrected MRI through the PROMO[9] method. We model $\boldsymbol{\alpha}^0, \ldots, \boldsymbol{\alpha}^t, \ldots$ as a first-order random walk process with zero-mean Gaussian process innovation $\boldsymbol{a}$ with covariance matrix $\boldsymbol{Q}$: $\boldsymbol{\alpha}^t = \boldsymbol{\alpha}^{t-1} + \boldsymbol{a}^t$. The PROMO method uses navigators to estimate motion, as they are relatively insensitive to time-varying signals like task activations. Using the sparse residual model described before, we propose extending this filtering approach to motion estimation from functional MRI images.

Augmenting the sparse residual optimization problem with the Kalman filter prediction model yields our proposed method:

$$\{\boldsymbol{x}_{\text{meas}}^t, \boldsymbol{a}^t\} \in \underset{\boldsymbol{x},\boldsymbol{a}}{\arg\min} \, \frac{1}{2\sigma^2}\|\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{x} - \boldsymbol{y}^t\|^2 + \lambda\|\boldsymbol{x} - \boldsymbol{T}(\boldsymbol{x}_{\text{reg}}^{t-1}; (\boldsymbol{\alpha}^{t-1})^{(-1)}, \boldsymbol{\alpha}^{t-1} + \boldsymbol{a})\|_1 + \frac{1}{2}\|\boldsymbol{a}\|_{(\boldsymbol{P}_\alpha^{t|t-1})^{-1}}^2, \quad (3)$$

where $\boldsymbol{P}_\alpha^{t|t-1} = \boldsymbol{P}_\alpha^{t-1} + \boldsymbol{Q}$ is the predicted covariance, and $\boldsymbol{P}_\alpha^{t-1}$ is the posterior covariance for the previous frame ($\boldsymbol{P}_\alpha^0 = \boldsymbol{0}$).

To solve this problem efficiently, we linearize the transformation $\boldsymbol{T}(\boldsymbol{x}_{\text{reg}}^{t-1}; (\boldsymbol{\alpha}^{t-1})^{(-1)}, \boldsymbol{\alpha}^{t-1} + \boldsymbol{a})$ around the previous motion estimate ($\boldsymbol{a} = \boldsymbol{0}$), as is done in PROMO:

$$\{\boldsymbol{x}^t, \boldsymbol{a}^t\} \in \underset{\boldsymbol{x},\boldsymbol{a}}{\arg\min} \, \frac{1}{2\sigma^2}\|\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{x} - \boldsymbol{y}^t\|^2 + \lambda\|\boldsymbol{x} - (\boldsymbol{J}_T\boldsymbol{a} + \boldsymbol{x}_{\text{reg}}^{t-1})\|_1 + \frac{1}{2}\|\boldsymbol{a}\|_{(\boldsymbol{P}_\alpha^{t|t-1})^{-1}}^2, \quad (4)$$

where $\boldsymbol{J}_T$ is the Jacobian matrix of $\boldsymbol{T}(\boldsymbol{x}_{\text{reg}}^{t-1}; (\boldsymbol{\alpha}^{t-1})^{(-1)}, \boldsymbol{\alpha}^{t-1} + \boldsymbol{a})$ around $\boldsymbol{a} = \boldsymbol{0}$. Next, we introduce the auxiliary variable $\boldsymbol{s} = \boldsymbol{x} - (\boldsymbol{J}_T\boldsymbol{a} + \boldsymbol{x}_{\text{reg}}^{t-1})$ and solve the resulting constrained problem using the alternating direction method of multipliers (ADMM):[15]

$$\{\boldsymbol{x}_{\text{meas}}^t, \boldsymbol{a}^t, \boldsymbol{s}\} \in \underset{\boldsymbol{x},\boldsymbol{a},\boldsymbol{s}}{\arg\min} \, \frac{1}{2\sigma^2}\|\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{x} - \boldsymbol{y}^t\|^2 + \lambda\|\boldsymbol{s}\|_1 + \frac{1}{2}\|\boldsymbol{a}\|_{(\boldsymbol{P}_\alpha^{t|t-1})^{-1}}^2 + \frac{\mu_s}{2}\|\boldsymbol{x} - \boldsymbol{J}_T\boldsymbol{a} - \boldsymbol{s} - \boldsymbol{x}_{\text{reg}}^{t-1} + \boldsymbol{u}_s\|_2^2$$

$$\text{s.t. } (\boldsymbol{x} - \boldsymbol{J}_T\boldsymbol{a}) - \boldsymbol{s} = \boldsymbol{x}_{\text{reg}}^{t-1}, \quad (5)$$

where $\boldsymbol{u}_s$ and μ_s are the scaled dual vector and ADMM tuning parameter, respectively. The subproblems iterating among solving for $\boldsymbol{x}$ and $\boldsymbol{a}$ jointly, solving for $\boldsymbol{s}$, and updating $\boldsymbol{u}_s$ are:

$$\{\boldsymbol{x}, \boldsymbol{a}\} \leftarrow \underset{\boldsymbol{x},\boldsymbol{a}}{\arg\min} \, \frac{1}{2\sigma^2}\|\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{x} - \boldsymbol{y}^t\|^2 + \frac{1}{2}\|\boldsymbol{a}\|_{(\boldsymbol{P}_\alpha^{t|t-1})^{-1}}^2 + \frac{\mu_s}{2}\|\boldsymbol{x} - \boldsymbol{J}_T\boldsymbol{a} - (\boldsymbol{s} + \boldsymbol{x}_{\text{reg}}^{t-1} - \boldsymbol{u}_s)\|_2^2; \quad (6)$$

