The University of
Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

Ward, Wil O.C. and Wilkinson, Paul B. and Chambers, Jon E. and Nilsson, Henrik and Kuras, Oliver and Bai, Li (2016) Tracking tracer motion in a 4-D electrical resistivity tomography experiment. Water Resources Research, 52 (5). pp. 4078-4094. ISSN 1944-7973

Access from the University of Nottingham repository:
http://eprints.nottingham.ac.uk/34202/1/WardEtAl-WRR2016-TrackingIndependentTracerMomentsInA4DElectricalResistivityTomographyExperiment.pdf

Correspondence to:
W. O. C. Ward,
wcw@cs.nott.ac.uk

Tracking tracer motion in a 4-D electrical resistivity tomography experiment

W. O. C. Ward[1,2], P. B. Wilkinson[2], J. E. Chambers[2], H. Nilsson[1], O. Kuras[2], and L. Bai[1]

[1]School of Computer Science, University of Nottingham, Nottingham, UK, [2]British Geological Survey, Keyworth, UK

Abstract A new framework for automatically tracking subsurface tracers in electrical resistivity tomography (ERT) monitoring images is presented. Using computer vision and Bayesian inference techniques, in the form of a Kalman filter, the trajectory of a subsurface tracer is monitored by predicting and updating a state model representing its movements. Observations for the Kalman filter are gathered using the maximally stable volumes algorithm, which is used to dynamically threshold local regions of an ERT image sequence to detect the tracer at each time step. The application of the framework to the results of 2-D and 3-D tracer monitoring experiments show that the proposed method is effective for detecting and tracking tracer plumes in ERT images in the presence of noise, without intermediate manual intervention.

1. Introduction

Applying computational imaging techniques is a relatively novel approach to the interpretation and analysis of electrical resistivity tomography (ERT) data [*Ward et al.*, 2014]. Prior research has investigated automated identification of subsurface boundaries in static images, using steepest gradient [*Chambers et al.*, 2012], Laplacian [*Hsu et al.*, 2010], or watershed [*Elwaseif and Slater*, 2010, 2012] edge detection. A common difficulty with such approaches is that ERT images tend to exhibit smooth boundaries between different regions. The cause may be artifacts of the smoothness constraints typically used to regularize the inversion, but can also reflect true gradational changes in the resistivity structure of the subsurface. Attempts have been made to mitigate the effects of smoothness constraints on edge detection by incorporating detected boundaries into the inversion process, to some success [*Elwaseif and Slater*, 2012; *Zhou et al.*, 2014].

Another approach is to use clustering techniques to automatically detect distinct regions in static ERT images [*Ward et al.*, 2014; *Audebert et al.*, 2014], as well as combining resistivity models created using multiple electrode arrays [*Ishola et al.*, 2014]. One particular method, distribution-guided fuzzy clustering, has been used previously on time-lapse ERT data [*Ward et al.*, 2014; *Chambers et al.*, 2015].

In a more complex situation, for example, where there is a need to track a particular region against a variable background (caused by noise or the influence of other processes), automated identification is a nontrivial problem. Clustering the overall change, for example to monitor quarry dewatering as in *Chambers et al.* [2015], is one of the methods for identifying regions-of-interest. However, that would require both a baseline image to compare with, and also manual intervention to choose the correct cluster associated with the region of interest (similar to the problem faced by *Elwaseif and Slater* [2013]).

In the case of tracking and monitoring a plume or tracer movement, clustering techniques are less suitable since they do not necessarily distinguish features, but rather fully segment an image. A more promising approach is to consider monitoring tracer movement as a visual tracking problem, incorporating the need for local feature detection [*Tuytelaars and Mikolajczyk*, 2007]. Such is the approach taken in research presented in this paper. The detections at each time step should be matched across the time series to form consistent tracks of objects throughout the time-lapse model [*Yilmaz et al.*, 2006]. Due to the effects of noise and variable spatial resolution in ERT imaging, this is not a simple task. Approaches to track objects would need to deal with the loss of tracked regions that either have not been correctly resolved in the current time step, or have moved out of the survey area (temporarily or permanently). The tracking should be able to handle artifacts in the model, as well as quantify uncertainty or variance in the results.
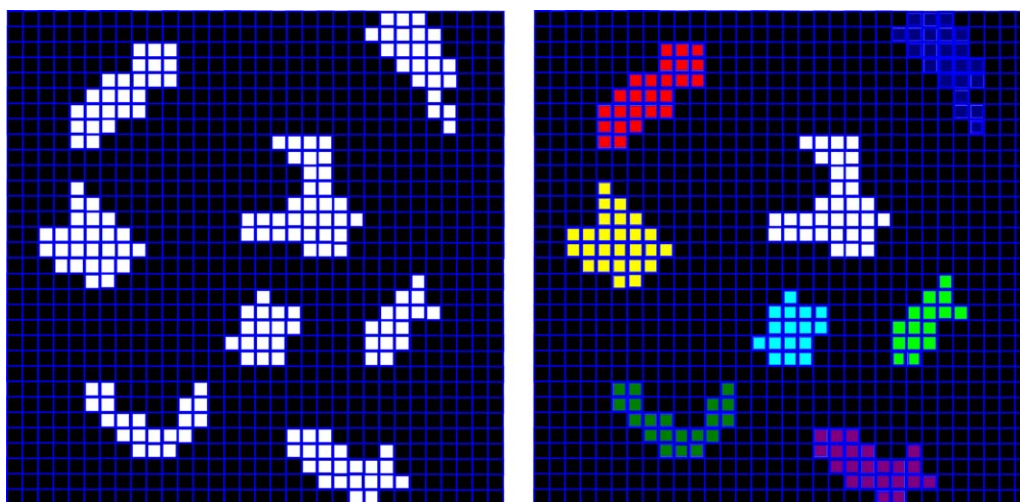
Figure 1. (right) A binary image with independent connected components individually colored.

The paper introduces a framework for tracking tracers in 2-D and 3-D time-lapse ERT images using a Kalman filter [*Kalman*, 1960], with promising results. The Kalman filter uses a measure of statistical correspondence to predict the next state in time sequence given information from all previous time steps. Updating each prediction step with observations, e.g., detected features from the model, allows for improved estimation of the true state while not relying on detection results being absolutely correct [*Yilmaz et al.*, 2006]. Multiple regions can be simultaneously tracked using assignment algorithms, to maintain and/or prune tracks as the time-series progresses.

A technique called maximally stable volumes (MSV) is used to detect features at each time step. The MSV approach involves converting an image into binary using thresholding (known as "binarising"), performed at several thresholds to create a tree of disconnected regions of the image at each threshold level [*Matas et al.*, 2002; *Donoser and Bischof*, 2006]. A measure of stability is then used to extract regions that remain relatively constant in size over a small set of consecutive thresholds. The method may be considered a localized adaptive thresholding method [*Matas et al.*, 2002]. Information from the identified MSVs is used to guide the Kalman filter by providing observations at the current time step.

In the rest of the manuscript, the background theory and algorithms are outlined, with a complete description of the framework. Details of the implementation and results of a 2-D and a 3-D ERT monitoring experiment are given. Following this, an evaluation of the technique and limitations are discussed, and recommendations for further work are outlined.

2. Methodology

2.1. Maximally Stable Volumes

The maximally stable volumes algorithm is a generalization of the maximally stable extremal regions (MSER) method [*Matas et al.*, 2002], used to find connected regions that are considered "stable." An input volume is binarized at multiple successive thresholds and a tree is created using the spatially connected components [*Donoser and Bischof*, 2006; *Vedaldi*, 2009]. MSER is widely regarded as one of the most effective means of local feature detection in images [*Tuytelaars and Mikolajczyk*, 2007]. In 3-D, MSV has been shown to perform well in automatic selection of points-of-interest, and has been found to be especially robust in the presence of noise [*Yu et al.*, 2012].

First, a set of thresholds, Δ, is specified, typically spanning the range of unique values in the input data and discretized over some predefined intervals. At each threshold, the input image is binarized to 0 or 1 based on whether values are below or above the threshold, and the individual connected components are identified—by nature they are spatially *disconnected* from each other. An example in 2-D is shown in Figure 1, with the binary image on the left, and connected components individually colored on the right. The algorithm for calculating connected components is simply finding an unlabeled voxel with value 1 and using a
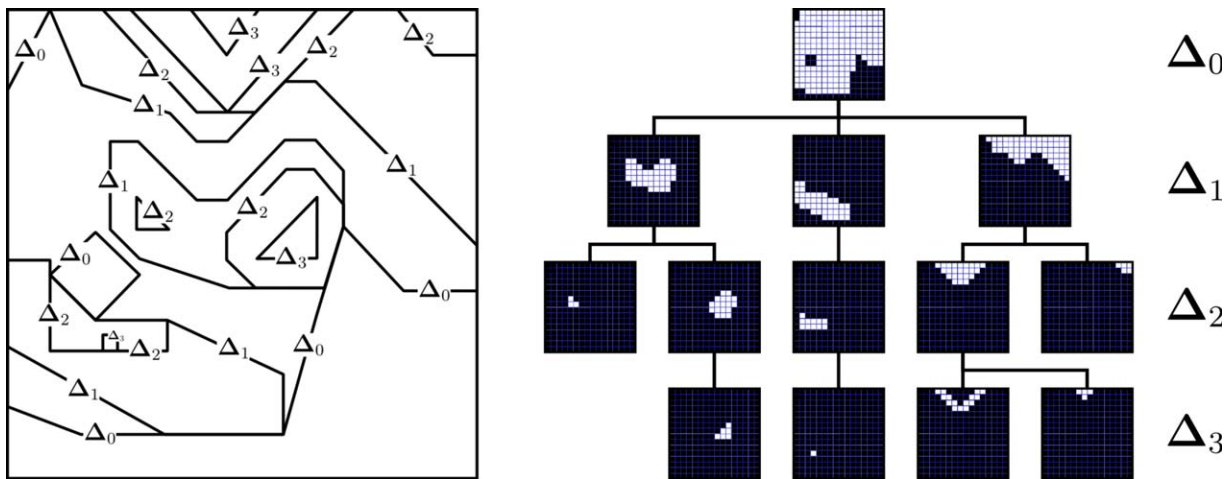
Figure 2. Demonstration of the connected component tree used in MSV. The nodes in the tree represent the connected components obtained by thresholding the input model with Δ_i, with contour plot corresponding to input data.

flood-fill to label all the voxels in the connected region containing that voxel. The process is repeated for every remaining unlabeled 1 value voxels in the model. In calculating connectivity, a voxel is adjacent both along principal coordinate axes and diagonally—8 connected for 2-D, and 26 connected for 3-D.

