Published in final edited form as:

Signal Processing. 2019 April ; 157: 170–179. doi:10.1016/j.sigpro.2018.12.001.

Real-Time Filtering with Sparse Variations for Head Motion in Magnetic Resonance Imaging

Daniel S. Weller^{a,*}, Douglas C. Noll^b, and Jeffrey A. Fessler^b

^aUniversity of Virginia, Charlottesville, VA, USA 22904

^bUniversity of Michigan, Ann Arbor, MI, USA 48109

Abstract

Estimating a time-varying signal, such as head motion from magnetic resonance imaging data, becomes particularly challenging in the face of other temporal dynamics such as functional activation. This paper describes a new Kalman filter-like framework that includes a sparse residual term in the measurement model. This additional term allows the extended Kalman filter to generate real-time motion estimates suitable for prospective motion correction when such dynamics occur. An iterative augmented Lagrangian algorithm similar to the alternating direction method of multipliers implements the update step for this Kalman filter. This paper evaluates the accuracy and convergence rate of this iterative method for small and large motion in terms of its sensitivity to parameter selection. The included experiment on a simulated functional magnetic resonance imaging acquisition demonstrates that the resulting method improves the maximum Youden's J index of the time series analysis by 2–3% versus retrospective motion correction, while the sensitivity index increases from 4.3 to 5.4 when combining prospective and retrospective correction.

Keywords

image processing; registration; Kalman filtering; sparsity; magnetic resonance imaging

1. Introduction

Prospective head motion correction [1–5] is an example of real-time estimation. Classical approaches to real-time motion tracking frequently involve a linearized (extended or unscented) Kalman filter and assume time series variations follow a Gauss-Markov model. This model efficiently shares information through time, favoring smooth transitions consistent with known (or estimated) autocorrelation matrices for both measurement noise and the random process of innovations influencing state transitions. In magnetic resonance

*Corresponding author: dsw8c@virginia.edu (Daniel S. Weller), dnoll@umich.edu (Douglas C. Noll), fessler@umich.edu (Jeffrey A. Fessler).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Declarations of interest: none.

imaging (MRI), PROMO [6] is a widely used example of such a motion correction framework, applying an extended Kalman filter to track a state vector composed of rigid motion (translation and rotation) measured from acquired “navigator data”. In addition to improving the statistical reliability of time series analysis in functional magnetic resonance imaging [7], prospective motion correction methods aid volumetric morphometry and other forms of quantitation on very high resolution brain maps [8–10] acquired in vivo. Unlike retrospective methods such as image registration, prospective motion tracking mitigates significant higher-order effects on the acquisition, such as spin history variation caused by motion disrupting the steady state between MRI excitation and relaxation. Accurate and precise motion tracking is necessary in such applications. However, real-time estimation of motion from the raw frequency-domain data (“k-space”) used to obtain functional or dynamic MRI’s is complicated by overlapping variations from motion and other physiological signals.

1.1. Our Contribution

This paper describes a new Kalman-filter-like real-time framework based on a combination of sparse and linearized Gauss-Markov signal and measurement models. While obtaining k-space data, applying this approach facilitates nearly continuous tracking of rigid motion in three dimensions. An appropriate real-time feedback mechanism in an MRI scanner can update the scan prescription over time to follow the estimated motion, mitigating both misalignments and higher-order distortions of the time series signal. We simulate our proposed motion correction and demonstrate improvement in sensitivity and specificity over existing registration corrections when detecting simulated brain activity.

This paper extends our previous work [11] by introducing slice-by-slice motion correction, by studying parameter selection effects on convergence of our iterative estimation algorithm, and by accounting for spin history and other effects encountered in real MRI scanning environments.

1.2. Related Work

See Refs. [1, 4, 12–14] for a comprehensive review of motion correction and image registration in MRI. The linearized motion model used in prospective correction appears in both PACE [15] and PROMO [6]. PACE estimates motion directly from functional images, while PROMO forms Kalman-filter-like estimates from special navigator data interspersed throughout the acquisition. Our technique combines PACE and PROMO, correcting motion directly from functional images, and using a real-time Kalman-filter-like design. Thus, additional navigator sequences are not needed to use our algorithm. Also, our motion correction algorithm improves over PACE by estimating motion continuously, rather than having to wait to obtain a complete volume. To accomplish this, we model temporally sparse changes in the image time series, enabling us to separate motion-related changes from dynamics such as functional activation or blood perfusion. Some retrospective motion correction techniques like k-t FOCUSS [16] and MASTeR [17] employ a similar sparse time residual, but they process the entire time series post-acquisition to suppress non-rigid motion via a block matching approach without any motion model. Unlike those techniques, and other recent sparsity-based reconstruction approaches, the proposed method is designed to

correct motion during the acquisition rather than after the fact. As the quality of the motion estimates matters more than the quality of the intermediate images reconstructed in this prospective setting, we do not consider accelerated high-resolution imaging techniques like parallel imaging or nonlinear or data-driven time series modeling approaches.

Another approach would treat the image at each time frame as a dynamic state and rigid motion as unknown parameters of the state transition model. The nonlinear dynamical system would include an extra sparse, non-Gaussian state innovation term to capture unknown temporal dynamics. A linearized Kalman filter can recursively update the motion parameters and the image state estimates, assuming the process variance is fixed [18], even with time delay in the measurements [19] or autoregressive moving average (ARMA) process noise [20]. The time delay is related to slice-by-slice motion estimation, but our approach must handle different time delays for each slice. Beyond Kalman filtering, other examples include nonlinear Wiener filtering [21] and H_∞ filters that constrain worst-case performance [22] and can handle randomly occurring faults, modeled as outliers [23]. Alternatively, after linearizing the effect of motion on the system, the problem of estimating both motion and image states becomes bilinear, leading to maximum likelihood, hierarchical identification, and stochastic gradient methods [24–26]. The proposed framework is distinct from all these approaches, as we do not impose a stochastic model on the image time series, only on the rigid motion parameters.

While this work describes the first addition of a sparse image difference model with a linearized Gauss-Markov model for motion state estimation, the idea of integrating sparsity and Kalman filtering is not new. Some methods replace the Gaussian assumption on the state innovation process with a heavy-tailed sparsity-promoting distribution leading to a 1-norm penalty on the process variation [27, 28]. Those methods [29–34] add sparse modeling of the state vector directly, improving estimation quality when prior information about the state's sparsity is known. That approach can be useful for dynamic or functional MRI reconstruction, where the time series images are known to be compressible via a sparsifying transform [35]. While integrating signal or innovation sparsity with the Kalman filter is essentially solved for when the filter's process autocorrelation matrix is known, online estimation of the process autocorrelation matrix becomes more challenging when accounting for sparse innovations. Various applications, such as tracking asset returns in high frequency market data, motivate new covariance estimation approaches [36]. However, none of these articles consider the role of sparse changes in the images being observed, as distinct from the motion parameters being tracked. Furthermore, adapting the process autocorrelation matrix remains unsolved in such settings.

In this paper, the sparse innovations appear in the measurement model instead of the state transition model, so extensions involving non-Gaussian noise and outliers in the measurements are also of interest. Unscented Kalman filters [37] can adapt to measurement uncertainty, but they cannot absorb sudden significant changes from sparse innovations. Adaptive filters also exist for measurement disturbances combining Gaussian noise and outliers, such as those that arise while measuring the state of pneumatic servo actuators [38]. A modified extended Masreliez-Martin filter is proposed for such nonlinear systems [39]. An important result here is that the method is consistent even when the measurements follow

a non-Gaussian distribution. For missing measurements, a probabilistic distribution on the missing data enables robust estimation in settings like complex sensor networks with real, physical constraints [40]. While these measurement outlier and missing data models are similar to our sparse innovations model, the sparse innovations due to functional brain activity occur in the image domain, not the measurement domain, requiring a different solution. For instance, Monte Carlo methods can adapt state estimates to abruptly changing environments, given a prior model for the sparse innovations [41]. However, a prior on functional activity in brain imaging may not be available.

