

University of Nottingham
School of Computer Science

THESIS SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Development of Machine Learning Techniques for Characterising Changes in Time-Lapse Resistivity Monitoring

Doctoral Dissertation of:
Wil O. C. Ward *MSci*

Principle Supervisors:
Dr Li Bai[1]

Dr Paul B. Wilkinson[2]

Secondary Supervisors:
Dr Henrik Nilsson[1]

Dr Jonathan E. Chambers[2]

[1] *School of Computer Science, University of Nottingham, UK*

[2] *Geophysical Tomography Group, British Geological Survey, UK*

in memory of Dorothy Margaret Ward,
and of Roy Ward,

my supporters, unconditionally

Abstract

Electrical resistivity tomography (ERT) is a geophysical technique for modelling the properties of the shallow subsurface. The technique provides a powerful tool for a volumetric representation of the spatial properties and spatio-temporal systems below the ground by indirectly measuring electrical properties. ERT has wide-reaching applications for imaging and monitoring in fields such as mineral exploration, infrastructure, and groundwater modelling. Developing tools that can perform predictions and analysis on the resistivity models with limited intervention will allow for ERT systems to be deployed remotely so that they might serve as an alert system, for example, in areas at risk of landslides, or groundwater contamination. However, the nature of indirect observation in ERT imaging means that there is a high degree of uncertainty in the resolved models, resulting from systematic artefacts that occur in inversion processes and from the fact that the underlying structures and processes cannot be directly observed.

This thesis presents a number of developments in automating the analysis and prediction of directly and indirectly observed uncertain systems, both static and dynamic. Drawing from principles in both fuzzy logic and probability, particularly Bayesian statistics, the different representations of uncertainty are exploited and utilised to make meaningful estimates of properties and parameters in noisy systems. The key contributions of the research presented include the unique combination of fuzzy inference systems in a recursive Bayesian estimator to resolve systems under the influence of multiple uncertain dynamic processes. Furthermore, frameworks for robustly isolating features with quantified certainty and for automatically tracking tracer moments in hydrodynamic systems are proposed and applied to a number of real-world case studies.

Acknowledgement

This thesis would not be without the help and support I have received. I would like to thank my supervisors, Dr Bai Li, Dr Paul Wilkinson, Dr Henrik Nilsson, and Dr Jon Chambers for their guidance and support while working through the challenges faced throughout my research. I would also like to acknowledge my family, and the friends and colleagues who have given me the moral support, providing advice and assistance as I have worked to become a good researcher and academic. In alphabetic order of family name, special mention to: Dr Mauricio Alvarez; Duan Jinming; James Harrison; Prof Natasa Milic-Frayling; Luke Sibbett; Kathryn Sillince; Tom Smith; and Dr Mike Smith.

Contents

List of Figures i

List of Tables vi

List of Acronyms viii

Preface x

1 Introduction 1

1.1 Motivation . . . 1
1.2 Aims and Objectives . . . 3
1.3 Methodology and Roadmap . . . 6
1.4 Contributions . . . 8
1.5 Dissemination . . . 9

2 Near Surface Geophysics 12

2.1 Electrical Properties of Subsurface Material . . . 12
2.2 Electrical Resistivity Tomography . . . 13
2.3 Principles of Groundwater and Hydrogeology . . . 15
2.3.1 Fluid Flow in Porous Media . . . 15
2.3.2 Relationship to Electrical Resistivity . . . 17
2.4 Hydrodynamic Solute Modelling . . . 18
2.4.1 Fick's Laws . . . 18
2.4.2 Diffusion Equation . . . 19
2.4.3 Advection-Dispersion Model . . . 19
2.4.4 Random Walk Particle Tracking . . . 21

3 Optimal Bayesian Filtering 24

3.1 State Space Models . . . 24
3.2 Bayesian Inference . . . 25
3.3 Kalman Filter . . . 27
3.3.1 Assumptions . . . 27

3.3.2 Linear Discrete Prediction . . . 28
3.3.3 Gain . . . 29
3.3.4 Correction . . . 30
3.4 Extensions of Kalman Filter . . . 31
3.4.1 Non-Linear Process and Observation Models . . . 31
3.4.2 Linearisation . . . 32
3.4.3 Unscented Transform . . . 33
3.4.4 Unscented Kalman Filter . . . 34
3.5 Smoothing . . . 37
3.5.1 Backward Recursive Smoothing . . . 37
3.5.2 Rauch-Tung-Striebel Smoother . . . 38
3.6 Multi-process systems . . . 39
3.6.1 Markov Switching SSM . . . 39
3.6.2 Interacting Multiple Models . . . 40

4 Fuzzy Logic and Control **44**
4.1 Fuzzy Sets . . . 44
4.2 Fuzzy Inference . . . 46
4.3 Type-2 Fuzzy Logic . . . 47
4.3.1 Type-2 Fuzzy Sets . . . 47
4.3.2 Interval Type-2 Fuzzy Logic . . . 49
4.3.3 Type Reduction . . . 52
4.4 Applications . . . 53
4.4.1 Fuzzy Clustering . . . 53
4.4.2 Predictive Fuzzy Control . . . 58

5 State of the Art **60**
5.1 Spatio-temporal Modelling . . . 60
5.2 Data Processing in Near Surface Geophysics . . . 63

6 Feature Detection in Electrical Resistivity Tomography **67**
6.1 Introduction . . . 67
6.2 Distribution-Guided Fuzzy c-Means . . . 68
6.2.1 Feature Detection for ERT . . . 68
6.2.2 Kernel Density Estimation . . . 68
6.2.3 Distribution-Weighted Distance Metric . . . 70
6.3 GFCM in Near-Surface Geophysics Applications . . . 72
6.3.1 Subsurface Interface Detection . . . 72
6.3.2 Groundwater Drawdown and Rebound Monitoring . . . 76

6.3.3 Discussion . . . 80
6.4 Maximally Stable Volumes in ERT . . . 81
6.4.1 Algorithm . . . 81
6.4.2 Feature Extraction . . . 83
6.5 Discussion . . . 84

7 Linear Filtering for Tracer Moment Monitoring 86
7.1 Problem Definition . . . 86
7.1.1 Plume Tracking . . . 86
7.1.2 Pit Experiment (2-D time-lapse) . . . 88
7.1.3 Tank Experiment (3-D time lapse) . . . 90
7.2 Tracking Centre-of-Mass . . . 91
7.2.1 Model Setup . . . 91
7.2.2 Results . . . 93
7.3 Moment Tracking . . . 97
7.3.1 Model Setup . . . 97
7.3.2 Results . . . 100
7.4 Discussion . . . 100

8 Fuzzy Rules for Filtering in Multi-Process Systems 104
8.1 Motivation . . . 104
8.2 Type-1 Fuzzy Kalman Filter . . . 105
8.2.1 Formulation . . . 105
8.2.2 Linear Models . . . 106
8.2.3 Estimate Error Covariances . . . 106
8.2.4 Component Gain . . . 108
8.2.5 Type-1 Fuzzy Inference . . . 110
8.3 Illustrative Example . . . 114
8.4 Unscented Fuzzy Kalman Filter . . . 117
8.4.1 Formulation . . . 117
8.4.2 Unscented Fuzzy Inference . . . 121
8.5 Interval Type-2 Fuzzy Kalman Filter . . . 122
8.5.1 Capturing the Uncertainty of State Estimation . . . 122
8.5.2 Credible Region . . . 123
8.5.3 Membership Functions . . . 125
8.5.4 Inference and Type-Reduction . . . 128
8.6 Discussion . . . 129

9 Experimental Evaluation 132

9.1 Method Comparisons and Metrics 132
9.2 Coordinated Turning Models 133
9.2.1 Planar System 133
9.2.2 3-D System 139
9.3 Separating Physical Processes 143
9.4 Electrical Resistivity Tomography 148
9.4.1 Process Models 148
9.4.2 Injection Types 149
9.4.3 Density-Driven Sinking 151

10 Discussion **154**
10.1 Contributions 154
10.2 Future Work 156
10.3 Conclusions 158

References **161**

List of Figures

1.1 Distribution of world water supply 1

2.1 Direct current electrode configuration with current flow lines from point A to point B and equipotential surfaces measured at points M and N 14
2.2 Pathways for particles detailing possible causes of mechanical dispersion of a solute in porous medium 16

3.1 Graphical model representing a discrete-interval state space model . . 25
3.2 Discrete Kalman filter process, showing prediction and correction cycle over increments of t, with regular measurements 31
3.3 A sampled distribution (left) transformed over some non-linear function (right), showing actual (blue) and estimated (red) distribution mean (cross) and covariance (ellipse). Original and transformed sigma points are denoted by red discs 35
3.4 Example of discrete-interval Kalman filter estimation and subsequent Rauch-Tung-Striebel smoothing on basic linear acceleration model with 100 noisy measurements. The smoother prior includes all observations 38
3.5 Markov switching state space model interpretation used by IMM and its extensions, with r SSMs, $\{M_1, M_2, \ldots, M_r\}$. Red arrows denote the switching between models, with probability $p_{ij} := \mathrm{P}(M_j \mid M_i)$ 40

4.1 An example of standard fuzzy operations, on triangular membership functions u_A and u_B, showing membership functions for complement, $u_\neg$, intersection, $u_\cap$, and union, $u_\cup$ 46
4.2 Type-1 fuzzy inference system 46
4.3 Type-2 fuzzy inference system 48

4.4 Interval type-2 fuzzy sets with uncertain mean (top) and uncertain variance, showing the upper (UMF) and lower (LMF) membership functions, and footprint of uncertainty (FOU) of the IT2 set. Comparatively, a type-1 fuzzy set is underlaid, along with the membership distributions corresponding to the extreme mean values in the membership interval (green, dotted) 49
4.5 Example clustering problem over domain $\mathbf{x} \in [0, 50] \times [0, 50]$, with four partially overlapping populations (top left). The result of k-means hard clustering gives non-overlapping regions for each population (top right), whereas a fuzzy approach like fuzzy c-means returns membership functions for each population giving handling the uncertainty (bottom). 56

6.1 Static electrical resistivity model and corresponding clustering. Intensity of clusters corresponds to certainty of clustering, proportional to each resistivity cell's membership to its primary allocated cluster 71
6.2 Density of resistivity in ERT model in Figure 6.1 calculated by KDE with bandwidth $\sigma = 0.25$. Fitted Gaussians are overlaid and used to calculated the corresponding GFCM distance metric and membership functions for each cluster 71
6.3 Membership functions (top) and clustering by k-means, iterative fuzzy c-means and distribution-guided fuzzy c-means (ordered respectively) on the ERT image in Figure 6.1 73
6.4 ERT model of the Willington site (top) with 2-clustering obtained by GFCM, with bandwidth 0.25. The interface between the two clusters is taken to be the bedrock surface, and an elevation map is shown (bottom) 75
6.5 3-D ERT survey design for Willington site, showing survey areas (red shading), lines (red lines) and borehole positions (black dots) 76
6.6 Setting of the Rock Common survey, with geological map with Ordnance Survey base layer (© Crown copyright and database rights 2015 Ordnance Survey 100021290), showing quarry boundary (dashed black line) and the study site location, with site location at national scale inset (left). Schematic field layout showing monitoring array, piezometers and well-point system (right) 77
6.7 Resistivity monitoring for pumping interruption between t_0 and $t_1 3$ (left), and log resistivity ratio plots showing change relative to base line t_0, and corresponding principle change cluster (right). Water level records for peizometer 4 and 6 are shown, and time is given in hours and minutes (hh:mm) 78

6.8 Principle change cluster area versus the water level records for piezometers 4 and 6 during the pumping experiment. The shaded region indicates the period during which pumping was reduced 79
6.9 A binary image with indendent connected components individually coloured (right) 81
6.10 2-D resistivity model with feature segmentation results from distribution-guided fuzzy clustering (middle) and maximally stable volumes (bottom). Individually coloured regions represent distinctly seperate detected regions, which, in the MSV case, represent nesting 84

7.1 Cross-sectional view of 2-D pit experiment setup, and demonstration of flow field and tracer injection 88
7.2 Time-lapse resistivity model for 2-D pit experiment over 37 time-steps showing the low resistivity tracer plume advection (ordered chronologically left to right, top to bottom) 91
7.3 Arrangement of vertical electrodes in the 3-D laboratory tank experiment, showing measurement array pairs and tracer injection point . . 92
7.4 Time-slices of the 2-D time-lapse ERT model of the pit tracer experiment, with Kalman filter tracks showing the tracked position of the centre-of-mass of the tracer plume up to the current time-slice 94
7.5 Plots showing horizontal (top) and vertical (bottom) velocity state estimates (m/day) for the three tracks shown in Figure 7.4. Calculated expected velocities are marked with dashed black lines: 4-5 and -0.4 m/day respectively. Shaded regions represent the estimation error to ± 1 standard deviation 95
7.6 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 97
7.7 Plots showing the lateral and vertical velocity state estimates for the tracks filtered in the 3-D tank experiment 98
7.8 Time-slices of the 2-D time-lapse ERT model of the pit tracer experiment, with detected and filtered ellipsoids representing the automatically tracked second-order moments of tracer at the current time-step . 101

8.1 Type-1 fuzzy Kalman filter process, showing fuzzy inference and correction cycle over increments of t, with regular measurements 114

8.2 Multi-process system with three component processes (*forward*, *left*, *right*) with sinusoidal switching, estimated with linear-discrete Kalman filter and type-1 fuzzy Kalman filter (top), with firing strengths of each component rule plotted against true system switching (dashed) . . . 116
8.3 Membership functions for each rule at $t = 50$, where the true case should result in a firing strength of $\hat{\omega}_1 = \hat{\omega}_2 = \hat{\omega}_3 = \frac{1}{3}$. The crisp input is highlighted by the dashed line. Calculated firing strengths were $\hat{\boldsymbol{\omega}} = \{0.3520, 0.3163, 0.3316\}$ for the minimum t-norm, and $\hat{\boldsymbol{\omega}} = \{0.3666, 0.3101, 0.3233\}$ for the product t-norm 118
8.4 Estimates by linear-discrete Kalman filter and type-1 fuzzy Kalman filter for state velocity parameters for each component motion model, *forward* ($x_t(3)$), *left* ($x_t(4)$), and *right* ($x_t(5)$). The dotted lines represents the 95% confidence interval for the estimates 118
8.5 Sample generated Gaussian membership functions for IT2 fuzzy set, with upper and lower membership functions generated by (8.36) using Gaussian kernels on the propagated confidence interval. Interior and exterior switching roots are highlighted with black discs 127

9.1 State observation and noisy measurements for the planar coordinated turn experiment using the coordinate and bearings-only observation models. The sensor locations used to calculate bearings are also plotted. Bearings are given in radians 134
9.2 Estimates of the $(\vec{x}, \vec{y})$ components of the object tracking experiment by six methods for non-linear filtering, using two different observation models. The left plot shows the filtered estimates based on Cartesian coordinate measurements; the right plot shows the estimates based on bearing-only measurements, based on four sensors 137
9.3 Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the planar tracking experiment using the Cartesian coordinate observation model. Turning angle, θ, was not estimated by UKF and SIR. Individual markers indicate outliers that fall outside 1.5±IQR 138
9.4 Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the planar tracking experiment using the bearings-only observation model 138
9.5 State observation and noisy measurements for the 3-D coordinated turn experiment using the coordinate and bearings-only observation models. The sensor locations used to calculate bearings are also plotted. Elevation and azimuth are given in radians 141