The connected component tree built for a simple 2-D image using four threshold levels is demonstrated in Figure 2. Each connected component forms a tree node at its threshold level, as a child of the component in the level above to which its voxels also belong. By the nature of the thresholding and connectivity criterion, there are no overlapping regions. The contour plot shows the embedded connected components in the equivalent input data.

From the complete tree, a measure of "stability" can be calculated based on the change in size of components as the threshold changes. Each connected component $\mathbf{c}$ has a corresponding stability $\Psi(\mathbf{c})$ that takes the difference in size between its parent node and its largest child divided by the size of $\mathbf{c}$. The root node (with no parent) and leaf nodes (with no children) have 0 stability. Traversing the tree depth-first to calculate local maxima in stability returns the nodes that correspond to maximally stable volumes. While some branches in the tree will have monotonically decreasing stability, it is also not uncommon for MSVs to be nested, as can be seen in Figure 3. The full algorithm is detailed in Algorithm 1 in Appendix A.

Constraints can be applied to remove MSVs that are too large or small as relative to the size of the image. Stability can also be calculated based on the sizes of the connected components $h(\geq 1)$ threshold levels above and below the focused component. Figure 3 shows a comparison of distribution-guided fuzzy clustering (as described by *Ward et al.*, [2014] and *Chambers et al.*, [2015]) and MSV applied to a time step from a 2-D ERT tracer monitoring experiment. It can be seen here that while the clustering approach is effective in picking out distinct regions, it is not as effective at pinpointing the local region of the tracer as MSVs. The issue is largely due to the complete segmentation nature where every point is given an arbitrary label, including "background" regions.

MSVs are well suited for the purpose of detecting tracers in ERT images, due both to their success at identifying distinct regions in a background of noise, and also the presence of nested volumes. The nested volumes mark steps of stability in the gradual shift in resistivity between the tracer and the background, the boundaries of which give discrete contours describing the contained tracer. Detected volumes are used to extract localized information at each time step, which is described in detail in a later section. The information is used for tracking purposes, incorporated into an implementation of the Kalman filter.

2.2. Kalman Filter

Object tracking is the process of identifying and locating a moving object or region in a sequence of images over time. In the cases studied in the presented work, the objects tracked are regions of increased conductivity, associated with the passage of a saline tracer in a time-lapse sequence of ERT images.
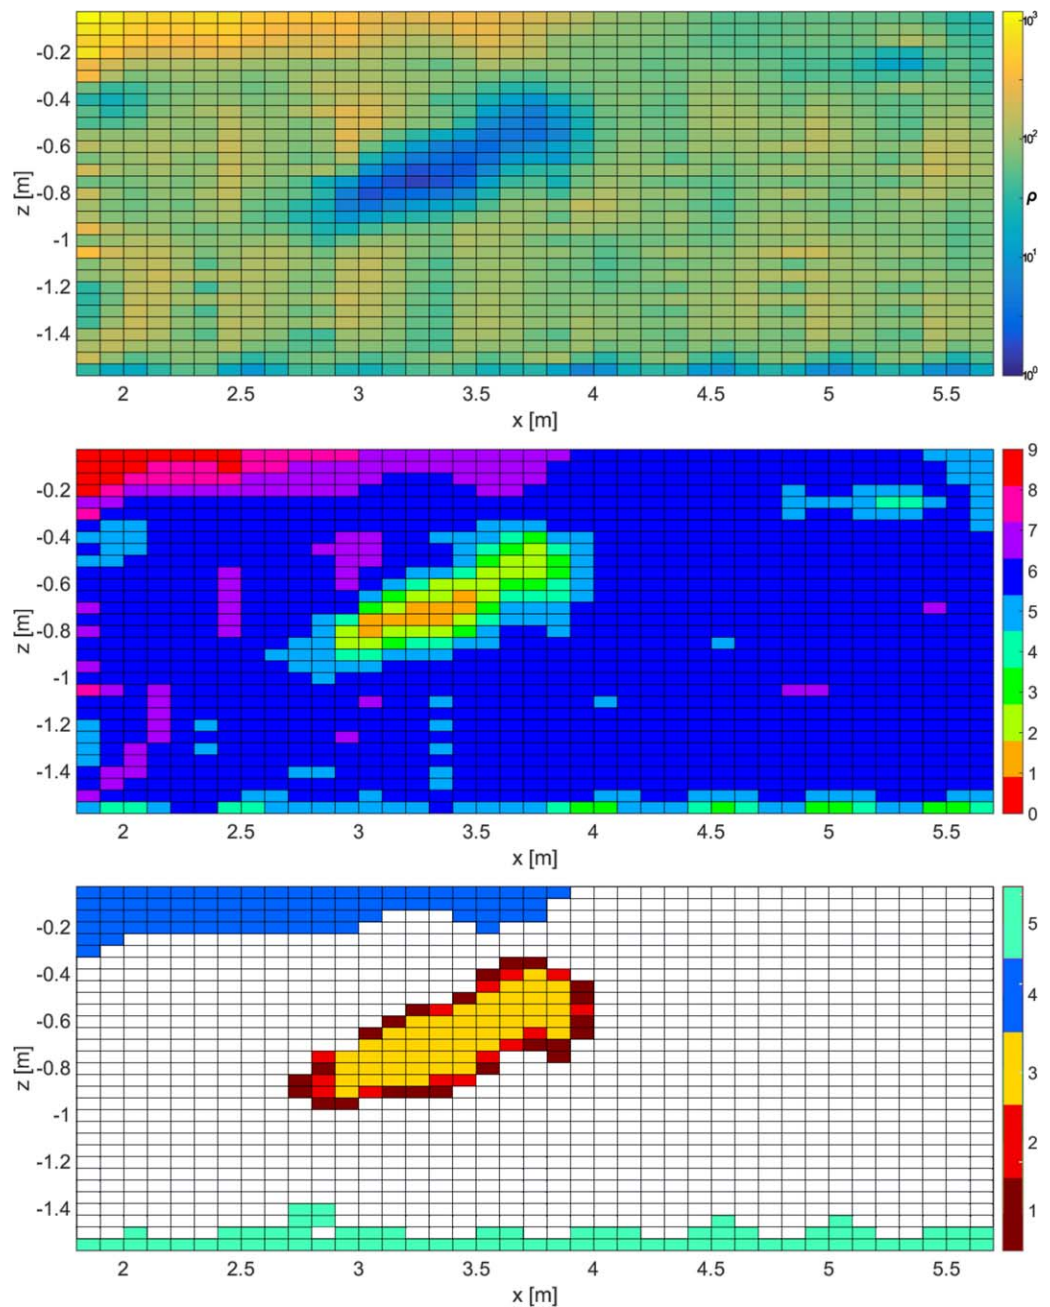
Figure 3. 2-D resistivity model of pit experiment at t_{20}, with results from (middle) distribution-guided fuzzy clustering and (bottom) maximally stable volumes. Individually colored regions represent distinctly separate detected regions which, in the MSV case, exhibit nesting.

Tracking algorithms are often split into two types characterized by the relations used to identify objects between time-slices. The first, deterministic correspondence, uses a cost function that is minimized between a current object and equivalent points of reference identified in a consecutive time-slice [*Yilmaz et al.*, 2006].

The alternative to deterministic approaches is to use statistical methods to identify correspondence of objects between time-slices. The basis of such tracking algorithms is to include model and detection errors into the algorithm. Errors are given in the form of covariance matrices that dictate the degree of uncertainty in the transition between two time-slices. Using a state model that is consistent with current knowledge of the underlying problem, and then updating this model as more data is observed, is the basis for many

approaches [*Masreliez and Martin*, 1977]. While there are a variety of such approaches that can be taken, requiring a different set of assumptions, such as Markov chain Monte Carlo [*Andrieu et al.*, 2003], or the particle filter [*Liu and Chen*, 1998], the Kalman filter tracker is used.

The Kalman filter is useful to approximate a linear system and gives an optimal estimate of the current state of the system, if the underlying errors are strictly normally distributed. The system is described by some state model, which is generated from observed features and dictates the rules of correspondence between two time-slices [*Kalman*, 1960]. An object state, X_t, at some time t, may be described by the following dynamic system:

$$X_t = \mathbf{F}X_{t-1} + \mathbf{Bu_t} + \mathbf{W}_t \tag{1}$$

where the matrix $\mathbf{F}$ is a system of linear equations representing the change between two states over some fixed time step; $\mathbf{B}$ and $\mathbf{u}_t$ form a control system that can be used to constrain the linear system; and $\mathbf{W}_t \sim (0, \mathbf{Q}_t)$ is some unknown Gaussian noise, with zero mean and known covariance $\mathbf{Q}_t$.

The relation between detected measurements z_t and the state representation is specified by an observation function:

$$Z_t = \mathbf{H}X_t + \mathbf{V}_t \tag{2}$$

where $\mathbf{H}$ is the mapping of state into measurement, and $\mathbf{V}_t \sim (0, \mathbf{R}_t)$ is unknown Gaussian noise, with zero mean and known covariance $\mathbf{R}_t$. It is referred to as "measurement noise" and is independent of $\mathbf{W}_t$.