Treating the motion as unknown parameters of a nonlinear dynamical system, metaheuristic techniques can mitigate the nonconvexity of the overall estimation problem. For instance, differential evolution evaluates several candidate parameter values in parallel [42]. However, these methods would have to incorporate the unknown sparse innovations in the cost function used to compare solutions to the parameter estimation problem. Related heuristics like parallel particle swarming, grey wolf pack-based optimization, or humpback whale-based spiraling optimization over the entire search space, can attempt to solve for these parameters directly, or in combination with conventional Kalman-based techniques [43, 44]. Another such heuristic, the cuckoo search uses a heavy-tailed step length to explore the parameter space more efficiently [45]. The firefly algorithm models the patterns of light flashes used by fireflies to attract mates, and leads to very efficient exploration of this space as well [46]. These meta-heuristic methods can jointly estimate state information and optimize the parameters simultaneously, such as for controlling parallel robotic platforms [47]. An improved metaheuristic search inspired by bats can improve the worst-case performance of such platforms, even over the cuckoo search and other common approaches [48]. Hypothetically, such metaheuristic search algorithms can tune other implementation parameters, such as the regularization or penalty parameters in the proposed motion estimation implementation, even though such parameters are not directly part of the state space model.

As a final alternative, artificial neural networks can learn to produce state estimates from a series of observations, incorporating nonlinearity without having to solve a difficult optimization problem in real time (training can be performed off-line if desired). For instance, a neural network derived for battery charge state estimation automatically adapts to the nonlinearity of the state estimation and can accommodate constraints in combination with a particle swarming algorithm [49]. In other work, a generative convolutional neural network can recover the latent dynamical state of nonlinear brain oscillations in magnetoencephalography (MEG) data without being constrained by the tractability of complicated noise models or a limited set of available training data [50]. Such methods can outperform extended or unscented Kalman smoothers in such applications, but the network complexity, and the amount of training data, would have to grow significantly to accommodate the larger problem size.

1.3. Outline

Section 2 introduces the Kalman filter formulation with sparse signal changes over time and describes a slice-by-slice motion correction framework for the same. Section 3 describes the

simulation design, evaluation criteria, and the algorithm parameters used. Section 4 presents experimental results regarding convergence analysis, parameter selection, motion estimation, and time series analysis. Section 5 discusses these results.

2. Theory

Consider complex vector-valued functions of time $\mathbf{x}(t)$. Without loss of generality, assume the time t are scaled in units of sampling frames, so the time series consists of vectors $\mathbf{x}_0 = \mathbf{x}(0)$, $\mathbf{x}_1 = \mathbf{x}(1)$, and so on. For BOLD functional MRI [51], these vectors represent the brain anatomy with T_2^* -weighted BOLD contrast for each of N_{xy} pixels in N_s slices acquired in a usually interleaved sequence repeatedly over time, so $\mathbf{x}_{n,t}$ contains the pixel values for the n th slice at time t . The spatial Fourier transform domain k-space data obtained via a Cartesian grid sampling pattern such as echo planar imaging (EPI) are related to these images via the 2D discrete Fourier transform (DFT) $\mathbf{F}$:

$$\mathbf{d}_{n,t} = \mathbf{F}\mathbf{x}_{n,t} + \boldsymbol{\xi}_{n,t}, \quad (1)$$

where $\boldsymbol{\xi}_{n,t}$ is a vector of iid circularly symmetric complex Gaussian noise with variance σ^2 . If k-space is undersampled, then the matrix $\mathbf{F}$ would exclude the corresponding rows of the DFT matrix. In general, real MRI acquisitions agree with this noise model, although occasional fluctuations due to radiofrequency interference or obvious spikes at a few individual k-space frequencies do occur. Due to the obviousness of such disturbances in the data, simple thresholding could exclude such disturbances from the DFT matrix used to reconstruct $\mathbf{x}_{n,t}$.

To describe the evolution of these images, we introduce an auxiliary vector $\mathbf{a}_{n,t}$ that describes how the corresponding signal relates to the reference frame, $\mathbf{x}_{:,ref}$, where the “:” notation indicates the vector contains the entire image stack (not just one slice). In the case of tracking head motion, we use three translation and three rotation parameters to describe the rigid motion throughout the brain. The notation $\mathbf{a}^{(-1)}$ describes suitable parameters of the inverse motion transformation, which exists when the motion is rigid. While this is an oversimplification near the neck and base of the cerebellum, it is reasonable in the brain regions that are usually imaged. We describe the process of applying a motion transformation to a volume of stacked slice images as $\mathbf{T}(\mathbf{a})\mathbf{x}_{:,t}$. Then, let $\mathbf{T}_n(\mathbf{a})\mathbf{x}_{:,t}$ denote extracting the n th slice from the transformed volume. In addition to motion, we allow other dynamics to occur over the course of imaging, such as functional brain activations. These dynamics are reflected in a residual image $s_{n,t}$ added to the transformed reference image:

$$\mathbf{x}_{n,t} = \mathbf{T}_n(\mathbf{a}_{n,t})\mathbf{x}_{:,ref} + s_{n,t}. \quad (2)$$

In our slice-by-slice motion correction, slices are acquired at slightly different times, meaning that the motion vectors $\mathbf{a}_{1,t}$, $\mathbf{a}_{2,t}$... can differ a bit. This possibility even applies to

the initial frame, where the first slice is consistent with the reference volume ($\mathbf{x}_{1,0} = \mathbf{x}_{1,\text{ref}}$), but the others follow the model in Eq. (2) with $\mathbf{s}_{n,0} = \mathbf{0}$ by definition.

The motion parameter vectors $\mathbf{a}_{n,t}$ are unknown. To aid estimating the motion, we assume they follow their own state evolution model. A Gauss-Markov model is a reasonable approximation for generating a motion vector $\mathbf{a}$ from the previous vector

$\mathbf{a}_{\text{prev}} : p(\mathbf{a} | \mathbf{a}_{\text{prev}}) = \mathcal{N}(\mathbf{a}_{\text{prev}}, \mathbf{Q})$ with symmetric positive definite covariance matrix $\mathbf{Q}$.

Indexing the slices in the order they are acquired, $\mathbf{a}_{\text{prev}} = \mathbf{a}_{n-1,t}$ for $n > 1$, and $\mathbf{a}_{N_s,t-1}$ for $n = 1$. By the Markov property, we can relate all the motion parameters $\mathbf{a}_{1,t}, \dots, \mathbf{a}_{N_s,t}$ for a time frame to $\mathbf{a}_{N_s,t-1}$ via the conditional distribution:

$$p\left(\begin{bmatrix} \mathbf{a}_{1,t} \\ \vdots \\ \mathbf{a}_{N_s,t} \end{bmatrix} \middle| \mathbf{a}_{N_s,t-1}\right) = \mathcal{N}\left(\begin{bmatrix} \mathbf{1}_{N_s \times 1} \otimes \mathbf{I} \end{bmatrix} \mathbf{a}_{N_s,t-1}, \mathbf{Q} \triangleq \begin{bmatrix} 1 & 1 & \dots & 1 \\ 1 & 2 & \dots & 2 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & \dots & N_s \end{bmatrix} \otimes \mathbf{Q}\right) \quad (3)$$

This enlarged covariance matrix $\mathbf{Q}$ is symmetric positive definite, and its inverse is a block tridiagonal matrix with closed form (for $N_s > 1$)

$$\mathbf{Q}^{-1} = \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & -1 & 1 \end{bmatrix} \otimes (\mathbf{Q}^{-1}). \quad (4)$$

For $N_s = 1$, $\mathbf{Q}^{-1} = \mathbf{Q}^{-1}$. Let $\mathbf{a}_{n,t}$ denote the difference between $\mathbf{a}_{n,t}$ and $\mathbf{a}_{N_s,t-1}$. Combining these models yields the state and measurement equations:

$$\mathbf{a}_{n,t} = \mathbf{a}_{N_s,t-1} + \mathbf{a}_{n,t}, \quad n = 1, \dots, N_s, \quad (5)$$

$$\mathbf{d}_{n,t} = \mathbf{F}\left(\mathbf{T}_n(\mathbf{a}_{n,t}) \mathbf{x}_{:, \text{ref}} + \mathbf{s}_{n,t}\right) + \boldsymbol{\xi}_{n,t}, \quad n = 1, \dots, N_s. \quad (6)$$

A challenge here is that the transformation $\mathbf{T}_n(\mathbf{a})$ is a nonlinear function of $\mathbf{a}$. The sequel describes a linearized Kalman-filter-like implementation to update $\mathbf{a}$ as new k-space measurements arrive, facilitating prospective motion correction.