9.6 Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the 3-D tracking experiment using the Cartesian coordinate observation model 144
9.7 Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the 3-D tracking experiment using the bearings-only observation model 144
9.8 60 frames of a video showing a ball falling and dropping, overlaid with bounding boxes denoting the filtered estimates of the ball's position at the current frame by LKF (blue) and UIT2-FKF (red). The green dashed bounding boxs are the measurements provided to the filters . 146
9.9 Estimates of the vertical position of the ball, $\vec{y}$ at frame t, filtered by the LKL and UIT2-FKF (top) and an area plot of corresponding firing strength of each component rule per frame 147
9.10 Measurement error covariance for each resistivity cell, taken from model resolution inversion output. The uncertainty is inversely proportional to the distance of each cell from measurement electrodes 149
9.11 Firing strength for component processes of the continuous and instantaneous injection models. The dashed lines represent the expected firing strengths 151
9.12 Velocity estimates for mixed initial condition UT1-FKF applied to the pit test cell ERT experiment. Shaded regions represent 95% confidence intervals, and the black dashed lines denote calculated expected values 152
9.13 Filtered centre-of-mass of the tracer plume using separable physical processes, horizontal velocity and density-driven sinking 153

List of Tables

1.1 Index of challenges addressed in this thesis by the relevent chapter . 9

5.1 Summary of state of the art 66

6.1 Common kernel definitions for use with kernel density estimation . . 69
6.2 Detected bedrock elevation by GFCM at boreholes in Willington survey area, compared with ground truth. Error is given as the absolute difference between actual and approximated bedrock depths, in metres 77
6.3 Contributions in methods for feature detection in ERT models 85

7.1 2-D Pit Model Setup 93
7.2 3-D Tank Model Setup 93
7.3 2-D second-order moment tracking model setup 99
7.4 Contributions of the tracer tracking framework 103

8.1 Kernel definitions for generative fuzzy membership functions in the fuzzy Kalman filter, scaled such that $\mathcal{G}(\boldsymbol{x} = \boldsymbol{m}) = 1$ 112
8.2 Model setup for type-1 fuzzy Kalman filter example system 115
8.3 Model setup for linear-discrete Kalman filter for comparison against T1-FKF illustrative example 116
8.4 Root mean-square error for estimation of element variables in the unified state for example multi-process system by linear-discrete Kalman filter and type-1 fuzzy Kalman filter 117
8.5 Contributions of the development of the fuzzy Kalman filters 129

9.1 Model setup for coordinated turn filtering on a plane with two component models, and two possible measurements: coordinates and angular bearings about some n_s sensors, s^k, where $s^k_{\vec{x}}$ are $s^k_{\vec{y}}$ are the coordinates of the k^{th} sensor position. Control parameter d_t denotes the discrete time interval between states, and ϵ is some very small ($< 10^{-15}$) value to force $\mathbf{Q}^j$ to be positive-definite 135

9.2 Mean (±1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for six non-linear filtering methods on planar coordinated turning experiment, using Cartesian coordinates as the observation model 136
9.3 Mean (±1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for six non-linear filtering methods on planar coordinated turning experiment, using bearings from four sensors as the observation model 136
9.4 Model setup for object tracking on in 3-D Cartesian space with three component models, and two possible measurements: coordinates and angular bearings about some n_s sensors, s^k, where $s^k_{\vec{x}}$, $s^k_{\vec{y}}$ and $s^k_{\vec{z}}$ are the coordinates of the k^{th} sensor position. Control parameter d_t denotes the discrete time interval between states, and ϵ is some very small ($< 10^{-15}$) value to force $\mathbf{Q}^j$ to be positive-definite 140
9.5 Mean (±1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for five non-linear filtering methods on 3-D coordinated turning experiment, using Cartesian coordinates as the observation model 142
9.6 Mean (±1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for five non-linear filtering methods on 3-D coordinated turning experiment, using bearings from four sensors as the observation model 142
9.7 Model setup for tracking the falling ball with three seperate models describing seperate process effects on the ball's state. 143

List of Acronyms

EIMM Extended Interacting Multiple Models
EKF Extended Kalman Filter
ERT Electrical Resistivity Tomography
FCM Fuzzy c-means [Clustering]
FIS Fuzzy Inference System
FKF Fuzzy Kalman Filters
FOU Footprint of Uncertainty [of an Interval Type-2 Fuzzy Set]
FS Fuzzy Set
GFCM Distribution-guided FCM
IMM Interacting Multiple Models
IQR Interquartile Region
IT2 Interval Type-2 [Fuzzy Set]
KDE Kernel Density Estimation
KM Karnik-Mendel [Methods]
LKF Linear-discrete Kalman Filter
LMF Lower Membership Function [of an Interval Type-2 Fuzzy Set]
MDE Mean Distance Error
MF Membership Function [of a Fuzzy Set]
ML Mamdani-Larson [Fuzzy Model]
RMSE Root Mean-Squared Error
ROI Region of Interest
RTS Rauch-Tung-Streibel [Smoother]
SIR Sequential Importance Resampling [Particle Filter]
SSM State Space Model
T1 Type 1 [Fuzzy Set]
T1-FKF Type 1 Fuzzy Kalman Filter
T2 Type 2 [Fuzzy Set]
TSK Takagi-Sugeno-Kang [Fuzzy Model]
UIMM Unscented Interacting Multiple Models
UIT2-FKF Unscented Interval Type 2 Fuzzy Kalman Filter

UKF Unscented Kalman Filter

UMF Upper Membership Function [of an Interval Type-2 Fuzzy Set]

UT Unscented Transform

UT1-FKF Unscented Type 1 Fuzzy Kalman

Preface

The nature of interdisciplinary projects means that the work presented herein is of significant interest to research communities in both machine learning and wider fields in computer science, and in geoelectrical monitoring and the areas to which it is applicable. With this in mind, this thesis is written such that it can be clearly understood if approached from either perspective. Extensive background is given to the topics in the respective fields: Chapter 2 gives an overview of the relevant techniques in near-surface geophysics and hydrogeological processes, and Chapter 3 and Chapter 4 cover introductions in recursive Bayesian inference for state space models, and fuzzy logic respectively. In each cases, relevant reviews and tutorials are cited.

A selected bibliography for background reading is presented:

Loke et al. (2013)	This survey outlines the major developments in geoelectrical imaging, predominantly focussing on ERT. The paper outlines areas of application and highlights the challenges and limitations of current imaging and inversion techniques.
Singha et al. (2015)	This paper reviews the application of ERT and other geoelectrical methods for modelling and imaging spatio-temporal physical processes, including hydrogeological. The study highlights contemporary development of methods for data collections and analysis.
Särkkä (2013)	This book comprehensively covers Bayesian filtering for linear and non-linear dynamics, in particular the Kalman filter and its extensions.
Mendel et al. (2006), Mendel (2014)	These tutorials introduce the concepts of type-2 fuzzy logic, and how it extends on type-1 fuzzy logic. In particular, they cover interval type-2 fuzzy logic, and limitations of representing general type-2 fuzzy sets.

Nomenclature

Drawing from different areas of computer science, probability theory and physics, there is some inconsistency in general notation. To present concepts and theory consistently, and avoid overlap of concepts with similar notational representation in their respective areas, some standard notation has been altered. The nomenclature and notation system in this thesis is outlined here.

Functions are typically defined with lower case, italic letters, and round brackets, e.g. $f(\cdot)$. Letters in common use to represent functions include f and h, which are typically used to denote dynamic processes and observation models in the state space model. In these cases, g is often a miscellaneous related function, either to distinguish from f or to serve as an extra step. Functionals, i.e. functions that operate over functions, are indicated using square brackets.

Non-bold, italic capital letters, e.g. A, B, are typically used to denote classical sets, with explicit definitions of sets denoted with curly brackets, defined with explicitly ordered elements, $A = \{a_1, a_2, \ldots\}$, or as conditional subsets, $B = \{b \in B' \,|\, b > 0\}$. Fuzzy sets are denoted similarly, but with a sub-set tilde, indicating their fuzziness, e.g. $\underset{\sim}{A}$. Vectors of numbers or functions are described with lower-case bold letters, e.g. $\mathbf{a}$, and matrices are denoted with upper-case bold letters, e.g. $\mathbf{A}$. Special cases of this are vectors and matrices that denote probability moments. A vector representing the mean or expected value of a random variable or distribution are italicised lower-case bold letters, e.g. $\boldsymbol{x}$, and covariance matrices use italicised matrix notation, $\boldsymbol{P}$.

For differentiation and antidifferentiation, Leibniz's notation is used, whereby $\frac{\mathrm{d}}{\mathrm{d}x}$ is the derivative with respect to x, and $\frac{\partial}{\partial x}$ is the partial derivative with respect to x. The gradient operator, ∇ is defined by $(\frac{\partial}{\partial x_i}, i = 1 \ldots n_x)$. Integrals denoted $\int_X \cdot \mathrm{d}x$ represent the definite integral over X with respect to x.

Notation

A, $\{\cdot\}$	Classic sets
$\underset{\sim}{A}$, (A, μ)	Fuzzy sets
$\mathbf{a}$, $[\ldots]^{\mathrm{T}}$	Vector
$\mathbf{A}$, $[\ddots]$	Matrix
$x(t)$	Hidden state at time t, also represented by x_t
$\tilde{x}$	Expression of the true (unknown) hidden state
$y(t)$	Observation state at time t, also represented by y_t
$\tilde{y}$	Expression of measured observation state
$\boldsymbol{x}$, $\tilde{\boldsymbol{x}}$	Expected value of hidden state, n_x-dimensional vector
$\boldsymbol{y}$, $\tilde{\boldsymbol{y}}$	Expected value of observation state, n_y-dimensional vector
$\boldsymbol{P}^{(x)}$	State estimation error covariance, $n_x \times n_x$ matrix
$\boldsymbol{P}^{(y)}$	Observation estimation error covariance, $n_y \times n_y$ matrix
$\boldsymbol{P}^{(xy)}$	State and observation error cross covariance, $n_x \times n_y$ matrix
$\boldsymbol{Q}$, $\boldsymbol{R}$	Covariance matrices for hidden state and observation noise
$\mathbb{E}[\cdot]$	Expectation
$\mathrm{cov}[\cdot, \cdot]$	Covariance
$\cdot \mid \cdot$	Conditional indicator
$\mathrm{tr}[\cdot]$	Trace operator
$\cdot_t$	Instance at discrete-time step t
$\cdot_i$	i^{th} element of set or vector (special case: $\cdot(k)$ is k^{th} element of observation state)
$\cdot_{i;t}$	i^{th} element of vector at discrete-time step t
$\cdot^j$	Instance pertaining to the j^{th} component or fuzzy rule
$\cdot^{\mathrm{T}}$	Transpose of a matrix or vector
$\vec{\mathbf{x}}$	Cartesian coordinate vector, $\vec{\mathbf{x}} = (\vec{x}, \vec{y}, \vec{z})$
$v_{\vec{x}}$	Velocity in Cartesian coordinate direction
$\sim$	"is distributed as"
$\propto$	"is proportional to"
$:=$	"is defined"
$\mathcal{N}(\mu, \Sigma)$	(Multivariate) Gaussian with mean μ and covariance Σ
$w_i^{(x)}$, $w_i^{(p)}$	i^{th} unscented transform weights
$\omega^j(\cdot)$	Firing strength of j^{th} fuzzy rule
$\bar{\cdot}$, $\underline{\cdot}$	Upper and lower parameters in interval type-2 fuzzy logic

1

Introduction

1.1 Motivation

Electrical resistivity monitoring is a technique increasingly used to investigate hydrogeophysical processes. The physical property of electrical resistivity is the quantification of how a material opposes electric current. Measurement of the apparent geoeletrical resistivity allows for the creation of spatial models of resistivity structure beneath the ground's surface which can give insight into the petrophysical nature of the local subsurface. The collection and analysis of this data can be a time consuming and labour intensive activity. Resistivity depends on many properties of interest, including rock-type, porosity, temperature and groundwater. Of the causes of the changes in resistivity, changes in groundwater and quality are among the most important.

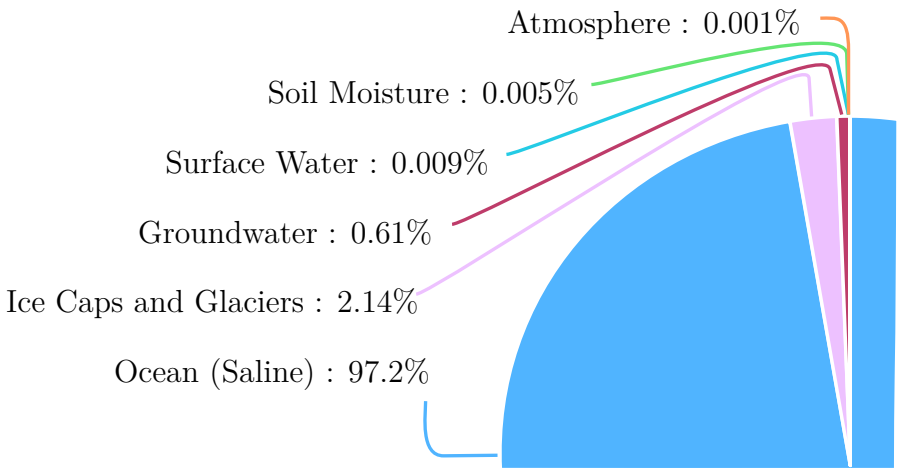


Figure 1.1: Distribution of world water supply (Fetter, 2001)

Groundwater is the largest source of freshwater on the earth, forming over 98% of available fresh water (Figure 1.1). In many parts of the world, it is the main source of drinking water. Groundwater feeds rivers and oceans and the ecosystems that they support. Groundwater plays a critical role in sustaining agriculture; it can also influence the structural properties of the ground and can lead to devastating widespread flooding. Understanding the behaviour of groundwater, and how natural and human-influenced disturbances impact its availability and quality, requires knowledge of the fundamental hydrological and hydrogeophysical processes. This, in turn, relies on adequate observational methods with clear and defined measures of uncertainty.

Using electrical resistivity tomography (ERT), it is possible to visualise underground processes occurring within tens of metres below the surface, while only carrying out a series of minimally invasive measurements. The technique works by injecting direct electrical current between two electrodes and measuring the resulting electrical potential differences between other electrode pairs. Applying these measurements across multiple pairs of electrodes, and changing the spacing between them, which allows for data to be obtained across the survey space that are sensitive to structure over a range of depths. The resulting data measure the apparent resistivity which is observed directly based on electrode separation. To create a volumetric model that shows the true resistivity with spatial positions at increasing distance from the measurement electrodes, an inversion process must be undertaken. However, there are many fundamental issues with models retrieved in this manner.