The operation of the Kalman filter can be considered in three distinct phases: initialization, prediction, and updating. While the initialization only occurs once, the prediction and correction phases will alternate for the duration of the tracking. Since each iteration of X_t is dependent on its prior state, its approximation is made with prior information of all previous measurements up to the current time step.

2.2.1. Initialization
To initialize a Kalman filter, some initial state X_0 and its corresponding error covariance P_0 need to be generated. Taking some initial measurement, Z_0, X_0 is simply the inverse of $\mathbf{H}$ multiplied by Z_0. The initialization of P_0 can be simply the identity matrix, or some estimate based on measurement noise covariance $\mathbf{R}$. If the initial state is unknown or the approximation is imperfect, the value of the initial covariance can be set to some large value L to minimize the effect of the initial information over the first few states: $P_0 = L \cdot \mathbf{I}$, where $\mathbf{I}$ is the identify matrix with dimension equal to the length of X_0.

2.2.2. Prediction
The prediction step creates an approximation of the state, $\hat{X}_{t|t-1}$, at time t using the dynamic state transition system in (1). Calculating the state error covariance involves updating the prior error with $\mathbf{F}$ and the process covariance $\mathbf{P}_t$, which can then used to calculate the weighting of the input measurements on the state model during the update phase.

Using the predicted state, a forecast of the present measurement $\hat{Z}_{t|t-1}$ can be attained given no information about the current time-step. The measurement can be found simply by applying the observation mapping $\mathbf{H}$ to $\hat{X}_{t|t-1}$, as in (2).

2.2.3. Updating
The correction phase corrects the predictions using observed measurements obtained from the tracked data, Z_t. The measurement is used both to better approximate any unobservable parameters in the state model, and to correct observations to fit the given state model, assuming some noise is present in the measurement.

The first step is to calculate the optimal "Kalman gain," K_t based on the estimated error covariance and observation model. The Kalman gain is a measure dictating the certainty of measurements with respect to the state estimate, with added measurement covariance $\mathbf{R}_t$. Given the current observed measurement Z_t, the Kalman gain is used to weight the difference between Z_t and approximation $\hat{Z}_{t|t-1}$. The result is added to the predicted state to get the updated estimation $\hat{X}_{t|t}$ at time t. Similarly, the gain is used to weight state error information that is subtracted from the covariance $\hat{P}_{t|t-1}$ to get $\hat{P}_{t|t}$. With a corrected state now available in the model, the observation function (2) can now be applied to get an updated measurement estimate, $\hat{Z}_{t|t}$.

Where no measurement Z_t is available, the Kalman filter loop can continue for the next time step $t + 1$ using only the approximations gained in the prediction step. The full equations for each phase are given in

Algorithm 2 in Appendix A, which detail the appropriate matrix multiplications for each calculation. Choices of noise parameters and setting up the state model are discussed in a later section, based on the context of the tracer tracking framework for time-lapse ERT data.

2.3. Selecting Kalman Filter Parameters

The creation of the state model requires an understanding of the underlying mathematics of the object being tracked. The aim of the research is to identify a saline tracer moving under gravity in an approximately uniform horizontal flow field. A simple state model for the Kalman filter is to represent the position and velocity of the tracer at a given time. The resistivity of the tracer is not expected to change significantly for the duration of the time-lapse model, $\rho V : \mathbb{R}^d \times \mathbf{t} \rightarrow \mathbb{R}$. With the mean log resistivity of the identified tracer region, $\bar{\rho}$, included, the state model is constructed as such: $X = (\bar{\mathbf{s}}, \bar{\rho}, d\mathbf{s})^T$.

The velocity of a tracer will be used to update the prediction of $\bar{\mathbf{s}}$, such that $\bar{\mathbf{s}}_{t+1} = \bar{\mathbf{s}}_t + d\mathbf{s}_t$, whereas the values of $\bar{\rho}$ and $d\mathbf{s}$ will be dependent only on their respective prior values. The assumed additive noise is not included, $\mathbf{W}_t$, which is instead considered in the calculation of P_t. Based on the assumptions outlined, the dynamic state transition function (1) can be set up as follows:

$$\mathbf{F} = \begin{pmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \; X_t = \begin{pmatrix} x_t \\ z_t \\ \bar{\rho}_t \\ dx_t \\ dz_t \end{pmatrix} \tag{3}$$

The above is for the 2-D case, with a horizontal x by vertical z imaging plane. For the 3-D case, X_t will have parameters for the y coordinate and its respective velocity component, dy. Since the velocity was expected to be only approximately constant, it was selected as a state variable, rather than a control parameter—hence its use in X instead of $\mathbf{u}$—allowing the velocity to be calculated adaptively, rather than relying on it being a known value. There are no other controls in the defined state system, so $\mathbf{B}=\mathbf{0}$. and any components containing $\mathbf{B}$ or $\mathbf{u}$ are omitted from calculations.

Mapping the state to the observations derived from the ERT images is a simple case of taking the position and resistivity components of the state X. $\mathbf{H}$ can therefore be defined as the following (again for a 2-D case):

$$\mathbf{H} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{pmatrix} \tag{4}$$

Observations, Z, are thus of the form $(\bar{\mathbf{s}}, \bar{\rho})^T$. We can extract such values though simple averaging, with the coordinates and distribution of resistivity in a measurement Z contributing to the geometric center of a detected MSV. For each detected MSV, a measurement will consist of the center-of-mass, $\bar{\mathbf{s}} = (x, y[\text{in 3D}], z)$, and mean log resistivity, $\bar{\rho}$. Center-of-mass can be calculated using a weighted mean, such that each spatial cell in a MSV is weighted by its conductivity change from some baseline, where available: $\bar{\mathbf{s}} = \sum \theta(\mathbf{s})\mathbf{s} / \sum \theta(\mathbf{s})$, with $\theta(\mathbf{s}) = \rho V(\mathbf{s}, t)^{-1} - \rho V(\mathbf{s}, 0)^{-1}$ and $\mathbf{s}$ the set of coordinates in the current MSV. Values for $\mathbf{Q}$ and $\mathbf{R}$ are setup as square matrices, created with diagonal elements containing predefined values of noise. The values chosen dictate the results of tracking, and directly affect the distance measure used in the assignment of tracks and observations.

2.4. Tracer Tracking Framework

The tracer tracking framework (TTF) introduced incorporates the previously described concepts of MSVs and the Kalman filter. The main principle of the TTF is to iterate through the time-slices of a time-lapse ERT model and at each state identify all MSVs. Rather than just managing a single Kalman filter that follows one detected region, the TTF uses a deterministic cost function to enable tracking of multiple objects. Taking this approach allows the creation of a "track" for each distinctly identified region and assigns an independent Kalman filter to each.

The use of MSVs may lead to multiple detected regions that could correspond to the tracer, but also to stationary regions, and random or systematic artifacts in the model. Additionally, it is very likely that nested volumes will occur: where a region is considered stable at different thresholds. Using a multitrack approach avoids the need to identify a single prominent region in each time-slice. Rather, minimizing the cost function $\hat{Z}_{t|t-1}$ in each track and the observed measurements allows the assignment of the detected regions to the independent Kalman filters.

The information taken from each individual MSV from an image $\rho V(\mathbf{S}, t)$ at time step t forms the observed measurements Z_t input to the Kalman filters. The cost function used to associate predictions with observations is based on a variant of the Hungarian assignment method [*Munkres*, 1957]. For some given matrix of distances between all combinations of predictions and observations, the method will find the optimal combination of assignments by minimizing the sum of distances. The distance measure for a given observation and track prediction is as follows:

$$d(Z_t, \hat{Z}_{t|t-1}) = \left(Z_t - \hat{Z}_{t|t-1}\right)^T \Sigma^{-1} \left(Z_t - \hat{Z}_{t|t-1}\right) \tag{5}$$

where $\Sigma = \mathbf{H}\hat{P}_{t|t-1}\mathbf{H}^T + \mathbf{R}_t$, with the error covariance of the Kalman filter associated to the given track.

Costs for nonassignment for both tracks and observed measurements are also incorporated into the assignment algorithm. If a track or prediction has no associated distance lower than the nonassignment cost, it will not be assigned to a prediction or track, respectively. For the case of tracks with assigned observations, the following step is simply a case of updating the Kalman filters with the given observed measurement. Measurements that do not correspond to any tracks are used to initialize a new Kalman filter, thereby starting a new track. Where a track has not been assigned, the predicted state and covariance will be used in the next time step to make the new prediction. If a track remains unassigned for a predetermined N consecutive time-slices, it will be considered "lost." Lost tracks will consist of detected regions that were simply noisy artifacts, as well as objects that have left the imaging region. Once a track is lost, it will no longer be updated or used in the assignment function.

The TTF is described explicitly in Algorithm 3 in Appendix A. The basic outline of the algorithm for a time-lapse ERT model, $\rho V(\mathbf{S}, \mathbf{t})$, is as follows. For the first time-slice, identify the MSVs. From each stable volume, obtain a measurement vector and initialize a Kalman filter. The Kalman filters will then form the initial set of tracks. For consecutive time-slices, the algorithm will continue by predicting the existing tracks, and identifying the MSVs in this slice. Then, the assignment method will be used for the tracks and measurements (taken from the MSVs). As described above, tracks will be updated and new tracks created. The process is then iterated for the whole time series.

The output of the TTF consists of the list of active tracks and lost tracks at the end of the time series, representing all the tracked objects, and can be post processed. Extraneous tracks or detected noise may be removed automatically, for example, based on age (i.e., the number of time-steps tracked), expected velocity, or the magnitude of the error covariance.

3. Results

The examples described in the presented work demonstrate the results of the tracer tracking framework when applied to two time-lapse cross-borehole ERT tracer experiments—one 2-D and one 3-D.