2.1. Slice-by-Slice Extended Kalman Filter

Here, we linearize the measurement model around the present motion estimate. Call $\mathcal{J}\{T_n(\cdot)\mathbf{x}\}(\boldsymbol{\alpha}_{n,t})$ the Jacobian matrix of first derivatives with respect to $\boldsymbol{\alpha}$ of the transform $T_n(\boldsymbol{\alpha})\mathbf{x}$ for the n th slice, evaluated at $\boldsymbol{\alpha} = \boldsymbol{\alpha}_{n,t}$. For convenience, we write the Jacobian matrix for this slice and time frame as just $\mathbf{J}_{n,t}$. Using this matrix, the first-order Taylor series expansion of $T_n(\boldsymbol{\alpha})\mathbf{x}_{:, \text{ref}}$ around $\boldsymbol{\alpha} = \boldsymbol{\alpha}_{N_s, t-1}$ is

$T_n(\boldsymbol{\alpha}_{N_s, t-1})\mathbf{x}_{:, \text{ref}} + \mathbf{J}_{n,t}(\boldsymbol{\alpha}_{n,t} - \boldsymbol{\alpha}_{N_s, t-1}) + o(\|\boldsymbol{\alpha}_{n,t} - \boldsymbol{\alpha}_{N_s, t-1}\|^2)$. When motion does not change too rapidly, the $o(\|\boldsymbol{\alpha}_{n,t} - \boldsymbol{\alpha}_{N_s, t-1}\|^2)$ term decays to zero. The linearized measurement model becomes

$$\mathbf{d}_{n,t} = F(T_n(\boldsymbol{\alpha}_{N_s, t-1})\mathbf{x}_{:, \text{ref}} + \mathbf{J}_{n,t}(\boldsymbol{\alpha}_{n,t} - \boldsymbol{\alpha}_{N_s, t-1}) + \mathbf{s}_{n,t}) + \boldsymbol{\xi}_{n,t}. \quad (7)$$

Let $\bar{\mathbf{x}}_{n,t-1} = T_n(\boldsymbol{\alpha}_{N_s, t-1})\mathbf{x}_{:, \text{ref}}$ with the appropriate estimate of $\boldsymbol{\alpha}_{N_s, t-1}$. Then, define the measurement residual $\mathbf{r}_{n,t} = \mathbf{d}_{n,t} - F(\bar{\mathbf{x}}_{n,t-1} + \mathbf{s}_{n,t})$ to be the difference between $\mathbf{d}_{n,t}$ and its prediction assuming no additional motion.

Since the measurements (and the images) are complex-valued, but the motion parameter vectors are real-valued, we must account for the real $[\cdot]_R$ and imaginary $[\cdot]_I$ parts of these variables in the Kalman filter expressions that follow. Given the previous state estimate $\hat{\boldsymbol{\alpha}}_{:, t-1|t-1}$ and previous error covariance estimate $\mathbf{P}_{t-1|t-1}$, the Kalman filter prediction step is straightforward: $\hat{\boldsymbol{\alpha}}_{n,t|t-1} = \hat{\boldsymbol{\alpha}}_{N_s, t-1|t-1}$, and

$\mathbf{P}_{t|t-1} = \left(\mathbf{1}_{N_s \times 1} \otimes \mathbf{I} \right) \mathbf{P}_{N_s, t-1|t-1} \left(\mathbf{1}_{N_s \times 1}^T \otimes \mathbf{I} \right) + \mathcal{Q}$, where $\mathbf{P}_{N_s, t-1|t-1}$ is the part of the error covariance for the motion of the last slice (since predictions for all the slices are with respect to that last slice's value).

To compute the Kalman filter update steps, following the incorporation of new data $\mathbf{d}_{:,t}$, we first consider $\mathbf{s}_{n,t}$ to be fixed and known. Through some matrix algebra, we have that the motion innovations $\mathbf{a}_{:,t}$ for all slices are estimated as

$$\hat{\mathbf{a}}_{:,t|t} = \left(\mathbf{P}_{t|t-1}^{-1} + \frac{1}{\sigma^2} \left[\mathbf{J}_{:,t}^H \left(\mathbf{I}_{N_s \times N_s} \otimes \mathbf{F}^H \mathbf{F} \right) \mathbf{J}_{:,t} \right]_R \right)^{-1} \cdot \frac{1}{\sigma^2} \left[\mathbf{J}_{:,t}^H \left(\mathbf{I}_{N_s \times N_s} \otimes \mathbf{F}^H \right) \mathbf{r}_{:,t} \right]_R. \quad (8)$$

Then, $\hat{\boldsymbol{\alpha}}_{:,t|t} = \hat{\boldsymbol{\alpha}}_{:,t|t-1} + \hat{\mathbf{a}}_{:,t|t}$. This $\hat{\mathbf{a}}_{:,t|t}$ also solves the least squares problem

$$\hat{\mathbf{a}}_{:,t|t} = \arg \min_{\mathbf{a}} \frac{1}{2\sigma^2} \left\| \left(\mathbf{I}_{N_s \times N_s} \otimes \mathbf{F} \right) \mathbf{J}_{:,t} \mathbf{a} - \mathbf{r}_{:,t} \right\|^2 + \frac{1}{2} \mathbf{a}^T \mathbf{P}_{t|t-1}^{-1} \mathbf{a}. \quad (9)$$

Using some more algebra, the posterior error covariance estimate $\mathbf{P}_{t|t}$ satisfies

$$\mathbf{P}_{t|t}^{-1} = \mathbf{P}_{t|t-1}^{-1} + \frac{1}{\sigma^2} \left[\mathbf{J}_{:,t}^H (\mathbf{I}_{N_s \times N_s} \otimes \mathbf{F}^H \mathbf{F}) \mathbf{J}_{:,t} \right]_R. \quad (10)$$

When the residual $s_{n,t}$ is not known, we impose a prior model on it. For dynamic MRI, a reasonable model is that the residual is sparse, assuming the dynamic changes not due to motion are localized in space and time. Thus, we expand the optimization to solve for both $\hat{\mathbf{a}}_{:,t|t}$ and $\hat{\mathbf{s}}_{:,t}$ together:

$$\begin{aligned} \{\hat{\mathbf{a}}_{:,t|t}, \hat{\mathbf{s}}_{:,t}\} = \arg \min_{\mathbf{a}, \mathbf{s}} \Psi(\mathbf{a}, \mathbf{s}) \triangleq & \frac{1}{2\sigma^2} \left\| \left(\mathbf{I}_{N_s \times N_s} \otimes \mathbf{F} \right) \left(\mathbf{J}_{:,t} \mathbf{a} + \mathbf{s} + \bar{\mathbf{x}}_{:,t-1} \right) - \mathbf{d}_{:,t} \right\|^2 \\ & + \frac{1}{2} \mathbf{a}^T \mathbf{P}_{t|t-1}^{-1} \mathbf{a} + \lambda \|\mathbf{s}\|_1. \end{aligned} \quad (11)$$

This convex optimization problem is the key component of our sparse Kalmanfilter-like algorithm. While one may solve for $\mathbf{a}$ in terms of $\mathbf{s}$ and plug back into Eq. (11), the resulting 1-norm-regularized least squares problem involves time-consuming matrix operations. We propose a variable-splitting approach that separates and simplifies the updates for $\mathbf{a}$ and $\mathbf{s}$. Since the estimate of $\mathbf{a}_{:,t}$ is no longer linear, the posterior error covariance estimate $\mathbf{P}_{t|t}$ becomes a loose approximation of the true posterior error covariance.

2.2. Initial Conditions and Covariance Estimation

The Kalman filter framework requires a reasonable estimate for $\mathbf{a}_{:, -1}$ and $\mathbf{P}_{-1/-1}$ consistent with the initial measurement information to perform the initial prediction step. We designate the initial head position as the reference $\mathbf{x}_{:, \text{ref}}$, so $\hat{\mathbf{a}}_{:, -1} = \mathbf{0}$, and $\mathbf{P}_{-1/-1} = \mathbf{0}$. The initial image $\mathbf{x}_{:, \text{ref}}$ can be obtained directly from the initial measurements without considering motion during this time. Thus, the initial conditions are known for this motion estimation problem.

Furthermore, the process and measurement noise covariances $\mathbf{Q}$ and $\sigma^2 \mathbf{I}$ are not necessarily known a priori. In MRI, we can estimate σ^2 by obtaining some additional measurements with the excitation radiofrequency field turned off. However, the process covariance $\mathbf{Q}$ is much harder to access in practice, so we adapt it from a sliding window of motion estimates. Autocovariance least squares (ALS) [52–54] also would work in the linear case. Recent efforts target the cases of nonlinear measurements [55] and sparse innovations [36]. Instead of adopting highly complex methods not necessarily amenable to rapid processing, we constrain $\mathbf{Q}$ to be diagonal and form simple estimates from the sample variances of the motion innovations $\hat{\mathbf{a}}_{n,t|t}$. This approach ignores cross-correlations across parameters and differences in motion sensitivities for different slices, but prospective correction would dynamically reorient the coordinate system to track the principal direction of the detected motion. So this constraint likely is not a major limitation on the overall algorithm's accuracy.