$$\boldsymbol{s} \leftarrow \underset{\boldsymbol{s}}{\arg\min} \, \lambda\|\boldsymbol{s}\|_1 + \frac{\mu_s}{2}\|\boldsymbol{s} - (\boldsymbol{x} - \boldsymbol{J}_T\boldsymbol{a} - \boldsymbol{x}_{\text{reg}}^{t-1} + \boldsymbol{u}_s)\|_2^2; \text{ and} \quad (7)$$

$$\boldsymbol{u}_s \leftarrow \boldsymbol{u}_s + (\boldsymbol{x} - \boldsymbol{J}_T\boldsymbol{a} - \boldsymbol{s} - \boldsymbol{x}_{\text{reg}}^{t-1}). \quad (8)$$

The first subproblem is a least-squares problem, so its corresponding normal equations are

$$\begin{bmatrix} \frac{1}{\sigma^2}\boldsymbol{\mathcal{F}}'\boldsymbol{M}'\boldsymbol{M}\boldsymbol{\mathcal{F}} + \mu_s\boldsymbol{I} & -\mu_s\boldsymbol{J}_T \\ -\mu_s\boldsymbol{J}_T' & (\boldsymbol{P}_\alpha^{t|t-1})^{-1} + \mu_s\boldsymbol{J}_T'\boldsymbol{J}_T \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{a} \end{bmatrix} = \begin{bmatrix} \frac{1}{\sigma^2}\boldsymbol{\mathcal{F}}'\boldsymbol{M}'\boldsymbol{y}^t + \mu_s(\boldsymbol{s} + \boldsymbol{x}_{\text{reg}}^{t-1} - \boldsymbol{u}_s) \\ -\mu_s\boldsymbol{J}_T'(\boldsymbol{s} + \boldsymbol{x}_{\text{reg}}^{t-1} - \boldsymbol{u}_s) \end{bmatrix}, \quad (9)$$

where $[\cdot]'$ is the conjugate-transpose operator. Denote the block matrix on the left-hand-side $\left[\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B}' & \boldsymbol{D} \end{smallmatrix}\right]$, and the block vector on the right-hand-side $\left[\begin{smallmatrix} \boldsymbol{f} \\ \boldsymbol{g} \end{smallmatrix}\right]$. Then, since $\boldsymbol{A}$ and $\left[\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B}' & \boldsymbol{D} \end{smallmatrix}\right]$ are invertible, for $\mu_s > 0$, we can compute $\boldsymbol{a} = (\mathcal{R}e\{\boldsymbol{D} - \boldsymbol{B}'\boldsymbol{A}^{-1}\boldsymbol{B}\})^{-1}(\mathcal{R}e\{\boldsymbol{g} - \boldsymbol{B}'\boldsymbol{A}^{-1}\boldsymbol{f}\})$, where $\mathcal{R}e\{\cdot\}$ constitutes the real part of a matrix or vector, and $\boldsymbol{x} = \boldsymbol{A}^{-1}(\boldsymbol{f} - \boldsymbol{B}\boldsymbol{a})$. In general, computing $\boldsymbol{A}^{-1}$ is easily accomplished by taking advantage of its circulant structure:

$$\boldsymbol{A}^{-1} = \boldsymbol{\mathcal{F}}^{-1}(\tfrac{N}{\sigma^2}\boldsymbol{M}'\boldsymbol{M} + \mu_s\boldsymbol{I})^{-1}\boldsymbol{\mathcal{F}}, \quad (10)$$

where N is the number of image voxels in $\boldsymbol{x}$. When $\boldsymbol{M}$ fully samples k-space, this expression simplifies further, to $\frac{1}{(\mu_s + N/\sigma^2)}\boldsymbol{I}$. We plug our expression for $\boldsymbol{A}^{-1}$ into

$$\boldsymbol{D} - \boldsymbol{B}'\boldsymbol{A}^{-1}\boldsymbol{B} = (\boldsymbol{P}_\alpha^{t|t-1})^{-1} + \mu_s\boldsymbol{J}_T'\boldsymbol{J}_T - \mu_s^2(\boldsymbol{\mathcal{F}}\boldsymbol{J}_T)'(\tfrac{N}{\sigma^2}\boldsymbol{M}'\boldsymbol{M} + \mu_s\boldsymbol{I})^{-1}(\boldsymbol{\mathcal{F}}\boldsymbol{J}_T). \quad (11)$$

Since the dimensions are small ($\boldsymbol{D}$ is a six-by-six matrix), direct inversion is possible. It is possible to reduce computations by precomputing $\boldsymbol{\mathcal{F}}\boldsymbol{J}_T$ and $\boldsymbol{\mathcal{F}}'\boldsymbol{M}'\boldsymbol{y}^t$ once per frame, so the only significant computation involves computing the DFT's associated with $\boldsymbol{A}^{-1}\boldsymbol{f}$ and $\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{a}$ once per ADMM iteration. Again, the computation simplifies in the fully sampled case.