The tracer tracking framework was applied to both the 2-D and 3-D experiments. For each time step, MSVs were calculated using 50 thresholds, with fixed intervals at 2% of the total range of log resistivity values. The list of thresholds was then reversed, such that it ran high to low, and MSVs were recalculated and combined with the initial results. The two-way thresholding is performed to obtain MSVs regardless of whether they shrink or grow as the resistivity threshold increases as a result of the contrast difference with the background. MSVs are removed automatically if they were outside the range of tolerated sizes, in both cases fixed between 2.5% and 25.0% of the total size of the image, inclusive.

3.1. Test Cell Experiment
The 2-D time-lapse experiment consists of a saline tracer monitored using time-lapse ERT in a reduced field scale hydraulic test cell [*Kuras et al.*, 2009]. The experimental area was an unlined elongated trench,
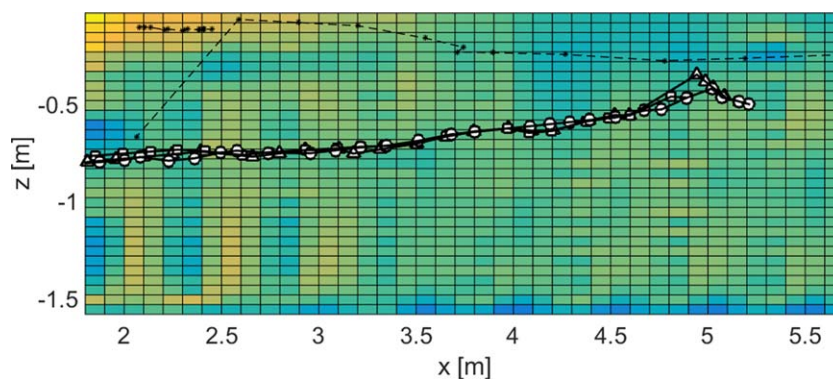
Figure 4. Kalman filter tracks showing the pathway and tracked positions of detected regions across the whole time series, plotted on the final ERT time-slice, t_{37}.

measuring 6.0 m × 1.5 m × 1.6 m, excavated from made ground comprising low permeability Cropwell Bishop Formation mudstones, and filled with fine washed and graded high silica sand (Chelford 52, average grain size 0.267 mm). Pumps and level sensors were used to produce a well-controlled horizontal flow field, with a relatively low flow rate (4–5 m/d) compared to the rate of ERT image acquisition to avoid blurring on the horizontal scale of the model cells (0.1 m) [*Wilkinson et al.*, 2010; *Rucker*, 2014].

The cell was instrumented with 19 vertical electrode arrays, 9 of which, spaced at 0.5 m intervals, formed a 2-D imaging plane along the long axis of the trench. Each array consisted of 16 stainless steel electrodes at 0.1 m separations. Data were acquired using cross-borehole bipole-bipole measurements between adjacent arrays, with a full data set being measured every 40 min over a 24 h period. The data were inverted using a 2.5-D least squares smoothness-constrained algorithm [*Loke and Barker*, 1995] with an L2 model smoothness constraint and an L1 measure of data misfit. The saline tracer had a concentration of 20 g/L and 25 L of tracer was injected into the flow field over a period of 1.5 h at a depth of 0.4 m and a distance of 0.5 m from the right-hand side of the imaging plane. The expected vertical Darcy velocity of the tracer, due to sinking under gravity, was 0.4 m/d. Further technical details of the experiment can be found in *Kuras et al.* [2009].

The results at the end of tracking are shown in Figure 4, highlighting the individually tracked regions over the whole time series. At the end of tracking, post processing was carried out to automatically remove extraneous tracks that did not meet the expected criteria: any tracks with horizontal velocity below −2 m/d relative to the flow field were considered a noisy response and removed. In addition, tracks that were only detected for 2 (or fewer) time-slices before being lost are also removed. A lost track was counted as one that was unassigned to a detected MSV for four consecutive time-slices.

The paths of the tracks in Figure 4 are represented by connecting solid and dashed lines, with point markers indicating the location of the tracked center-of-mass of a region, predicted or corrected by some detected MSV. Notable tracks of interest can be seen specifically across the central area of the imaging plane (solid with triangle, square, and circles markers), but also in the top and lower left corners of the image (around $(2.2, -0.1)$ and $(2.1, -1.2)$, respectively). The latter are indicative of stationary regions with a slight resistivity contrast from surrounding areas, detected as MSVs in multiple time series. Their velocities are relatively constant at 0 m/d and so are tracked successfully by the framework despite not being due to the tracer. It is simple then to identify and isolate these in post processing.

The three tracks representing the tracer path are highlighted in the image and shown progressing across the time series in Figure 5. This is demonstrated at 12 time-steps at evenly spaced intervals in the 2-D time-lapse model. The tracks show assignment and detection to the tracer plume (blue region of low resistivity, $\sim 1-30$ Ω m). The tracked horizontal and vertical velocity of these tracks are plotted against time in Figure 6. The plot in Figure 6a shows the tracked velocity reaching the expected region of 4−5 m/d, consistent with the flow field. Toward the end of the tracking period, the measured velocity is slightly higher. Despite this, the error bars representing standard deviation include the upper velocity limit for each time step.
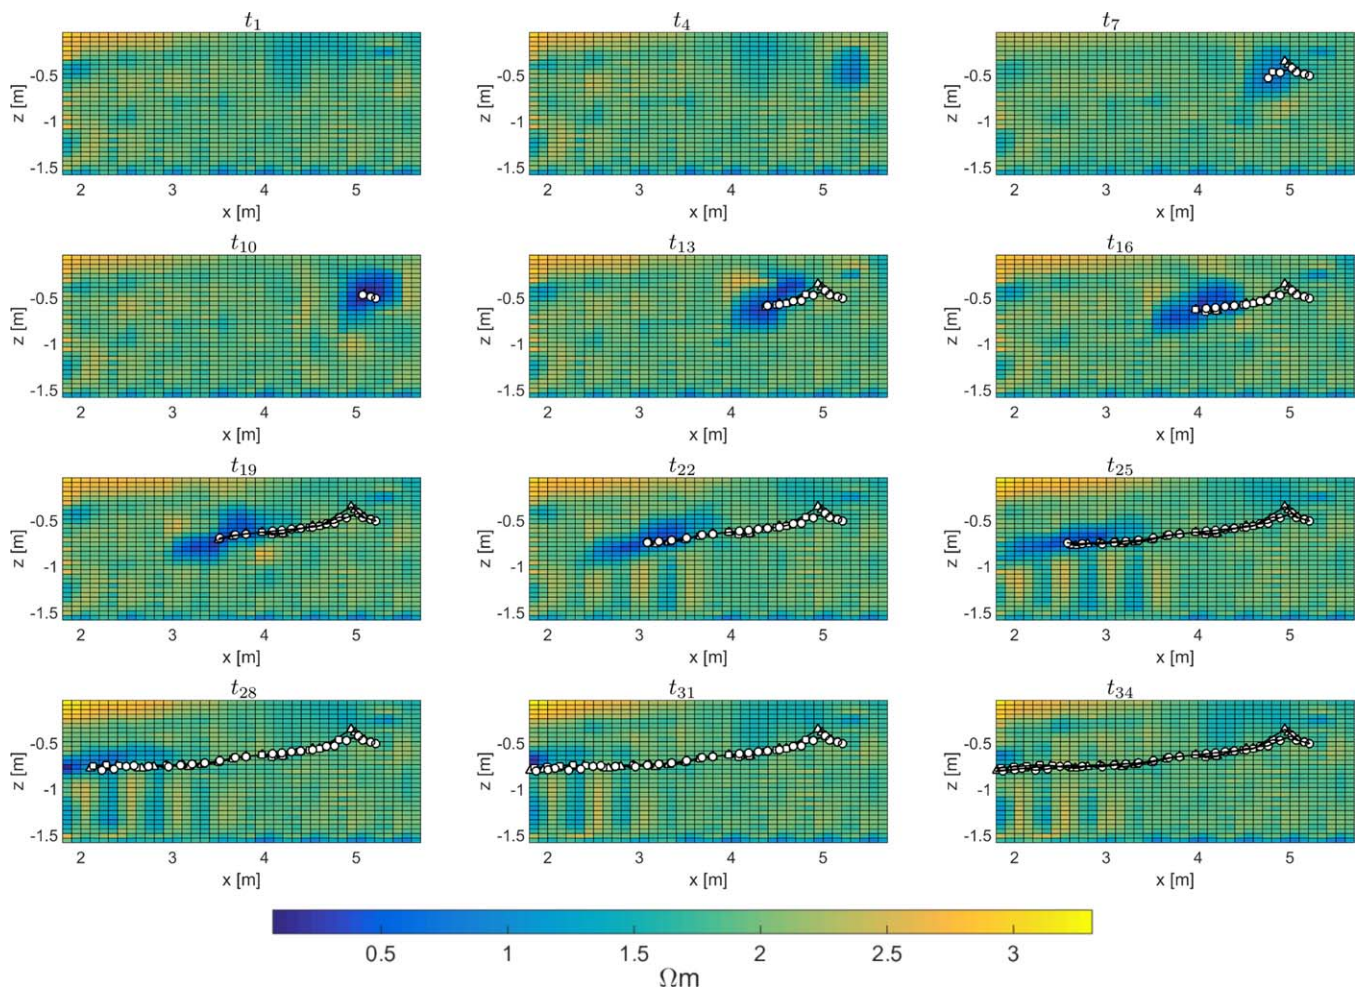
Figure 5. Time-slices of 2-D time-lapse ERT model of test cell tracer experiment, with Kalman filter tracks showing tracked positions of tracer plume up to the current time-slice.

Density-driven sinking shown in Figure 6b demonstrates the tracked vertical velocity decreasing to a value consistent with the calculated Darcy velocity of -0.4 m/d. This pattern is present in each of the tracer tracks and shows the reliability of the tracking model.