By the nature of inversion, the construction of the model is limited by noise and sensitivity that varies spatially. Enforced smoothness and constraints are typically added to assist in solving the inverse problem. As a result, boundaries between objects within the survey site can appear smoothed and gradational. In addition, variations in moisture content in the subsurface, which directly affect resistivity, tend to be gradual, meaning there may not be a distinct interface of resistivity to resolve.

Practical use of the technology ranges from mineral investigation, with static ERT, to time-lapse analysis for infrastructure maintenance and hydrogeological monitoring for imaging contaminant transport and other petrophysical changes. ERT can provide a useful visualisation tool to relevant research and industrial sectors. Creating automated computational tools that provide quantitative analysis with defined uncertainty should help the rapid uptake of this technology by reducing

the time and expert input required for full interpretation.

Time-lapse ERT is simply the process of repeating surveys, in 2-D or 3-D, in a fixed position at discrete times over some period. This enable near real-time measurements of geoelectric and hydrogeologic properties and processes. Data made available from these tools allows for early warning systems that can provide information about the ongoing physical processes. Expert ERT users are aware of these limitations, but there is a need for automated tools that can detect and classify changes to aid interpretation both by expert and non-expert users. The current state of research in this area is limited, with little overlap between the understanding and pitfalls of ERT data and numerical and predictive tools.

Any such tool for the automated interpretation and analysis of ERT data must handle the significant amount of uncertainty. Uncertainty can come from noisiness in data, which in the case of ERT is the inverted resistivity model subject to constraints, and smoothness; and in the parameters that are used to explain the underlying system. Uncertainty can be resolved through a variety of means. Among the most common is the concept of probability, that measures certainty as a degree of belief in a given prediction. Distributions of probability can be utilised here, whereby the potential value of a given event may manifest as different values with varying frequency or likelihood. This gives rise to metrics of expectation and confidence, and the nature of subjectivity. For example, a the probability of one event given another can be resolved given approximate prior knowledge.

An alternatives approach to uncertainty is the measurement of subjective truth of a particular value. A key example is fuzzy logic, whereby the class of a datum may belong to a number of possible groups. Fuzzy logic deals with the vagueness of the classification by a metric of membership, and an axiomatic logic to resolving uncertainty. It is considered distinct from probability theory amongst theoreticians in both fields, but there is significant overlap in the concepts and both the qualification and quantification of uncertainty.

1.2 Aims and Objectives

In the fields of machine learning and pattern recognition, the general principles seek to achieve a conversion of real-world data into some form of abstract understanding. The development of ideas began with the aim of emulating human vision, intelligence and reasoning in solving problems based on the perception and understanding of data. The problem of estimation is one that describes the

methods of identifying values that affect a given system based on some available empirical measurements. Probabilistic approaches treat these measurements as random variables with a distribution relating to the estimable parameters. Quantitative predictions of underlying processes require an understanding of the basic system and must draw on observed information to correct and guide assumptions. Estimators are related to inversion and regression problems, whereby some function of latent random variables $f(\mathbf{x})$ returns an observation $\mathbf{y}$:

$$f(\mathbf{x}) = \mathbf{y}$$
$$\Rightarrow \mathbf{x} = f^{-1}(\mathbf{y})$$

The functional f is presumed known, along with some amount of $\mathbf{y}$, but the latent state $\mathbf{x}$ and inverse model f^{-1} are unknown, or do not exist in an exact form. The predominant driving motivation for pattern recognition and machine learning algorithms is to identify $\mathbf{x}$, so that predictions of $\mathbf{y}$ might be made.

In geoelectrical monitoring there are a great number of assumptions, pitfalls and generalisations in the generation of images. Regardless, there is a general understanding of the influencing processes and the theoretical geophysical and hydrogeological models that affect subsurface properties and groundwater flows. To make reliable quantitative predictions, computational tools are required to convert measurements and visual information to quantified data and reliably interpret the data to identify model parameters in a robust manner. The model state may be propagated through the underlying model system to make data-driven predictions, and incorporate measurement and model uncertainty.

Examples of predictions that can be made from electrical resistivity include, but are not limited to, mineral volume estimation, structural integrity based on hydrological build-up, and contaminant monitoring. Engineering applications rely on robust interpretation and known levels of uncertainty to make appropriate judgement.

Typically ERT models comprise irregularly spaced or completely unstructured grids. When applying image analysis techniques to ERT, many of the assumptions such as regular pixel spacing or well-defined edges may not hold. Feature detectors, algorithms that make abstractions on image data and provide localised information of notable regions, must extend to unstructured grids and handle the smoothness and irregular resolution of ERT. The measurements taken from these detectors may be used as observations for estimations of the underlying system models.

Hydrogeological and petrophysical mathematical models describe theoretical groundwater processes, and how they interact with the subsurface structure. These can model the process of general groundwater flow through aquifers and porous media, such as with Darcy's law, or describe contaminant movement in pre-existing hydrological systems. Petrophysical models exist that can be used to relate electrical conductivity to hydrological properties, such as concentration ratio, and in turn be adapted to map between these and electrical resistivity. The models themselves rely on many assumptions and known properties, but in real-world systems these may not hold, or may be unknown.

Filtering with Bayesian inference has developed from the desire to understand hidden process through some observable form. For a system dictated by a transitioning process, with some observation or measurement, Bayesian filtering seeks to provide an optimal estimation of the system at a given state. The recursive filter forms a cycle of prediction and updating an estimation. Bayesian estimators differ from Monte Carlo-based ones, such as the particle filter, in that they are deterministic in their approach.

Early optimal filters were developed in the '50s and '60s to solve the problem of Gaussian and linear systems through a least squares process. The Weiner filter and the Kalman filter were seminal developments towards what became the family of Bayesian estimators. From these methods, further filters were developed to extend on functionality and theory: from continuous-time modelling in the Kalman-Bucy filter to non-linear processes with the extended and unscented Kalman filters.

This thesis seeks to address the problem of interpretation in near-surface electrical resistivity tomography. The core focus is to prototype automated systems for the identification and inference of credible patterns and features describing the underlying systems in the shallow subsurface. The research questions, and the objectives for answering them, that form the focus of this research are outlined below.

Can features-of-interest in electrical resistivity tomography be robustly identified, with little-to-no expert intervention?

i) Existing detection methods will be reviewed and assessed for suitability on detecting features
ii) Novel techniques will be developed and adapted for isolating relevant features in ERT data

iii) Context-specific requirements need to be assessed and, where suitable, appropriate detection methods should be applied

Can automatically detected features, extracted from ERT models, be used to estimate and make predictions of underlying systems dynamics?

iv) Time-lapse estimators need to be identified and developed for monitoring ERT systems

v) Methodology need to be addressed to handle complex predictive and estimation problems

vi) The developed methods should be tested on general problems and applied to ERT models

How certain are the predictions made, and how can the certainty be quantified?

vii) The uncertainty in both observations and modelling should be quantified and presented in a way that is understandable for application purposes

viii) The developed techniques should recognise uncertainty in their system and be able to use this information to make meaningful predictions and updates to the models

1.3 Methodology and Roadmap

Computer science is not just science. It is engineering; it is mathematics. In a paper by Tedre and Sutinen (2008), computer science is discussed as a distinctly unique combination of traditions. Building on the report by Denning et al. (1989) on computing as a discipline, the work outlines the intrinsically entwined paradigms that build on the traditions of science, engineering and mathematics. Computing as theory follows from mathematics in the characterisation of a study, hypothesising relationships and interpreting and determining the truth in them. Abstractions come from traditional science, whereby the hypotheses are formulated and modelled; predictions are tested against some designed experiment and analysed. From the perspective of computer science as engineering, a system might designed, and be built up of requirements, specifications and an implementation. Design is validated through testing. Both Tedre and Sutinen (2008) and Denning et al. (1989) conclude that the paradigms are distinct but so too inseparable from the

perspective of computing. Computer science blends theory, abstraction and design; the commonality between each is the value in experimentation and in providing computational support for its stages.

With this in mind, the approach taken in this study draws from each paradigm to different degrees. The problem statement is in essence an engineering problem: a problem to solve with a set of requirements, namely robust predictions with interpretability, automation and quantified uncertainty. The interdisciplinary nature tends towards providing computing tools to solve a problem, but the body of work led to the hypotheses defined above. To solve the problem in a generalised way, theory was developed, predicated on the exploration of Bayesian probability and fuzzy logic and their co-representation of uncertainty. The treatment of the developed methods to the application of time-lapse ERT monitoring served as one aspect of the interpretation and analysis of results. Abstracting the hypotheses, experiments were designed to compare with toy problems in a different, but closely related, domain, namely target tracking. All of this relates back to the development of computational tools for an applied problem through means of experimentation, but the direction of enquiry led to further contributions in both theory and generalised tools.

Being a jointly supervised project, this thesis has significant interdisciplinary components to it. The key contributions have great relevance to the relatively young field of geoelectrical monitoring, and rationalise and motivate for its use in engineering solutions in mineral exploration, infrastructure and civil engineering solutions. Indeed, Chapter 6 gives practical examples of the how machine learning and pattern recognition can be applied to solving problems of resolution and automation in ERT monitoring. The methods developed in Chapter 6 and Chapter 7 can be considered at the forefront of automated processing and interpretation in geoelectrical monitoring. The techniques are not just applicable to ERT, but would easily extend to related imaging technologies of electrical impedance tomography and capacitative resistivity imaging. The background required for interpreting ERT data and the underlying systems is given in Chapter 2.

In quantifying and characterising uncertainty for ERT data, the relationship between fuzzy logic and probability theory is examined. Backgrounds of optimal Bayesian filtering and fuzzy logic are discussed in Chapter 3 and Chapter 4, respectively. A literature review is discussed in Chapter 5, covering the state-of-the-art in spatio-temporal problems, including probabilistic filtering and predictive fuzzy control (Section 5.1), and how machine intelligence has been applied to near

surface geophysics (Section 5.2). Initially, the tools for approximating probability are applied in guiding the fuzzy membership problem (Section 6.2) with focus on application, but as Bayesian probability is utilised, this serves as inspiration for the core theoretical and methodological contributions. A unison of fuzzy inference and recursive Bayesian inference is exploited to solve a problem of probability using the concept of truth and membership that comes with fuzzy inference. The derivation is motivated by error minimisation over different possibilities of a system, and analogues are identified between the proposed methods under specific assumptions and existing probabilistic modelling tools. The resulting fuzzy Kalman filter and its variants are developed and discussed in Chapter 8, with an illustrative toy example.

To validate the proposed filtering methods a series of experiments were designed and implemented. Looking beyond the application of near-surface geophysics, the problem of target tracking serves as an analogue to demonstrate the general capabilities of the proposed methods. In Chapter 9, a series of demonstrative experiments are performed, with the proposed filtering methods compared against existing methods. Following this, the theoretical developments are applied to the a groundwater modelling problem, using the ERT as an observation model. The contributions and conclusions of the thesis are discussed in Chapter 10.

1.4 Contributions

In Table 1.1, the problems and associated challenges addressed in this thesis are summarised. Each problem is based on the objective described, and the contributions in this thesis, particularly algorithmic solutions, address multiple challenges.

The following list outlines the contributions in terms of the aforementioned contributed algorithms:

Distribution-guided fuzzy c-means The GFCM method is adapted from fuzzy c-means to exploit density information as a prior for cluster initialisation and assignment. It avoids the need for an time-complex iterative scheme by defining distance based on the density information and is therefore $\mathcal{O}(n)$

Tracer tracking framework This framework was defined for tracking and modelling tracer movement in a subsurface hydrodynamic system, using time-lapse ERT as an observational model. It combines different feature representations with multi-object Bayesian filtering methods and is shown to accurate predictions

Problem	Challenges	Chapters
Characterising subsurface features using ERT	Smoothness and noise in the inversion leads to indistinct boundaries	2, 6
Distinguishing regions in stationary ERT models	Regions of interest vary between cases, so method requirements could differ, especially without user intervention	4, 6
Detecting change in time-lapse models	Solutions to stationary images require additional considerations when extending to time-lapse data	6, 7
Combining mechanistic models with data	Exploiting known or assumed mechanistic models to better recognise and predict the underlying system and its features	2,3, 7, 9
Handling complex prediction problems	Physical systems and noisy models can lead to systems that are nonlinear and difficult to resolve, especially when there is great uncertainty both in the observation and system model	8, 9
Quantifying the uncertainty in the system and making use of this in predictions	Propagation of uncertain estimates can lead to catastrophic failure of an estimator if some metric of that uncertainty is not maintained within the system estimation	3, 6, 7, 8
Presenting the uncertainty in the system	In a system with uncertainty, an point estimate may be one of many possibilities; this should be presented accordingly	3, 4, 6, 7, 8, 9

Table 1.1: Index of challenges addressed in this thesis by the relevent chapter

of system parameters

Fuzzy Kalman filters A series of extensions to the Kalman filter are developed in this thesis, using fuzzy inference systems to resolve state estimation over multiple dynamic processes. The type-1 fuzzy (linear) Kalman filter is derived, and extended to non-linear models using the unscented transform. The representation of uncertainty in process choice is extended using type-2 fuzzy logic and the proposed filters are compared experimentally with common alternatives to filtering approaches, performing favourably in all cases

1.5 Dissemination

The following papers have resulted from work produced in this thesis (Chapter 6 and Chapter 7):

- W. O. C. Ward, P. B. Wilkinson, J. E. Chambers, L. S. Oxby, and L. Bai. Distribution-based fuzzy clustering of electrical resistivity tomography images for interface detection. *Geophysical Journal International, 197(1):310-321, Feb. 2014*
- W. O. C. Ward, P. B. Wilkinson, J. E. Chambers, H. Nilsson, O. Kuras, and L. Bai. Tracking tracer motion in a 4-D electrical resistivity tomography experiment. *Water Resources Research, 52(5):4078-4094, 2016*

- J. E. Chambers, P. B. Wilkinson, S. S. Uhlemann, J. P. R. Sorensen, C. Roberts, A. J. Newell, W. O. C. Ward, A. M. Binley, P. J. Williams, D. C. Gooddy, G. Old, and L. Bai. Derivation of lowland riparian wetland deposit architecture using geophysical image analysis and interface detection. *Water Resources Research, 50:5886-5905, 2014*
- J. E. Chambers, P. I. Meldrum, P. B. Wilkinson, W. O. C. Ward, C. Jackson, B. Matthews, P. Joel, O. Kuras, L. Bai, S. Uhlemann, and D. Gunn. Spatial monitoring of groundwater drawdown and rebound associated with quarry dewatering using automated time-lapse electrical resistivity tomography and distribution guided clustering. *Engineering Geology, 193:412-420, 2015*

The following paper based on parts of Chapter 8 and Chapter 9 has been submitted:

- W. O. C. Ward, H. Nilsson, and L. Bai. A type-1 fuzzy inference system for the kalman filter in switching models. *IEEE Transactions on Signal Processing, submitted to*

Work from this thesis has also been presented in the following conferences:

- W. O. C. Ward, P. B. Wilkinson, J. E. Chambers, and L. Bai. Comparative investigation of guided fuzzy clustering and mean shift clustering for edge detection in electrical resistivity tomography images of mineral deposits. *In EGU AGM Geophysical Research Abstracts, volume 16, 2014a*
- W. O. C. Ward, P. B. Wilkinson, J. Chambers, and L. Bai. Predicting and tracking spatiotemporal moments in electrical resistivity tomography. *In AGU Fall Meeting Abstracts, 2015*
- W. O. C. Ward, P. B. Wilkinson, J. Chambers, and L. Bai. Tracer moment tracking and forecasting in time-lapse electrical resistivity tomography. *In 3rd International Workshop on Geoelectrical Monitoring (GELMON), 2015*
- J. Chambers, P. Meldrum, D. Gunn, P. Wilkinson, S. Uhlemann, O. Kuras, and R. Swift. Proactive infrastructure monitoring and evaluation (PRIME): a new electrical resistivity tomography system for remotely monitoring the internal condition of geotechnical infrastructure assets. *In 3rd International Workshop on Geoelectrical Monitoring (GELMON) 2015*
- S. Uhlemann, J. Chambers, P. Wilkinson, M. Kirkham, and M. Beamish. Archaeological and resources assessment of sand and gravel deposits using 3-D

ERT and automated interface detection. *In Near Surface Geoscience - 22nd European Meeting of Environmental and Engineering Geophysics, 2016*

Other co-authored journal and conference papers, published during the period of study of this PhD, but not contributing directly to the thesis:

- J. Duan, W.O.C. Ward, L. Sibbett, Z. Pan, L. Bai. Introducing diffusion tensor to high order variational model for image reconstruction. *Digital Signal Processing, 2017*
- Y. Ding, M.C. Pardon, A. Agostini, H. Faas, J. Duan, W.O.C. Ward, F. Easton, D. Auer, L. Bai. Novel Methods for microglia segmentation, feature extraction and classification. *IEEE/ACM Trans. Comput. Biol. Bioinform, 2016*
- C.R. Nortje, W.O.C. Ward, B.P. Neuman, L. Bai. Spherical harmonics for surface parametrisation and remeshing. *Mathematical Problems in Engineering, 2015*
- J. Duan, B. Haines, W.O.C. Ward, L. Bai. Surface reconstruction from point clouds using a novel variational model. *SGAI: Research and Development in Intelligent Systems XXXII, 2015*
- Y. Ding, W.O.C. Ward, J. Duan, D.P. Auer, P. Gowland, L. Bai. Retinal vasculature classification using novel multifractal features. *Physics in Medicine and Biology, 2015*
- Y. Ding, W.O.C. Ward, T. Wästerlid, P.A. Gowland, A.M. Peters, J. Yang, S. Nakagawa, L. Bai. Three-dimensional vessel segmentation using a novel combinatory filter framework. *Physics in Medicine and Biology, 2014*

2

Near Surface Geophysics

2.1 Electrical Properties of Subsurface Material

Applied geophysics is used for the mapping and exploration of the subsurface by means of remote measurements. Some physical property is estimated and results are used for the interpretation and analysis. Successful geophysical surveying is dependent on measureable contrasts in the physical properties of the targeted medium. Such properties commonly investigated include density, magnetic susceptibility, electrical conductivity and radioactivity (Milsom and Eriksen, 2011).

Advancements in near-surface geophysics has largely been driven by oil and mineral industries, especially in the early 1900's. Instrumentation is designed typically to measure between the ranges of 0.5 m to 100 m below the surface. Beyond these resource industries, geophysical surveying is used in archaeological, environmental and engineering surveys. The basic objective of geoelectrical monitoring is to inject electric current into the ground and measure the effects that subsurface material have on voltage. How the current is conducted and resisted is dependent on the electrical properties of subsurface material (Milsom and Eriksen, 2011).

Electrical conduction is the movement of charged particles through a medium and resistivity is the measure of how this is inhibited. Different rock types exhibit varying degrees of resistivity depending on factors such as their chemical make-up, density, and porosity (Loke, 2012). Hard materials such as igneous or metamorphic rocks tend to have highly resistive properties, while sedimentary rocks such as sandstone have lower resistivity values. Unconsolidated sands and gravels are most likely to have the lowest base resistivity, but this too can vary. For example, soils with high clay content are more conductive than sandy soils (Chambers et al.,

2011).

Resistivity in porous or fractured material is largely dependent on the saturation of pores. Where the water content is high, and a material is fully saturated, the most porous medium will be a greater conductor. If there are many air-filled fractures, the resistivity will be high as electrical charge does not pass easily through air.

Groundwater affects the electrical properties similarly based on its ionic make-up. In saturated porous media, the presence of contaminants such as saline solutes, or dense non-aqueous phase liquids can change local resistivity. For example, in regions of high saline concentration, the electrical conductivity is high due to the presence of charged ions (e.g. sodium and chloride ions). As a result of this, saline tracers are well resolved in electrical resistivity imaging (Singha and Gorelick, 2005; Kuras et al., 2009).

2.2 Electrical Resistivity Tomography

Geophysical methods that rely on measured voltages or the magnetic fields associated with electric current can either rely on passive current or that which is artificially generated. Electrical resistivity tomography (ERT) relies on current passed through the ground by direct injection, normally using direct current (DC). This current is reversed at short intervals, normally to cancel background potentials and certain types of electronic noise effects. ERT produces an estimate of the electrical resistivity in the survey imaging area (Loke et al., 2013).

The initial means of measuring ground resistivity is by passing electric current between a pair of electrodes. Measurements of the voltage are taken between a separate pair of electrodes, to avoid such issues as contact resistance in the current pair. There are many ways to arrange these arrays, each with different sensitivities to structure in the subsurface. For example, the differences in electrode spacing between the current and potential electrode, a, and between potential electrodes, b, affects the depth of measurable resistivity. Along with this, different combinations can affect the measurement quality under various conditions (Milsom and Eriksen, 2011).

To create resistivity cross-sections, ERT uses electrode arrays and passes current between increasing separations to generate a pseudo-section of apparent resistivity against position at depth. This pseudo-section only gives a measure of apparent resistivity, and does not reflect the true resistivity of the subsurface, since it relies

on homogeneous characteristics that may not be present. To convert the resistivity section into a *true* resistivity model, an inversion process must be undertaken to identify the model properties that would produce the observed apparent resistivity pseudo-section (Loke, 2012).

Geophysical inversion seeks to identify a model (potentially one of many) that gives a response that matches the measured values to the level of noise in the data (Menke, 1984). This model is a mathematical representation of the resistivity properties of the imaged section of earth. Solving the inverse problem is typically done using finite element method solutions on structured and unstructured meshing, and least-squares to find a resistivity image as a function of depth and position: $V_\rho(\vec{\mathbf{x}})$. Since the problem is often ill-posed, non-unique and mixed determined, prior information is required, often in the form of constraints that yield a stable, unique solution. Typically some Tikhonov regularisation constraint is used, using either L1 or L2 norms (Loke et al., 2013). The best solution is considered as the best fit based on a metric of closeness, but may suffer from both over- and under-determinedness (Menke, 1984).

Different types of electrode array setups can be used to vary the resolution at different locations in the subsurface. For example, a cross-borehole electrode setup is suitable for time-lapse monitoring of processes at depth. However, regardless of electrode configuration and array type, the resistivity resolution in the inverted

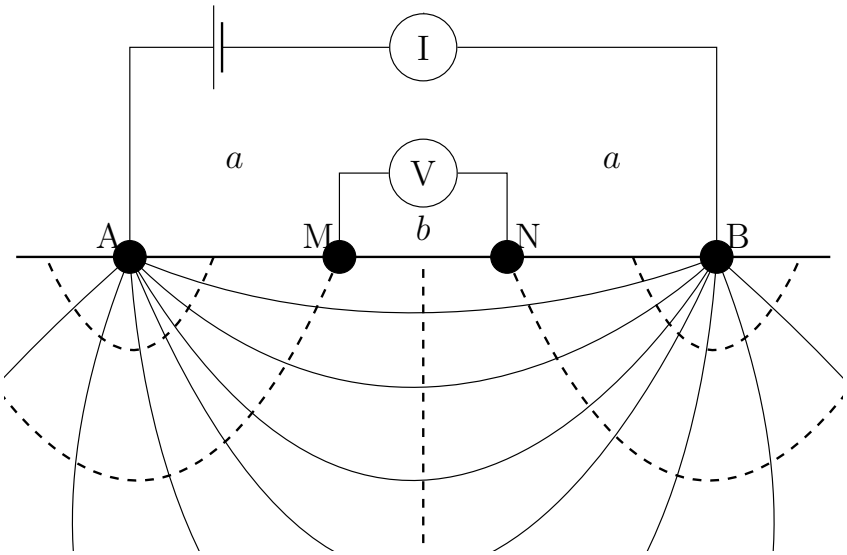


Figure 2.1: Direct current electrode configuration with current flow lines (solid) from point A to point B and equipotential surfaces (dashed) measured at points M and N. Electrode distances a and b are array parameters used in inversion

model decreases as a function of distance from the electrodes. Model resolution and sensitivity are both metrics of the distribution of information in a given ERT model (Tsourlos, 1995). Singha and Gorelick (2005) discusses the effects of smoothness constraints on field experiments of resolving solute concentration using ERT. It was demonstrated that decreased sensitivity of the ERT to detect features with increasing distance from electrodes, and smoothing as a result of regularisation were predominant factors in the difficulties by ERT to resolve petrophysical relationships in space and time. This relationship was also discussed in the works by Gharibi and Bentley (2005) and by Rucker (2014).

2.3 Principles of Groundwater and Hydrogeology

2.3.1 Fluid Flow in Porous Media

Groundwater processes are largely dominated by the laws of physics and thermodynamics. The major forces that affect groundwater are gravity, external pressure, surface tension and friction (Fetter, 2001). The former is the downward pull that exists in most physical system, and draws the water down. The second describes any force that affects groundwater as a result of processes outside the aquifer it is contained in: for example, rain-caused refill or atmospheric pressure. Surface tension as a result of molecular attraction also describes the adsorption and adherence of water particles to physical surfaces and the capillary mechanism. Frictional forces resist the flow of the fluid through the subsurface.

Groundwater flow describes the movement of fluid in porous media beneath the ground. Flow can be described as a function of space and time, describing the differential between two measured points. The rate of flow through a unit parallelepiped is termed the flux. The average flux across some defined area describes the velocity of a flow field. Flow is affected by a variety of physical properties, and can be modelled by *Darcy's law*.

Darcy's law is the mathematical relation that describes the flow of fluid in porous media. In its simplest terms, it can be stated as the negative hydraulic pressure gradient, ∇h, multiplied by the hydraulic conductivity, K, of the porous medium. The units of both K and velocity $\vec{\mathbf{v}}$ are given in terms of distance over time, and ∇h is unitless.

$$\vec{\mathbf{v}} = -K\nabla h \tag{2.1}$$

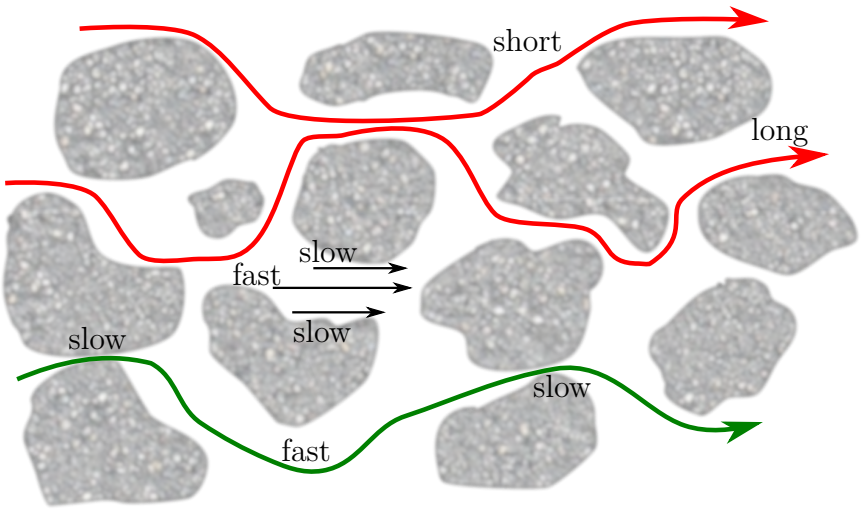


Figure 2.2: Pathways for particles detailing possible causes of mechanical dispersion of a solute in porous medium. The figure demonstrates short versus long paths a particle may take; the effect of pore size on the speed of a particle; and the effects of friction and adsorption

---

Hydraulic conductivity describes the ease with which water particles can pass through gaps in the media. K is high for very porous media, such as gravel and course sand, as wells as fractured rock, and lower for finer sands. Some materials such as clay and granite have very low hydraulic conductivities, meaning that they hinder or even prevent groundwater flow (Bear, 1972).

When assessing quality and contamination of groundwater, a basic measure of water quality is to calculate the dissolved solids, which is defined as the total amount of solids after a water sample is fully evaporated. In real-time monitoring, however, this is not always enough: studies on contamination often also consist of modelling solute movement in groundwater systems (Fetter, 2001).

Consider a solute, i.e. a substance of particles dissolved in fresh water, moving through a porous media in a gradient flow field. The underlying physical system describes both chemical interactions and mechanical effects. As with any molecular or ionic solute, it is expected that the particles will experience chemical diffusion as they move from regions of high concentration to regions of lower concentration. Along with this, the solute will experience advective forces, i.e. predominantly horizontal convection, as a result of the bulk fluid movement.

In porous media, systems experience advective and diffusive processes and also experience mechanical dispersion as a result of interaction with the surrounding

rock or soil. The dispersion causes increased mixing with the solvent groundwater. There are three basic contributories to mechanical dispersion: where fluid will move faster in the middle of the pores than along the edges; when particle pathways are of varying length; and where fluid passes through larger pores it will travel faster than when moving through smaller pores. This is visualised in Figure 2.2. Dispersion can also occur vertically, perpendicular to the flow direction, as flow paths bifurcate and branch into pores and fractures.

Hydrodynamic dispersion of a solute describes the combination of chemical diffusion and mechanical dispersion. In most cases, mechanical dispersion is predominant in the system:

$$D = D^* + \underbrace{\vec{\mathbf{v}} \cdot a(\vec{\mathbf{x}})}_{\text{mechanical dispersion}} \tag{2.2}$$

Here, D^* denotes the diffusion coefficient of particles in the solution: the slowest of any anions and cations in ionic solutes. The mechanical dispersion term describes the aforementioned processes, and is formed of a function of velocity field $\vec{\mathbf{v}}$ and dynamic dispersivity $a(\vec{\mathbf{x}})$. The dynamic dispersivity is a function of space describing a logarithmic relation to the length of the flow path (Xu and Eckstein, 1995). The mechanical dispersion component forms the diagonal elements of a diffusion tensor. The dynamic dispersivity function is defined:

$$a(\vec{\mathbf{x}}) = \alpha \cdot (\log \|\vec{\mathbf{x}} - \vec{\mathbf{x}_0}\|)^{\beta} \tag{2.3}$$

The parameters α and β are defined based on experimental relationships, commonly reported as 0.83 and 2.414 respectively as calculated through a weighted least-squares fitting mapping over the units of metres to metres.