The $\boldsymbol{s}$-update consists of soft-thresholding the vector $\boldsymbol{x} - \boldsymbol{J}_T\boldsymbol{a} - \boldsymbol{x}_{\text{reg}}^{t-1} + \boldsymbol{u}_s$ with shrinkage threshold λ/μ_s. Ideally, μ_s would be chosen so that λ/μ_s thresholds an appropriate fraction of that vector's coefficients. In our simulations, we choose μ_s and λ empirically. We initialize the image using the direct reconstruction $\frac{1}{N}\boldsymbol{\mathcal{F}}'\boldsymbol{M}'\boldsymbol{y}^t$
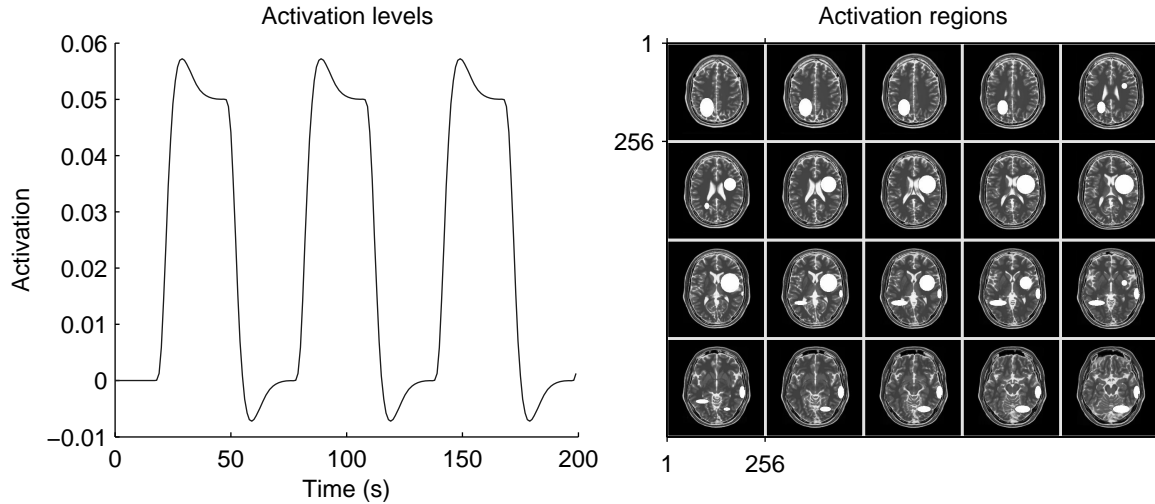
Figure 1. The activation pattern was generated by convolving a rough hemodynamic response function (area 0.05) with a block design with a 60 second period and a 50% duty cycle. The 20 (sampled) axial slices of the high-resolution T2-weighted BrainWeb phantom contain parts of five ellipsoidal activation regions (white) that share this activation pattern.

and start with $\boldsymbol{a} = \boldsymbol{0}$. We also begin with $\boldsymbol{s}$ set to $\boldsymbol{x} - \boldsymbol{J}_T \boldsymbol{a} - \boldsymbol{x}_{\text{reg}}^{t-1}$ and $\boldsymbol{u}_s = \boldsymbol{0}$. We then iterate through solving the subproblems until the next time frame is ready to begin.

Once the iterations are complete, we update the posterior covariance $\boldsymbol{P}_\alpha^t$. However, finding the exact posterior covariance for the nonlinear estimator $\widehat{\boldsymbol{\alpha}}^t$ is challenging. Instead, we approximate the posterior covariance with that of the linear Kalman filter with the sum-of-squared-differences registration cost function $\frac{1}{2\sigma^2}\|\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{J}_T \boldsymbol{a} - (\boldsymbol{y}^t - \boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{x}_{\text{reg}}^{t-1})\|_2^2$. So $(\boldsymbol{P}_\alpha^t)^{-1} = (\boldsymbol{P}_\alpha^{t|t-1})^{-1} + \frac{1}{\sigma^2}(\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{J}_T)'(\boldsymbol{M}\boldsymbol{\mathcal{F}}\boldsymbol{J}_T)$, where $\boldsymbol{x}_{\text{reg}}^{t-1}$ and $\boldsymbol{J}_T$ are treated as fixed (nonrandom) quantities, ignoring their dependence on past values of $\boldsymbol{\alpha}$.

3. METHODS

To evaluate our method, we simulated sparse activations on a T2-weighted normal BrainWeb phantom[16] available online at http://www.bic.mni.mcgill.ca/brainweb/. The phantom volume initially measured $181 \times 217 \times 60$ voxels with resolution $1 \times 1 \times 3$ mm. The volume was cropped at the very top of the head and padded in the axial plane to measure $256 \times 256 \times 50$ voxels with the same high resolution.

The sparse activations were simulated using five ellipsoids placed arbitrarily in the brain sharing a single activation pattern characterized by the convolution of a rough hemodynamic response function with a block design with 30 seconds on-task and a 60 second cycle. We resampled the canonical hemodynamic response function[17] included in SPM8, which is available online at http://www.fil.ion.ucl.ac.uk/spm/. The amplitude of the activation pattern was chosen to be 5% of the maximum anatomical amplitude. The activation time series and ellipsoidal activation regions are shown in Fig. 1.