3.2. Tank Experiment

The 3-D experiment was similar to the 2-D test cell, but was carried out in a laboratory tank. A volume of 1.30 m $\times$ 0.80 m $\times$ 0.95 m was filled with Chelford 52 sand and tap water, and a hydraulic gradient was established along the long dimension of the tank, giving an approximately uniform horizontal flow field with a velocity of 1.70 m/d. The saline tracer had the same concentration as in the 2-D experiment (20 g/L) and therefore the same expected density-driven vertical Darcy velocity of 0.4 m/d. A tracer volume of 1.30 L was injected over 0.5 h at a depth of 0.20 m and a distance of 0.10 m upstream of an ERT imaging volume. The volume, of dimensions 0.80 m $\times$ 0.40 m $\times$ 0.75 m, was defined by 8 vertical electrode arrays (A1-A8) situated within the tank as shown in Figure 7, each comprising 16 electrodes at 0.05 m spacing between depths of 0.10 m and 0.85 m. As above, cross-borehole bipole-bipole measurements were carried out between pairs of arrays shown by the numbered dashed lines in Figure 7. Full data sets were acquired every hour over a 16 h period and were inverted using a 3-D algorithm with the same types of constraints used in the 2-D experiment.

The results of the maximally stable volumes algorithm (after automatically removing MSVs greater than or less than 40% and 1% of the total image size, respectively) for each time step are shown in Figure 8. The surfaces representing volume boundaries are shaded corresponding to the mean log resistivity of the MSV. From these plots, it is clear that the MSV is effective at clearly isolating distinct regions, such as those
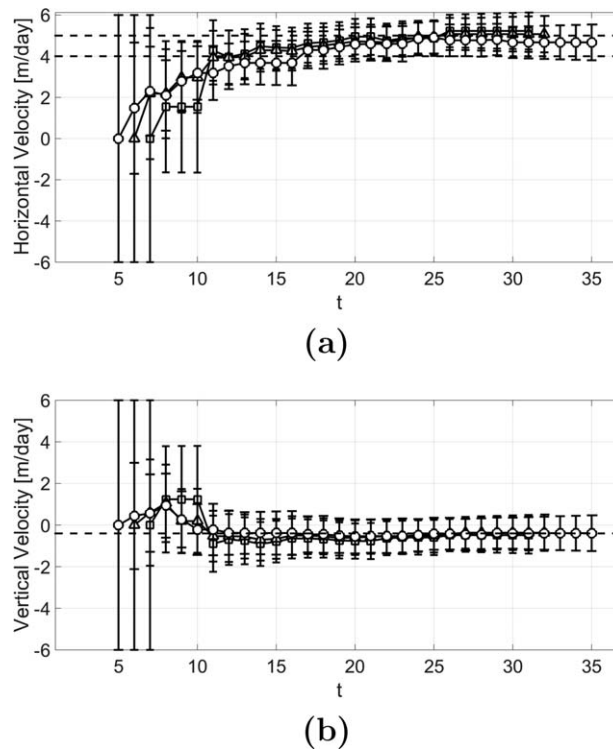
Figure 6. Plots showing (a) horizontal and (b) vertical velocity (m/d) for the three tracks shown in Figure 5 based on Kalman filter states in the test cell tracer experiment. Calculated expected velocities are marked with dashed black lines: $4-5$ and -0.4 m/d, respectively. Error bars represent one standard deviation either side of the estimation.

corresponding to the tracer plume, in 3-D ERT data. While the tracer is clearly defined, there are numerous detected regions that do not correspond to areas-of-interest. These are likely caused by artifacts in the ERT data, such as those resolved in t_1.

The tracks shown in Figure 8 represent the best fits for the tracer, found by automatically removing those that were deemed too short, were stationary, or had negative horizontal velocity with respect to the expected flow field. The plots of both lateral ($||(dx_t, dy_t)||_2$) and vertical velocity (dz_t) of these tracks are shown in Figures 9a and 9b, respectively. Each path shown clearly corresponds to the tracer plume, which at later time-steps is represented by a set of nested MSVs, and show both horizontal and vertical velocities that match the expected values (denoted by the dashed lines in Figure 9). This is especially the case for the vertical velocity. Notable deviations from the expected parameters occur toward the end of the monitoring period, where the horizontal velocities for the circle and triangle tracks begin to decrease in the final four time steps. This corresponds to the tracer plume moving

outside the imaging volume, which occurs over multiple time steps. Detections of the tracer in this plume are thus incomplete and the detected center-of-mass begins to affect the Kalman state for these tracks. The phenomenon does not affect the diamond track, as this is no longer dictated by detection and is driven only by predictions based on the prior states. However, this is a clear indicator that the tracking and estimation of velocity parameters is effective for forecasting the physical progression where no measurements are available.
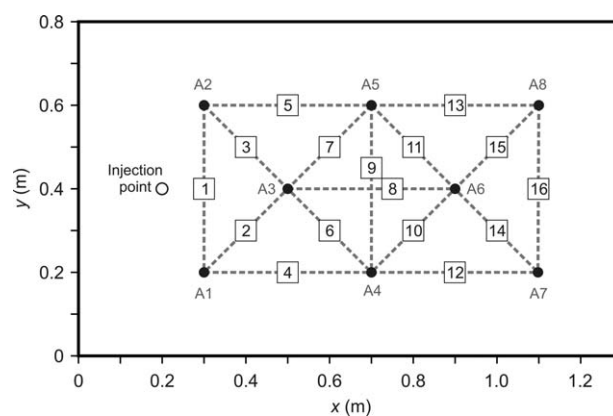
4. Limitations

Figure 7. Arrangement of vertical electrode arrays (black dots A1-8) in laboratory tank experiment, showing measurement array pairs and tracer injection point.

Despite the many benefits of the technique for the process of tracking tracer movement in ERT time-lapse data, there are a number of limitations to the framework in its current form. In using ERT data, there are greater assumptions required for the detection certainty. Due to the inversion process, it is always the case that image resolution varies throughout the imaging space [*Gharibi and Bentley*, 2005]. A degree of uncertainty is present that is nonuniform in the imaging survey, something that is not explicitly dealt with in the tracking framework.

Additionally, the detection step of the process (i.e., MSV), relies on local connectivity between resistivity cells to calculate
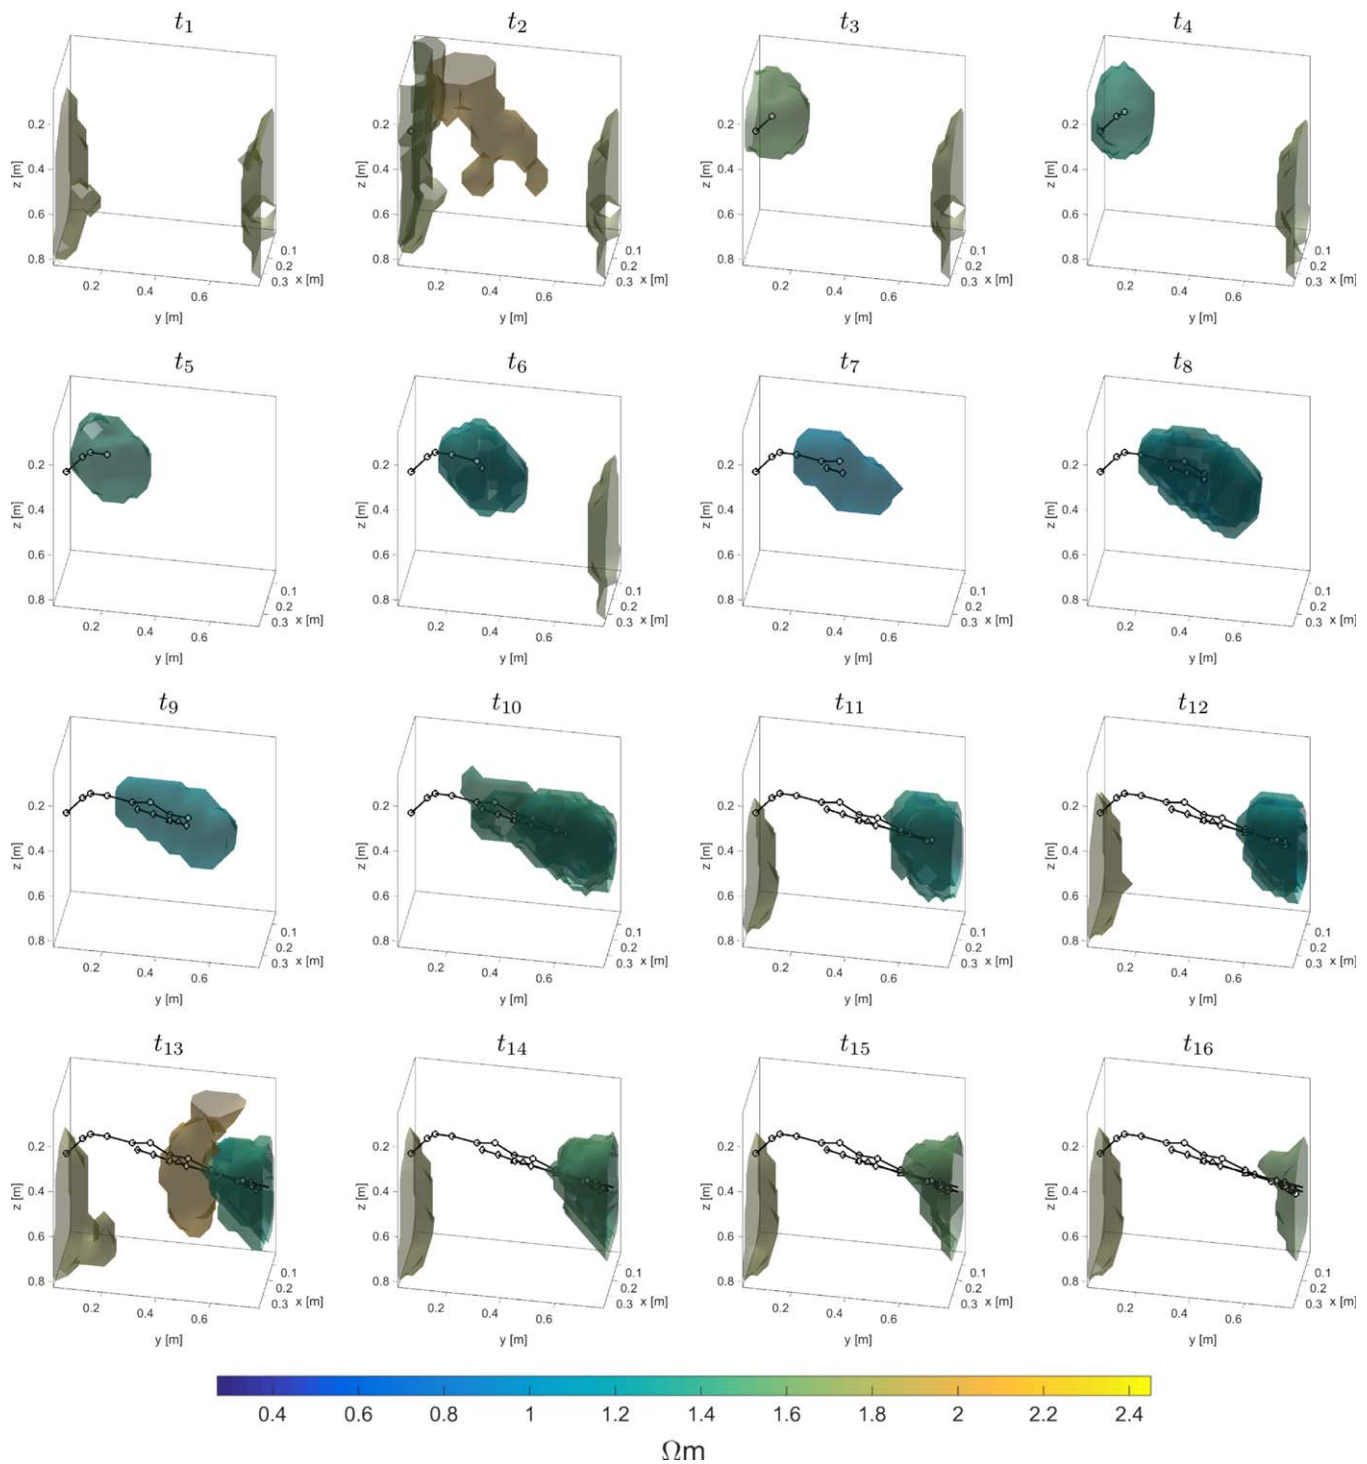
Figure 8. Time-slices of 3-D time-lapse ERT model of tank experiment, showing detected MSVs and Kalman filter tracks of the tracked positions up to the current time-slice. MSV surfaces are shaded to represent log mean resistivity.