2.3.2 Relationship to Electrical Resistivity

It is possible to use the petrophysical relations between electrical conductivity and porosity and saturation using *Archie's law* (Singha and Gorelick, 2005). Archie's law is an equation describing the expected electrical conductivity, κ, of the saturated rock as a function of media porosity, ϕ, saturation, θ_s, and electrical conductivity of the solute, κ_s:

$$\kappa = a^{-1}\phi^{-m}\kappa_s\theta_s^{-n} \tag{2.4}$$

The parameters a, m, and n respectively represent tortuosity, cementation and saturation factors of the porous media and are often determined experimentally. However, since conductivity is inversely related to resistivity: $\kappa = \rho^{-1}$, a version of Archie's law exists for electrical resistivity.

To determine the saturation of the porous media, the following rule applies, with baseline ρ_0 denoting the resistivity of the saturated media without contamination.

$$\theta_s = \left(\frac{\rho}{\rho_0}\right)^n \tag{2.5}$$

2.4 Hydrodynamic Solute Modelling

2.4.1 Fick's Laws

The general models of diffusion are described by a pair of rules describing how the flux of a substance relates to concentration in a steady state system, and how this changes over change. The spatial gradient of the concentration multiplied by the diffusivity of an area derives the diffusion flux, J. The flux describes the amount of substance that can flow through a unit volume during a unit time interval. Fick's first law describes this relationship, for diffusion coefficient, D and concentration $C(\vec{\mathbf{x}}, t)$. Note that dispersivity may be a function of space, in the case of heterogeneous media, $D := D(\vec{\mathbf{x}})$.

$$J = -D\nabla C \tag{2.6}$$

The second of Fick's laws predicts the change in concentration as the divergence of the diffusive flux, J, such that the following equality holds:

$$\frac{\partial C}{\partial t} + \nabla \cdot J = 0 \tag{2.7}$$

2.4.2 Diffusion Equation

Relating the two equations that form Fick's laws, the calculation of flux, in (2.6) may be factored into the partial differential equation of concentration over time in (2.7). The resulting differential equation describes the diffusion equation, for non-constant dispersitivty D.

$$\frac{\partial}{\partial t}C(\vec{\mathbf{x}},t) = \nabla \cdot (D(\vec{\mathbf{x}})\nabla C(\vec{\mathbf{x}},t)) \tag{2.8}$$

The diffusion equation is a generalised form of the heat equation, and as such, the resulting solution is typically Gaussian in nature. An example of this can be seen in the analytic solution for the diffusion equation in 1-D, with constant dispersivity and an initial condition that is an instantaneous point injection $C(\vec{x},0) = C_0\delta(\vec{x})$. The resolved concentration is proportional to the general solution to the heat equation at time t. Here, C_0 refers to the initial concentration of solute injected, and $\delta(\cdot)$ is the Dirac delta function.

$$C(\vec{x},t) = \frac{C_0}{\sqrt{4\pi Dt}}\exp\left(-\frac{\vec{x}^2}{4Dt}\right) \tag{2.9}$$

Here, the analytic solution is a normal distribution of solute particles with variance as a function of time, $C \sim \mathcal{N}(0,\sigma^2(t))$, where $\sigma^2(t) = 2Dt$. This result allows for the general phenomenon that where a transport system can be considered Fickian, i.e. it follows Fick's laws, the general dispersion of particles can be approximately assumed to be proportionally Gaussian. Where dispersivity is non-constant, there may be some convolution of the shape of the Gaussian, most notably with the shift in porosity in heterogeneous media.

2.4.3 Advection-Dispersion Model

While the diffusion equation describes how a substance will spread, in systems under the influence of flow fields, there can be some velocity component describing convection of particles. This generally can be assumed as the shift of concentration at a rate approximately equal to the velocity $\vec{\mathbf{v}}(\vec{\mathbf{x}})$ of the underlying flow field. As a differential equation, this is an additive term combined with the diffusion equation defined in (2.8).

$$\frac{\partial}{\partial t}C(\vec{\mathbf{x}},t) = \underbrace{\nabla\cdot(D(\vec{\mathbf{x}})\nabla C(\vec{\mathbf{x}},t))}_{\text{diffusion}} + \underbrace{\nabla\cdot\vec{\mathbf{v}}(\vec{\mathbf{x}})C(\vec{\mathbf{x}},t)}_{\text{convection}} \tag{2.10}$$

In solute transport systems, in a homogeneous porous medium that has some uniform and steady flow through it, the component advective velocity dominates, and the diffusion tensor and velocity vector can be defined respectively:

$$D(\vec{\mathbf{x}}) = \begin{bmatrix} D_{\vec{x}} & 0 & 0 \\ 0 & D_{\vec{y}} & 0 \\ 0 & 0 & D_{\vec{z}} \end{bmatrix}, \qquad \vec{\mathbf{v}}(\vec{\mathbf{x}}) = \begin{bmatrix} v_{\vec{x}} \\ 0 \\ 0 \end{bmatrix} \tag{2.11}$$

Here, each component diagonal elements of D refer to the longitudinal and traverse dispersion coefficients, and $v_{\vec{x}}$ denotes the advection velocity. The general principle of the advection-dispersion equation (ADE), so-called when the system is under predominantly advective flow and mechanical dispersion, is that the centre of the mass of a solute is moving at the rate of the flow field, while dispersing mechanically and chemically, approximately normally. The ratio between advective velocity and dispersivity is termed the Péclet number, $\text{Pe} = v_{\vec{x}}/|D|$. The ratio property is are useful in the choice of modelling the ADE. In the scenarios covered in this topic, the solutions can be considered to have large values of Pe, denoting advection-dominant systems, $v_{\vec{x}} \gg |D|$. Note that Pe is dimensionless.

The analytical solution of an instantaneous point injection at the origin, with the steady velocity and dispersion terms in (2.11) gives a a normally distributed plume, now centered on the centre of mass as it has been advected over time. With initial conditions, $C(\vec{\mathbf{x}},0) = C_0\delta(\vec{\mathbf{x}})$ in a d-dimensional system, the concentration evolution is defined:

$$C(\vec{\mathbf{x}},t) = \frac{C_0}{(4\pi|D|t)^{\frac{d}{2}}}\exp\left(-\left(\vec{\mathbf{x}}-\vec{\mathbf{v}}t\right)^{\mathrm{T}} D^{-1}\left(\vec{\mathbf{x}}-\vec{\mathbf{v}}t\right)\right) \tag{2.12}$$

The aforementioned solution holds only for the conditions where the solute is injected instantaneously and travels through the porous medium as a consolidated slug. Where there is continuous injection, or the velocity field and dispersivity are not uniform or constant, the analytical solution is much harder to resolve, and in

some cases is intractable. In cases such as these, the analytical solution can be approximated using the complementary function, $\text{erfc}(x)$, where

$$\begin{aligned}\text{erfc}(x) &= 1 - \text{erf}(x) \\ &= \frac{2}{\sqrt{\pi}} \int_x^\infty \exp\left(-t^2\right) \mathrm{d}t\end{aligned} \tag{2.13}$$

2.4.4 Random Walk Particle Tracking

2.4.4.1 Space-based Random Walk

Numerical solutions to approximate the advection dispersion equation can be performed using a variety of computational tools, including finite element and finite different techniques. Another approach to modelling hydrodynamic processes, including advection-dispersion, is to represent the solute as a number of particles and simulate their movement through the medium, under the forces of the velocity field and mechanical dispersion.

In space-based random walk particle tracking, the particles are each represented by a coordinate in the modelling space, $\vec{\mathbf{x}}_i$, forming a set $\vec{P} = \{\vec{\mathbf{x}}_i \,|\, i = 1, \ldots, n_p\}$. The trajectory of the particles are modelled to simulate the underlying system. As seen in in the analytic solution of the diffusion equation, in (2.9), the distribution of the particles can be generally be assumed to be normal. As such, some stochastic process is applied to each particle, with the random component sampled from a normal distribution.

Given a particle, in an advective-dispersive system, its trajectory from one time interval to the next can be described by the following random process:

$$\vec{\mathbf{x}}_i(t + \delta_t) = \vec{\mathbf{x}}_i(t) + \underbrace{\left(\vec{\mathbf{v}}(\vec{\mathbf{x}}_i(t)) + \nabla \cdot D(\vec{\mathbf{x}})|_{\vec{\mathbf{x}}_i(t)}\right)}_{\text{flux component}} \delta_t + \underbrace{\sqrt{2dD(\vec{\mathbf{x}})}}_{\text{dispersivity}} \delta_t^{\frac{1}{2}} \xi \tag{2.14}$$

The additive components of the trajectory propagation form the flux and dispersivity parts of the underlying model, each scaled by the time interval, δ_t. The random process ξ is independently sampled from some d-dimensional normal distribution, $\xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, and applies only to the dispersion of the particle. Note that the where the dispersivity is constant, the divergence of the dispersion tensor, D is a vector of zeros, with length d.

The moments of the tracer can be approximated by the particles using the discrete integrals. The first and (central) second order moments corresponding to the centroid and spread of the tracer plume at time t, can be approximated respectively, based on the zeroth moment, i.e. mass, being equal to C_0.

$$\begin{aligned}
\bar{\mathbf{x}}(t) &= \int \vec{\mathbf{x}}(t) C(\vec{\mathbf{x}}, t) \mathrm{d}\vec{\mathbf{x}}(t) \\
&\approx \frac{1}{n_p} \sum_{i=1}^{n_p} \vec{\mathbf{x}}_i(t) \\
\bar{P}(t) &= \int \left(\vec{\mathbf{x}}(t) - \bar{\mathbf{x}}(t)\right) \left(\vec{\mathbf{x}}(t) - \bar{\mathbf{x}}(t)\right)^{\mathrm{T}} C(\vec{\mathbf{x}}, t) \mathrm{d}\vec{\mathbf{x}}(t) \\
&\approx \frac{1}{n_p} \sum_{i=1}^{n_p} \left(\vec{\mathbf{x}}_i(t) - \bar{\mathbf{x}}(t)\right) \left(\vec{\mathbf{x}}_i(t) - \bar{\mathbf{x}}(t)\right)^{\mathrm{T}}
\end{aligned} \tag{2.15}$$

The concentration can also be approximated over a spatial domain by calculating the density over the whole domain, using appropriate estimation properties. For example, a histogram can be used, binning particle locations in spatial cells at a time-step and normalising the bin sizes and scaling the histogram such that its area is equal to the zeroth moment, C_0. The histogram will give a solution to concentration over a set of discrete intervals. Alternatively, kernel density estimation can be used for finer resolution meshes and a smooth concentration function.

2.4.4.2 Time-domain Random Walk

Space-based random walk methods allow for particles to move freely, and spatial discretisation can be performed once all particles are propagated. An alternative approach is to use a time-domain approach which can, instead of solving for the PDE in (2.10), model the concentration-time curve at a set of observation points, mimicking the corresponding analytical solution. For simulating concentration in a solute transport problem, Bodin (2015) separates analytical solutions into two parameters: the density function of particle travel times; and a recovery function that accounts for dispersion occurring at a given observation point.

The analytical solution of an equation can be split into these two parts, for initial concentration, C_0 and injection rate Q.

$$C(\vec{\mathbf{x}}, t) = Q C_0 \Gamma(\vec{\mathbf{x}}) \Psi(\vec{\mathbf{x}}, t) \tag{2.16}$$

The recovery ratio function, $\Gamma(\vec{\mathbf{x}})$, is an infinite-time asymptotic value of the concentration at some point, $\vec{\mathbf{x}}$, where the boundary conditions are fixed. $\Psi(\vec{\mathbf{x}}, t)$ is a cumulative density function at $\vec{\mathbf{x}}$ as a function of t, as opposed to the probability density function approach used in SBRW. In the time-domain random walk approach (TDRW), this is formed of a threshold function that counts the number of particles with travel time less than t or, for the instantaneous injection boundary condition, the number of particles between the interval defined by $t \pm 0.5\delta_t$.

For the continuous injection of solute in an advection-only flow field with dispersivity and velocity as defined in (2.11) , the recovery ratio is defined:

$$\begin{aligned} \Gamma(\vec{\mathbf{x}}) &= \frac{1}{4\pi\vec{\gamma}(\vec{\mathbf{x}})\sqrt{D_{\vec{y}}D_{\vec{z}}}} \exp\left(\frac{v_{\vec{x}}\left(\vec{x} - \vec{\gamma}(\vec{\mathbf{x}})\right)}{2D_{\vec{x}}}\right) \\ \text{where } \vec{\gamma}(\vec{\mathbf{x}}) &= \sqrt{\vec{x}^2 + \frac{D_{\vec{x}}}{D_{\vec{y}}}\vec{y}^2 + \frac{D_{\vec{x}}}{D_{\vec{z}}}\vec{z}^2} \end{aligned} \tag{2.17}$$

Where the injection is instantaneous, $\Gamma(\vec{\mathbf{x}}) = 1/Q$, to simply cancel out the injection rate in (2.16). Thus the function described by Ψ, which in TDRW is a t-thresholded cumulative density function, must be resolved for some n_p particles. In the time-domain approach, a particle is defined by its time to a particular observation point, $t_i(\vec{\mathbf{x}})$, and the set is defined $\vec{P} = \{t_i(\vec{\mathbf{x}}) \,|\, i = 1, \ldots, n_p\}$. A particle is resolved over a measurement space by the following term, for advection-dispersion.

$$\begin{aligned} t_i(\vec{\mathbf{x}}) &= \frac{\vec{\gamma}(\vec{\mathbf{x}})}{v_{\vec{x}}} + \frac{\zeta_i^2 + \zeta_i\sqrt{\zeta_i^2 + 4v_{\vec{x}}\vec{\gamma}(\vec{\mathbf{x}})}}{2v_{\vec{x}}^2} \\ \text{where } \zeta_i &= 2\sqrt{D_{\vec{x}}}\,\text{erfc}^{-1}\left(2\xi_u\right) \end{aligned} \tag{2.18}$$

Here, the stochastic aspect of the particle comes from its initialisation, ζ_i, where ξ_u is some random number, independently sampled from a *uniform* distribution, within the interval $[0, 1]$.

3

Optimal Bayesian Filtering

3.1 State Space Models

A state space model (SSM) describes a probabilistic model over which a sequence of hidden states and observations are linked through some continuous or discrete dynamic processes (Särkkä, 2013). Given a continuous variable, t, the hidden state $x(t)$ and its observations $y(t)$ can be described by the following respective process and observation models:

$$\begin{aligned} x'(t) &= f(t, x(t), u(t); \xi(t)) \\ y(t) &= h(t, x(t), u(t); \nu(t)) \end{aligned} \tag{3.1}$$

The process and observation models may be dynamic where they are dependent on t, or static over all t, e.g. $f(t, x(t), u(t)) := f(x(t), u(t))$. Control parameters, $u(t)$ are independent of the state $x(t)$; $\xi(t)$ and $\nu(t)$ represent noise in the state system process and observation models respectively.