We generated rigid-body frame motion according to the first-order random walk process described in Section 2. During each time frame, we simulated iid zero-mean Gaussian innovations with variance 0.125^2 units^2 per second, where translational units are mm, and rotation is measured with respect to the major axes, in degrees. We enforced a cap of ± 5 units maximum motion. Figure 2 depicts the generated true motion parameters for 200 seconds of simulation time.

Before sampling, the high-resolution volume was transformed to measurement coordinates using b-spline interpolation. We then generated low-resolution k-space data for the same twenty 3 mm axial slices of each time frame using a single-shot echo planar imaging (EPI) trajectory (full field of view, 64×64 matrix size). We added iid zero-mean additive complex Gaussian noise with 40 dB SNR. We did not consider EPI ghosting

Figure 2. Six-parameter rigid-body motion generated by a Gaussian-stepping first-order random walk. These motion parameters were the true motion used in the simulation. The translations Δx, Δy, and Δz are in the right-left, anterior-posterior, and superior-inferior directions, respectively. The major-axis rotations α, β, and γ are in the axial, coronal, and sagital planes, respectively.

effects, intra-frame motion, spin saturation, field inhomogeneity, and respiratory and cardiac noise to simplify our simulation. We collected 200 frames with a repetition time (TR) of 1.0 s, corresponding to 50 ms per slice. Real and imaginary components of noisy data were quantized to 16 bits and "transmitted" in real time to the reconstruction code written in Matlab running on the same machine. We did not model network latency effects, which we would expect to be minimal, since the reconstruction machine would be directly connected to the scanner backplane.

Our preliminary experiments compare the time series correlations to the expected task-based activations with and without prospective motion correction (using the proposed method). We avoid using a general linear model with motion estimates as regressors because of the potential for task-related motion being correlated with our stimulus. We examine the correlation maps before and after retrospective realignment and re-slicing of the acquired time series using the standard tool found in SPM8 (run with default parameters). We compute the time series correlation coefficient r between an image $\boldsymbol{x}$ at voxel n and the activation sequence $a^0, \ldots, a^{N_F-1}$ using the standard formula

$$r[n] = \frac{\frac{1}{N_F-1}\sum_{t=0}^{N_F-1}(\boldsymbol{x}^t[n] - \bar{\boldsymbol{x}}[n])(a^t - \bar{a})}{s_{\boldsymbol{x}}[n]s_a}, \tag{12}$$

where $\bar{\boldsymbol{x}}$ and $\bar{a}$ are the sample means of the image and activation sequence over time, and $s_{\boldsymbol{x}}$ and s_a are the (unbiased) sample standard deviations of the same.

4. RESULTS

First, we demonstrate the effectiveness of our real time motion estimation algorithm using the sparse residual and Kalman filter tracking. Figure 3 shows that the motion estimate errors from our proposed prospective correction method have been reduced to within ±0.3 units for most frames, nearly an order of magnitude lower than the uncorrected motion. The axial-plane motion estimates (Δx, Δy, and α) appear to have smaller errors than the other motion, with errors of less than ±0.1 units.

The mean absolute residual translational motion for the prospectively corrected data is only 0.220 mm (for comparison, a voxel measures $4 \times 4 \times 3$ mm), versus 1.24 mm total translational motion for the simulated acquisition without prospective correction. The mean absolute residual rotation angle (in axis-angle form) is only

Figure 3. Motion errors after prospective motion correction. The translational motion errors are shown in black lines, and the rotational motion errors are in gray lines.

0.264 degrees after prospective correction, versus 1.73 degrees without it. The residual motion after prospective correction is more than five times smaller than the overall motion, on average.

In the context of functional MRI, one uses such time series acquisitions to measure task-related correlations between voxels in the brain and the known block design, identifying activation regions like those shown in the ideal map in Fig. 1. The low-resolution reconstructed images portrayed in Fig. 4 for the 100th time frame all appear very similar with or without prospective correction, and with or without retrospective registration. However, small differences correlated with the stimulus can have a greater impact on our statistical analysis. First, we plot the correlation coefficients for the acquired reconstructed images with and without prospective motion correction, before any further processing of the data. At this stage, no further motion correction has been performed, so both time series are misaligned, although the misalignment is expected to be minimal for the prospectively corrected data. Figure 5 displays the correlation coefficient maps. Note how clean the activation regions are in the prospectively corrected image. The errors due to motion are much more noticeable in the map without prospective correction, especially in the difference images.

To address the effects of residual motion, we process both simulated acquisitions using the standard tools for retrospective rigid motion correction found in SPM8. The spline-based interpolator introduces local spatial smoothing among the realigned time frames, and we observe the effects in increased background correlations in both data sets. The correlation coefficients for the realigned and re-sliced time series are portrayed in Fig. 6. The outermost slices were masked out by the re-slicing process, since no additional padding slices were acquired in our simulation. The quality of the activation regions signified by bright ellipses in both processed data sets nearly matches the original prospectively corrected data, but the background correlations are much more noticeable, especially in the case without prospective motion correction.

