Beneficial Techniques for Spatio-Temporal Imaging in Electrical Impedance Tomography

Alistair Boyle[1], Kirill Aristovich[2], Andy Adler[3]

[1] School of Human Kinetics, University of Ottawa, Ottawa, Canada

[2] Department of Medical Physics and Biomedical Engineering, UCL, London, UK.

[3] Systems and Computer Engineering, Carleton University, Ottawa, Canada

E-mail: `alistair.boyle@uottawa.ca`

April 2020

Abstract. *Objective:* Electrical Impedance Tomography (EIT) typically reconstructs individual images from electrical voltage measurements at pairs of electrodes due to current driven through other electrode pairs on a body. EIT images have low spatial resolution, but excellent temporal resolution. There are four methods for integrating temporal data into an EIT reconstruction: filtering over measurements, filtering over images, combined spatial and temporal (spatio-temporal) regularization, and Kalman filtering. These spatio-temporal methods have not been directly compared, making it difficult to evaluate relative performance and choose an appropriate method for particular use cases.

Approach: We (1) develop a common framework, (2) develop comparison metrics, (3) perform simulation and tank studies which directly compare algorithms, and (4) report on relative advantages of the different algorithms.

Main Results: Temporal filtering is well understood, but often not considered as part of the imaging process despite a direct impact on image reconstruction quality. Spatio-temporal regularized techniques are not yet efficient but offer tantalizing advantages. Kalman filtering enables adaptive filtering for time-varying measurement/image noise at the cost of often over-regularized (sub-optimal) images which can now be understood in the same framework as the other techniques. Further research into efficient implementations of Gauss-Newton spatio-temporal regularization will allow temporal and spatial covariance to be explicitly defined for longer time series ($n > 10$ frames) where temporal regularization can be more effective. For the immediate analysis of temporally varying images, we recommend the use of adaptive (time-varying) temporal filtering of measurements followed by adaptive spatial regularization (hyperparameter selection) as the most computationally efficient and effective approach currently available.

Significance: The analysis of variation within regions of an EIT image to extract physiological measures (functional imaging), has become an important EIT technique where temporal and spatial aspects of analysis are tightly integrated. This work gives guidance on available methods and suggests directions for future research.

1. Introduction

Spatio-temporal refer to those items "belonging to both space and time" (Stevenson and Waite 2011). The spatial qualities of an image are often considered paramount, though sequences of images are frequently used to select a representative moment. Close interactions exist between the temporal and spatial changes in a signal measured by an instrument. Longer temporal capture periods can be used to improve measured signal quality and an image's spatial resolution at the risk of blurring fast changing images. Suitable time series measurement data enable a systematic investigation of the exchange between spatial and temporal qualities of an image when the reconstruction method allows explicit, even regional, spatio-temporal choices.

Electrical Impedance Tomography (EIT) is a biomedical technique where low frequency electrical stimulus current is applied to surface electrodes and the resulting voltage distribution is measured at other electrode locations. Electrical Resistivity Tomography (ERT) is the same mathematical problem with geophysical applications. The stimulus and measurements are repeated many times at a comprehensive set of electrode locations, forming a "frame" of data used to reconstruct the interior conductivity. Many frames of data may be collected while monitoring a system for interesting changes: some biomedical EIT instruments can collect data at 100 frames per second, while geophysics ERT instruments typically collect data at lower frame rates (for example, one frame per day over many years) due to using many more electrodes and lower stimulus frequencies. In either case, large sets of time-series data are collected. These datasets capture changes in both spatial and temporal dimensions: three-dimensional changes in conductivity, boundary shape, and contact impedance over time (Boyle and Adler 2011). Most modern reconstruction algorithms for EIT use regularization, with a penalty to impose spatial smoothness. Many publications focus on the spatial characteristics of images at particular time points (Boone and Holder 1996, Friedel 2003, Grychtol *et al* 2012, Crabb *et al* 2014, Mewes *et al* 2017) including a growing array of techniques focused on functional imaging (Adler and Boyle 2017, and references there-in). Multiple frames of data may be reconstructed together and have regularization (smoothness) applied across them, leading to spatio-temporal regularization. Spatio-temporal reconstruction can be formulated as

(a) combining spatio-temporal regularization via an augmented Gauss-Newton reconstruction matrix (Adler *et al* 2007, Boyle 2017, Adler and Aristovich 2018),

(b) as a Kalman-based smoother (Vauhkonen *et al* 1998),

(c) filtering images: image reconstruction, followed by filtering over images, or

(d) filtering measurements: filtering over measurements, then image reconstruction (Yerworth and Bayford 2013).

Measurements are often filtered by averaging data frames (either at acquisition, or post-hoc) to improve Signal-to-Noise (SNR) which smears regions of rapidly changing conductivity. Images are sometimes filtered after reconstruction, though this is often
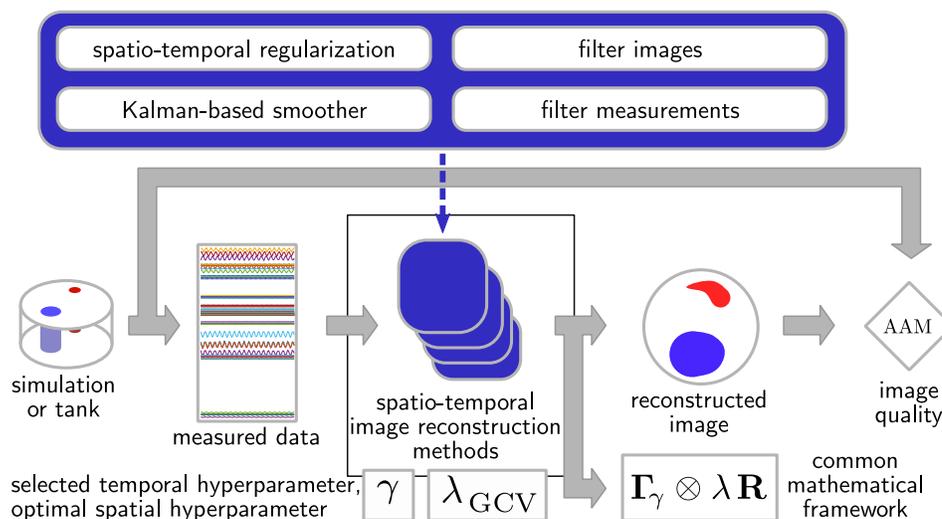
Figure 1: Four methods for spatio-temporal regularization are examined: spatio-temporal regularized Gauss-Newton, Kalman filter, filtering images and filtering measurements

perceived negatively as a manipulation of the resulting images. The alternatives, Kalman filtering and spatio-temporal Gauss-Newton methods, are algorithmically more complex and computationally demanding. On the other hand, these alternatives afford the opportunity to make better, even "optimal," use of available data. We consider image reconstruction methods using spatio-temporal regularized Gauss-Newton, Kalman filter, filtering images and filtering measurements (Fig. 1). How are these spatio-temporal methods related and when is each method appropriate?

These methods have not been directly compared, making it difficult to evaluate relative performance and choose an appropriate method for particular use cases. We (1) develop a common framework, (2) develop comparison metrics, (3) perform simulation and tank studies which directly compare algorithms, and (4) report on relative advantages of the different algorithms. We reformulate each of these methods in a common Kronecker product-based notation to illustrate their similarities and introduce conceptual tools to understand their performance. Despite initial appearances, we find these methods to be closely related. To evaluate the performance of these approaches, we simulate under a range of SNR conditions with stationary and non-stationary models (Fig. 2). Stationary simulations using filtering illustrate that an optimal balance between spatial regularization and temporal filtering exists (Fig. 4). This optimal balance is restricted in the non-stationary case by the rate of conductivity change. We reflect on the strengths and weaknesses of each technique to understand when each method is best employed and present some sample observations (Fig. 5, Fig. 6, Fig. 7).

2. Gauss-Newton on single frames

Gauss-Newton methods are generally used to reconstruct a conductivity image for a single frame of data (Holder 2005, Adler and Boyle 2017). Regularized EIT image reconstruction seeks an optimum image $\mathbf{x}$ to minimize the norm

$$\arg \min_{\hat{\mathbf{x}}} \|\mathbf{y} - \mathcal{F}(\mathbf{x})\|_{\Sigma_\eta^{-1}} + \|\mathbf{x}_* - \mathbf{x}\|_{\Sigma_x^{-1}} \tag{1}$$

for data $\mathbf{y}$, and a forward model $\mathcal{F}$ which generates an estimate of the data based on the image $\mathbf{x}$. Measurement noise is assumed Gaussian $\Sigma_\eta \sim \mathcal{N}(0, \sigma_\eta) = \mathbf{W}^{-1}$, and image noise is also assumed Gaussian $\Sigma_x \sim \mathcal{N}(0, \sigma_x) = \mathbf{R}^{-1}$ centred at a prior image estimate $\mathbf{x}_*$. The Gauss-Newton iterative update (GN-update) for a single frame of data, solving (1), is

$$\delta\mathbf{x}_k = (\mathbf{J}_k^\mathsf{T}\mathbf{W}\mathbf{J}_k + \lambda^2\mathbf{R})^{-1}\left(\mathbf{J}_k^\mathsf{T}\mathbf{W}\mathbf{y}_k + \lambda^2\mathbf{R}(\mathbf{x}_* - \mathbf{x}_{k-1})\right)$$
$$\mathbf{x}_k = \mathbf{x}_{k-1} + \alpha_n\,\delta\mathbf{x}_k \tag{2}$$

Updated parameters $\mathbf{x}_k$ are calculated using results from the previous iteration $\mathbf{x}_{k-1}$ and a prior estimate $\mathbf{x}_*$. The Jacobian $\mathbf{J}_k$ is calculated at each iteration using the forward model with parameters $\mathbf{x}_{k-1}$. The measurements misfit $\mathbf{y} = \mathbf{y}_a - \mathcal{F}(\mathbf{x}_{k-1})$ is the difference between the measurements $\mathbf{y}_a$ and simulated measurements, using the forward model $\mathcal{F}$ and most recent estimate of the parameters $\mathbf{x}_{k-1}$. The misfit $\mathbf{y}_k$ is weighted by an inverse noise covariance matrix $\mathbf{W}$, and the reconstruction is regularized by the image's inverse covariance $\mathbf{R}$ with a hyperparameter λ controlling regularization strength. At each step $\delta\mathbf{x}_k$, a line search may be executed to find the best scaling α_k, or for sufficiently linear problems the line search may be skipped ($\alpha_k = 1$). The initial guess $\mathbf{x}_0$ is usually taken to be the prior estimate $\mathbf{x}_* = \mathbf{x}_0$.