connected components. The assumption here is that the cells are sized and distributed uniformly—which is the case for the data used, but would not necessarily be so for e.g., surface monitoring experiments. The current implementation of MSV relies on structured grids for calculating connectivity. In theory, the extension of MSVs to unstructured polygonal and polyhedral meshes is feasible, by using a measure of connectivity based
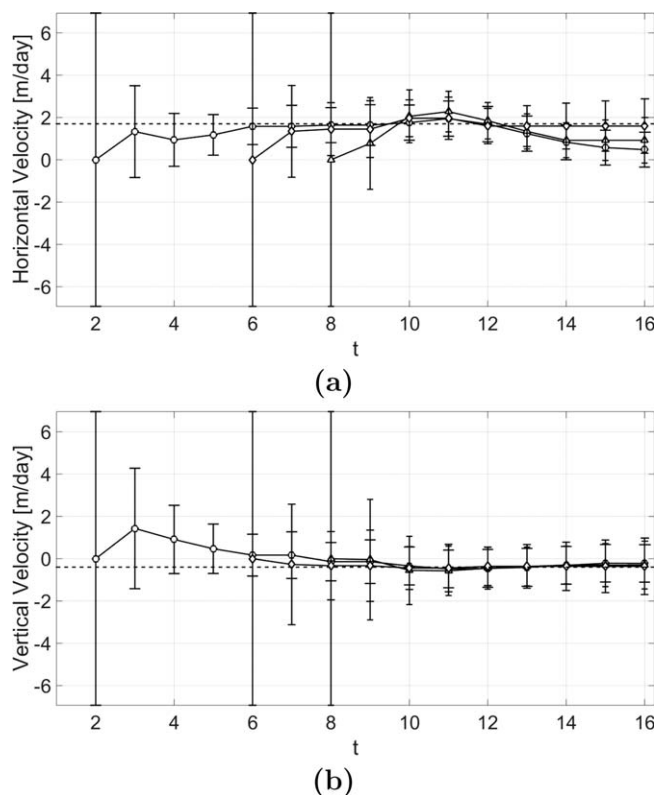
Figure 9. Plots showing (a) horizontal and (b) vertical velocity (m/d) for the three tracks shown in Figure 8 based on Kalman filter states in the 3-D tank tracer experiment. Calculated expected velocities are marked with dashed black lines: 1.7 and $-0.4\,\text{m/d}$, respectively. Error bars represent one standard deviation either side of the estimation.

on shared edges, vertices, or faces (as 2-D quadric grid uses 8-connectivity, and 3-D cubic grid uses 26-connected in the presented work). However, some investigation should be made into the effects on computational efficiency and implementation.

The estimation of tracer center-of-mass is based on a pseudomeasure of concentration that weights cells based on their observed change from some baseline. Where no convenient baseline exists, e.g., in monitoring a region where the tracer plume occurs in the initial time step(s), the assumptions behind this weighting may not hold. An alternative weighting using solely the current conductivity would additionally allow a focused estimation of the plume distribution. However, using only electrical conductivity is not as reliable in representing tracer distribution, as it does not account for background artifacts that may be removed with use of a baseline.

Despite these issues, a comparison of the two approaches, demonstrated in Figure 10, shows that there is little deviation in the overall values given, and that the estimated velocity is reliably predicted in both cases. The results give positive indications that conductivity may be used as an appropriate alternative where there is no reliable baseline.

In the use of the Kalman filter, to calculate the optimal gain, all noise is assumed to be normally distributed, whereas the data noise in ERT is strongly voltage-dependent and the model covariances are influenced by the type of smoothing constraint. Similarly, we are limited to only a linear system for state transition. Alternative methods based on Kalman filter and other Bayesian approaches are available and will be discussed in section 6.

5. Conclusions

In summary, the tracer tracking framework presented is an effective means of automatically identifying tracer movement with minor user intervention in identifying specific tracks. The results take into account inaccuracies in detection methods and variance from the expected model.

Results have clearly shown the capabilities of tracking a moving object, demonstrating simple post processing steps that can be both automated and manually controlled. Namely, short tracks and stationary tracks can easily be identified for the given state transition parameters.

The detection method can effectively identify and isolate maximally stable volumes, seen in Figure 8. The capabilities for nested volumes allow for multiple tracks to be assigned to the same tracer. Such an approach may be useful for observing different levels of concentration in the tracer—one track following the movement of high concentration region, and another initialized and tracking the larger region as the tracer disperses into the outlying environment as in the 3-D tank experiment.

Using a Bayesian approach such as Kalman filter is effective in dealing with uncertainties involved with ERT imaging. Taking into account potential errors in both the detection method and expected trajectory allows an
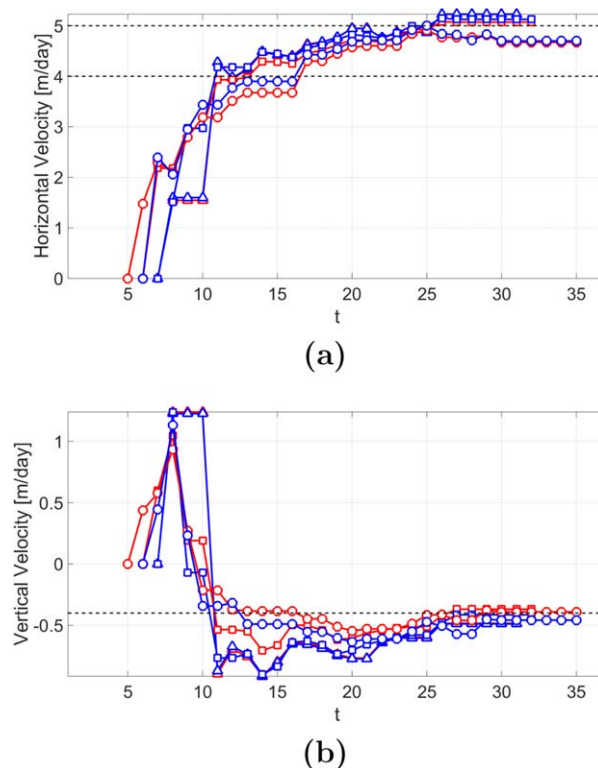
Figure 10. Plots showing (a) horizontal and (b) vertical velocity (m/d) comparing tracking results for center-of-mass calculations weighted with resistivity ratio (red) and conductivity (blue) for tracks corresponding to those in Figure 5. Calculated expected velocities are marked with dashed black lines: 4−5 and −0.4 m/d, respectively.

appropriate level of response to changes in the system. The Kalman filter in particular is useful as the calculation of Kalman gain, K_t, allows an optimal weighting of the expectation versus observation based on the covariance of state tracking. It is evident that the tracer tracking framework described here can provide useful functionality in the detection and tracking of subsurface hydrogeophysical processes, both in assisting and automating observation processes.