To evaluate the quality of these correlation maps quantitatively, we compute receiver operator characteristic (ROC) curves in Fig. 7, by computing sensitivity and specificity values for absolute value correlation coefficient thresholds from zero to one. The sensitivity and specificity rates exclude the first and last slices in the retrospectively corrected cases to avoid bias from missing data. The sensitivity represents the fraction of correctly detected activations, and the false alarm rate represents the fraction of inactive voxels incorrectly marked as active. Aiming for the top-left corner of the plot, we observe the prospectively corrected reconstructions yield better ROC curves than their uncorrected counterparts. Retrospective registration appears to improve the sensitivity of the correlation statistic, although not as effectively as prospective correction.
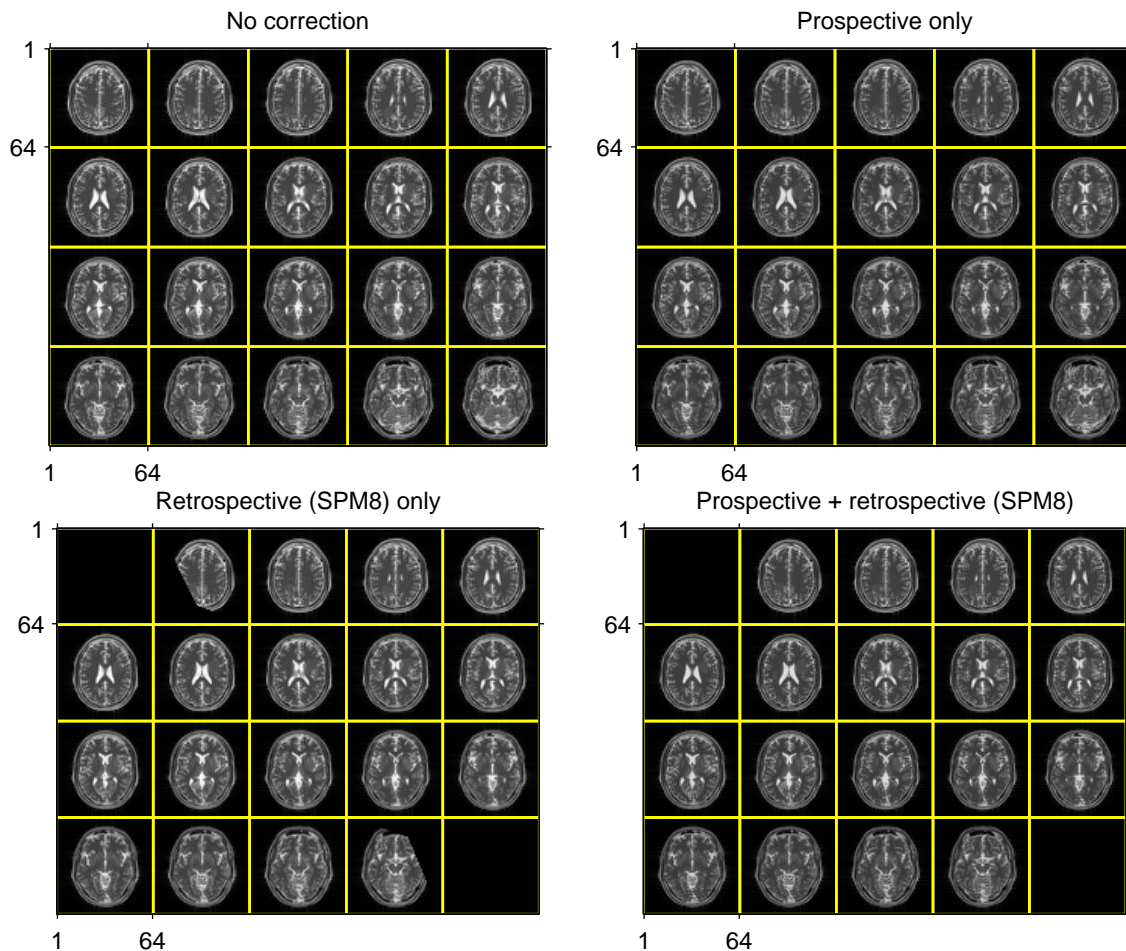
Figure 4. Reconstructed images are shown with and without prospective correction, before and after retrospective registration (using SPM8). Note that the retrospective processing masks our outermost slices since we did not acquire additional slices for padding in our simulation.

5. DISCUSSION

The proposed joint image reconstruction and motion estimation method combines both a Kalman-filter-like motion tracking component and a sparse residual image registration cost function to yield high-quality motion estimates suitable for prospective motion correction on a frame-by-frame basis. The quality of the time series correlation coefficients estimated from the prospectively corrected data exceeds that of the retrospectively registered data. The cause of this improvement is most likely the reduced range of residual motion experienced after prospective correction, as reported by the SPM8 retrospective rigid-body registrations. We emphasize that this high quality was achieved in real (simulated) time, while the slices were being "acquired" by the simulation and transferred to the reconstruction and motion estimation software, without using navigator sequences or anything else that would inflate the repetition time of the time series.