Spatial regularization is often expressed in terms of Tikhonov regularization $\mathbf{R} = \mathbf{I}$ which suppresses large image values, or the Laplacian discrete smoother as a first- or second-order difference operator across adjacent voxels in two or three dimensions.

For time difference EIT, the true change in conductivity $\hat{\mathbf{x}}_\delta = \mathbf{x}_b - \mathbf{x}_a$ is a function of the measurement misfit, which is the difference in measurements $\mathbf{y}_\delta = \mathbf{y}_b - \mathbf{y}_a$ between two times a and b, where

$$\mathbf{y}_{\Delta,k} = (\mathbf{y}_b - \mathbf{y}_a) - (\mathcal{F}(\mathbf{x}_{b,k-1}) - \mathcal{F}(\mathbf{x}_{a,k-1})) = \mathbf{y}_\delta - (\mathcal{F}(\mathbf{x}_a + \mathbf{x}_{\Delta,k-1}) - \mathcal{F}(\mathbf{x}_a))$$
$$\delta\mathbf{x}_{\Delta,k} = (\mathbf{J}_k^\mathsf{T}\mathbf{W}\mathbf{J}_k + \lambda^2\mathbf{R})^{-1}\left(\mathbf{J}_k^\mathsf{T}\mathbf{W}\mathbf{y}_{\Delta,k} + \lambda^2\mathbf{R}(\mathbf{x}_* - \mathbf{x}_{\Delta,k-1})\right)$$
$$\mathbf{x}_{\Delta,k} = \mathbf{x}_{\Delta,k-1} + \alpha_k\,\delta\mathbf{x}_{\Delta,k} \tag{3}$$

and after a number of iterations k, the solution $\mathbf{x}_{\Delta,k}$ should converge near the true solution $\hat{\mathbf{x}}_\delta$. For small changes in conductivity, it is usually sufficient to use (3) for a single-step, rather than in an iterative solution. A single-step with no line search ($\alpha_k = 1$) considerably simplifies the update when $\mathbf{x}_* = \mathbf{x}_{\Delta,0} = 0$ (assume no change in conductivity), which gives $\mathbf{y}_\Delta = \mathbf{y}_\delta$ and reduces (3) to a single linear

(spatially regularized) reconstruction matrix $\mathbf{Q}_s$ for conductivity change $\mathbf{x}_\delta = (\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J} + \lambda^2\mathbf{R})^{-1}\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{y}_\delta$ so that

$$\mathbf{x} = \mathbf{Q}_s\mathbf{y} \qquad \text{for } \mathbf{y} = \mathbf{y}_\delta \text{ and } \mathbf{x} = \mathbf{x}_\delta \tag{4}$$

where the forward model $\mathcal{F}$ no longer plays a direct role in the reconstruction: the effect of modelling errors are significantly reduced. The Jacobian $\mathbf{J}$ is usually calculated based on an estimated (and often homogeneous) conductivity, for example $\mathbf{x}_0 = 1$. We focus on difference imaging ($\mathbf{x}_* = 0$) in the following simulations.

3. Gauss-Newton over multiple frames

The original optimization problem (1) can be extended to express an expanded reconstruction for optimum images $\tilde{\mathbf{x}}$ which minimize the norm

$$\arg\min_{\hat{\tilde{\mathbf{x}}}} \|\tilde{\mathbf{y}} - \mathcal{F}(\tilde{\mathbf{x}})\|_{\tilde{\Sigma}_\eta^{-1}} + \|\tilde{\mathbf{x}}_* - \tilde{\mathbf{x}}\|_{\tilde{\Sigma}_x^{-1}} \tag{5}$$

with the inverse measurement $\tilde{\Sigma}_\eta^{-1}$ and image $\tilde{\Sigma}_x^{-1}$ covariances now endowed with information about inter-frame, as well as the original intra-frame, correlations. Mathematically, there are many choices for how to represent these covariances; in this work we have selected the Kronecker product (Loan 2000) as a useful rhetorical device which reduces cleanly (and exactly) to the original single frame imaging problem when inter-frame correlations are ignored. The conceptual framework of a covariance amongst frames and within frames enables complex relationships to be expressed: for example, the cyclical, regional variations of the heart or variations in measurement noise over time, though these opportunities are not explored further in this paper.

For many frames of data, the images can be reconstructed independent of temporal relationships as

$$\mathbf{X} = \mathbf{Q}_s\mathbf{Y} \quad \text{for} \quad \mathbf{X} = \begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_n \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} \mathbf{y}_1 & \mathbf{y}_2 & \cdots & \mathbf{y}_n \end{bmatrix} \tag{6}$$

when the n measurement frames are arranged as columns in matrix $\mathbf{Y}$ with a column vector of p measurements $\mathbf{y}_i = [y_{i,1}, y_{i,2}, \cdots y_{i,p}]^\mathsf{T}$ for each measurement frame i. Similarly, the conductivity images are arranged in a matrix $\mathbf{X}$ with a column vector of q voxel conductivities $\mathbf{x}_i = [x_{i,1}, x_{i,2}, \cdots x_{i,q}]^\mathsf{T}$ for each measurement frame i. The first image $\mathbf{x}_1$ corresponds to the first frame of data $\mathbf{y}_1$.

The vectorization operator vec() stacks columns of a matrix converting a matrix to a column vector, and the Kronecker product $\otimes$ resizes the reconstruction matrix to agree

$$\tilde{\mathbf{x}} = \mathrm{vec}(\mathbf{X}) = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_n \end{bmatrix} \quad \tilde{\mathbf{y}} = \mathrm{vec}(\mathbf{Y}) = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \vdots \\ \mathbf{y}_n \end{bmatrix} \quad \mathbf{I} \otimes \mathbf{Q}_s = \begin{bmatrix} \mathbf{Q}_s & 0 & & \\ 0 & \mathbf{Q}_s & & \\ & & \ddots & 0 \\ & & 0 & \mathbf{Q}_s \end{bmatrix}$$

$$\tilde{\mathbf{x}} = (\mathbf{I} \otimes \mathbf{Q}_s)\tilde{\mathbf{y}} = (\mathbf{I} \otimes \mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J} + \mathbf{I} \otimes \lambda^2\mathbf{R})^{-1}(\mathbf{I} \otimes \mathbf{J}^\mathsf{T}\mathbf{W})\tilde{\mathbf{y}} \tag{7}$$

When we assume intra-frame measurement and image covariance to be stationary, $\tilde{\Sigma}_\eta^{-1} = \boldsymbol{\Upsilon} \otimes \mathbf{W}$ and $\tilde{\Sigma}_x^{-1} = \boldsymbol{\Gamma} \otimes \mathbf{R}$, for constant forward model and Jacobian $\mathbf{J}_k \to \mathbf{I} \otimes \mathbf{J}_k$. Our estimate of measurement noise $\Sigma_\eta^{-1} = \mathbf{W}$ and choice of spatial regularization $\Sigma_x^{-1} = \mathbf{R}$ have inverse covariances $\boldsymbol{\Upsilon}$ and $\boldsymbol{\Gamma}$ respectively across frames. By expansion from (2), the spatio-temporal Gauss-Newton solution (Adler *et al* 2007, Dai *et al* 2008) takes the form

$$\delta\tilde{\mathbf{x}}_k = ((\mathbf{I} \otimes \mathbf{J}_k^\mathsf{T})(\boldsymbol{\Upsilon} \otimes \mathbf{W})(\mathbf{I} \otimes \mathbf{J}_k) + (\boldsymbol{\Gamma} \otimes \mathbf{I})(\mathbf{I} \otimes \lambda^2\mathbf{R}))^{-1}$$
$$\left((\mathbf{I} \otimes \mathbf{J}_k^\mathsf{T})(\boldsymbol{\Upsilon} \otimes \mathbf{W})\,\tilde{\mathbf{y}} + (\boldsymbol{\Gamma} \otimes \mathbf{I})(\mathbf{I} \otimes \lambda^2\mathbf{R})\,(\tilde{\mathbf{x}}_* - \tilde{\mathbf{x}}_{k-1})\right) \tag{8}$$

by applying the Kronecker identity $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC} \otimes \mathbf{BD})$ to collect terms

$$\delta\tilde{\mathbf{x}}_k = (\boldsymbol{\Upsilon} \otimes \mathbf{J}_k^\mathsf{T}\mathbf{W}\mathbf{J}_k + \boldsymbol{\Gamma} \otimes \lambda^2\mathbf{R})^{-1}\left(\boldsymbol{\Upsilon} \otimes \mathbf{J}_k^\mathsf{T}\mathbf{W}\,\tilde{\mathbf{y}} + \boldsymbol{\Gamma} \otimes \lambda^2\mathbf{R}\,(\tilde{\mathbf{x}}_* - \tilde{\mathbf{x}}_{k-1})\right)$$
$$\tilde{\mathbf{x}}_k = \tilde{\mathbf{x}}_{k-1} + (\alpha_k \otimes \mathbf{I})\,\delta\tilde{\mathbf{x}}_k \tag{9}$$

which will result in very large sparse matrices with dense blocks matching the non-zero entries of $\boldsymbol{\Upsilon}$ for $\boldsymbol{\Upsilon} \otimes \mathbf{J}_k^\mathsf{T}\mathbf{W}\mathbf{J}_k$ and a large sparse matrix for $\boldsymbol{\Gamma} \otimes \mathbf{R}$ which leads to a large dense matrix inverse (Cholesky or LU decomposition). The Wiener filter form, as suggested in (Dai *et al* 2008), will be a better choice when there are fewer measurements per frame than image voxels, though it still results in similar large dense matrices when handling many frames. As a result, direct implementations over many frames are exceptionally inefficient, often limiting this approach to fewer than 10 frames.