The approach is distinct from focused inversion [*Cardiff and Kitanidis*, 2009], in that it is reliant only on the inverted data. As a result of the predict/update cycle, with estimators accounting for the assumed uncertainty in the ERT models, the framework is largely invariant to inversion parameters. While it has been shown that predictions of the center-of-mass from smoothness-constrained ERT of tracers is relatively robust [*Singha and Gorelick*, 2005; *Doetsch et al.*, 2012], some bias exists due to spatially variable resolution. Adjusting the choices of inversion and associated constraints to optimize the solution would go some way to improve the results, but the presented method mitigates the effects by means of Bayesian inference.

In contrast to some other feature detection frameworks for ERT, which return pseudo-optimal segmentations based on multiple inversions with changing parameters [*Audebert et al.*, 2014], the detection results are corrected by the dynamic state model.

6. Further Research

The framework outlined demonstrates the effectiveness of Kalman filters applied to maximally stable volumes for hydraulic feature tracking in ERT monitoring. As discussed previously, the current state model involves only the tracer velocity and resistivity. With additional parameters in the state model, it may be possible to develop a more accurate forecasting technique.

Introducing more statistical measures, such as covariance models to represent stable volumes, will provide increased information to track. Similarly, ancillary data may be incorporated into the model as additional observations, such as local temperature measurements. By increasing the dimensionality of potential state models, both by improved information and the use of second-order moments, the complexity of state transition will be greater which, while relying on more computation, will provide a more detailed result based on increased understanding of the underlying system.

It may also be possible to use more complex state equations to model the tracer, especially if incorporating more spatial information such as higher spatial moments (e.g., those quantifying the extend of the tracer). For example, incorporating a convection-dispersion model to describe the tracer plume in a flow [*Kemna et al.*, 2002]. Such a model may be incorporated into use with the Kalman filter, however, the state system may now be nonlinear. In modeling tracers that experience decay, or in the presence of continuous injection, it would be possible to model known physical effects on the tracer.

Adaptations to the basic Kalman filter exist, designed to handle a nonlinear state transition function. Indeed, approximating a linear model derived from the underlying nonlinear system is possible by using a first-order Taylor expansion, as with the extended Kalman filter [*Ljung*, 1979].

The extended Kalman filter requires calculations of derivatives in the form of a Jacobian, which may lead to high computational expense. An alternative is to use the unscented transform to numerically sample the linear predictors [*Julier and Uhlmann*, 2004]. This unscented Kalman filter is similar to particle filter, except that the sample scheme is deterministic instead of stochastic [*Ahn and Chan*, 2014].

Regardless of the filtering technique used, alternative relations between multiple tracks may be explored. Instead of using a deterministic cost minimization approach, statistical measures may be used in the form of multiple hypothesis testing [*Cox and Hingorani*, 1996]. The method works by maintaining all permutations of track assignment and uses a likelihood function to assign final tracks.

Appendix A: Algorithms

Algorithm 1 Maximally Stable Volumes

---

inputs
$V : \mathbb{R}^d \to \mathbb{R}$ d-dimensional volume, indexed by coordinates $\mathbf{s} \in \mathbf{S} \subset \mathbb{R}^d$
$n \in \mathbb{N}$ number of threshold levels
$\Delta_{[1,n]} \subset \mathbb{R}$ ascending (or descending) vector of thresholds, indexed by $i \in [1,n] \cap \mathbb{N}$
variables
$B : \mathbb{R}^d \to \{0,1\}$ binary d-dimensional volume indexed by coordinates $\mathbf{s}$
$\mathbf{C}_{[0,n]} \subset \mathcal{P}(\mathbf{S})$ set of connected components for each threshold level, indexed by $i \in [0,n] \cap \mathbb{N}$
$T : [0,n] \times \mathcal{P}(\mathbf{S})$ tree containing connected components with $n+1$ levels
$\Psi : [0,n] \times \mathcal{P}(\mathbf{S}) \to \mathbb{R}$ stability measure for pairs $(i,C) \in T$
algorithm
$\mathbf{C}_0 \leftarrow \{\mathbf{S}\}$ (or $\{\varnothing\}$ if Δ is descending)
$root(T) \leftarrow (0, \mathbf{C}_0)$
for $i = 1 \dots n$ **do**
    for all $\mathbf{s} \in \mathbf{S}$ **do**
        $B(\mathbf{s}) \leftarrow V(\mathbf{s}) \geq \Delta_i$
    end for
    $\mathbf{C}_i \leftarrow$ set of connected components in B
    for all $\mathbf{c} \in \mathbf{C}_{i-1}$ **do**
        $children(T, (i-1, \mathbf{c})) \leftarrow \{(i, \mathbf{c}')|\mathbf{c}' \in \mathbf{C}_i, c' \subseteq \mathbf{c}\}$
    end for
    for all $(i, \mathbf{c}) \in T$ **do**
        if $i = 0$ **or** $children(T, (i, \mathbf{c})) = \varnothing$ **then**
            $\Psi(i, \mathbf{c}) \leftarrow 0$
        else
            $\Psi(i, \mathbf{c}) \leftarrow |\mathbf{c}|^{-1}(|parent(T, (i, \mathbf{c}))| - \max\{|\mathbf{c}'||(\_, \mathbf{c}') \in children(T(i, \mathbf{c}))\})$
        end if
    end for
end for
return
$\{\mathbf{c}|(i, \mathbf{c}) \in T, \Psi(i, \mathbf{c}) \geq \Psi(parent(T, (i, \mathbf{c})))$
        and $\Psi(i, \mathbf{c}) \geq \Psi(i+1, \mathbf{c}'), \forall (i+1, \mathbf{c}') \in children(T, (i, \mathbf{c}))\}$

Algorithm 2 Kalman Filter

---

inputs

Z_0 initial observations

$\mathbf{Q}_t \in \mathbb{R}^x \times \mathbb{R}^x$ state estimation noise, with dimension $x = |X|$

$\mathbf{R}_t \in \mathbb{R}^z \times \mathbb{R}^z$ measurement noise with dimension $z = |Z|$

algorithm

<u>*initialize*</u>

$\hat{X}_{0|0} = \mathbf{H}^{-1}Z_0$

$\hat{P}_{0|0} = L \cdot \mathbf{I}$

$t \leftarrow 0$: time-step

loop

    $t \leftarrow t+1$

    <u>*predict*</u>

        $\hat{X}_{t|t-1} \leftarrow \mathbf{F}\hat{X}_{t-1|t-1} + \mathbf{B}u_t$

        $\hat{P}_{t|t-1} \leftarrow \mathbf{F}\hat{P}_{t-1|t-1}\mathbf{F}^T + \mathbf{Q}_t$

        $\hat{Z}_{t|t-1} \leftarrow \mathbf{H}\hat{X}_{t|t-1}$

    input Z_t: current observations

    <u>*update*</u>

        $K_t \leftarrow \hat{P}_{t|t-1}\mathbf{H}^T[\mathbf{H}\hat{P}_{t|t-1}\mathbf{H}^T + \mathbf{R}_t]^{-1}$

        $\hat{X}_{t|t} \leftarrow \hat{X}_{t|t-1} + K_t[Z_t - \hat{Z}_{t|t-1}]$

        $\hat{P}_{t|t} \leftarrow \hat{P}_{t|t-1} - K_t\mathbf{H}\hat{P}_{t|t-1}$

        $\hat{Z}_{t|t} \leftarrow \mathbf{H}\hat{X}_{t|t}$

end loop

return $\hat{Z}_{t|t}$ observation estimate at time-step t

Algorithm 3 Tracer Tracking Framework

---

inputs

$\rho V : \mathbb{R}^d \times \mathbf{t} \to \mathbb{R}$ ERT volume, indexed by coordinates $\mathbf{s} \in \mathbf{S} \subset \mathbb{R}^d$ and time $t \in \mathbf{t}$

$\varepsilon \in \mathbb{R}$ cost of non-assignment

$M \in \mathbb{N}$ minimum age of a track before it can be considered lost

$N \in \mathbb{N}$ maximum consecutive steps a track is invisible before considered lost

variables

$\mathcal{T} \subseteq \varnothing$ set of active tracks

$\mathcal{L} \subseteq \varnothing$ set of lost tracks

algorithm

for all $t \in \mathbf{t}$ **do**

    $v \leftarrow \rho V(\mathbf{S}, t)$

    for all $\kappa \in \mathcal{T}$ **do**

        $\kappa(t) \leftarrow predict(\kappa)$

        $age(\kappa) \leftarrow age(\kappa) + 1$

        $invisible(\kappa) \leftarrow invisible(\kappa) + 1$

    end for

    $\mathbf{C}_v \leftarrow msv(v)$

    $\mathbf{Z}_v \leftarrow \{(\bar{\mathbf{s}}, \bar{\rho})|\mathbf{c} \in \mathbf{C}_v, \mathbf{s} \in \mathbf{c}, \bar{\mathbf{s}} \leftarrow \sum \theta(\mathbf{s})\mathbf{s}/\sum \theta(\mathbf{s}), \theta(\mathbf{s}) = v(\mathbf{s})^{-1} - \rho V(\mathbf{s}, 0)^{-1},$

        $\rho \leftarrow v(\mathbf{s}), \bar{\rho} \leftarrow \sum \rho/|\rho|\}$

    for all $Z \in \mathbf{Z}_v$ **do**

        $\kappa^* \leftarrow \arg\min\{cost(Z, \kappa(t))|\kappa \in \mathcal{T}\}$

        if $cost(Z, \kappa^*(t)) < \varepsilon$ **then**

            $\kappa^*(t) \leftarrow update(k^*, Z)$

            $invisible(\kappa^*) \leftarrow 0$

        else

$$\kappa' \leftarrow initialize(Z)$$
$$\mathcal{T} \leftarrow \mathcal{T} \cup \{\kappa'\}$$
$$age(\kappa\prime) \leftarrow 1$$
$$invisible(\kappa\prime) \leftarrow 0$$
     end if
   end for
$$\mathcal{L} \leftarrow \mathcal{L} \cup \{\kappa | \kappa \in \mathcal{T}, age(\kappa) > M \ \textbf{and} \ invisible(\kappa) > N\}$$
$$\mathcal{T} \leftarrow \mathcal{T} \setminus \mathcal{L}$$
end for
return $\mathcal{T}, \mathcal{L}$

Notation

s	spatial coordinates, e.g., (x, y, z).
S	spatial domain: set of coordinates $\{\mathbf{s}\}$.
t	time series of ERT slices.
ρ	log resistivity in Ωm.
ρV	electrical resistivity tomography model.
θ	electrical conductivity change.
$\lvert \cdot \rvert$	cardinality, i.e. number of elements.
$\mathbf{C}_i$	set of connected components at threshold Δ_i.
c	a connected component in $\mathbf{C}_i$.
Ψ	stability measure of a region.
Z	data observation.
X	state representation.
P	state error covariance.
$d(\cdot, \cdot)$	Kalman weighted distance function.
$(\hat{\cdot})_{a\vert b}$	estimation for time a given prior b.
$\leftarrow$	assignment.
$\cup$	set union.
$\{a\vert b\}$	set of values a that meet condition(s) b.
$a \setminus b$	set a without contents of set b.
$\mathcal{P}(a)$	power set: the set of all combinations of elements in set a.
$\varnothing$	empty set.