2.3. AL-Based Implementation

Majorization-minimization [56–58] could solve Eq. (11), but guaranteeing a small Lipschitz constant (see supplementary material) for fast convergence to the solution $\{\hat{\mathbf{a}}, \hat{\mathbf{s}}\}$ is challenging. Instead, we create an auxiliary variable $\mathbf{z}$ for the motion-transformed image and use an augmented Lagrangian (AL) scheme similar to the alternating direction method of multipliers (ADMM) [59–62].

First, let $\mathbf{z}$ denote the unknown image stack at time t : $\mathbf{z} = \bar{\mathbf{x}}_{:,t-1} + \mathbf{J}_{:,t}\mathbf{a} + \mathbf{s}$. We define the following modified augmented Lagrangian:

$$\begin{aligned} \mathcal{L}_A(\mathbf{a}, \mathbf{s}, \mathbf{z}; \mathbf{u}) \triangleq & \frac{1}{2\sigma^2} \|(\mathbf{I} \otimes \mathbf{F})\mathbf{z} - \mathbf{d}_{:,t}\|_2^2 + \frac{1}{2} \mathbf{a}^T \mathbf{P}_{t|t-1}^{-1} \mathbf{a} + \lambda \|\mathbf{s}\|_1 + \frac{\mu}{2} \|(\mathbf{J}_{:,t}\mathbf{a} + \mathbf{s} + \bar{\mathbf{x}}_{:,t-1}) \\ & - \mathbf{z} + \mathbf{u}\|_2^2, \end{aligned} \quad (12)$$

where $\mathbf{u}$ and μ are the scaled Lagrange dual vector and augmented Lagrangian penalty parameter, respectively.

Alternating minimizing $\mathbf{a}$, $\mathbf{s}$, and $\mathbf{z}$ would correspond to ADMM. Instead, we update $\mathbf{a}$ and $\mathbf{z}$ jointly, and separately from $\mathbf{s}$:

$$\mathbf{s}^{i+1} \leftarrow \arg \min_{\mathbf{s}} \lambda \|\mathbf{s}\|_1 + \frac{\mu}{2} \|(\mathbf{J}_{:,t}\mathbf{a}^i + \mathbf{s}) - \mathbf{z}^i + \bar{\mathbf{x}}_{:,t-1} + \mathbf{u}^i\|_2^2, \quad (13)$$

$$\begin{aligned} \{\mathbf{a}^{i+1}, \mathbf{z}^{i+1}\} \leftarrow \arg \min_{\mathbf{a}, \mathbf{z}} & \frac{1}{2\sigma^2} \|(\mathbf{I} \otimes \mathbf{F})\mathbf{z} - \mathbf{d}_{:,t}\|_2^2 + \frac{1}{2} \mathbf{a}^T \mathbf{P}_{t|t-1}^{-1} \mathbf{a} + \frac{\mu}{2} \|(\mathbf{J}_{:,t}\mathbf{a} + \mathbf{s}^{i+1}) - \mathbf{z} \\ & + \bar{\mathbf{x}}_{:,t-1} + \mathbf{u}^i\|_2^2, \end{aligned} \quad (14)$$

$$\mathbf{u}^{i+1} \leftarrow \mathbf{u}^i + (\mathbf{J}_{:,t}\mathbf{a}^{i+1} + \mathbf{s}^{i+1}) - \mathbf{z}^{i+1} + \bar{\mathbf{x}}_{:,t-1}. \quad (15)$$

Updating $\mathbf{s}$ soft-thresholds $\mathbf{z}^i - \bar{\mathbf{x}}_{:,t-1} - \mathbf{J}_{:,t}\mathbf{a}^i - \mathbf{u}^i$, which element-wise shrinks the complex-valued argument (call it s_n) according to the threshold $\frac{\lambda}{\mu}$: $s_n \leftarrow \text{sign}(s_n) \cdot \max\{0, |s_n| - \frac{\lambda}{\mu}\}$. The normal equations from Eq. (14) yield $\mathbf{a}$ and $\mathbf{z}$:

$$\begin{aligned}
& \left(\mathbf{P}_{t|t-1}^{-1} + \mu [\mathbf{J}_{:,t}^H \mathbf{J}_{:,t}]_R \right) \mathbf{a} - \mu [\mathbf{J}_{:,t}^H \mathbf{z}]_R = -\mu [\mathbf{J}_{:,t}^H (\bar{\mathbf{x}}_{:,t-1} + \mathbf{u}^i + \mathbf{s}^{i+1})]_R, \\
& -\mu \mathbf{J}_{:,t} \mathbf{a} + \left[\frac{1}{\sigma^2} (\mathbf{I} \otimes (\mathbf{F}^H \mathbf{F})) + \mu \mathbf{I} \right] \mathbf{z} = \frac{1}{\sigma^2} (\mathbf{I} \otimes \mathbf{F}^H) \mathbf{d}_{:,t} + \mu (\bar{\mathbf{x}}_{:,t-1} + \mathbf{u}^i + \mathbf{s}^{i+1}). \quad (16)
\end{aligned}$$

Schur complements solve this system of equations efficiently (see supplementary material). The final step of our algorithm updates $\mathbf{u}$ via vector addition.

As the problem scales in the image size N_{xy} , number of slices N_s , or number of frames, the update steps of our ADMM-like implementation remain reasonable computationally. As N_{xy} increases, the DFT operation grows like $N_{xy} \log N_{xy}$, and all the other operations involved in solving the normal equations in Eq. (16) grow linearly. Similarly, the matrix-vector inversions scale linearly with N_s , since the matrices involved are either block diagonal or block tridiagonal. Matrix-vector products involving $\mathbf{I} \otimes \mathbf{F}$ or $\mathbf{J}_{:,t}$ also scale linearly in N_s , as they are either stacked or block diagonal matrices. Finally, the computational and storage complexity remains constant as the number of time frames grows, as opposed to methods that process the entire time series jointly.

2.4. Initialization and Parameter Selection

The convergence rate of the ADMM-like method proposed here is affected by both the initialization and the choice of penalty parameter μ . We initialize $\mathbf{a}^0 = \mathbf{0}$ by setting the motion for the current slice is the same as the previous one. To initialize $\mathbf{z}$ and $\mathbf{s}$, we perform a direct reconstruction of the data $\mathbf{d}_{:,t}$ to form $\mathbf{z}^0$, and assign the residual image $\mathbf{z}^0 - \bar{\mathbf{x}}_{:,t-1} - \mathbf{J}_{:,t} \mathbf{a}^0$ to $\mathbf{s}^0$.

The choice of penalty parameter μ that leads to fastest possible convergence is expected to vary with the noise and process innovation covariances as well as the choice of λ . AL methods like ADMM converge, even with finite μ [61]. The parameter selection experiment explores how a fixed value of μ can affect convergence for a range of motions. Heuristics for fixing μ based on the numerical conditioning of the quadratic subproblem and the sparse thresholding fraction of the shrinkage subproblem can achieve consistent results across reconstructions [63]. At the same time, adaptive heuristic methods [62] allow μ to vary to balance minimizing the objective function and preserving the equality constraint ($\mathbf{z} = \bar{\mathbf{x}}_{:,t-1} + \mathbf{J}_{:,t} \mathbf{a} + \mathbf{s}$). Optimizing the choice of μ is also feasible for quadratic problems [64]; this method could approximate non-quadratic problems as well. Thus, we can select fixed or adaptive μ automatically.

The other tuning parameter introduced in our implementation, λ , controls the sparsity of the image innovations $\mathbf{s}$, so λ can be considered an implicitly described parameter of the state space model. We propose using a fixed λ during prospective motion correction and adapt the process covariance matrix to abrupt motion changes in real time. Alternatively, automatic parameter selection methods like Stein's unbiased risk estimator [65–67] can help pick λ . Derivative-free metaheuristic search algorithms described in Sec. 1 can optimize λ to make the residual energy in the data consistency term consistent with the expected noise variance

of the data (consistent with the discrepancy principle [68]). However, tuning λ during the motion estimation process would involve calculating and comparing many motion estimates in parallel, so parameter selection ideally would be performed only during the initial stimulus phases.