As before with difference imaging (4), we can considerably simplify (9) with $\tilde{\mathbf{x}}_* = \tilde{\mathbf{x}}_k = 0$ and $\alpha_k = 1$, to arrive at the difference imaging single-step Gauss-Newton update (1-step GN-update) with spatio-temporal regularization

$$\tilde{\mathbf{x}} = (\boldsymbol{\Upsilon} \otimes \mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J} + \boldsymbol{\Gamma} \otimes \lambda^2\mathbf{R})^{-1}(\boldsymbol{\Upsilon} \otimes \mathbf{J}^\mathsf{T}\mathbf{W})\tilde{\mathbf{y}} \tag{10}$$

In this work, we reconstruct difference images assuming stationary uncorrelated noise $\boldsymbol{\Upsilon} = \mathbf{I}$.

The application of electrical resistivity tomography (ERT) to geophysics problems is mathematically the same problem as EIT: similar algorithms are often employed. A number of spatio-temporal approaches in the geophysics literature have made use of additional terms in the minimization $\arg\min_{\mathbf{x}} \|\tilde{\mathbf{y}} - \mathcal{F}(\tilde{\mathbf{x}})\|_2^2 + \|\mathbf{L}_s\tilde{\mathbf{x}}\|_2^2 + \|\mathbf{L}_t\tilde{\mathbf{x}}\|_2^2$ (Kim *et al* 2009, Kim *et al* 2010, Hayley *et al* 2011, Loke *et al* 2014). Structuring additional regularization in this form is equivalent to constructing a tall regularization matrix $\mathbf{L} = [\mathbf{L}_s\ \mathbf{L}_t]^\mathsf{T}$ so that $\|\mathbf{L}_s\tilde{\mathbf{x}}\|_2^2 + \|\mathbf{L}_t\tilde{\mathbf{x}}\|_2^2 = \|\mathbf{L}\tilde{\mathbf{x}}\|_2^2$ (Hansen 2010). The temporal regularization $\mathbf{L}_t$ introduced in these ERT approaches is uniform between frames and reduces to our Kronecker expansion (10). In summary, similar techniques have been studied in the geophysics setting over a few frames of data ($n < 4$).

Naïve implementations will be expensive to compute as the number of frames grows, and rapidly run out of storage. Are there approaches that can achieve similar outcomes efficiently?

4. Kalman filtering

Using Kalman filtering, one can build a similar formulation for spatially regularized EIT image reconstruction (Vauhkonen *et al* 1998). The Kalman model is

$$\mathbf{x}_k = \mathbf{A}\mathbf{x}_{k-1} + \mathbf{w} \qquad \text{prediction model}$$
$$\mathbf{y}_k = \mathbf{J}_k\mathbf{x}_k + \mathbf{n} \qquad \text{observation model} \qquad (11)$$

where $\mathbf{w}$ is state noise, $\mathbf{A}$ is the state transition matrix and $\mathbf{n}$ is measurement noise. For Kalman time-difference EIT image reconstruction, the current image $\mathbf{x}_k$ is based on the previous image $\mathbf{x}_{k-1}$ ($\mathbf{x}_0 = 0$), and new measurements $\mathbf{y}_k$ at each iteration k

$$\mathbf{x}_k^- = \mathbf{A}\mathbf{x}_{k-1} + \mathbf{B}\mathbf{u}_{k-1} \qquad \text{state estimate}$$
$$\mathbf{x}_k = \mathbf{x}_k^- + \mathbf{K}_k(\mathbf{y}_k - \mathbf{J}_k\mathbf{x}_k^-) \qquad \text{state}$$
$$\mathbf{C}_k^- = \mathbf{A}\hat{\mathbf{C}}_{k-1}\mathbf{A}^\mathsf{T} + \mathbf{P} \qquad \text{error covariance estimate}$$
$$\mathbf{K}_k = \mathbf{C}_k^-\mathbf{J}_k^\mathsf{T}(\mathbf{J}_k\mathbf{C}_k^-\mathbf{J}_k^\mathsf{T} + \mathbf{V})^{-1} \qquad \text{Kalman gain}$$
$$\hat{\mathbf{C}}_k = (\mathbf{I} - \mathbf{K}_k\mathbf{J}_k)\mathbf{C}_k^- \qquad \text{error covariance} \qquad (12)$$

where $\mathbf{P} = \mathbf{R}^{-1}$ is image covariance, $\mathbf{K}_k$ is the Kalman gain, calculated from the error covariance estimate $\mathbf{C}_k^-$ having an initial value $\hat{\mathbf{C}}_0 = 0$, $\mathbf{V} = \lambda^2\mathbf{W}^{-1}$ is the measurement noise covariance, and we don't know the input (assume $\mathbf{u} = 0$) so the control model $\mathbf{B}$ is moot.

Smoothing the state-space, rather than filtering measurements, leads to the idea of Kalman inter-frame temporal regularization as a "fixed-lag smoother" (Moore 1973). Kalman methods have been used in the past to account for intra-frame measurement delays in EIT (Kim *et al* 2001, Kim *et al* 2006), and changing conductivity for known interior boundaries (Kim *et al* 2004).

In general for offline state estimation, the Rauch-Tung-Striebel fixed interval smoother provides an optimal estimate using all measurements in the interval (Rauch *et al* 1965). The smoother takes a sequence of states produced from the Kalman filter in a forward pass and runs them through a correction on a backward pass to adjust the state rather than only using measurements in the past. The backward pass is computed

$$\tilde{\mathbf{x}}_k = \mathbf{x}_k + \mathbf{H}_k(\tilde{\mathbf{x}}_{k+1} - \mathbf{x}_{k+1}^-) \qquad \text{for } \mathbf{H}_k = \hat{\mathbf{C}}_k\mathbf{A}^\mathsf{T}(\mathbf{C}_{k+1}^-)^{-1}$$
$$\tilde{\mathbf{C}}_k = \hat{\mathbf{C}}_k + \mathbf{H}_k(\tilde{\mathbf{C}}_{k+1} - \mathbf{C}_{k+1}^-) \qquad (13)$$

to give spatially and temporally smoothed image reconstruction.

5. Filtering over measurements or images

Many frames of data are sometimes averaged together to produce a single image to improve measurement SNR and thereby reduce artifacts in the resulting image. Windowing is often used to modify the finite impulse response filter characteristics

and avoid frequency aliasing in discrete time systems by weighting samples for digital signal processing (Lyons 2011). One can choose to produce multiple images by "sliding" a window function across the data frames. The result is then sometimes "decimated" to produce frames at a reduced frame-rate or at relevant instants. One can use this technique to interpolate between frames, allowing alignment of (interpolated) frames to gating signals such as a heartbeat (Ross 2010). Averaging the data $\mathbf{Y}$ by using a rectangular window $\mathbf{\Gamma}_{\bar{x}}^{-\mathsf{T}}$, or filtering using an exponential window $\mathbf{\Gamma}_{\gamma}^{-T}$, using a single-step reconstruction $\mathbf{Q}_s$ can be written

$$\bar{\mathbf{x}} = \mathbf{Q}_s \mathbf{Y} \mathbf{\Gamma}_{\bar{x}}^{-\mathsf{T}} \qquad \text{for} \qquad \mathbf{\Gamma}_{\bar{x},i}^{-\mathsf{T}} = \frac{1}{n} \tag{14}$$

$$\mathbf{X} = \mathbf{Q}_s \mathbf{Y} \mathbf{\Gamma}_{\gamma}^{-T} \qquad \text{for} \qquad \mathbf{\Gamma}_{\gamma,i,j}^{-\mathsf{T}} = \frac{\gamma^{|i-j|}}{\sum_j \gamma^{|i-j|}}; \ 0 < \gamma \leq 1$$

with filter $\mathbf{\Gamma}_{\bar{x}}^{-\mathsf{T}}$ (a column vector) or $\mathbf{\Gamma}_{\gamma,i,j}^{-T}$ (a dense matrix) over n frames, with measurement frame i and image frame j. Averaging is a special case of the exponential filter with $\gamma = 1$ and in such a case all image frames will be identical.

For the exponential filter, a value for γ can be selected based on sampling and cutoff frequencies using the z-transform. The discrete exponential filter can be calculated based on a $3\,\mathrm{dB}$ cutoff frequency f_c (Hz) using the instrument frame rate f_s (Hz) as

$$\gamma = 2 - \cos(\omega) - \sqrt{\cos^2(\omega) - 4\cos(\omega) + 3} \tag{15}$$

for frequency $\omega = 2\pi \frac{f_c}{f_s}$. Noise is reduced as $(1-\gamma)/(1+\gamma)$ or approximately $(1-\gamma)/2$ as $\gamma \to 1$ (Lyons 2011, §11.6.1) and, in fact, very little filtering occurs until γ approaches one because the effective width of the filter n is small ($n < 10$) until $\gamma > 0.999$.