References

Ahn, K. W., and K.-S. Chan (2014), Approximate conditional least squares estimation of a nonlinear state-space model via an unscented Kalman filter, *Comput. Stat. Data Anal.*, *69*, 243–254, doi:10.1016/j.csda.2013.07.038.

Andrieu, C., N. de Freitas, A. Doucet, and M. I. Jordan (2003), An introduction to MCMC for machine learning, *Mach. Learn.*, *50*(1–2), 5–43, doi:10.1023/A:1020281327116.

Audebert, M., R. Clément, N. Touze-Foltz, T. Günther, S. Moreau, and C. Duquennoi (2014), Time-lapse ERT interpretation methodology for leachate injection monitoring based on multiple inversions and a clustering strategy (MICS), *J. Appl. Geophys.*, *111*, 320–333, doi: 10.1016/j.jappgeo.2014.09.024.

Cardiff, M., and P. K. Kitanidis (2009), Bayesian inversion for facies detection: An extensible level set framework, *Water Resour. Res.*, *45*, W10416, doi:10.1029/2008WR007675.

Chambers, J. E., et al. (2012), Bedrock detection beneath river terrace deposits using three-dimensional electrical resistivity tomography, *Geomorphology*, *177–178*, 17–25, doi:10.1016/j.geomorph.2012.03.034.

Chambers, J. E., et al. (2015), Spatial monitoring of groundwater drawdown and rebound associated with quarry dewatering using automated time-lapse electrical resistivity tomography and distribution guided clustering, *Eng. Geol.*, *193*, 412–420, doi:10.1016/j.enggeo.2015.05.015.

Cox, I. J., and S. L. Hingorani (1996), An efficient implementation of Reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking, *IEEE Trans. Pattern Anal. Mach. Intell.*, *18*(2), 138–150, doi:10.1109/34.481539.

Doetsch, J., N. Linde, T. Vogt, A. Binley, and A. G. Green (2012), Imaging and quantifying salt-tracer transport in a riparian groundwater system by means of 3D ERT monitoring, *Geophysics*, *77*(5), B207–B218, doi:10.1190/geo2012-0046.1.

Donoser, M., and H. Bischof (2006), 3D segmentation by maximally stable volumes (MSVs), in *18th International Conference on Pattern Recognition (ICPR'06)*, vol. 1, edited by Y.Y. Tang, et al., pp. 63–66, IEEE, Hong Kong, doi:10.1109/ICPR.2006.33.

Elwaseif, M., and L. D. Slater (2010), Quantifying tomb geometries in resistivity images using watershed algorithms, *J. Archaeol. Sci.*, *37*(7), 1424–1436, doi:10.1016/j.jas.2010.01.002.

Elwaseif, M., and L. D. Slater (2012), Improved resistivity imaging of targets with sharp boundaries using an iterative disconnect procedure, *J. Environ. Eng. Geophys.*, *17*(2), 89–101.

Elwaseif, M., and L. D. Slater (2013), Reconstruction of discrete resistivity targets using coupled artificial neural networks and watershed algorithms, in *Near Surface Geophysics*, vol. 11, edited by J. Simms, and U. Yaramanci, pp. 517–530, EAGE, doi:10.3997/1873-0604.2013045.

Gharibi, M., and L. R. Bentley (2005), Resolution of 3-D electrical resistivity images from inversions of 2-D orthogonal lines, *J. Environ. Eng. Geophys.*, *10*(4), 339–349, doi:10.2113/JEEG10.4.339.

Hsu, H.-L., B. J. Yanites, C.-c. Chen, and Y.-G. Chen (2010), Bedrock detection using 2D electrical resistivity imaging along the Peikang River, central Taiwan, *Geomorphology*, *114*(3), 406–414, doi:10.1016/j.geomorph.2009.08.004.

Ishola, K. S., M. N. M. Nawawi, and K. Abdullah (2014), Combining multiple electrode arrays for two-dimensional electrical resistivity imaging using the unsupervized classification technique, *Pure Appl. Geophys.*, *172*(6), 1615–1642, doi:10.1007/s00024-014-1007-4.

Julier, S., and J. Uhlmann (2004), Unscented filtering and nonlinear estimation, *Proc. IEEE*, *92*(3), 401–422, doi:10.1109/JPROC.2003.823141.

Kalman, R. E. (1960), A new approach to linear filtering and prediction problems, *J. Basic Eng.*, *82*(D), 35–45, doi:10.1115/1.3662552.

Kemna, A., J. Vanderborght, B. Kulessa, and H. Vereecken (2002), Imaging and characterisation of subsurface solute transport using electrical resistivity tomography (ERT) and equivalent transport models, *J. Hydrol.*, *267*, 125–146, doi:10.1016/S0022-1694(02)00145-2.

Kuras, O., J. D. Pritchard, P. I. Meldrum, J. E. Chambers, P. B. Wilkinson, R. D. Ogilvy, and G. P. Wealthall (2009), Monitoring hydraulic processes with automated time-lapse electrical resistivity tomography (ALERT), *C. R. Geosci.*, *341*(10–11), 868–885, doi:10.1016/j.crte.2009.07.010.

Liu, J. S., and R. Chen (1998), Sequential Monte Carlo methods for dynamic systems, *J. Am. Stat. Assoc.*, *93*(443), 1032–1044, doi:10.1080/01621459.1998.10473765.

Ljung, L. (1979), Asymptotic behavior of the extended Kalman filter as a parameter estimator for linear systems, *IEEE Trans. Autom. Control*, *24*(1), 36–50, doi:10.1109/TAC.1979.1101943.

Loke, M. H., and R. D. Barker (1995), Least-squares deconvolution of apparent resistivity pseudosections, *Geophysics*, *60*(6), 1682–1690, doi:10.1190/1.1443900.

Masreliez, C., and R. Martin (1977), Robust Bayesian estimation for the linear model and robustifying the Kalman filter, *IEEE Trans. Autom. Control*, *22*(3), 361–371, doi:10.1109/TAC.1977.1101538.

Matas, J., O. Chum, M. Urban, and T. Pajdla (2002), Robust wide-baseline stereo from maximally stable extremal regions, in *British Machine Vision Conference*, vol. 22, edited by D. Marshall, and P. L. Rosin, pp. 761–767, BMVA, Cardiff, U. K. doi:10.5244/C.16.36.

Munkres, J. (1957), Algorithms for the assignment and transportation problems, *J. Soc. Ind. Appl. Math.*, *5*(1), 32–38, doi:10.1137/0105003.

Rucker, D. (2014), Investigating motion blur and temporal aliasing from time-lapse electrical resistivity, *J. Appl. Geophys.*, *111*, 1–13, doi:10.1016/j.jappgeo.2014.09.010.

Singha, K., and S. M. Gorelick (2005), Saline tracer visualized with three-dimensional electrical resistivity tomography: Field-scale spatial moment analysis, *Water Resour. Res.*, *41*, W05023, doi:10.1029/2004WR003460.

Tuytelaars, T., and K. Mikolajczyk (2007), Local invariant feature detectors: A survey, *Found. Trends Comput. Graph. Vis.*, *3*(3), 177–280, doi:10.1561/0600000017.

Vedaldi, A. (2009), An implementation of multi-dimensional maximally stable extremal regions, UCLA, Los Angeles, Calif.

Ward, W. O. C., P. B. Wilkinson, J. E. Chambers, L. S. Oxby, and L. Bai (2014), Distribution-based fuzzy clustering of electrical resistivity tomography images for interface detection, *Geophys. J. Int.*, *197*(1), 310–321, doi:10.1093/gji/ggu006.

Wilkinson, P. B., P. I. Meldrum, O. Kuras, J. E. Chambers, S. Holyoake, and R. D. Ogilvy (2010), High-resolution electrical resistivity tomography monitoring of a tracer test in a confined aquifer, *J. Appl. Geophys.*, *70*(4), 268–276, doi:10.1016/j.jappgeo.2009.08.001.

Yilmaz, A., O. Javed, and M. Shah (2006), Object tracking: A survey, *ACM Comput. Surv.*, *38*(4), Article 13, 45 pp., doi:10.1145/1177352.1177355.

Yu, T.-H., O. J. Woodford, and R. Cipolla (2012), A performance evaluation of volumetric 3D interest point detectors, *Int. J. Comput. Vis.*, *102*(1–3), 180–197, doi:10.1007/s11263-012-0563-2.

Zhou, J., A. Revil, M. Karaoulis, D. Hale, J. Doetsch, and S. Cuttler (2014), Image-guided inversion of electrical resistivity data, *Geophys. J. Int.*, *197*(1), 292–309, doi:10.1093/gji/ggu001.