In some instances, it is desirable to discretise independent parameter t, for example in the event that observations are obtainable at discrete intervals. A discrete process model of the SSM described in (3.1) can be defined as follows:

$$x(t + \delta(t)) = f(t, x(t), u(t); \xi(t), \delta(t)) \tag{3.2}$$

The value $\delta(t)$ is the interval of t. The interval is commonly fixed, such that it is constant and independent of t: $\delta(t) = \delta$. In terms of notation, for discrete-interval

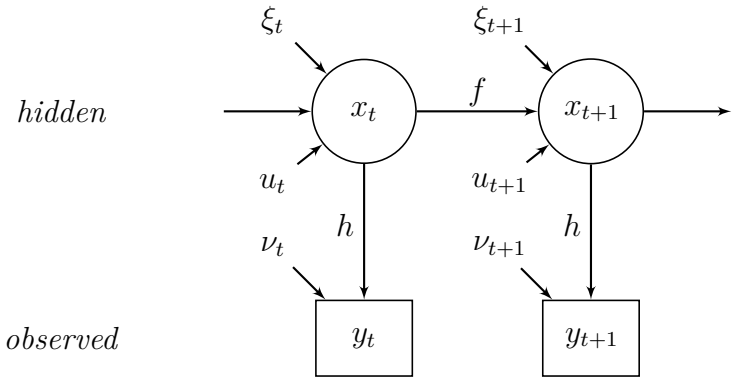


Figure 3.1: Graphical model of Hidden Markov model representing a discrete-interval state space model. The subscript t represents discrete time, where $x_t := x(t)$ and $x_{t+1} := x(t + \delta)$

state uses subscript t with increments of 1 representing δ, i.e. $x_t = x(t)$ and $x_{t+1} = x(t + \delta)$, where $t \in \mathbb{N}$.

A discrete-interval SSM can be considered *Markovian*, such that the state x is a hidden Markov process with y representing the observable aspects of the model. The discrete state space model visualised as such in Figure 3.1, shows a general case snapshot of the state process and observation models for t and $t + 1$. Due to the nature of the Markov process, the snapshot represents any two states in the SSM when $t > 0$.

3.2 Bayesian Inference

As opposed to frequentist probability, which infers the probability of some random event based on the relative frequency of success from a sample, Bayesian probability is based on the concept that probability can be inferred based on the likelihood of the proposed hypothesis (Särkkä, 2013). Where a sample may suggest that a given hypothesis is correct, a Bayesian approach uses *a priori* probability, i.e. the likelihood of an event without any supporting or contradicting evidence. The *a posteriori* probability represents the likelihood of the event given the sample evidence. The relation between *a posteriori* and *a priori* is defined by Bayes' rule.

$$\mathrm{P}(A \mid B) = \frac{\mathrm{P}(B \mid A)\,\mathrm{P}(A)}{\mathrm{P}(B)} \tag{3.3}$$

The component probabilities of Bayes' rule are described as follows:

$\mathrm{P}(A \mid B)$	*a posteriori* probability of hypothesis A given sample evidence B
$\mathrm{P}(A)$	*a priori* probability of the A given no evidence
$\mathrm{P}(B)$	probability of evidence B occurring, regardless of the hypothesis
$\mathrm{P}(B \mid A)$	probability, or likelihood, of B were the hypothesis true

While the above interpretation of Bayes' rule describes the probability of a hypothesis occurring, it is also possible to apply the same rule for a probability distribution. In the case of an SSM, the process model can be described through a predictive distribution $p(x'(t))$. A prior and posterior predictive distribution are defined by calculating marginal likelihood over $x(t)$, giving probabilistic reference of the state to process model.

$$\begin{aligned} p\left(x'(t)\right) &= \int p\left(x'(t) \mid x(t)\right) p\left(x(t)\right) \mathrm{d}x(t) \\ p\left(x'(t) \mid y(t)\right) &= \int p\left(x'(t) \mid x(t)\right) p\left(x(t) \mid y(t)\right) \mathrm{d}x(t) \end{aligned} \tag{3.4}$$

In the discrete-interval case of the SSM, as described by (3.2), the Markovian nature of the process model allows certain assumptions to be made. The primary assumption is that the conditional probability distributions of $x(t)$ given all previous discrete states, and of observations, $y(t)$ given the current and previous states, are equivalent to the distribution given only the previous and current state respectively:

$$\begin{aligned} p\left(x(t) \mid x(t-\delta), x(t-2\delta), \ldots, x(0)\right) &\equiv p\left(x(t) \mid x(t-\delta)\right) \\ p\left(y(t) \mid x(t), x(t-\delta), \ldots, x(0)\right) &\equiv p\left(y(t) \mid x(t)\right) \end{aligned} \tag{3.5}$$

These equations hold intuitively, given that prior states are dependent on their respective priors up to the initial state $x(0)$.

If the hidden state is considered to have an unknown probability distribution, Bayesian inference can be used to recursively estimate the state at a given time. This estimator uses the process model to formulate an *a priori* state distribution, along with measurements of the state observations at a given value of t as a sample. By Bayes' rule (3.3), the *a posteriori* estimate of the state probability distribution is defined below:

$$
\begin{aligned}
p\left(x_t \mid y_t\right) &= \frac{p\left(y_t \mid x_t\right) p\left(x_t \mid y_{t-1}\right)}{p\left(y_t \mid y_{t-1}\right)} \\
&= \frac{\overbrace{p\left(y_t \mid x_t\right)}^{\text{observation}} \int \overbrace{p\left(x_t \mid x_{t-1}\right)}^{\textit{a priori}} p\left(x_{t-1}\right) \mathrm{d}x_{t-1}}{\underbrace{\int p\left(y_t \mid x_t\right) p\left(x_t \mid y_{t-1}\right) \mathrm{d}x_t}_{\text{normalisation term}}}
\end{aligned} \tag{3.6}
$$

The observation model can be described as the probability distribution of y_t given state x_t, and the remainder of the components in (3.6) are be marginalised by the previous state x_{t-1}, as in (3.4). The numerator of the resulting function describes the observation weighted by an *a priori* estimation of x_t. The denominator becomes a normalisation term: the integral of the numerator.

The formulation of the Bayes estimator can be used, for a given SSM, to approximate the state probability at a given time using progressive measurements. Numerically, there are many approaches to state estimation using recursive Bayesian inference. Where the process and observation models are linear, and the probability distribution of state, observation and respective noise are normal, the Kalman filter exactly describes an optimal recursive Bayesian estimator. Where these assumptions do not hold, there are extensions to the Kalman filter, as well as stochastic estimators. Such filters include Markov Chain Monte Carlo (MCMC) filters (Andrieu et al., 2003), and Sequential Monte Carlo filters, also known as particle filters (Liu and Chen, 1998).

3.3 Kalman Filter

3.3.1 Assumptions

For a state system model, the Kalman filter seeks to approximate the unknown variables of the expected state (Kalman, 1960; Särkkä, 2013). To provide an optimal estimator, the filter relies on certain assumptions. Where these assumptions hold, the Kalman filter has been proved to be optimal. In cases where the assumptions are approximately true, the Kalman filter can still provide a good approximation of the state.

If $x(t)$ is an unknown probability distribution of the state given t, and $y(t)$ is the unknown probability distribution of the corresponding observation, then the

Kalman filter seeks to estimate the expected value of the state and observation.

$$\mathbb{E}[x(t)] := \tilde{\boldsymbol{x}}_t \qquad \mathbb{E}[y(t)] := \tilde{\boldsymbol{y}}_t$$

Where the distribution is normal, the *a priori* state and observation distributions can be approximated using the SSM in (3.1):

$$\begin{aligned} p\left(x_t \mid x_{t-1}\right) &= \mathcal{N}\left(f(t, \boldsymbol{x}_{t-1}, u_t), \boldsymbol{Q}_t\right) \\ p\left(y_t \mid x_t\right) &= \mathcal{N}\left(h(t, \boldsymbol{x}_t, u_t), \boldsymbol{R}_t\right) \end{aligned} \tag{3.7}$$

The input parameter $\boldsymbol{x}$ denotes the Kalman approximation of the state expectation, $\tilde{\boldsymbol{x}}$; the covariance $\boldsymbol{Q}$ and $\boldsymbol{R}$ are respectively the known, or approximated, covariance matrices of the unknown model noise ξ and ν, both assumed to be zero-mean normally distributed.

$$\underbrace{\xi_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q}_t)}_{\text{process noise}} \qquad \underbrace{\nu_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_t)}_{\text{measurement noise}}$$

As an optimal filter, the Kalman filter seeks to minimise the residual between estimation and actual value. As a result, the expected residual is zero. The actual values of the residuals correspond to the estimation errors in the filter.

$$\mathbb{E}[\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t] = \boldsymbol{0} \qquad \mathbb{E}[\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_t] = \boldsymbol{0}$$

Here, the measured observation given t is denoted $\tilde{\boldsymbol{y}}_t$ and model estimation is $\boldsymbol{y}_t = h(t, \boldsymbol{x}_t, u_t)$. For the purposes of weighting the effects of model parameter estimations in filtering, it is necessary to maintain estimations of the covariance between the state and observation errors, as well as of their joint distribution.

$$\begin{aligned} \boldsymbol{P}_t^{(x)} &:= \operatorname{cov}[\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t, \tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t] \\ \boldsymbol{P}_t^{(y)} &:= \operatorname{cov}[\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_t, \tilde{\boldsymbol{y}}_t - \boldsymbol{y}_t] \\ \boldsymbol{P}_t^{(xy)} &:= \operatorname{cov}[\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t, \tilde{\boldsymbol{y}}_t - \boldsymbol{y}_t] \end{aligned} \tag{3.8}$$

3.3.2 Linear Discrete Prediction

The first step in a recursive Bayesian estimator is to calculate the *a priori* estimation of the state, i.e. the prediction given the previous state. Given the state space models, it is possible to propagate the state via the process model to estimate the

next state. The corresponding measurement of this predicted state can also be approximated, as a result of the observation model on the *a priori* state (Kalman, 1960).

Where the systems are linear-discrete, the process and observation models can be considered systems of linear equations and represented in matrix form:

$$\begin{aligned} f(t, \boldsymbol{x}_t, u_t) &:= \mathbf{F}_t \boldsymbol{x}_t + \mathbf{U}_t^f u_t \\ h(t, \boldsymbol{x}_t, u_t) &:= \mathbf{H}_t \boldsymbol{x}_t + \mathbf{U}_t^h u_t \end{aligned} \tag{3.9}$$

The propagation of the state by these equations forms the predictive step of the Kalman filter. The *a priori* estimation is thus denoted with subscript "$t|t-1$". The covariance of the state estimation error can also be explicitly propagated, by the linear system $\mathbf{F}$, incorporating the process noise covariance.

$$\begin{aligned} \boldsymbol{x}_{t|t-1} &= f(t, \boldsymbol{x}_{t-1}, u_t) \\ \boldsymbol{y}_{t|t-1} &= h(t, \boldsymbol{x}_{t|t-1}, u_t) \\ \boldsymbol{P}_{t|t-1}^{(x)} &= \mathbf{F}_t \boldsymbol{P}_{t-1}^{(x)} \mathbf{F}_t^{\mathrm{T}} + \boldsymbol{Q}_t \end{aligned} \tag{3.10}$$

3.3.3 Gain

To correct the state approximation, during the calculation of an *a posteriori* estimation, the weighting of measured observations must be calculated. This weight is referred to as the *gain*, and in this case the *Kalman gain*, and is designed to optimally weight the uncertainty and error in the state estimation and measurement criteria.

The covariance of joint process and estimation error is approximated using linear observation model $\boldsymbol{H}$ and the *a priori* process error covariance. This can form an intermediate step in the approximation of the observation error covariance, which incorporates the measurement uncertainty.

$$\begin{aligned} \boldsymbol{P}_t^{(xy)} &= \boldsymbol{P}_{t|t-1}^{(x)} \mathbf{H}_t^{\mathrm{T}} \\ \boldsymbol{P}_t^{(y)} &= \mathbf{H}_t \boldsymbol{P}_t^{(xy)} + \boldsymbol{R}_t \end{aligned} \tag{3.11}$$

The matrix division of the joint covariance by the observation error covariance results in the gain. This has been shown, under the appropriate conditions, to minimise the square difference between the state error residual when factored into the calculation of *a posteriori* estimate. The calculation of the gain, $\mathbf{K}_t$, is therefore:

$$\mathbf{K}_t := \boldsymbol{P}_t^{(xy)} \left(\boldsymbol{P}_t^{(y)}\right)^{-1} \tag{3.12}$$

The gain weights the observation residual, such that where the uncertainty in the process model is high, but the measurements are relatively reliable, the gain will by large and apply significant correction to the *a priori* state based on the residual between estimated and measured observations. In contrast, where the process model is reliable, and the measurements noisy, the effect of the gain will be minimised.

3.3.4 Correction

As discussed, the calculation of the *a posteriori* state approximation is the weighted combination of the *a priori* state and the residual between estimated and measured observations. The correction terms are described:

$$\begin{aligned} \boldsymbol{x}_{t|t} &= \boldsymbol{x}_{t|t-1} + \mathbf{K}_t \left(\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_{t|t-1}\right) \\ \boldsymbol{P}_{t|t}^{(x)} &= \boldsymbol{P}_{t|t-1}^{(x)} - \mathbf{K}_t \boldsymbol{P}_t^{(y)} \mathbf{K}_t^{\mathrm{T}} \end{aligned} \tag{3.13}$$

Note that where the estimated observation and the measurement are exactly, or near, equal, the *a posteriori* and *a priori* state estimates are close. However, this does not necessarily imply that the estimate is correct, due to the known presence of both process and measurement noise. Such uncertainty is thus captured in the *a posteriori* estimation error covariance, $\boldsymbol{P}_{t|t}^{(x)}$.

By the recursive nature of the filter *a posteriori* state forms the basic for the *a priori* estimate of the next state in a discrete-interval SSM. Where a measurement is not available of perform correction, it is possible to use the *a priori* estimate of a previous state in (3.10). Figure 3.2 demonstrates the recursive process of the Kalman filter.

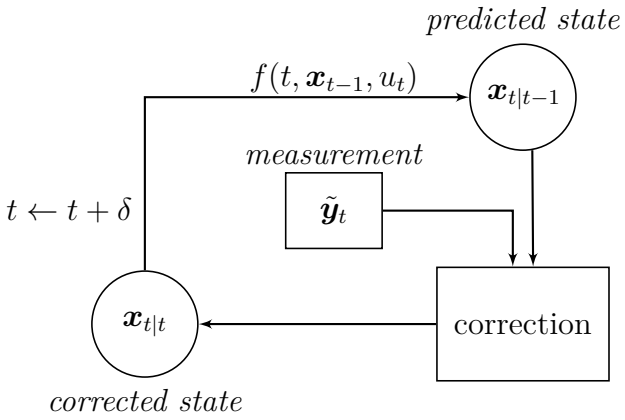


Figure 3.2: Discrete Kalman filter process, showing prediction and correction cycle over increments of t, with regular measurements

3.4 Extensions of Kalman Filter

3.4.1 Non-Linear Process and Observation Models

The basic Kalman filter, as described above, is defined to optimally estimate the state in a strictly linear SSM. The linear nature of the underlying process and observation models allow for explicit calculation of the estimate error covariance matrices. Where either of the process or observation models are non-linear, the calculation of estimation errors, and therefore the gain, is non-trivial.