In some EIT instruments, the measurements are taken sequentially about the circumference on a single-ring electrode belt. The sequential delay between measurements in a single frame can result in apparent phase differences between left and right lung filling. Corrections for this delay which estimate the signal as if all measurements had been instantaneous, can be implemented though an FFT-based phase correction or a predictive filter (Yerworth and Bayford 2013). This intra-frame filter $\mathbf{F}^{-1}$ (filtered measurements within each frame) can be combined with inter-frame filtering $\mathbf{\Gamma}^{-\mathsf{T}}$ (filtering across frames), as

$$\mathbf{X} = \mathbf{Q}_s \mathbf{F}^{-1} \mathbf{Y} \mathbf{\Gamma}^{-\mathsf{T}} \tag{16}$$

though for most of this work we ignore the intra-frame filter and set $\mathbf{F} = \mathbf{I}$.

6. Simulations and Tank Measurements

Each of the four spatio-temporal methods was simulated at a variety of SNR and with varying temporal and spatial filtering/regularization. Measured tank data (32-electrode Swisstom Pioneer Set EIT system; Swisstom AG, Switzerland) was also collected and reconstructed.
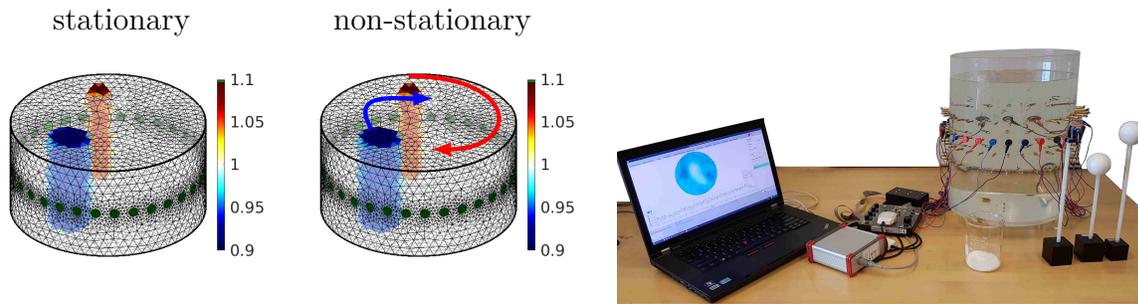
Figure 2: *(above)* Simulations: 32-electrode 3D tank model containing conductive and insulating targets; non-stationary model targets rotating clockwise; *(below)* Tank Measurements: 32-electrode Swisstom Pioneer Set EIT system (Swisstom AG, Switzerland; configured for 47.68 frames/sec, 195 kHz 3 mA stimulation, skip-4 stimulation) and saline tank (CSEM, Switzerland; salinity adjusted to an average 257 Ω electrode impedance) with resistive plastic targets.

A three-dimensional tank forward model (EIDORS `ng_mk_cyl_models`, radius $r = 1$ m, 267,325 elements, 55,309 nodes) with a single ring of 32 electrodes was simulated at a background conductivity of 1 S/m (Fig. 2). Two inclusions were inserted, a conductive (1.10 S/m, 0.1 m diameter) and resistive (0.91 S/m, 0.2 m diameter) object. For the non-stationary model, the objects were rotated around the centre of the model at 1000 measurement frames per clockwise rotation ($f/f_s = 10^{-3}$). Both stationary and non-stationary models were used to generate 256,000 frames with Additive White Gaussian Noise (AWGN) at eight SNR (5–40 dB) relative to the measured signal prior to calculating difference measurements. The difference measurements represented at most 0.63% of the individual measurements (high common mode signal content). Noise was high-pass filtered (6th-order Butterworth) $f_c/f_s > 2 \times 10^{-3}$ and mixed (added) to broadband AWGN noise at a ratio of 90% filtered to 10% broadband noise. This mixed noise was important to demonstrate the different performance between spatial $f/f_s = 10^{-3}$ and temporal noise filtering in the following simulations.

Difference images were reconstructed on a "dual model" using a three-dimensional forward tank model (fwd: 229,905 elements, 47,317 nodes) projected onto a two-dimensional reconstruction model (rec: 1658 elements, 924 nodes). Stimulus and measurement sequences for simulations were configured to match tank measurements: a skip-4 sequence dropping all measurements which saturate the hardware (> 1.5 mV) and any measurements within 2 electrodes of the stimulus electrode, leaving 640 of 928 possible measurements.

Using the models for stationary and non-stationary conditions, simulations at each SNR and range of temporal cut-off frequencies ($10^{-6} < f_c/f_s < 10^{-1}$) were evaluated for each of the four reconstruction methods: spatio-temporal Gauss-Newton, Kalman filtering, filtering measurements, and filtering images (§III–V). Gauss-Newton reconstructions used single-step difference imaging ($\mathbf{x}_* = \mathbf{x}_0 = 0$). Spatial Tikhonov

regularization $\mathbf{R} = \mathbf{I}$ was used with a hyperparameter λ selected via Generalized Cross Validation (GCV) (Hansen 2010). An appropriate λ was selected heuristically for the small subset of cases where GCV failed. The resulting image was compared to an expected image, the forward model. The difference between expected $\mathbf{x}_a$ and reconstructed $\mathbf{x}_b$ images was quantified using a voxel volume corrected Artifact Amplitude Measure (AAM) as

$$\text{AAM} = \sqrt{\sum_i v_i \left(\frac{\mathbf{x}_{a,i}}{\max |\mathbf{x}_a| + \epsilon} - \frac{\mathbf{x}_{b,i}}{\max |\mathbf{x}_b| + \epsilon} \right)^2} \tag{17}$$

for the i-th voxel volume v_i, so that very similar images have low AAM. A small $\epsilon = 10^{-9}$ was supplied to manage situations where the regularization λ becomes overly strong and results in a zero or near zero image. Differences in amplitude between the images were ignored: these are expected when filtering and regularization are performed on a linear one-step reconstruction.

We considered the GREIT consensus "performance figures of merit:" amplitude response, position error, resolution, shape deformation, ringing, and noise figure (Adler *et al* 2009). These measures use the quarter-amplitude response of the reconstruction compared to a single ideal target to evaluate the performance of a reconstruction algorithm over the imaged area. Noise figure, a measure of how measurement noise is amplified in the reconstruction, was recommended by the GREIT authors for selecting the spatial hyperparameter λ, where we have elected to use GCV for this purpose with satisfactory results. In our initial tests using the GREIT metrics, movement artifacts were jointly captured by amplitude response, position error, resolution, shape deformation, and ringing. All five types of artifacts tend to occur simultaneously as movement artifacts reach a critical threshold. We did not see a noticeable benefit in breaking out the different types of image artifacts. (Sample plots for the GREIT figures of merit have been supplied in Appendix A.) In this work, we have elected to use AAM as a generalized summary statistic which accurately captures the artifacts observed in our images.

7. Results

We explored proxies for AAM which might allow an optimal selection of temporal filtering γ when the "true" image is not know (Fig. 3). Changes in the relative norm of the conductivity $\|\mathbf{x}\|$ or the data misfit $\|\mathbf{Jx} - \mathbf{b}\|$ may be a way forward. Correlation scatter plots showed no reliable global replacement for AAM.

The variation of AAM and GCV-selected spatial regularization λ_{GCV} with temporal filtering f_c/f_s for reconstruction of filtered measurements is plotted in Fig. 4. Some illustrative sample reconstructions using stationary and non-stationary simulation data and measured tank data, are shown in Fig. 5 (25 dB SNR), Fig. 6 (40 dB SNR), and Fig. 7 (tank data). (Red dots on Fig. 4 indicate samples presented in Fig. 5a–f. Blue

dots indicate samples presented in Fig. 6b–f.) An AAM > 7.0 tended to indicate images that were no longer recognizably similar to the expected image.

For stationary simulations, more temporal filtering reduced artifacts and the need for spatial regularization which in turn improved spatial resolution (Fig. 5abc) (reducing f_c/f_s reduced λ_{GCV}). As expected, there is a noticeable drop in λ_{GCV} around $2 \times 10^{-3} < f/f_s < 10^{-3}$ for stationary reconstructions where temporal filtering removes 90% of the measurement noise; the high-pass filtered AWGN noise. This drop is less reliably present in non-stationary reconstructions. We observed the GCV criteria occasionally fail in a well understood way (Hansen 2010) by under-regularizing images such as Fig. 6a particularly under very low SNR conditions.

Non-stationary images experienced similar improvements when compared to the stationary reconstructions in spatial resolution and AAM up to the limit of the rate of movement ($f/f_s = 10^{-3}$). Beyond this movement threshold the spatial regularization selected by GCV continued to fall, as with the stationary case, but images became significantly distorted due to smearing of the moving targets (reflected in an increasing AAM) (Fig. 5def). This is an important observation: that λ_{GCV} cannot be used as a mechanism to detect an optimal temporal-spatial regularization trade-off for non-stationary images. An optimal λ exists for a specified γ. Optimal γ might be selected for entire frames by using adaptive temporal filtering techniques leading to a frame-by-frame optimal γ and λ, though this approach was not explored further in this work.

For filtering images, GCV failed to produce usable images for poor SNR (SNR $\leq$ 25 dB) (see lines exiting the top of Fig. 4's plots) because large λ_{GCV} suppressed the reconstructed image so that there was nothing remaining to filter at the next stage. In such cases (Fig. 5ghi), a fixed hyperparameter was selected to match Fig. 5def. As noted in §7.3, with a fixed λ filtering by measurements or images is exactly equivalent: compare SNR=25 dB Fig. 5def to Fig. 5ghi. When λ is found for each frame before filtering images the results are similar to filtering measurements but not identical: compare SNR=40 dB Fig. 6def to Fig. 6ghi. Where SNR was sufficient to allow successful GCV without temporal filtering (SNR > 25 dB) very similar results were often observed when comparing filtering measurements or images. Occasionally, filtering images gave odd "smeared" cyclic oscillations due to harmonic variations in λ at the rate of rotation: the image intensity varied as spatial regularization changed (Fig. 6g). This non-uniform smearing could lead to misinterpretation of reconstructed images by giving apparent structure where none should exist.