As an alternative to adjusting λ for large motion, we temporarily scale our estimate of the process innovation covariance $\mathbf{Q}$ by a large number whenever we detect large impulsive motion. Doing so temporarily decouples the last motion estimate from the current one, permitting rapid adjustment to the large motion. To detect this large motion, we use a threshold on the total energy of the latter half of slices s , since those slices are most affected by the new motion. When the motion is consistent with the random walk model, so the change in motion is well-modeled by the extended Kalman filter, the sparse innovations s for those slices should include only the functional activity and some small residual from matching the measurement noise. The 1-norm sparsity penalty suppresses the latter. Both of these are on the order of a few percent of the total signal energy, so a much larger change can only be explained by a failure of the motion model. Thus, a relatively conservative threshold would be a reliable change detector, and a relatively large scaling of $\mathbf{Q}$ would effectively reset the Kalman filter to handle large motion. The process covariance estimation then naturally decreases $\mathbf{Q}$ over time as the motion returns to being consistent with the model.

2.5. Sensing Matrix for Echo Planar Imaging

Echo planar imaging (EPI) is used often for functional MRI, and is highly susceptible to systematic k-space offsets caused by eddy current effects. These produce aliasing artifacts (ghosts) offset by $N_y/2$ in the phase encode direction. These ghosts disrupt motion estimation, making ghost correction necessary in our reconstruction. We estimate the in-plane k-space shifts through a calibration pre-scan [69, 70]. Since prospective motion correction rotates the acquisition readout gradient axis, correcting phase ghosts requires calculating oblique-plane k-space offsets [71]. Shifting the odd and even phase encode lines of our data corrects the ghost artifacts before performing motion estimation. When we include the EPI phase shifts in our forward model, we can still implement the sensing matrix $\mathbf{F}$ efficiently using the fast Fourier transform by treating the odd and even EPI lines as two half transforms.

3. Materials and Methods

To analyze the proposed Kalman-filter-like framework incorporating a sparse residual model in the images, we implemented our method for prospective motion estimation in MATLAB (The Mathworks, Natick, MA), using the Michigan Image Reconstruction Toolbox, which is available online from <http://web.eecs.umich.edu/~fessler/code/index.html>, along with MEX files previously developed for B-spline-based interpolation [72, 73]. We first evaluate the effect of varying the sparse regularization parameter λ and including $\mathcal{Q}$ -adaptation. Then, we vary the AL penalty parameter μ and measure how the objective function convergence changes. Our real-time experiments simulate a realistic BOLD functional MRI acquisition described next.

3.1. Simulated Data

Our simulation of BOLD functional MR imaging uses a high-resolution T_2^* -weighted Brainweb phantom [74], shown in Fig. 1 with active regions high-lighted. This phantom contains $181 \times 217 \times 181$ voxels, each 1 mm isotropic resolution, with no noise or inhomogeneity effects. From this volume we extract twenty slices, with a uniform slice profile 3-mm thick, covering a 6 cm slab with no gaps. We also construct a high-resolution T_1 map of this phantom for modeling spin history effects, by tracking tissue-specific longitudinal relaxations, corresponding to the time between successive excitations, of each high-resolution voxel. Our simulated acquisition includes spatial variations in intensity caused by different longitudinal relaxations, or “spin histories.” We interleave the slice order in “bit-reversed” fashion, with greater distance between the initial slices of each frame, to facilitate three-dimensional motion estimation. We simulate a single-shot EPI acquisition, including small phase variations reflecting the presence of eddy current effects. To account for partial volume effects, k-space samples from this high-resolution phantom are reconstructed at a lower resolution, measuring 4×4 mm in-plane.

To this brain phantom, we add five 3D ellipsoidal regions of varying dimensions to represent sets of activated voxels. All of the activated voxels have added to their amplitudes an activation time series (shown in Fig. 2) corresponding to a block design task repeated three times over 200 seconds (200 frames with TR=1 second), with 30 seconds off, and 30 seconds on. This block design is convolved with a canonical hemodynamic response function [75] to emulate the vascular response to activation. These activations produce signal changes of approximately 3% of the maximum image intensity.

We simulate motion for each slice acquisition in every frame using a combination of two motion models, producing variability not expected by our Kalman filter. The first is the random-walk Markov model discussed in Sec. 2, with zero-mean Gaussian innovations with a standard deviation of 0.05 mm or degrees, per second (unknown to our motion estimation system). We add to this random walk impulsive motion with a larger magnitude of 1 mm or degree over a second; impulse times are exponentially distributed with a mean occurrence time of 50 seconds. Each of the six motion parameters are simulated independently. The generated motion trajectories had no more than ± 5 mm or degrees of motion in each direction. When simulating prospective correction, these simulated motions are applied after adjusting the volume coordinates for the motion estimated for the last slice of the previous time frame ($\hat{\alpha}_{N_s, t-1}$).

3.2. Parameter Selection

Before running our complete simulation, we study the effects of choosing λ and μ on the performance of our motion estimation. Recognizing that activation patterns may change between scans, and desiring to determine broadly applicable choices of parameters, we exclude functional activations from the parameter selection tests. For simplicity, we also ignore spin history, EPI phase ghosts, quantization, and other effects for the parameter selection experiments. We run a short time series based on the acquisition strategy with realistic slice-by-slice motion for $\lambda \in \{1, 10, 100, 1000\}$. For each λ , we run our method for

a range of μ 's, retaining the result that achieved the lowest objective function value. To test our adaptive $\mathbf{Q}$ scheme, we scale $\mathbf{Q}$ by a factor of ten whenever the latter half of slices of s on average contain greater than unit energy (our threshold). We aim to show that adapting $\mathbf{Q}$ (with λ held fixed) can provide some of the same advantages of adapting λ in the presence of large, impulsive motion.

Next, we vary the penalty parameter μ to determine a value that yields rapid convergence with both small and large motion. For large impulsive motions, using the adaptive scheme just described would increase $\mathbf{Q}$ by several orders of magnitude, yielding much smaller precision matrices $\mathbf{P}_{t|t-1}^{-1}$. However, we hold μ fixed to simplify real-time implementation.

We run our AL-based method for $\mu \in \{1, 10, 100, 1000, 10^4\}$ and compare the objective function convergence against each other. We aim to show that even a fixed $\mu = 100$ can achieve relatively consistent convergence in both settings (small and large motion).

4. Experimental Results

4.1. Parameter Selection

Figure 3 illustrates the effects of increasing λ when large motion, such as from an impulse, is present. The predicted error covariance $\mathbf{P}_{t|t-1}$ attempts to keep $\mathbf{a}$, the change in motion, small to be consistent with the random walk model. The error due to underestimating the motion appears in the sparse residual s . As λ increases, more of this error appears in the image estimate $\mathbf{x}$ instead of in the sparse residual, but the translational and rotational errors in the estimated motion decrease much more rapidly. Figure 3(c) shows that temporarily increasing $\mathbf{Q}$ achieves the desired effect of reducing the magnitude and duration of errors due to large residual motion without altering λ . Also, this strategy does not introduce additional errors into the image estimates, unlike increasing λ .

Figure 4 plots the objective function $\Psi(\mathbf{a}, s)$, relative to the long-term-optimal value Ψ^* , for our ADMM-like method with the three fastest-converging values of μ . We depict two cases: frame #20, containing only random walk motion, and frame #34, featuring the aftereffects of large impulsive motion. In both cases, the ADMM-like method converges rapidly in objective function value over a wide range of μ , and the AL penalty parameter $\mu = 100$ performs reasonably in both motion settings. Since the computation time is about the same throughout (9 ms/iteration each), plots of objective function convergence versus time would appear similar to those in Fig. 4.

4.2. Real-Time Simulated Experiments

Our real-time experiments jointly simulate BOLD functional MRI acquisitions and prospective motion estimation on two separate MATLAB instances communicating with each other in real time. Figure 5 shows that the motion estimates generated by the slice-by-slice prospective correction algorithm, appear very similar to the true motion. The absolute errors for the translational and rotational motion estimates measure 0.063 ± 0.10 mm and 0.085 ± 0.17 degrees, respectively (mean $\pm$ std. dev.). These translational errors are computed as the 2-norm of the x , y , and z errors together. The rotational motion errors are computed as the magnitude of the angle in the axis-angle representation of the difference

between the true and estimated rotations, where rotations are relative to the center of the 6 cm slab.