The proposed method has multiple limitations that we plan to address in the near term. We acquired fully-sampled data in our simulations, so our reconstructions did not require regularization. Using regularization would require performing DFT's in each iteration of the proposed algorithm, slowing the method slightly. Also, the tracking component of our algorithm is highly sensitive to the estimate of the motion process' innovation variance. In reality, we do not know this variance a priori, and we can expect the variance to change over time. We aim to adjust our algorithm to estimate and adaptively tune our innovation variance based on the motion estimates actually observed. Also, the ADMM penalty parameter and sparsity parameter are set manually, which
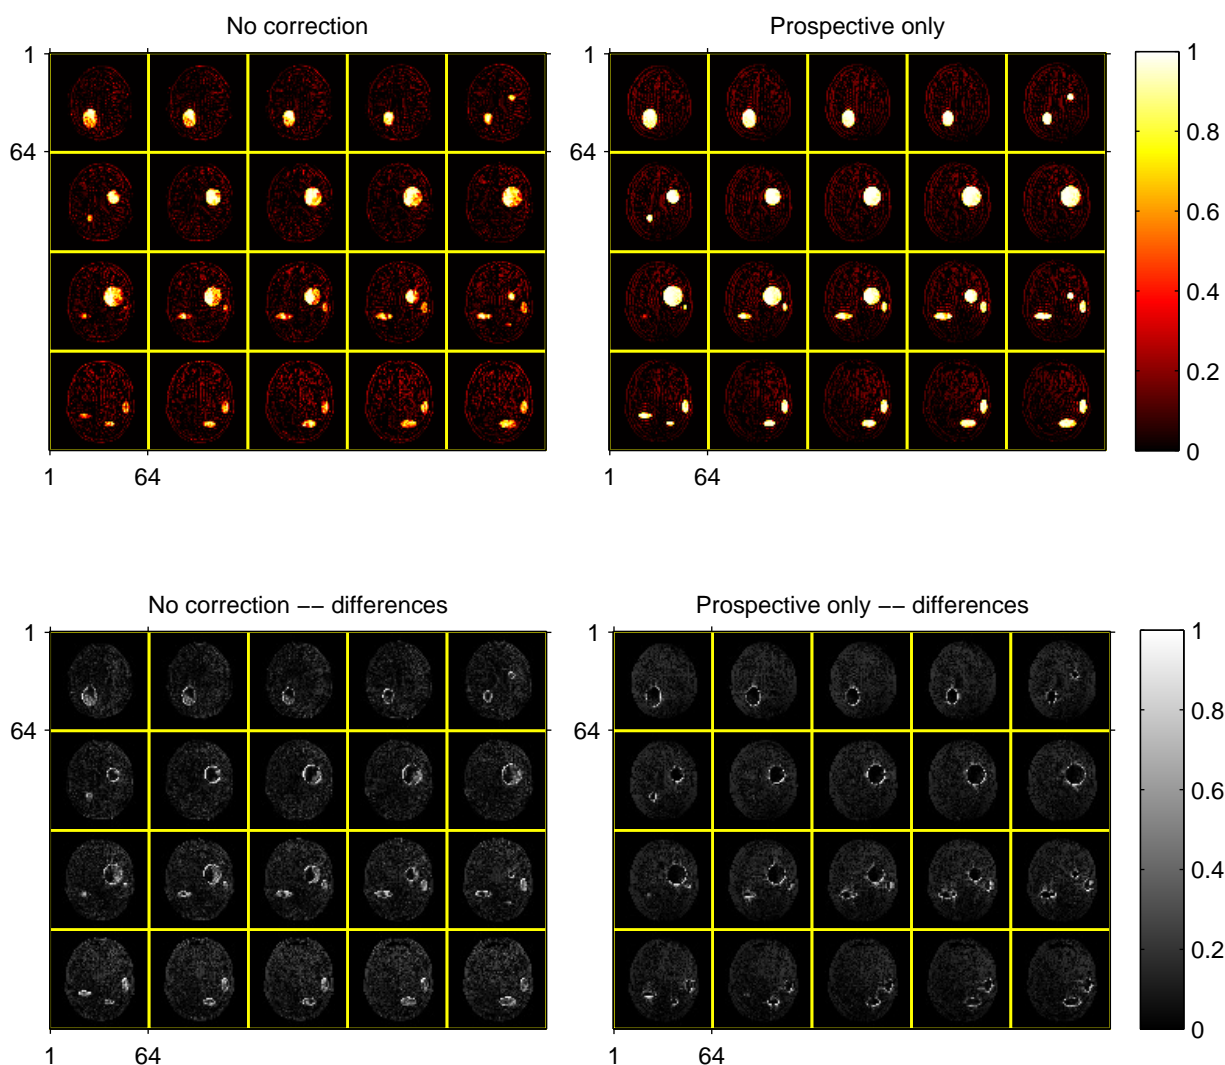
Figure 5. Correlation coefficients for the 200-frame time series reconstructed with and without prospective motion correction, before retrospective processing. Absolute difference images relative to the ideal activation mask also are shown.

may not be optimal in practice. We plan to experiment with automatic tuning of the sparsity parameter using Stein's unbiased risk estimate[18] to predict the error in the motion estimates.