Reconstructions using some methods were much more time consuming (most notably Kalman filtering) but showed similar trends across methods: increasing filtering beyond the rate of target movement resulted in significant smearing of the image. Kalman filtering (Fig. 5j), was a notable exception which we discuss later (§7.2). Prior to this filtering threshold $f_c/f_s > 10^{-3}$, reconstructed images were quite similar, with the exception of Gauss-Newton Kronecker reconstructions which we discuss in the following section (§7.1). Individual artifacts and how smearing artifacts appeared differ in their particulars for each method. Reconstructions with measured tank data (Fig. 7) showed
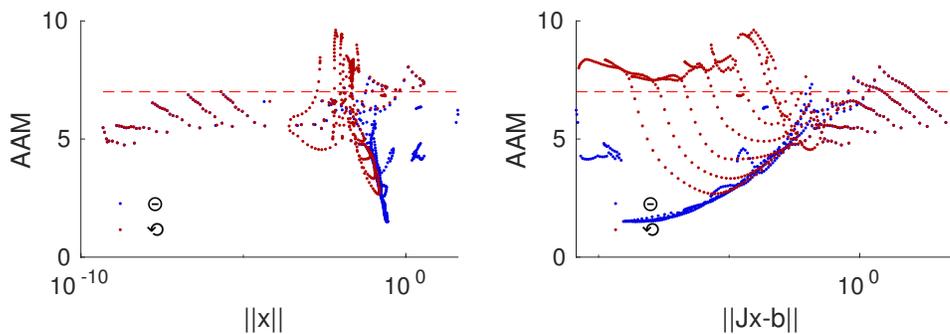
Figure 3: No correlation between AAM (known "true" image) and $\|\mathbf{x}\|$ or $\|\mathbf{Jx} - \mathbf{b}\|$ ("true" image unknown) for stationary (● blue dot) and moving (● red dot) images; AAM > 7.0 (– – red dashed line) are no longer recognizable reconstructions

similar trends to the simulations, though images differed due to using only resistive targets. Kalman filtering reconstructed worse images than in simulation. Spatio-temporal Gauss-Newton regularized reconstruction with limited frames ($n < 10$) did manage to reconstruct reasonable images due to having sufficient SNR to reconstruct good images with very little temporal filtering.

7.1. Gauss-Newton

In Fig. 5mno, Fig. 6mno, and Fig. 7yzaa, Kronecker-based spatio-temporal Gauss-Newton reconstructions appear largely independent of the temporal cut-off frequency (f_c/f_s) due to the general limitations of a naïve Kronecker approach, as we explain in the following.

Two observations arise from manipulating the spatio-temporal Gauss-Newton equation (10). First, the augmented time difference spatio-temporal reconstruction matrix contains no temporal regularization if the noise covariances are equal

$$\tilde{\mathbf{Q}}_{st} = \left(\mathbf{\Upsilon} \otimes \mathbf{J}^\mathsf{T}\mathbf{WJ} + \mathbf{\Gamma} \otimes \lambda^2\mathbf{R} \right)^{-1} \left(\mathbf{\Upsilon} \otimes \mathbf{J}^\mathsf{T}\mathbf{W} \right)$$
$$\tilde{\mathbf{Q}}_s = \tilde{\mathbf{Q}}_{st} = (\mathbf{I} \otimes \mathbf{Q}_s) \qquad \text{when } \mathbf{\Upsilon} = \mathbf{\Gamma} \tag{18}$$

because the covariances $\mathbf{\Upsilon} = \mathbf{\Gamma}$ can be collected using Kronecker identities (Loan 2000). This is an unlikely scenario in practice since temporal correlations are 1-dimensional while spatial correlations are 2- or 3-dimensional.

Second, over many frames, a naïve Gauss-Newton Kronecker expansion will run out of storage. The Kronecker product of large matrices results in very large matrices, growing at $O(n^2)$ for n frames. It is not efficient or desirable to compute these large matrices directly. There needs to be some more efficient implementation since we are otherwise left with large sparse blockwise matrices with dense blocks distributed about the matrix. In our implementation for this work we were limited to ten frames (10 frames: 2.05 GB, 100 frames: 205 GB) which limited our ability to construct effective temporal filtering as seen in Fig. 5mno, Fig. 6mno, and Fig. 7yzaa, where images with
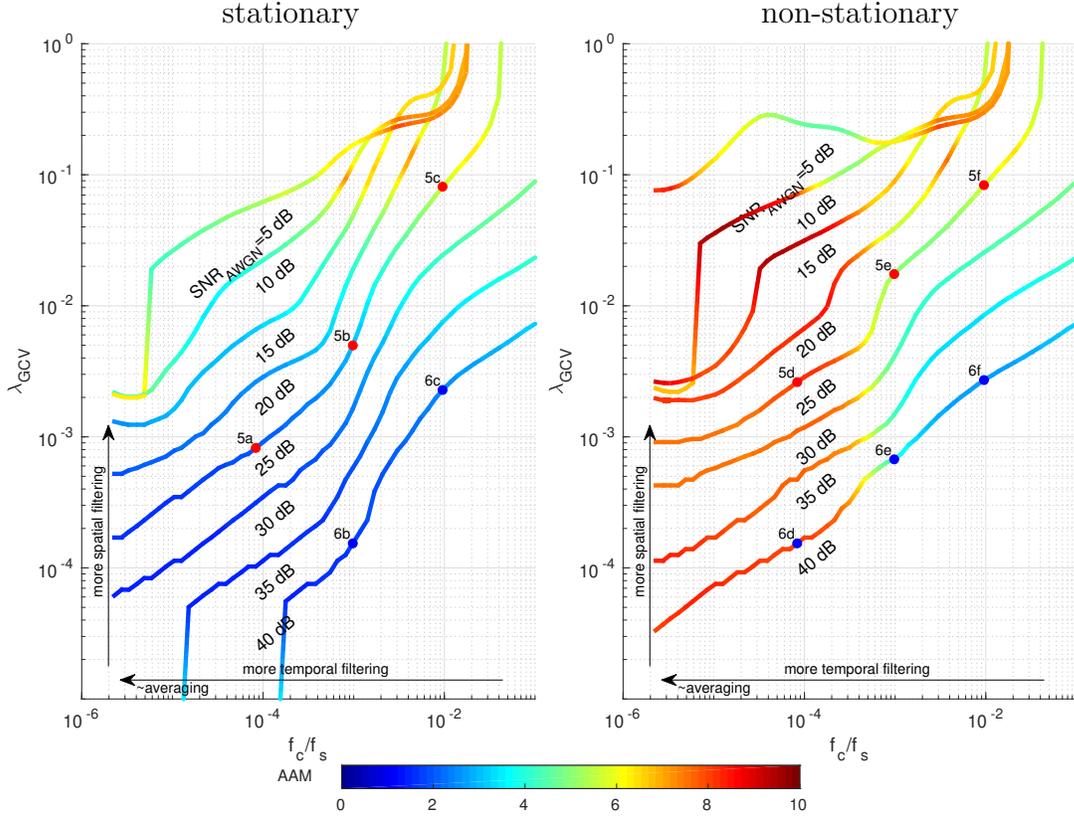
Figure 4: Simulated reconstructions (Fig. 2) over a range of SNR (5–40 dB) with AWGN (90% filtered $f_c/f_s = 2 \times 10^{-3}$, 10% broadband AWGN); *(left)* stationary; *(right)* non-stationary (1000 frames per rotation $= 10^{-3}$); optimal spatial regularization λ_{GCV} is reduced as temporal filtering f_c/f_s reduces measurement noise; temporal filtering is limited by movement in the non-stationary model; • and • dots indicate selected image samples in Fig. 5a–f and Fig. 6b–f, respectively

nominally greater temporal filtering look essentially the same. More frames of data need to be incorporated. Noise (AWGN) is reduced by $\sqrt{n}$: when averaging, a large number of frames are required to make a significant difference in measurement SNR. A more efficient implementation or an approximation of the Kronecker product approach is needed.

With a low frame limit ($n < 10$), spatio-temporal Gauss-Newton regularization using Kronecker products was not effective unless there was a well known prior image to use as a reference. Forcing a strong temporal correlation between a few frames is a common approach in geophysics settings where 2 or 3 images are used in this Kronecker Gauss-Newton-type solution (Kim *et al* 2009, Kim *et al* 2010, Hayley *et al* 2011, Loke *et al* 2014). In this "known prior image" case, a spatial prior $\mathbf{x}_*$ in (2) and using spatial-only regularization is a more efficient solution which will penalize the reconstruction as it moves away from the prior.