At least as important as the accuracy of the motion estimates is the impact of prospective correction using these measurements. Even if these motion estimates are exact, we expect some residual motion to remain since we apply the previous frame's last motion estimate to acquire the current frame. In prospective correction, the residual motion should remain small relative to the overall motion throughout the course of the scan, so effects like spin history should be minimized. In fact, the average residual translation and rotation decreases from 2.2 ± 1.3 mm to 0.11 ± 0.19 mm, and from 1.6 ± 1.0 degrees to 0.13 ± 0.25 degrees, after prospective correction.

The functional MRI analysis that follows compares prospective, retrospective (the standard realign-and-reslice tool in SPM12 available online at <http://www.fil.ion.ucl.ac.uk/spm/>), and both motion corrections against an ideal acquisition without head motion. The time series data for each voxel in the reconstructed/realigned volume are then correlated against the true activation time series in Fig. 2. Figure 6(a)–(d) features slices #9 and #11 with overlaid correlation maps for no motion, retrospective correction, prospective correction, and prospective and retrospective correction combined. These overlays are all thresholded according to a significance threshold of $\alpha = 0.01$, adjusted for multiple comparisons using Bonferroni correction by the number of voxels in the brain. Not shown, the motion-corrupted correlation map without any correction has very low quality. Figure 6(e) quantitatively compares these fMRI time series analyses over all twenty slices using receiver operating characteristic (ROC) curves to track the sensitivity and specificity of each method for detecting functional activations. The ground truth is generated by reducing the resolution of the high-resolution activation mask shown in Fig. 1. We focus on the low false-alarm-rate regime, as activations are usually rare events in task-based functional MRI studies. The area under the curve (AUC) is one summary statistic for quantifying the overall quality of the detection algorithm, but the AUC equally weights all prior probabilities for the true hypothesis, ignoring the scarcity of true activations. Alternatively, another statistical measure of the quality of detection is the maximum Youden's J index p' , which is the maximum difference between sensitivity and false alarm rate. This criterion is related to the sensitivity index d' according to the formula $d' = \Phi^{-1}(\text{sensitivity}) - \Phi^{-1}(1 - \text{specificity})$ for the J-index-maximizing sensitivity and specificity, where $\Phi^{-1}(\cdot)$ is the inverse cumulative density function of the standard Normal distribution. Table 1 lists the AUC, p' , and d' statistics for these ROC curves.

5. Discussion

Ensuring a reasonable convergence rate for this motion estimation method is essential for a nearly real-time application like prospective motion correction, but the fastest possible ADMM convergence is likely not necessary given the progressive nature of prospective motion correction permits the algorithm to make additional corrections to motion through time. As long as retrospective correction after the prospectively corrected acquisition demonstrates that the residual motion errors are roughly of the same scale as the motion changes from frame to frame, the benefits of further improvements in prospective motion

estimation accuracy become negligible. Given the motion errors in Fig. 3 are relatively small before and after the large motion, in both the cases of $\lambda = 1000$ and of $\lambda = 1$ with adaptive $\mathbf{Q}$, this condition appears to be satisfied by our method with a simple, fixed μ . That said, the adaptive heuristics described could further minimize the number of iterations between accurate motion parameter updates, helping the prospective estimation method to keep up with head motion during a more challenging scan. In our experiments, we keep the number of algorithm iterations between updates below 50 steps. Ensuring our ADMM-like optimization makes sufficient progress towards the optimum during this time is important. Observing in Fig. 4 rapid convergence within this many iterations for both small and big motions confirms the suitability of our optimization method for this simulation.

While we observe that the objective converges quickly using $\mu = 100$ regardless of whether the motion is small (with a realistic covariance $\mathbf{Q}$), or large (temporarily increasing the covariance $\mathbf{Q}$ by two or three orders of magnitude, to enable larger jumps), further study is needed to demonstrate the broad applicability of $\lambda = 1$, $\mu = 100$ across many data sets. Metaheuristic methods and other automatic parameter selection techniques described earlier may provide a means to tune λ during the earliest task activation, on a per-scan basis. This approach would automate selection of λ without overburdening the real-time estimation after the first task activation period is done (a functional MRI scan frequently runs for many such periods in succession). The sparse residual energy also appears to be a suitable criterion for detecting large impulsive motion, although other techniques like parallel imaging motion detection [76] may also be effective alternatives. Also, observing that $\lambda = 1$ generalizes from the no-activation experiment to our simulation with functional activations supports the practicality of using a fixed parameter without additional tuning.

Significant differences are apparent in the activation correlation maps as a result of using this motion correction method prospectively. The content of these ellipses are noticeably degraded with retrospective correction, likely due to the inability to mitigate spin history effects and the smoothing nature of retrospective interpolation. Prospective correction noticeably improves the uniformity of these ellipses. Prospective correction, either alone or in combination with retrospective correction, also improves the maximum Youden's J index (sensitivity and specificity) of detecting functional activations across a wide range of thresholds, especially in the small false positive regime, the important region for rare events like activations. Combining prospective and retrospective correction provides even greater improvement.

However, our conclusions come with several caveats. Our simulations cannot exactly replicate what we would observe with a real acquisition. Real functional MRI acquisitions are not explicitly task-specific, and data often contain unrelated activations. Other motion-related changes, such as to the magnetic field sensitivities of the receiver antennas used, are not modeled by our method. Our method also does not explicitly track and account for outliers in the measured k-space, and undersampling k-space to accelerate the acquisition also may make motion estimation less accurate. In general, our method assumes our measurements are Gaussian, which may not fit all practical applications. Non-Gaussian noise likely would produce non-sparse errors in the image domain, so a sparse regularizer (such as that applied to s) would likely mitigate the effects of such noise or outliers.

Furthermore, rigid motion estimation is low-dimensional on its own, so the effects of a few outliers may be small. Additional evaluation is necessary before applying this method to such environments where such noises are prevalent. Additionally, we do not incorporate time-varying physiological signals that may alias onto the frequencies of our task-related activations during a typical study. Other time-varying effects include slow temperature drift and magnetic field inhomogeneity changes, and these effects may add distortions that do not fit the sparse residual image model. Techniques like bandpass filtering and source separation (e.g., RETROICOR [77]) effectively mitigate some of these perturbations and can be combined with our phase-corrected EPI reconstruction, assuming nearly real-time implementations are available.

6. Conclusion

In conclusion, we developed a novel Kalman-filter-like real-time estimation framework incorporating a sparse residual term in the measurement model to accommodate other time-varying dynamics. We presented two implementations of this framework for prospective motion correction. Our experiments analyzed the tuning parameters for these methods and compared their convergence rates for different degrees of motion. We also demonstrated the effectiveness of this prospective motion correction on a time series analysis of a simulated functional MRI acquisition. We discussed our results and limitations of our experimental setup. In summary, our prospective motion correction method has the potential to significantly improve the sensitivity and specificity of functional MRI time series analysis versus existing retrospective methods.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Funding: This work was supported by the National Institutes of Health [grant numbers F32EB015914, R21EB022309, and R01EB023618] and the Thomas F. and Kate Miller Jeffress Memorial Trust, Bank of America, Trustee.

References

- [1]. Maclaren J, Herbst M, Speck O, Zaitsev M, Prospective motion correction in brain imaging: A review, *Mag. Res. Med* 69 (2013) 621–36.
- [2]. Power JD, Schlaggar BL, Petersen SE, Recent progress and outstanding issues in motion correction in resting state fMRI, *NeuroImage* 105 (2015) 536–551. [PubMed: 25462692]
- [3]. Goto M, Abe O, Miyati T, Yamasue H, Gomi T, Takeda T, Head motion and correction methods in resting-state functional MRI, *Magnetic Resonance in Medical Sciences* 15 (2016) 178–186. [PubMed: 26701695]
- [4]. Godenschweger F, Kägebein U, Stucht D, Yarach U, Sciarra A, Yakupov R, Lüsebrink F, Schulze P, Speck O, Motion correction in MRI of the brain, *Physics in Medicine & Biology* 61 (2016) R32. [PubMed: 26864183]
- [5]. Zaitsev M, Akin B, LeVan P, Knowles BR, Prospective motion correction in functional MRI, *NeuroImage* 154 (2017) 33–42. [PubMed: 27845256]