Further developments and applications of the proposed method include extending the measurement model to non-Cartesian sampling (such as spiral) and to accelerated parallel imaging with multiple receiver channels of data processed simultaneously. We also expect that our method can improve further if time-consuming components were optimized for parallel processing, since we could perform more iterations within a single TR. Reducing the computational cost of steps like computing the Jacobian matrix $\boldsymbol{J}_T$ would enable prospective motion compensation within time frames, which could alleviate spin-saturation effects.
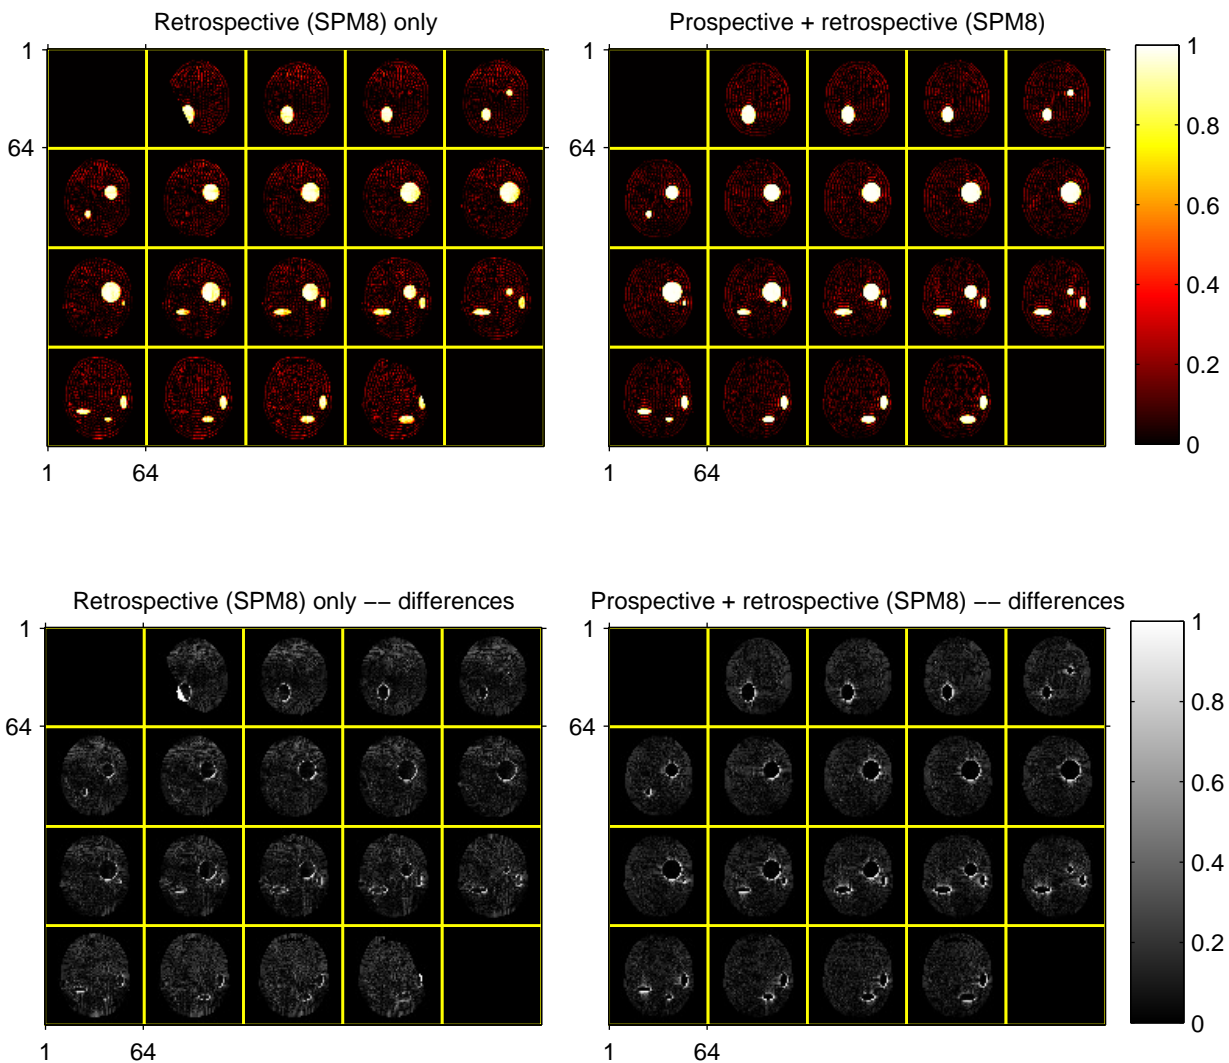
ACKNOWLEDGMENTS

Figure 6. Correlation coefficients for the 200-frame time series reconstructed with and without prospective motion correction, after retrospective processing using SPM8. Absolute difference images relative to the ideal activation mask also are shown. The retrospective registration and re-slicing masks out most of the outermost slices since we did not acquire additional slices for padding purposes.

REFERENCES

[1] Hajnal, J. V., Myers, R., Oatridge, A., Schwieso, J. E., Young, I. R., and Bydder, G. M., "Artifacts due to stimulus correlated motion in functional imaging of the brain," *Mag. Res. Med.* **31**, 283–91 (Mar. 1994).

[2] Friston, K. J., Ashburner, J., Frith, C. D., Poline, J. B., Heather, J. D., and Frackowiak, R. S. J., "Spatial registration and normalization of images," *Hum. Brain Map.* **3**(3), 165–89 (1995).

[3] Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S. J., and Turner, R., "Movement-related effects in fMRI time-series," *Mag. Res. Med.* **35**, 346–55 (Mar. 1996).

[4] Oakes, T. R., Johnstone, T., Ores Walsh, K. S., Greischar, L. L., Alexander, A. L., Fox, A. S., and Davidson, R. J., "Comparison of fMRI motion correction software tools," *NeuroImage* **28**(3), 529–43 (2005).