If the limitation in number of frames could be alleviated, the Gauss-Newton
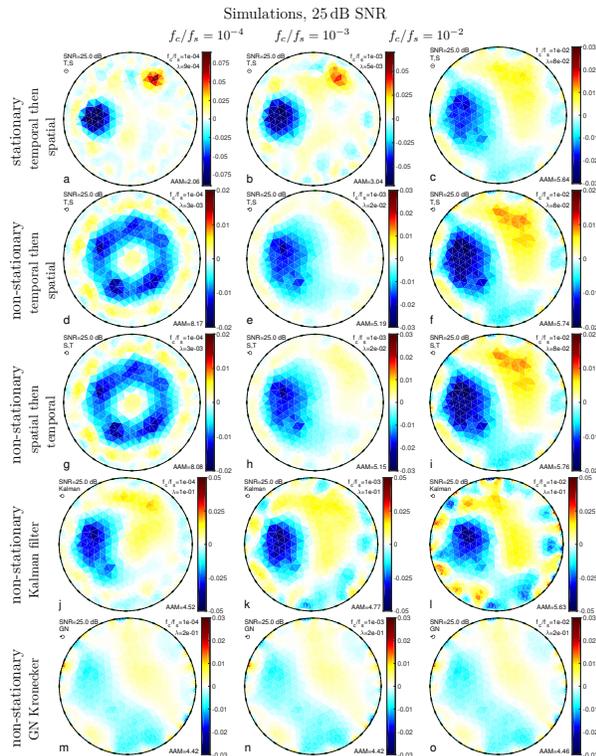
Figure 5: Sample conductivity reconstructions for SNR=25 dB for each reconstruction method; *(a–c)* stationary model where greater temporal filtering removes noise and improves resolution; *(d–f)* non-stationary model where temporal filtering causes blurring when beyond the rate of rotation; *(g–i)* filtering images gives similar results to filtering measurements (d–f); *(j–l)* Kalman filtering gives nice images but is computationally expensive and the Kalman gain sometimes fails to converge (j) possibly resulting in unanticipated regularization (see text for details); *(m–o)* cut-off frequency (f_c/f_s) has little impact on spatio-temporal Gauss-Newton reconstructions
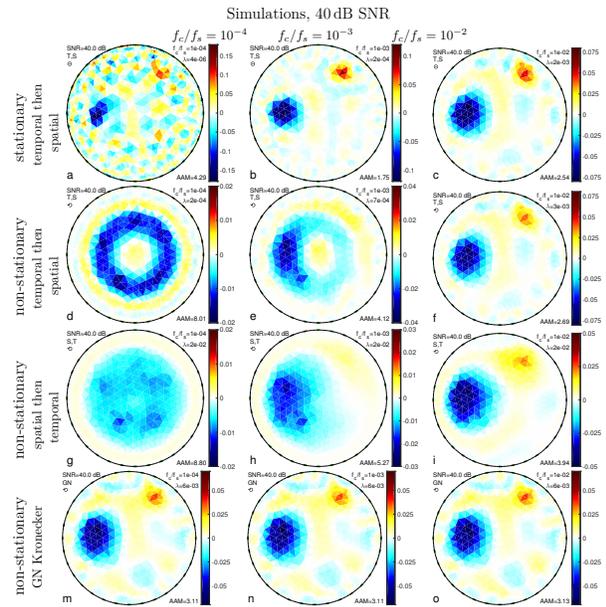
Figure 6: Sample conductivity reconstructions for SNR=40 dB for each reconstruction method; *(a–c)* stationary model where greater temporal filtering removes noise and improves resolution; *(d–f)* non-stationary model where temporal filtering causes blurring when beyond the rate of rotation; *(g–i)* filtering images gives similar results to filtering measurements (d–f); *(m–o)* as with Fig. 5mno, cut-off frequency (f_c/f_s) has little impact on spatio-temporal Gauss-Newton reconstructions
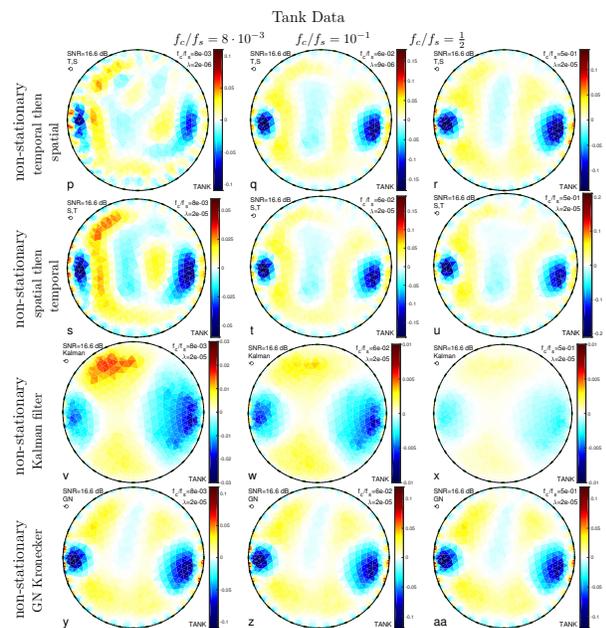
Figure 7: Tank conductivity reconstructions (resistive targets) for each reconstruction method; measured tank data gives similar results to simulations with the notable exceptions that Kalman filtering *(v–x)* gives poorer images and limited Gauss-Newton spatio-temporal filtering gives a reasonable result since filtering is not required to produce good images

Table 1: Kalman filter time (iterations) to convergence

elems (fwd/rec)	f_c/f_s		
	10^{-4}	10^{-2}	10^{-1}
21682/813	1:01:07.5 (10000)*	0:08:18.4 (1380)	0:00:58.4 (157)
54055/1478	6:04:31.6 (10000)*	0:50:21.9 (1406)	0:05:46.2 (160)
229905/1658	8:00:00.0* (9032)	1:11:25.0 (1411)	0:09:02.8 (160)

elems = reconstruction dual model elements (fwd/rec); time, iterations = *hh:mm:ss (iterations).*
* = Iteration or time limited, convergence terminated early.

approach would be appealing. The technique is algorithmically rich in potential applications due to the control afforded by complete freedom in explicitly selecting time-varying measurement and regional image covariance.

7.2. Kalman filter

A direct alternative to explicitly setting measurement and image covariance is to try to approximate them using the model and measurements. In Fig. 5jki, and Fig. 7vwx, Kalman filtering shows promising results but is slow to converge as stronger temporal filtering is applied (Tab. 1). The value of Kalman gain $\mathbf{K}$ and error covariance $\hat{\mathbf{C}}$ (12) converge to stable values independent of measurements for constant $\mathbf{A}, \mathbf{J}, \mathbf{P}, \mathbf{V}$ and can be precomputed to give asymptotic $\mathbf{K}_\infty$ and $\hat{\mathbf{C}}_\infty$ learned from the model (Kim *et al* 2006). This pre-computation is often the most time consuming part of the calculation (Tab. 1 on 3.3 Ghz 4-core i5-2500K with 32 GB memory avail.) and halting prior to convergence is similar to truncating a singular value decomposition or halting iterative optimizations early: it can have an uncontrolled regularizing effect as if stronger spatial/temporal regularization were applied.

Given stable gain and covariance $\mathbf{K}_\infty, \hat{\mathbf{C}}_\infty$, the smoother (13) also becomes stable $\mathbf{H}_\infty$. The Kalman gain does not converge with $\mathbf{A} = \mathbf{I}$ since the error covariance grows at each iteration when $\mathbf{P} > 0$. To obtain a converged $\mathbf{K}_\infty, \hat{\mathbf{C}}_\infty$, we set $\mathbf{A} = \gamma\mathbf{I}$ for $\gamma \to 1$ and $\mathbf{P} = (1 - \|\mathbf{A}\|)\mathbf{R}^{-1}$; $\mathbf{R}^{-1} = \mathbf{I}$: a state transition model predicting a voxel independent exponential decay with state uncertainty related to rate of change γ.

The state $\mathbf{x}_k$ for converged gain and covariance $\mathbf{K}_\infty, \hat{\mathbf{C}}_\infty$ simplifies (12) to

$$\mathbf{x}_k = \mathbf{A}\mathbf{x}_{k-1} + \mathbf{K}_\infty(\mathbf{y}_k - \mathbf{J}\mathbf{A}\mathbf{x}_{k-1}) \qquad \text{with } \mathbf{A} = \gamma\mathbf{I} \text{ and } \mathbf{K}_\infty = \hat{\mathbf{C}}_\infty\mathbf{J}^\mathsf{T}(\mathbf{J}\hat{\mathbf{C}}_\infty\mathbf{J}^\mathsf{T} + \mathbf{V})^{-1}$$

$$\mathbf{x}_k = \gamma\mathbf{x}_{k-1} + \hat{\mathbf{C}}_\infty\mathbf{J}^\mathsf{T}(\mathbf{J}\hat{\mathbf{C}}_\infty\mathbf{J}^\mathsf{T} + \mathbf{V})^{-1}(\mathbf{y}_k - \mathbf{J}\gamma\mathbf{x}_{k-1})$$

$$\mathbf{x}_k = (\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J} + \lambda^2\hat{\mathbf{R}})^{-1}(\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{y}_k + \lambda^2\hat{\mathbf{R}}\mathbf{x}_*) \tag{19}$$

for $\hat{\mathbf{R}} = \hat{\mathbf{C}}_\infty^{-1}$, $\mathbf{W} = \lambda^2\mathbf{V}^{-1}$, $\mathbf{x}_* = \gamma\mathbf{x}_{k-1}$ where we find the Wiener filter form (Wiener 1964) of the single-step Gauss-Newton update and recall that it is algebraically equivalent (Boyle 2010, appendix B.4) to (2) using the Sherman-Morrison-Woodbury formula. From this manipulation we see that the time-invariant Kalman filter has

learned a spatial regularization $\hat{\mathbf{R}}$ that is the specific spatial covariance of the model limited by noise estimates for state (conductivity parameters) and measurements.

It has generally been observed that Kalman filtering over-regularizes images, indicating that the "optimal Kalman gain" is not in fact the best choice for EIT images. There is no temporal regularization in this Kalman formulation until fixed lag smoothing is applied, so we note that Kalman reconstruction must construct a more heavily regularized set of images from many frames, as any reconstructed image noise is propagated into future images through the prior $\mathbf{x}_* = \gamma\mathbf{x}_{k-1}$ term. Given the numerical noise inherent in reconstructing many frames of data, it becomes apparent that greater "optimal" spatial regularization (more image smoothing) is required to control the additional cumulative noise from prior reconstructions. The key advantage of a Kalman filter is its ability to adapt to changing conditions (the error covariance estimate) and incorporate predictive information about the model (the state estimate), neither of which are leveraged in most EIT applications where the Jacobian is not updated, the system inputs $\mathbf{u}$ are unknown, and noise estimates are fixed. There are opportunities here for better application of Kalman filtering in EIT image reconstruction but the slow computation of the initial converged error covariance estimate is a major draw-back. One can argue that this is a "one time cost," but it must be paid at each modification of the reconstruction model. Potentially, model updates (modifications to the FEM mesh, rather than re-meshing) might converge rapidly from a previously converged solution.