- [6]. White N, Roddey C, Shankaranarayanan A, Han E, Rettmann D, Santos J, Kuperman J, Dale A, PROMO: Real-time prospective motion correction in MRI using image-based tracking, *Mag. Res. Med* 63 (2010) 91–105.
- [7]. Friston KJ, Williams S, Howard R, Frackowiak RSJ, Turner R, Movement-related effects in fMRI time-series, *Mag. Res. Med* 35 (1996) 346–55.
- [8]. Callaghan MF, Josephs O, Herbst M, Zaitsev M, Todd N, Weiskopf N, An evaluation of prospective motion correction (PMC) for high resolution quantitative MRI, *Frontiers in Neuroscience* 9 (2015) 97. [PubMed: 25859178]
- [9]. Stucht D, Danishad KA, Schulze P, Godenschweger F, Zaitsev M, Speck O, Highest resolution in vivo human brain MRI using prospective motion correction, *PLOS ONE* 10 (2015) e0133921. [PubMed: 26226146]
- [10]. Federau C, Gallichan D, Motion-correction enabled ultra-high resolution in-vivo 7T-MRI of the brain, *PLOS ONE* 11 (2016) e0154974. [PubMed: 27159492]
- [11]. Weller DS, Fessler JA, Prospective motion correction for functional MRI using sparsity and Kalman filtering, in: *Proc. SPIE 8858 Wavelets and Sparsity XV*, 2013, p. 885823.
- [12]. Zaitsev M, Maclaren J, Herbst M, Motion artifacts in MRI: A complex problem with many partial solutions, *Journal of Magnetic Resonance Imaging* 42 (2015) 887–901. [PubMed: 25630632]
- [13]. Schnabel JA, Heinrich MP, Papiez BW, Brady SJM, Advances and challenges in deformable image registration: From image fusion to complex motion modelling, *Medical Image Analysis* 33 (2016) 145–148. [PubMed: 27364430]
- [14]. Viergever MA, Maintz JA, Klein S, Murphy K, Staring M, Pluim JP, A survey of medical image registration - under review, *Medical Image Analysis* 33 (2016) 140–144. [PubMed: 27427472]
- [15]. Thesen S, Heid O, Mueller E, Schad LR, Prospective acquisition correction for head motion with image-based tracking for real-time fMRI, *Mag. Res. Med* 44 (2000) 457–65.
- [16]. Jung H, Sung K, Nayak KS, Kim EY, Ye JC, k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI, *Mag. Res. Med* 61 (2009) 103–16.
- [17]. Asif MS, Hamilton L, Brummer M, Romberg J, Motion-adaptive spatio-temporal regularization for accelerated dynamic MRI, *Mag. Res. Med* 70 (2013) 800–12.
- [18]. Ding F, Wang X, Chen Q, Xiao Y, Recursive least squares parameter estimation for a class of output nonlinear systems based on the model decomposition, *Circuits, Systems, and Signal Processing* 35 (2016) 3323–3338.
- [19]. Ding F, Wang X, Mao L, Xu L, Joint state and multi-innovation parameter estimation for time-delay linear systems and its convergence based on the kalman filtering, *Digital Signal Processing* 62 (2017) 211–223.
- [20]. Ding F, Meng D, Dai J, Li Q, Alsaedi A, Hayat T, Least squares based iterative parameter estimation algorithm for stochastic dynamical systems with arma noise using the model equivalence, *International Journal of Control, Automation and Systems* 16 (2018) 630–639.
- [21]. Hu Y, Liu B, Zhou Q, Yang C, Recursive extended least squares parameter estimation for wiener nonlinear systems with moving average noises, *Circuits, Systems, and Signal Processing* 33 (2014) 655–664.
- [22]. Chen C, Xiong R, Shen W, A lithium-ion battery-in-the-loop approach to test and validate multiscale dual h infinity filters for state-of-charge and capacity estimation, *IEEE Transactions on Power Electronics* 33 (2018) 332–342.
- [23]. Dong H, Wang Z, Ding SX, Gao H, On h-infinity estimation of randomly occurring faults for a class of nonlinear time-varying systems with fading channels, *IEEE Transactions on Automatic Control* 61 (2016) 479–484.
- [24]. Li M, Liu X, The least squares based iterative algorithms for parameter estimation of a bilinear system with autoregressive noise using the data filtering technique, *Signal Processing* 147 (2018) 23–34.
- [25]. Zhang X, Xu L, Ding F, Hayat T, Combined state and parameter estimation for a bilinear state space system with moving average noise, *Journal of the Franklin Institute* 355 (2018) 3079–3103.
- [26]. Wang X, Ding F, Recursive parameter and state estimation for an input nonlinear state space system using the hierarchical identification principle, *Signal Processing* 117 (2015) 208–218.

- [27]. Asif MS, Charles A, Romberg J, Rozell C, Estimation and dynamic updating of time-varying signals with sparse variations, in: Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2011, pp. 3908–3911.
- [28]. Asif MS, Romberg J, Sparse recovery of streaming signals using ℓ_1 -homotopy, IEEE Transactions on Signal Processing 62 (2014) 4209–4223.
- [29]. Angelosante D, Roumeliotis SI, Giannakis GB, Lasso-kalman smoother for tracking sparse signals, in: Proc. 43rd Asilomar Conference on Signals, Systems and Computers, 2009, pp. 181–185.
- [30]. Charles A, Asif MS, Romberg J, Rozell C, Sparsity penalties in dynamical system estimation, in: Proc. 45th Annual Conference on Information Sciences and Systems, 2011, pp. 1–6.
- [31]. Charles AS, Rozell CJ, Dynamic filtering of sparse signals using reweighted ℓ_1 , in: Proc. IEEE International Conference on Acoustics, Speech and Signal Processing, 2013, pp. 6451–6455. (2011) 183–196.
- [32]. Karseras E, Leung K, Dai W, Tracking dynamic sparse signals with kalman filters: Framework and improved inference, in: Proc. 10th Int. Conf. on Sampling Theory and Applications (SampTA), 2013, pp. 224–227.
- [33]. Liu H, Li Y, Zhou Y, Truong TK, Sparse kalman filter, in: Proc. IEEE China Summit and International Conference on Signal and Information Processing (ChinaSIP), 2015, pp. 1022–1026.
- [34]. Charles AS, Balavoine A, Rozell CJ, Dynamic filtering of time-varying sparse signals via ℓ_1 minimization, IEEE Transactions on Signal Processing 64 (2016) 5644–5656.
- [35]. Vaswani N, Kalman filtered compressed sensing, in: Proc. 15th IEEE International Conference on Image Processing, 2008, pp. 893–896.
- [36]. Ho M, Xin J, Sparse kalman filtering approaches to covariance estimation from high frequency data in the presence of jumps, 2016 URL: <http://arxiv.org/abs/1602.02185>, arXiv:1602.02185.
- [37]. Partovibakhsh M, Liu G, An adaptive unscented kalman filtering approach for online estimation of model parameters and state-of-charge of lithium-ion batteries for autonomous mobile robots, IEEE Transactions on Control Systems Technology 23 (2015) 357–363.
- [38]. Filipovic V, Nedic N, Stojanovic V, Robust identification of pneumatic servo actuators in the real situations, Forschung im Ingenieurwesen 75 (2011) 183–196.
- [39]. Stojanovic V, Nedic N, Joint state and parameter robust estimation of stochastic nonlinear systems, International Journal of Robust and Nonlinear Control 26 (2016) 3058–3074.
- [40]. Hu J, Wang Z, Liu S, Gao H, A variance-constrained approach to recursive state estimation for time-varying complex networks with missing measurements, Automatica 64 (2016) 155–162.
- [41]. Nemeth C, Fearnhead P, Mihaylova L, Sequential monte carlo methods for state and parameter estimation in abruptly changing environments, IEEE Transactions on Signal Processing 62 (2014) 1245–1255.
- [42]. Mehmood A, Aslam MS, Chaudhary NI, Zameer A, Raja MAZ, Parameter estimation for hammerstein control autoregressive systems using differential evolution, Signal, Image and Video Processing 12 (2018) 1603–1610.
- [43]. Sankaranarayanan S, Sivakumaran N, Radhakrishnan T, Swaminathan G, Metaheuristic-based approach for state and process parameter prediction using hybrid grey wolf optimization, Asia-Pacific Journal of Chemical Engineering 13 (2018) e2215.
- [44]. Srivastava A, Das DK, Rai A, Raj R, Parameter estimation of a permanent magnet synchronous motor using whale optimization algorithm, in: 2018 Recent Advances on Engineering, Technology and Computational Sciences (RAETCS), 2018, pp. 1–6.
- [45]. Liao Q, Zhou S, Shi H, Shi W, Parameter estimation of nonlinear systems by dynamic cuckoo search, Neural Computation 29 (2017) 1103–1123. [PubMed: 28181879]
- [46]. Nedic N, Stojanovic V, Djordjevic V, Optimal control of hydraulically driven parallel robot platform based on firefly algorithm, Nonlinear Dynamics 82 (2015) 1457–1473.
- [47]. Stojanovic V, Nedic N, Prsic D, Dubonjic L, Djordjevic V, Application of cuckoo search algorithm to constrained control problem of a parallel robot platform, The International Journal of Advanced Manufacturing Technology 87 (2016) 2497–2507.