In the general case, including where the underlying distributions are discrete or multimodal, a Monte Carlo approach could be used. However, where the probability distributions and noise terms can be reasonably approximated as Gaussian, there are extensions to the Kalman filter that can estimate the state in a non-linear SSM (Julier and Uhlmann, 2004). One approach is to linearise parts of the underlying models using Taylor expansions, and approximate the estimate errors using these linearised models. Linearised filters such as these are not optimal, and in highly non-linear models can lead to catastrophic divergence of the estimate from the true state.

One alternative to linearisation is the use of the unscented transform, which captures the first two moments of a distribution by a set of deterministic points and propagates these points individually through system equations to approximate the resulting probability distribution (Wan et al., 2000).

Algorithm 1 (Linear-discrete) Kalman filter

procedure PREDICT($\boldsymbol{x}_{t-1}$, u_t, $\boldsymbol{P}_{t-1}^{(x)}$)
  $\boldsymbol{x}_{t|t-1} \leftarrow \mathbf{F}_t \boldsymbol{x}_{t-1} + \mathbf{U}_t^f u_t$ ▷ If no input, the addition of u_t is omitted
  $\boldsymbol{P}_{t|t-1}^{(x)} \leftarrow \mathbf{F}_t \boldsymbol{P}_{t|t-1}^{(x)} \mathbf{F}_t^{\mathrm{T}} + \boldsymbol{Q}_t$
  return $\boldsymbol{x}_{t|t-1}$, $\boldsymbol{P}_{t|t-1}^{(x)}$
end procedure
procedure UPDATE($\boldsymbol{x}_{t|t-1}$, u_t, $\boldsymbol{P}_{t|t-1}^{(x)}$, $\tilde{\boldsymbol{y}}_t$)
  $\boldsymbol{y}_t \leftarrow \mathbf{H}_t \boldsymbol{x}_{t|t-1} + \mathbf{U}_t^f u_t$ ▷ If no input, the addition of u_t is omitted
  $P_t^{(xy)} \leftarrow \boldsymbol{P}_{t|t-1}^{(x)} \mathbf{H}_t^{\mathrm{T}}$
  $P_t^{(y)} \leftarrow \mathbf{H}_t P_t^{(xy)} + \boldsymbol{R}_t$
  $\mathbf{K}_t \leftarrow \boldsymbol{P}_t^{(xy)} \left(\boldsymbol{P}_t^{(y)}\right)^{-1}$
  $\boldsymbol{x}_{t|t} \leftarrow \boldsymbol{x}_{t|t-1} + \mathbf{K}_t \left(\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_t\right)$
  $\boldsymbol{P}_{t|t}^{(x)} \leftarrow \boldsymbol{P}_{t|t-1}^{(x)} - \mathbf{K}_t \boldsymbol{P}_t^{(y)} \mathbf{K}_t^{\mathrm{T}}$
  return $\boldsymbol{x}_{t|t}$, $\boldsymbol{P}_{t|t}^{(x)}$
end procedure
procedure KALMAN FILTER($\boldsymbol{x}_0$, $\boldsymbol{P}_0^{(x)}$; $[\mathbf{F}_t, \mathbf{U}_t^f, \boldsymbol{Q}_t, \mathbf{H}_t, \mathbf{U}_t^h, \boldsymbol{R}_t]_{t=1}^T$)
  $\boldsymbol{x}_{t-1}, \boldsymbol{P}_{t-1}^{(x)} \leftarrow \boldsymbol{x}_0, \boldsymbol{P}_0^{(x)}$
  for $t \in 1, \ldots, T$ **do**
    $\boldsymbol{x}_{t|t-1}, \boldsymbol{P}_{t|t-1}^{(x)} \leftarrow$ PREDICT($\boldsymbol{x}_{t-1}$, u_t, $\boldsymbol{P}_{t-1}^{(x)}$)
    if measurement $\tilde{\boldsymbol{y}}_t$ available **then**
      $\boldsymbol{x}_{t|t}, \boldsymbol{P}_{t|t}^{(x)} \leftarrow$ UPDATE($\boldsymbol{x}_{t|t-1}$, u_t, $\boldsymbol{P}_{t|t-1}^{(x)}$, $\tilde{\boldsymbol{y}}_t$)
      $\boldsymbol{x}_t, \boldsymbol{P}_t^{(x)} \leftarrow \boldsymbol{x}_{t|t}, \boldsymbol{P}_{t|t}^{(x)}$
    else
      $\boldsymbol{x}_t, \boldsymbol{P}_t^{(x)} \leftarrow \boldsymbol{x}_{t|t-1}, \boldsymbol{P}_{t|t-1}^{(x)}$
    end if
  end for
  return $\left[(\boldsymbol{x}_t, \boldsymbol{P}_t^{(x)})\right]_{t=1}^T$
end procedure

3.4.2 Linearisation

An early method proposed for estimation of non-linear state-space model is the extended Kalman filter, which uses linearisation methods to approximate the estimate error covariance models (Einicke and White, 1999). For non-linear, but differentiable, f and h, the Jacobian can be used to linearise the system about the expectations of state and observation models respectively. Based on Taylor series expansions, the equations for the estimation error covariance and cross-covariance matrices in (3.10) and (3.11) are predicted by using $\mathbf{F}$ and $\mathbf{H}$ as the respective Jacobians of f and h evaluated for $\boldsymbol{x}_{t-1}$ and $\boldsymbol{x}_{t|t-1}$ respectively:

$$\mathbf{F}_t = \left.\frac{\partial f}{\partial x}\right|_{\boldsymbol{x}_{t-1},u_t}, \qquad \mathbf{H}_t = \left.\frac{\partial h}{\partial x}\right|_{\boldsymbol{x}_{t|t-1},u_t} \tag{3.14}$$

By extending the higher order terms of the Taylor expansion of the covariance approximation, greater representation by the approximated state process model can be attained (Roth and Gustafsson, 2011). The difficulties in such problems lie both in the numerical stability of the Taylor expansion, and in the difficulties in analytically deriving the Jacobian or Hessian for higher order differentials (Julier and Uhlmann, 2004).

3.4.3 Unscented Transform

The unscented transform (UT) is a numerical method developed for the propagation of some probability distribution over some non-linear function. As opposed to a stochastic approach, which may propagate a large number of sample points through the function and estimate the posterior distribution from the transformed samples, the UT uses a small set of deterministic *sigma points*. The sigma points are defined such that they capture the mean and covariance of the original distribution (Wan et al., 2000; Julier and Uhlmann, 2004).

For a given non-linear system equation, mapping probability distributions x to y, $y = g(x)$, a finite set of sigma points can be defined, with the spread about the mean controlled by some parameter λ:

$$\begin{aligned} \chi_0 &= \mathbb{E}[x] \\ \chi_i &= \begin{cases} \chi_0 + \left[\sqrt{(n_x+\lambda)\mathrm{cov}[x,x]}\right]_i & i = 1,\ldots,n_x \\ \chi_0 - \left[\sqrt{(n_x+\lambda)\mathrm{cov}[x,x]}\right]_{i-n_x} & i = n_x+1,\ldots,2n_x \end{cases} \end{aligned} \tag{3.15}$$

The number of sigma points, $2n_x + 1$, is based on the dimensionality of distribution x: $n_x = |\chi_0|$. The matrix square root, $\sqrt{\mathbf{A}} = \mathbf{B} \Rightarrow \mathbf{A} = \mathbf{B}\mathbf{B}^{\mathrm{T}}$, is calculated by some numerically stable method, typically Cholesky decomposition. Matrix subscript $\mathbf{A}_i$ is defined as the i^{th} matrix column.

The sigma points are then individually propagated by g, and the resulting distribution, y, can be calculated in terms of its mean and covariance by the weighted sum of propagated sigma points.

$$\begin{aligned} \gamma_i &= g(\chi_i) \\ \mathbb{E}[y] &\approx \sum_{i=0}^{2n_x} w_i^{(x)} \gamma_i \\ \operatorname{cov}[y, y] &\approx \sum_{i=0}^{2n_x} w_i^{(p)} \left(\gamma_i - \mathbb{E}[y]\right)\left(\gamma_i - \mathbb{E}[y]\right)^{\mathrm{T}} \end{aligned} \tag{3.16}$$

The definition of the spread and weighting parameters define the degree to which the distribution is captured. The key controlling parameters are α, β and κ. The value of β allows incorporation of information on the potentially non-Gaussian distribution, and values α and κ are quadratic and constant controls of the spread, respectively. The spread of the sigma points about the mean are thus defined by $\lambda = \alpha^2 \left(n_x + \kappa\right) - n_x$. The inverse, in recombining sigma points, gives the weights defined, with $w^{(x)}$ and $w^{(p)}$ weighting the calculation of mean and covariance respectively.

$$\begin{aligned} w_0^{(x)} &= \lambda / (n_x + \lambda) \\ w_0^{(p)} &= w_0^{(x)} + \left(1 - \alpha^2 + \beta\right) \\ w_i^{(x)} = w_i^{(p)} &= \left(1 - w_0^{(x)}\right) / 2n_x \end{aligned} \tag{3.17}$$

Figure 3.3 demonstrates the effectiveness of the unscented transform in approximating a probability distribution over some non-linear function. In the figure, the actual distribution was calculated by Monte Carlo simulation with 10,000 samples, while the UT approximates the transformation with only five sigma points.

3.4.4 Unscented Kalman Filter

To avoid the pitfalls of linearisation in approximating an SSM, a version of the Kalman filter that utilises the UT was developed (Julier and Uhlmann, 2004; Särkkä, 2007). This *unscented Kalman filter* (UKF). The unscented transform is used to sample the state distribution $p(x_t)$, using sigma points. For a given state estimate, $\boldsymbol{x}_{t-1}$ and its error covariance $\boldsymbol{P}_{t-1}^{(x)}$, the corresponding sigma points are defined as in (3.15).

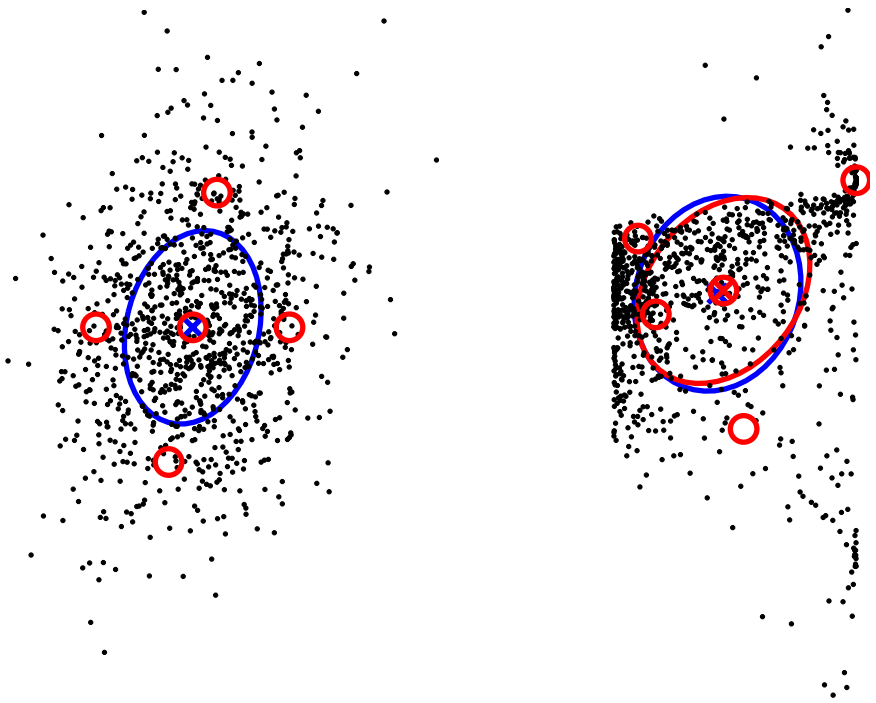


Figure 3.3: A sampled distribution (left) transformed over some non-linear function (right), showing actual (blue) and estimated (red) distribution mean (cross) and covariance (ellipse). Original and transformed sigma points are denoted by red discs

$$\chi_{0;t-1} = \boldsymbol{x}_{t-1}$$
$$\chi_{i;t-1} = \begin{cases} \chi_{0;t-1} + \left[\sqrt{(n_x+\lambda)\boldsymbol{P}_{t-1}^{(x)}}\right]_i & i = 1,\ldots,n_x \\ \chi_{0;t-1} - \left[\sqrt{(n_x+\lambda)\boldsymbol{P}_{t-1}^{(x)}}\right]_{i-n_x} & i = n_x+1,\ldots,2n_x \end{cases} \tag{3.18}$$

3.4.4.1 Prediction

As in the general unscented transform, the number of sigma points is based on the dimensionality of the state, i.e. $n_x = |\boldsymbol{x}_{t-1}|$. Where the process model f is non-linear, the *a priori* estimates of the state and covariance can be approximated by the propagation of sigma points through f. In the case of additive noise, the noise covariance can be incorporated into $\boldsymbol{P}^{(x)}$ directly through summation.

$$\begin{aligned}
\chi_{i;t|t-1} &= f(t, \chi_{i;t-1}, u_t) \\
\boldsymbol{x}_{t|t-1} &= \sum_{i=0}^{2n_x} w_i^{(x)} \chi_{i;t|t-1} \\
\boldsymbol{P}_{t|t-1}^{(x)} &= \sum_{i=0}^{2n_x} w_i^{(p)} \Big(\boldsymbol{x}_{t|t-1} - \chi_{i;t|t-1}\Big)\Big(\boldsymbol{x}_{t|t-1} - \chi_{i;t|t-1}\Big)^{\mathrm{T}} + \boldsymbol{Q}_t
\end{aligned} \tag{3.19}$$

While the *a priori* state can be used to approximate $\boldsymbol{y}_{t|t-1}$ by propagating it through the observation model, where h is non-linear the unscented transform should be applied again to calculate a set of sigma points in the observation domain. These sigma points can be used to get an estimate of the expected observation.

$$\begin{aligned}
\gamma_{i;t|t-1} &= h(t, \chi_{i;t|t-1}, u_t) \\
\boldsymbol{y}_{t|t-1} &= \sum_{i=0}^{2n_x} w_i^{(x)} \gamma_{i;t|t-1}
\end{aligned} \tag{3.20}$$

3.4.4.2 Correction

In terms of correction, the *a posteriori* estimate of the state in UKF calculated in the same way as (3.13) in the linear-discrete Kalman filter. The correction gain $\boldsymbol{K}_t = \boldsymbol{P}_t^{(xy)} \left(\boldsymbol{P}_t^{(y)}\right)^{-1}$, as in (3.12). However, like the covariance of the state estimation error, the required covariance matrices cannot be explicitly determined when the observation model is non-linear. As a result, the UKF approximates the observation and joint distribution covariances using the UT weighted sums of state and observation sigma points.

$$\begin{aligned}
P_t^{(xy)} &= \sum_{i=0}^{2n_x} w_i^{(p)} \Big(\boldsymbol{x}_{t|t-1} - \chi_{i;t|t-1}\Big)\Big(\boldsymbol{y}_{t|t-1} - \gamma_{i;t|t-1}\Big)^{\mathrm{T}} \\
P_t^{(y)} &= \sum_{i=0}^{2n_x} w_i^{(p)} \Big(\boldsymbol{y}_{t|t-1} - \gamma_{i;t|t-1}\Big)\Big(\boldsymbol{y}_{t|t-1} - \gamma_{i;t|t-1}\Big)^{\mathrm{T}} + \boldsymbol{R}_t
\end{aligned} \tag{3.21}$$

The general advantage of the UKF over linearisation, such as the EKF, is that it does not require formally differentiable functions or the calculation of closed form expected values. The UKF can capture higher order moments than at least the first-order EKF, and is simpler in terms of implementation because there is no need to calculate the Jacobian of the underlying models.