[5] Bhagalia, R. and Kim, B., "Spin saturation artifact correction using slice-to-volume registration motion estimates for fMRI time series," *Med. Phys.* **35**, 424–34 (Feb. 2008).
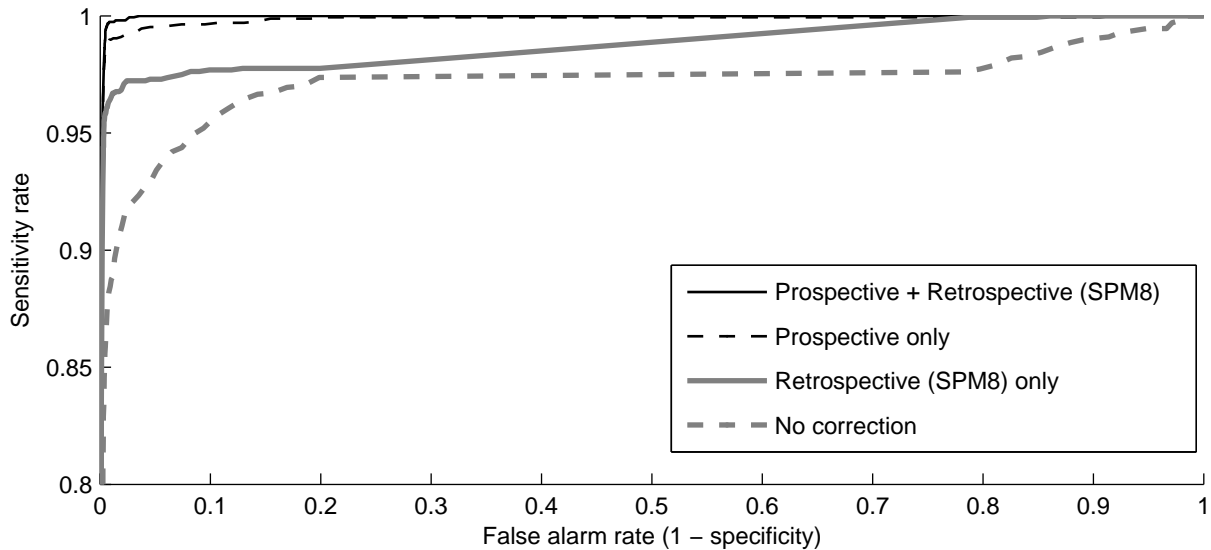
Figure 7. Receiver operator characteristic curves depicting correct detections (as activations), known as sensitivity, versus false alarm rates for reconstructions with and without prospective motion correction (using the proposed method) and with and without retrospective registration (using SPM8). For clarity, the sensitivity axis is zoomed to [0.8, 1].

[6] Maclaren, J., Herbst, M., Speck, O., and Zaitsev, M., "Prospective motion correction in brain imaging: A review," *Mag. Res. Med.* **69**(3), 621–36 (2013).

[7] Zaitsev, M., Dold, C., Sakas, G., Hennig, J., and Speck, O., "Magnetic resonance imaging of freely moving objects: Prospective real-time motion correction using an external optical motion tracking system," *NeuroImage* **31**(3), 1038–50 (2006).

[8] Thesen, S., Heid, O., Mueller, E., and Schad, L. R., "Prospective acquisition correction for head motion with image-based tracking for real-time fMRI," *Mag. Res. Med.* **44**, 457–65 (Sept. 2000).

[9] White, N., Roddey, C., Shankaranarayanan, A., Han, E., Rettmann, D., Santos, J., Kuperman, J., and Dale, A., "PROMO: Real-time prospective motion correction in MRI using image-based tracking," *Mag. Res. Med.* **63**, 91–105 (Jan. 2010).

[10] Jung, H., Sung, K., Nayak, K. S., Kim, E. Y., and Ye, J. C., "k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI," *Mag. Res. Med.* **61**, 103–16 (Jan. 2009).

[11] Asif, M. S., Hamilton, L., Brummer, M., and Romberg, J., "Motion-adaptive spatio-temporal regularization for accelerated dynamic MRI," *Mag. Res. Med.* (2013).

[12] Unser, M., Aldroubi, A., and Eden, M., "B-spline signal processing: Part I—theory," *IEEE Trans. Sig. Proc.* **41**, 821–33 (Feb. 1993).

[13] Unser, M., Aldroubi, A., and Eden, M., "B-spline signal processing: Part II—efficient design and applications," *IEEE Trans. Sig. Proc.* **41**, 834–48 (Feb. 1993).

[14] Fessler, J. A., "Model-based image reconstruction for MRI," *IEEE Sig. Proc. Mag.* **27**, 81–9 (July 2010). Invited submission to special issue on medical imaging.

[15] Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J., "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. & Trends in Machine Learning* **3**(1), 1–122 (2010).

[16] Kwan, R. K.-S., Evans, A. C., and Pike, G. B., "MRI simulation-based evaluation of image-processing and classification methods," *IEEE Trans. Med. Imag.* **18**, 1085–97 (Nov. 1999).

[17] Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M. D., and Turner, R., "Event-related fMRI: Characterizing differential responses," *NeuroImage* **7**(1), 30–40 (1998).

[18] Stein, C., "Estimation of the mean of a multivariate normal distribution," *Ann. Stat.* **9**, 1135–51 (Nov. 1981).