- [48]. Stojanovic V, Nedic N, A nature inspired parameter tuning approach to cascade control for hydraulically driven parallel robot platform, *Journal of Optimization Theory and Applications* 168 (2016) 332–347.
- [49]. Rahbari O, Omar N, Firouz Y, Rosen MA, Goutam S, Bossche PVD, Mierlo JV, A novel state of charge and capacity estimation technique for electric vehicles connected to a smart grid based on inverse theory and a metaheuristic algorithm, *Energy* 155 (2018) 1047–1058.
- [50]. Ambrogioni L, Güçlü U, Maris E, van Gerven M, Estimating nonlinear dynamics with the ConvNet smoother, 2017 URL: <http://arxiv.org/abs/1702.05243>, arxiv 1702.05243.
- [51]. Logothetis NK, What we can do and what we cannot do with fMRI, *Nature* 453 (2008) 869–878. [PubMed: 18548064]
- [52]. Mehra RK, On the identification of variances and adaptive Kalman filtering, *IEEE Trans. Auto. Control* 15 (1970) 175–84.
- [53]. Mehra RK, Approaches to adaptive filtering, *IEEE Trans. Auto. Control* 17 (1972) 693–98.
- [54]. Odelson BJ, Rajamani MR, Rawlings JB, A new autocovariance least-squares method for estimating noise covariances, *Automatica* 42 (2006) 303–8.
- [55]. Feng B, Fu M, Ma H, Xia Y, Wang B, Kalman filter with recursive covariance estimation–sequentially estimating process noise covariance, *IEEE Transactions on Industrial Electronics* 61 (2014) 6253–6263.
- [56]. Beck A, Teboulle M, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM J. Imaging Sci* 2 (2009) 183–202
- [57]. Taylor A, Hendrickx J, Glineur F, Exact worst-case performance of first-order methods for composite convex optimization, *SIAM Journal on Optimization* 27 (2017) 1283–1313.
- [58]. Kim D, Fessler J, Another look at the fast iterative shrinkage/thresholding algorithm (FISTA), *SIAM Journal on Optimization* 28 (2018) 223–250. [PubMed: 29805242]
- [59]. Glowinski R, Marrocco A, Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de dirichlet non linéaires, *Modélisation Mathématique et Analyse Numérique* 9 (1975) 41–76.
- [60]. Gabay D, Mercier B, A dual algorithm for the solution of nonlinear variational problems via finite-element approximations, *Comput. Math. Appl* 2 (1976) 17–40.
- [61]. Eckstein J, Bertsekas DP, On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators, *Mathematical Programming* 55 (1992) 293–318.
- [62]. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J, Distributed optimization and statistical learning via the alternating direction method of multipliers, *Found. & Trends in Machine Learning* 3 (2011) 1–122.
- [63]. Weller DS, Ramani S, Fessler JA, Augmented Lagrangian with variable splitting for faster non-Cartesian L1-SPIRiT MR image reconstruction, *IEEE Trans. Med. Imag* 33 (2014) 351–361.
- [64]. Ghadimi E, Teixeira A, Shames I, Johansson M, Optimal parameter selection for the alternating direction method of multipliers (ADMM): quadratic problems, *IEEE Trans. Auto. Control* 60 (2015) 644–58.
- [65]. Ramani S, Blu T, Unser M, Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms, *IEEE Trans. Im. Proc* 17 (2008) 1540–1554.
- [66]. Ramani S, Weller DS, Nielsen J-F, Fessler JA, Non-Cartesian MRI reconstruction with automatic regularization via Monte-Carlo SURE, *IEEE Trans. Med. Imag* 32 (2013) 1411–1422.
- [67]. Weller DS, Ramani S, Nielsen J-F, Fessler JA, Monte Carlo SURE-based parameter selection for parallel magnetic resonance imaging reconstruction, *Mag. Res. Med* 71 (2014) 1760–1770.
- [68]. Engl HW, Discrepancy principles for Tikhonov regularization of ill-posed problems leading to optimal convergence rates, *Journal of Optimization Theory and Applications* 52 (1987) 209–215.
- [69]. Hu X, Le TH, Artifact reduction in EPI with phase-encoded reference scan, *Mag. Res. Med* 36 (1996) 166–71.
- [70]. Chen N, Wyrwicz AM, Removal of EPI Nyquist ghost artifacts with two-dimensional phase correction, *Mag. Res. Med* 51 (2004) 1247–53.
- [71]. Reeder SB, Atalar E, Faranesh AZ, McVeigh ER, Referenceless interleaved echo-planar imaging, *Mag. Res. Med* 41 (1999) 87–94.

- [72]. Thevenaz P, Ruttimann UE, Unser M, A pyramid approach to subpixel registration based on intensity, *IEEE Trans. Im. Proc* 7 (1998) 27–41.
- [73]. Chun SY, Fessler JA, A simple regularizer for B-spline nonrigid image registration that encourages local invertibility, *IEEE J. Sel. Top. Sig. Proc* 3 (2009) 159–69.
- [74]. Kwan RK-S, Evans AC, Pike GB, MRI simulation-based evaluation of image-processing and classification methods, *IEEE Trans. Med. Imag* 18 (1999) 1085–97.
- [75]. Friston KJ, Fletcher P, Josephs O, Holmes A, Rugg MD, Turner R, Event-related fMRI: characterizing differential responses, *NeuroImage* 7 (1998) 30–40. [PubMed: 9500830]
- [76]. Kober T, Marques JP, Gruetter R, Krueger G, Head motion detection using FID navigators, *Mag. Res. Med* 66 (2011) 135–43.
- [77]. Glover GH, Li T-Q, Ress D, Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR, *Mag. Res. Med* 44 (2000) 162–7.

Highlights

- Sparse residual image modeling adapts Kalman filtering to temporal dynamics.
- Augmented Lagrangian optimization can produce motion estimates in real time.
- Adapting the process covariance estimate can improve handling of large motions
- Simulations depict improved brain activity detection with prospective correction.

Figure 1:
Active regions (red) are overlaid on the high-resolution T_2^* -weighted Brainweb phantom used in our simulations.

Figure 2:
The activation time series is typical for a block design with alternating off- and on-task blocks.

Figure 3:

The sparse residual (top), image error (middle), and motion error (bottom) are plotted for both (a) $\lambda = 1$ and (b) $\lambda = 1000$ with $\mathbf{Q}$ held fixed (at the true value for the random walk model). The sparse residual, image error, and motion error also are plotted for (c) $\lambda = 1$, resetting $\mathbf{Q}$ when the sparse residual energy is large, such as after the large motion around 33 seconds.

Figure 4:

The objective function $\Psi(a, s)$ convergence to a minimum value Ψ^* is plotted for our ADMM-like method with $\mu = \{10, 100, 1000\}$, for both small motion (a) and large motion (b) cases. In both cases, our algorithm rapidly reaches the optimal objective function value over a wide range of μ . Despite small motion and large motion cases using significantly different values of Q (according to the adaptive scheme), choosing $\mu = 100$ yields good convergence over the first 50 iterations for both cases.

Figure 5:

Ground truth (a) and estimated (b) translations x , y , z and rotations α , β , and γ (relative to the z -, y -, and x -axes, respectively) using the proposed sparse-residual Kalman filter estimation technique.

Figure 6:

Correlation maps ($\alpha = 0.01$) are overlaid on the low-resolution reconstructed slices #9 and #11 for (a) no motion (ideal), (b) retrospective correction (using SPM8) only, (c) slice-by-slice prospective correction only, and (d) both prospective and retrospective corrections. The receiver operating characteristic (ROC) curves (e) plot sensitivity versus false alarm rate (1-specificity) for these methods is shown.

Table 1:

ROC Curve Statistics for Functional MRI Correlation Analysis

	AUC	p'	d'
No correction	0.945	0.81	2.8
Retrospective only	0.994	0.96	4.3
Prospective only	0.997	0.98	4.8
Prospective + retrospective	> 0.999	0.99	5.4
No motion (ideal)	0.998	> 0.99	5.7