7.3. Spatial then temporal filtering

Fig. 5d matches Fig. 5g to within numerical noise, as do Fig. 5ef with Fig. 5hi, respectively. Notice that, because the solution is linear, we can perform inter-frame filtering on the images, rather than the data to achieve identical results (for fixed $\mathbf{F}$ and λ, leading to fixed $\mathbf{Q}_s$) without changing the filter.

$$\mathbf{X} = \mathbf{Q}_s\mathbf{F}^{-1}(\mathbf{Y}\mathbf{\Gamma}^{-\mathsf{T}}) = (\mathbf{Q}_s\mathbf{F}^{-1}\mathbf{Y})\mathbf{\Gamma}^{-\mathsf{T}} \tag{20}$$

Unlike in Fig. 5, in Fig. 6 the corresponding images (Fig. 6def to Fig. 6ghi, respectively) do not match because an optimal hyperparameter λ was selected at each frame. If one is selecting the hyperparameter λ via GCV, L-curve or another technique at each frame, then the results will differ because $\mathbf{Q}_s$ will be changing frame-to-frame. In our sample images (Fig. 6d–i), the hyperparameter λ differs by small amounts resulting in noticeably different reconstructions; particularly Fig. 6e and f. Despite this detail, in our exploration the images for temporal-then-spatial or spatial-then-temporal filtering are often similar, inclusive of any artifacts that may be present.

The fact that an identical image can be obtained by applying fixed filters to either image or measurements (20) seems to imply an assumption on the common spatial and temporal noise covariances, but this is not the case, as we shall see in the next section.

7.4. Temporal then spatial filtering

The similarities in Fig. 5def/ghi are striking. Are temporal and spatial filtering an equivalent and efficient replacement for Gauss-Newton spatio-temporal regularization? Measurement filtering (16) can be manipulated into something resembling the spatio-temporal Gauss-Newton equation (10) featuring Kronecker products. If we assume an (approximately) invertible inter-frame filter $\mathbf{\Gamma}^{-\mathsf{T}}$ such as the exponential filter for $\gamma < 1$, then by applying the Kronecker identities (Loan 2000)

$$(\mathbf{C}^{\mathsf{T}} \otimes \mathbf{A})\text{vec}(\mathbf{B}) = \text{vec}(\mathbf{ABC}) \qquad (\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1}$$

$$(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD}) \qquad \mathbf{A} \otimes (\mathbf{B} + \mathbf{C}) = (\mathbf{A} \otimes \mathbf{B}) + (\mathbf{A} \otimes \mathbf{C})$$

the expression (16) can be rearranged as

$$\begin{aligned}
\tilde{\mathbf{x}} &= (\mathbf{\Gamma}^{-1} \otimes \mathbf{Q}_s \mathbf{F}^{-1})\tilde{\mathbf{y}} = \left(\mathbf{\Gamma}^{-1} \otimes \left((\mathbf{J}^{\mathsf{T}}\mathbf{WJ} + \lambda^2 \mathbf{R})^{-1}\mathbf{J}^{\mathsf{T}}\mathbf{WF}^{-1}\right)\right)\tilde{\mathbf{y}} \\
&= \left(\mathbf{\Gamma}^{-1} \otimes (\mathbf{J}^{\mathsf{T}}\mathbf{WJ} + \lambda^2 \mathbf{R})^{-1}\right)(\mathbf{I} \otimes \mathbf{J}^{\mathsf{T}}\mathbf{WF}^{-1})\tilde{\mathbf{y}} \\
&= (\mathbf{\Gamma} \otimes \mathbf{J}^{\mathsf{T}}\mathbf{WJ} + \mathbf{\Gamma} \otimes \lambda^2 \mathbf{R})^{-1}(\mathbf{I} \otimes \mathbf{J}^{\mathsf{T}}\mathbf{WF}^{-1})\tilde{\mathbf{y}}
\end{aligned} \qquad (21)$$

Comparing (10) and (21) ($\mathbf{F}^{-1} = \mathbf{I}$), the inter-frame measurement noise and temporal image covariances are fixed to the same matrix $\mathbf{\Gamma}$ on the left-hand side of (10), while the right-hand side of (21) features an assumption of independence $\mathbf{I}$ between frames. Therefore, we can see that spatial and temporal filtering are not exactly equivalent to Gauss-Newton spatio-temporal regularization. This disconnect between covariances on the right and left-hand sides of (21) is striking, yet for many cases reconstructed images looked quite similar.

8. Discussion

We have examined four methods for making use of the temporal information in EIT signals.

- Two temporal filtering methods (on measurements and images) were shown to be exactly equivalent for fixed λ.
- Image filtering was shown to be more prone to failure for low SNR reconstructions.
- Combined spatio-temporal regularized reconstruction was shown to be nearly similar to filtering methods when the inverse of the filter matrix was used as the temporal regularization matrix, as is the case for common exponential filters. The one-dimensional Laplacian smoothing operator $[-1 \ +2 \ -1]$ is an approximate inverse for the exponential smoothing filter, and clarifies the functional behaviour of the temporal regularization term.
- Kalman filtering was shown to be a method that gives the spatio-temporal regularization form of reconstruction but uses the previous image as an image prior and finds an approximate covariance for the model through extensive iterations.

- The spatio-temporal Gauss-Newton formulation offers the greatest freedom in selecting a fixed prior across all images or, as with Kalman filtering, a prior based on previous images.

From the temporal filtering approaches, it becomes clear that temporal regularization across a few frames for controlling noise is not particularly effective, as it drastically limits the effectiveness of averaging across noisy measurements. We find that the spatio-temporal regularization using Kronecker products is not currently practical for large numbers of frames (computational time and storage rapidly become unreasonable for any useful quantity of temporal regularization), but is an excellent tool for understanding the inter-play between temporal and spatial regularization as generalized covariances. Recently, there has been some promising preliminary work in finding efficient implementations for the spatio-temporal regularized Gauss-Newton method (Boyle 2017, Adler and Aristovich 2018).

We find the temporal filtering methods using exponential smoothing to be easily understood and controlled. Performing spatial regularization last enables more reliable use of "optimal" hyperparameter selection methods such as GCV of L-curve. A key draw-back of the method is that it cannot be extended to more complex (inter/intra-frame) regional spatio-temporal covariance relationships because the operations are performed sequentially.

The Kalman filtering method takes its strength from the ability to adaptively find "optimal" approximations. Under conditions where the exact values for these approximations are known, this can only result in poorer quality reconstructions. The great promise of the Kalman filtering approach is its ability to adapt to changing noise and model. We feel that the ability to control the interplay between spatial and temporal regularization, and find globally optimal solutions for a particular set of data, means that the Gauss-Newton approach (combined with methods such as GCV or L-curve for evaluating the optimal trade-off and to adapt to varying noise levels or other measurement conditions over time) provide a generalized framework in which the highest quality reconstructed images might be produced.

9. Conclusions

This work explored four techniques for spatio-temporal regularized imaging in EIT: combining spatio-temporal regularization in a Gauss-Newton framework, a Kalman-based smoother, filtering over reconstructed images, and filtering over measurements followed by reconstruction of images.

We find that the equations describing these techniques, after translation to a common Kronecker-based structure, have differences which help explain why artifacts differ in low quality images. Practically we find that, despite these differences, these techniques give quite similar results in most of our tests.

In conclusion, we find great promise in spatio-temporal regularized Gauss-Newton reconstruction approaches for EIT. In particular, the ability to apply different

regularization and priors to regions or times in an image sequence is a powerful technique that could aid in, for example, separating lung and heart which operate at different frequencies and are spatially separated. To realize this promise, efficient implementations that facilitate the determination of optimal spatial and temporal hyperparameters must be developed. Adaptive filtering techniques which account for time-varying measurement noise are closely related to optimal hyperparameter selection for temporal regularization. We recommend temporal measurement filtering approaches, in the meantime, because they provide a good baseline method with well understood failure mechanisms, efficient implementation techniques, and the breadth of (temporal) digital signal processing literature available to support a variety of measurement filtering techniques.

Acknowledgement

Appendix A. Reconstruction metrics: AAM and GREIT

We examined the possibility of using the GREIT performance figures of merit in this work instead of the Artifact Amplitude Measure (AAM). The GREIT metrics are: amplitude response, position error, resolution, shape deformation, ringing, and noise figure (Adler *et al* 2009). Noise figure was proposed by the GREIT authors as a way of setting the spatial regularization λ and we have not considered it further here, as we have elected to use GCV. Each GREIT metric identifies types of image artifacts that typically occur in stationary EIT reconstruction.

Reconstructions were simulated at $50\,\text{dB}$ SNR, with temporal filtering and Gauss-Newton reconstruction using Tikhonov regularization and a hyperparameter selected by GCV. Circular movement of a single conductive target was simulated. (The GREIT metrics only support a single target in the reconstruction.) In Fig. A1, an optimal hyperparameter which was large relative to the singular values of the reconstruction matrix was a strong indicator that the image was being suppressed, as illustrated by a reduced GREIT amplitude response. A combination of the GREIT position error, resolution, and shape deformation identify when smearing of the image, due to greater temporal filtering, has occurred. GREIT's ringing metric did not identify interesting artifacts, as the amplitude response and optimal hyperparameter already identify suppressed (low amplitude) images. The same smearing artifact can be identified by setting a threshold for AAM, as we have done in this work.

The GREIT metrics were designed to isolate specific spatial reconstruction artifacts.
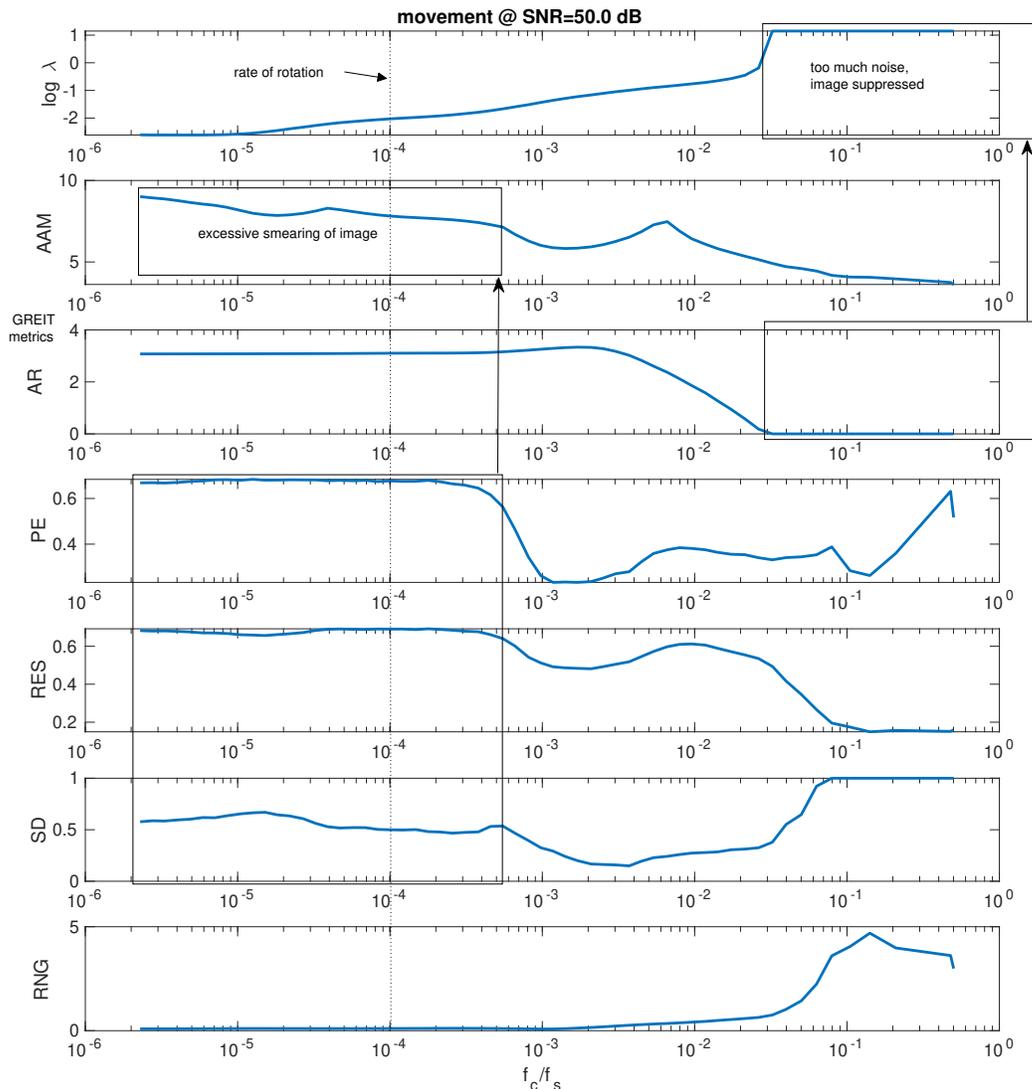
Figure A1: Movement reconstruction with SNR = 50 dB using temporal filtering; GREIT's amplitude response (AR), position error (PE), resolution (RES), shape deformation (SD), and ringing (RNG) jointly identify the same artifacts identified by examining the optimal spatial hyperparameter λ_{GCV} and Artifact Amplitude Measure (AAM).

Movement smearing was captured in four of the five GREIT metrics. We find that AAM is capable of identifying movement smearing, along with other artifacts, and have elected to use AAM in this work.

References

Adler A and Lionheart W R B 2006 Uses and abuses of EIDORS: an extensible software base for EIT *Physiological Measurement* 27(5) S25–S42.

Adler A, Dai T and Lionheart W R B 2007 Temporal image reconstruction in electrical impedance tomography *Physiological Measurement* 28(7) S1–S11.

Adler A, Arnold J, Bayford R, Borsic A, Brown B, Dixon P, Faes T, Frerichs I, Gagnon H, Gärber Y, Grychtol B, Hahn G, Lionheart W R B, Malik A, Patterson R, Stocks J, Tizzard A, Weiler N and Wolf G 2009 GREIT: a unified approach to 2D linear EIT reconstruction of lung images *Physiological Measurement* 30(6) S35–S55.

Adler A, Boyle A, Braun F, Crabb M G, Grychtol B, Lionheart W R B, Tregidgo H F J and Yerworth R 2017 EIDORS 3.9 *in* '18th International Conference on Biomedical Applications of Electrical Impedance Tomography' Dartmouth, USA.

Adler A and Boyle A 2017 Electrical impedance tomography: Tissue properties to image measures *IEEE Transactions on Biomedical Engineering* 64(11) 2494–2504.

Adler A and Aristovich K 2018 Accelerating space-time regularized reconstructions *in* '19th International Conference on Biomedical Applications of Electrical Impedance Tomography' Edinburgh, UK.

Boone K G and Holder D S 1996 Effect of skin impedance on image quality and variability in electrical impedance tomography: a model study *Medical and Biological Engineering and Computing* 34(5) 351–354.

Boyle A 2010 The effect of boundary shape deformation on two-dimensional electrical impedance tomography Master's thesis Carleton University Ottawa, Canada.

Boyle A and Adler A 2011 Impact of electrode area, contact impedance and boundary shape on EIT images *Physiological Measurement* 32(7) 745–754.

Boyle A 2017 Spatio-temporal regularization over many frames *in* '18th International Conference on Biomedical Applications of Electrical Impedance Tomography' Dartmouth, USA.

Crabb M G, Davidson J, Little R, Wright P, Morgan A, Miller C, Naish J, Parker G, Kikinis R, McCann H and Lionheart W R B 2014 Mutual information as a measure of image quality for 3D dynamic lung imaging with EIT *Physiological Measurement* 35(5) 863–879.

Dai T, Soleimani M and Adler A 2008 EIT image reconstruction with four dimensional regularization *Medical & Biological Engineering & Computing* 46(9) 889–899.

Friedel S 2003 Resolution, stability and efficiency of resistivity tomography estimated from a generalized inverse approach *Geophysical Journal International* 153(2) 305–316.

Grychtol B, Lionheart W R B, Bodenstein M, Wolf G K and Adler A 2012 Impact of model shape mismatch on reconstruction quality in electrical impedance tomography *IEEE Transactions on Medical Imaging* 31(9) 1754–1760.

Hansen P 2010 *Discrete Inverse Problems: Insight and Algorithms* SIAM.

Hayley K, Pidlisecky A and Bentley L 2011 Simultaneous time-lapse electrical resistivity inversion *Journal of Applied Geophysics* 75(2) 401–411.

Holder D, ed. 2005 *Electrical impedance tomography: Methods, history and applications* 1 edn IOP Publishing Ltd.

Kim B S, Kim K Y, Kao T J, Newell J C, Isaacson D and Saulnier G J 2006 Dynamic electrical impedance imaging of a chest phantom using the Kalman filter *Physiological Measurement* 27(5) S81–S91.

Kim J, Yi M, Park S and Kim J 2009 4-D inversion of DC resistivity monitoring data acquired over a dynamically changing earth model *Journal of Applied Geophysics* 68(4) 522–532.

Kim J, Yi M, Ahn H and Kim K 2010 4-D inversion of resistivity monitoring data using L1 norm minimization *in* 'Near Surface Geoscience 2010' Zurich, Switzerland.

Kim K Y, Kim B S, Kim M C, Lee Y J and Vauhkonen M 2001 Image reconstruction in time-varying electrical impedance tomography based on the extended Kalman filter *Measurement Science*

and Technology 12(8) 1032–1039.

Kim K Y, Kim B S, Kim M C and Kim S 2004 Dynamic inverse obstacle problems with electrical impedance tomography *Mathematics and Computers in Simulation* 66(4-5) 399–408.

Loan C V 2000 The ubiquitous Kronecker product *Journal of Computational and Applied Mathematics* 123(1-2) 85–100.

Loke M H, Dahlin T and Rucker D 2014 Smoothness-constrained time-lapse inversion of data from 3D resistivity surveys *Near Surface Geophysics* 12(1) 5–24.

Lyons R 2011 *Understanding Digital Signal Processing* 3 edn Prentice Hall.

Mewes B, Hilbich C, Delaloye R and Hauck C 2017 Resolution capacity of geophysical monitoring regarding permafrost degradation induced by hydrological processes *The Cryosphere* 11(6) 2957–2974.

Moore J 1973 Discrete-time fixed-lag smoothing algorithms *Automatica* 9(2) 163–173.

Rauch H, Tung F and Streiebel C 1965 Maximum likelihood estimates of linear dynamic systems *AIAA Journal* 3(8) 1445–1450.

Ross H 2010 Technical developments to enable electrical impedance tomography (EIT) measurement of blood flow to monitor cardiac output: Key engineering contributions Master's thesis Carleton University Ottawa, Canada.

Schöberl J 1997 NETGEN an advancing front 2D/3D-mesh generator based on abstract rules *Computing and visualization in science* 1(1) 41–52.

Stevenson A and Waite M, eds 2011 *Concise Oxford English dictionary* 12th edn Oxford University Press.

Vauhkonen M, Karjalainen P A and Kaipio J P 1998 A Kalman filter approach to track fast impedance changes in electrical impedance tomography *IEEE Transactions on Biomedical Engineering* 45(4) 486–493.

Wiener N 1964 *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications* MIT Press.

Yerworth R and Bayford R 2013 The effect of serial data collection on the accuracy of electrical impedance tomography images *Physiological Measurement* 34(6) 659–669.