In the instance of non-additive noise, the error covariance can be *augmented* into the sigma points, rather than incrementing the state and observation error estimates with the respective process and observation model noises (Ahn and Chan, 2014). The augmented state, $\breve{x}_t$, combines the expected value and covariance of the noise terms with that of the state in the respective sigma point estimations. The corresponding augmented estimate and error covariance are thus defined:

$$\breve{\boldsymbol{x}}_t = \begin{bmatrix} \boldsymbol{x}_t \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \qquad \breve{\boldsymbol{P}}_t^{(x)} = \begin{bmatrix} \boldsymbol{P}_t^{(x)} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{Q}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{R}_t \end{bmatrix} \tag{3.22}$$

3.5 Smoothing

3.5.1 Backward Recursive Smoothing

While the concept of a filter describes the estimation of a state, given observations up to and including the current state, an optimal *smoother* seeks to calculate an *a posteriori* distribution given observations of states consequent to the current state (Särkkä, 2013).

The *a posteriori* distribution is defined, given observations for state interval up to and including $T > t$ is defined as follows:

$$p\left(x_t \,|\, y_0, \ldots, y_T\right) = p\left(x_t \,|\, y_t\right) \int \left(\frac{p\left(x_{t+1} \,|\, x_t\right) p\left(x_{t+1} \,|\, y_0, \ldots, y_T\right)}{p\left(x_{t+1} \,|\, y_t\right)} \right) \mathrm{d}x_{t+1} \tag{3.23}$$

This distribution describes the relation between states, marginalised over the next state x_{t+1}. By the Markovian nature of the state process, the conditional distribution of x_t given x_{t+1} is independent of later measurements, $\{y_{t+1}, \ldots, y_T\}$. Thus, $p(x_t \,|\, x_{t+1}, y_0, \ldots, y_T) \equiv p(x_t \,|\, x_{t+1}, y_0, \ldots, y_t)$. By these properties, a backwards recursive approach can be used to smooth x_t, with the *a priori* smoothed state approximated by the proceeding state. The joint posterior distribution of the smoother is defined:

$$p\left(x_t, x_{t+1} \,|\, y_0, \ldots, y_T\right) = \frac{p\left(x_{t+1} \,|\, x_t\right) p\left(x_t \,|\, y_t\right) p\left(x_{t+1} \,|\, y_0, \ldots, y_T\right)}{p\left(x_{t+1} \,|\, y_t\right)} \tag{3.24}$$

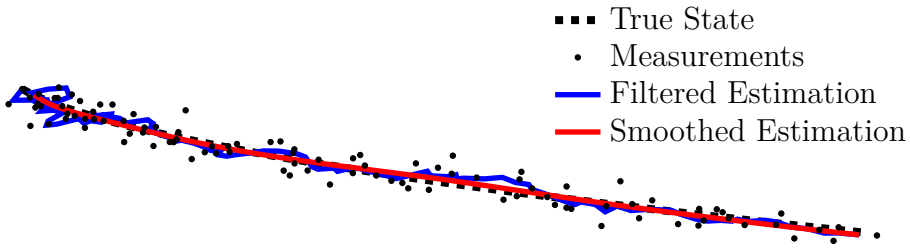


Figure 3.4: Example of discrete-interval Kalman filter estimation and subsequent Rauch-Tung-Striebel smoothing on basic linear acceleration model with 100 noisy measurements. The smoother prior includes all observations

The key differences between prediction, filtering and smoothing can be defined by the prior information available for the estimation of a given state. Where a predictor approximates a state given observations up to but not including the current point in the process, $x_{t|t-1}$, the filter is the combination of prediction and current observation, $x_{t|t}$. The smoothed estimation, $x_{t|T>t}$ differs as it can be based on observations from a fixed period ahead of the state, e.g. fixed-lag smoothing, or from an entire, complete time-series.

3.5.2 Rauch-Tung-Striebel Smoother

There are many implementations of smoothers using the principles of optimal smoothing. The Rauch-Tung-Striebel (RTS) smoother, also known as the discrete-time Kalman smoother, is one such example (Rauch et al., 1965). The RTS smoother is backward-recursive and uses the discrete-interval process model. For a given SSM, the prediction of the state is estimated and compared with the *a posterior* (smoothed) state. The estimate error covariance is also forward propagated, and a smoothing gain, similar to the Kalman gain, is calculated.

$$\begin{aligned} \boldsymbol{x}_{t|t-1} &= f\left(t, \boldsymbol{x}_{t-1}, u_t\right) \\ \boldsymbol{P}^{(x)}_{t|t-1} &= \mathbf{F}_t \boldsymbol{P}^{(x)}_{t-1} \mathbf{F}_t^{\mathrm{T}} + \boldsymbol{Q}_t \\ \mathbf{S}_t &= \boldsymbol{P}^{(x)}_t \mathbf{F}_t \left(\boldsymbol{P}^{(x)}_{t|t-1}\right)^{-1} \end{aligned} \tag{3.25}$$

The smoothing gain, $\mathbf{S}_t$ is used to weight the difference predicted and smoothed of the next forward state, as part of a correction step in calculating the smoothed a

posteriori state, $x_{t-1|t}$. There is no requirement for observations to be provided in the smoother, as the relevant information is contained within the filtering estimates of the states: $x_{T|T}$ is the final filtered state that is used as the initial state in the smoother.

$$\begin{aligned} \boldsymbol{x}_{t-1|t} &= \boldsymbol{x}_{t-1} + \mathbf{S}_t \left(\boldsymbol{x}_{t|t+1} - \boldsymbol{x}_{t|t-1} \right) \\ \boldsymbol{P}^{(x)}_{t-1|t} &= \boldsymbol{P}^{(x)}_{t-1} + \mathbf{S}_t \left(\boldsymbol{P}^{(x)}_{t|t+1} - \boldsymbol{P}^{(x)}_{t|t-1} \right) \mathbf{S}_t^{\mathrm{T}} \end{aligned} \tag{3.26}$$

Figure 3.4 demonstrates the result of Kalman filtering and the subsequent RTS smoothing of a discrete-interval SSM. The SSM is formed of a basic linear-acceleration velocity process with observations formed of the state position co-ordinates at a given time. The noisy measurements are taken at regular intervals and the RTS smoothing is performed using the final *a posteriori* state estimate of the Kalman filter. The figure shows that while the Kalman filter is powerful at approximating the state given process noise, the smoother provides an estimation much closer to the true state. Extensions to the RTS smoother have been made where the process model of the SSM is non-linear. Adaptations have been made based on linearisation, such as the extended RTS smoother; and using the UT to create an unscented RTS smoother (Särkkä, 2008).

3.6 Multi-process systems

3.6.1 Markov Switching SSM

Certain practical applications are not able to rely on a singular process model to describe a system, either due to the underlying state processes consisting of a set independent driving processes that may occur simultaneously or separately with varying likelihood; or the relative complexity of a joint model means it is difficult to construct or solve (Hamilton, 1996; Särkkä, 2013). In such a multi-process system, the system can be represented as a set of state space models, each representing an individual component model, M^j:

$$\begin{aligned} x^j_t &= f^j(t, x_{t-1}, u^j_t; \xi^j_t) \\ y^j_t &= h^j(t, x^j_t, u^j_t; \nu^j_t) \end{aligned} \tag{3.27}$$

A component process, f^j may be dependent on a single prior state x_{t-1} or a component representation of that state, x^j_{t-1}. For a component model, M_j, in a set

of r models representing the state at a given point, there exists a probability that the state process will *switch* to another component model, $\mathrm{P}(M_i \,|\, M_j)\,, i = 1, \ldots, r$. These switching probabilities represent a measure of uncertainty about the model, and filtering approaches, such as *interacting multiple models* (IMM) use these probabilities to weight their *a priori* distributions when estimating the state. Figure 3.5 represents the switching Markov model concept used in IMM, with each individual component model propagated individually, weighted by the respecting switching properties (Mazor et al., 1998).

3.6.2 Interacting Multiple Models

The interacting multiple models acts as an extension of the Kalman filter applied individually to component models, using a model probability, $\mu_t^j = \mathrm{P}(M^j \,|\, t)$, to weight the contribution of each component model's state estimation in the estimation of a unified *a posteriori* state estimate (Mazor et al., 1998). Alongside the unified state estimate, $\boldsymbol{x}_t$ and its error covariance, $\boldsymbol{P}_t^{(x)}$, state space representations of component models are maintained and propagated, with estimated state and

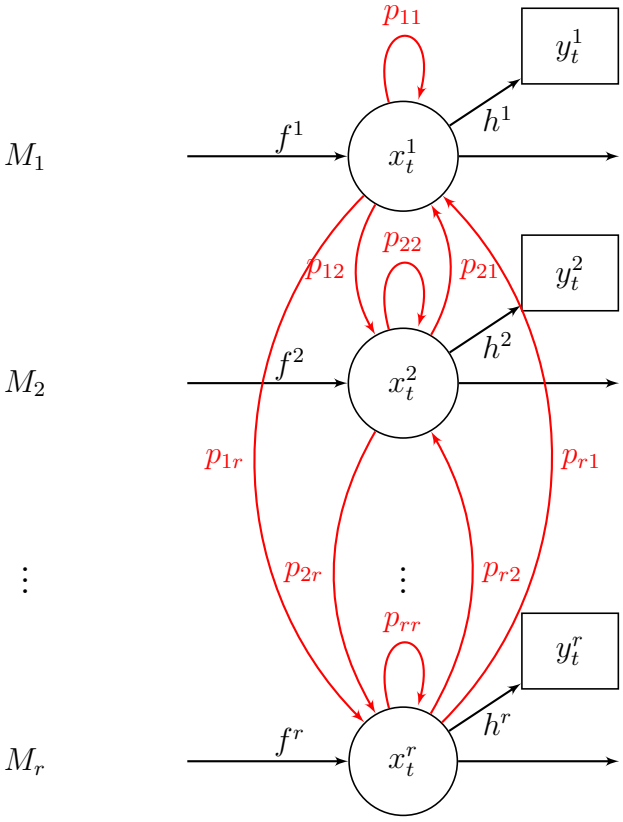


Figure 3.5: Markov switching state space model interpretation used by IMM and its extensions, with r SSMs, $\{M_1, M_2, \ldots, M_r\}$. Red arrows denote the switching between models, with probability $p_{ij} := \mathrm{P}(M_j \,|\, M_i)$

error covariance denoted $\boldsymbol{x}_t^j$ and $\boldsymbol{P}_t^{j(x)}$ respectively.

The process weights at interval, t, are calculated as a mixing probability describing the normalised likelihood of switching between each pair of component states, $\mu_t^{i|j}$. The calculation is based on the a prior mixing probability and the likelihood of switching between the two states:

$$\begin{aligned} c_t^j &= \sum_{k=0}^{r} \mu_{t-1}^k \mathrm{P}(M_k \mid M_j) \\ \mu_t^{k|j} &= \frac{\mu_{t-1}^k \mathrm{P}(M_k \mid M_j)}{c_t^j} \end{aligned} \tag{3.28}$$

The *a priori* state estimate for a component model is a mixed input combining all previous individual component states weighted by the respective mixing probabilities. The prior state estimate for model M_j, and its estimation error covariance, is thus defined:

$$\begin{aligned} \boldsymbol{x}_{t-1}^{0j} &= \sum_{k=0}^{r} \mu_t^{k|j} \boldsymbol{x}_{t-1}^k \\ \boldsymbol{P}_{t-1}^{0j(x)} &= \sum_{k=0}^{r} \mu_t^{k|j} \left(P_{t-1}^{k(x)} - \left(\boldsymbol{x}_{t-1}^k - \boldsymbol{x}_{t-1}^{0j}\right)\left(\boldsymbol{x}_{t-1}^k - \boldsymbol{x}_{t-1}^{0j}\right)^{\mathrm{T}} \right) \end{aligned} \tag{3.29}$$

Component *a priori* states are propagated through the Kalman filter (or its extensions) prediction and correction steps to produce an *a posteriori* estimate for each component model. Alongside this, the likelihood of each component observation is calculated from the measurement residual of the respective SSM filter output, and these are combined with the prior switching likelihood and probabilities to estimate the posterior mixing probability.

$$\begin{aligned} \Lambda_t^j &= \mathcal{N}\left(\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_{t|t-1}^j; \mathbf{0}, \boldsymbol{P}_{t|t-1}^{j(y)}\right) \\ \mu_t^j &= \frac{\Lambda_t^j c_t^j}{\sum_{k=1}^{r} \Lambda_t^k c_t^k} \end{aligned} \tag{3.30}$$

The observation estimate, $\boldsymbol{y}_{t|t-1}^j$, describes the Kalman predicted observation for SSM M_j, with $\boldsymbol{P}_{t|t-1}^{j(y)}$ denoting the corresponding estimate error covariance. Combining the component model state estimations into a single uniform estimate

is simply the sum of each component state weighted by its respective posterior mixing probability.

$$\begin{aligned} \boldsymbol{x}_{t|t} &= \sum_{j=1}^{r} \mu_t^j \boldsymbol{x}_t^j \\ \boldsymbol{P}_{t|t}^{(x)} &= \sum_{j=1}^{r} \mu_t^j \left(\boldsymbol{P}_t^{j(x)} - \left(\boldsymbol{x}_{t|t} - \boldsymbol{x}_t^j\right)\left(\boldsymbol{x}_{t|t} - \boldsymbol{x}_t^j\right)^{\mathrm{T}} \right) \end{aligned} \tag{3.31}$$

Note that the unified state $\boldsymbol{x}_{t|t}$ is not used in the recursive propagation of the filter, instead the individual component states are maintained, along with the mixing probabilities for IMM. The IMM also relies on the switching likelihood of the Markov-switching system, $\mathrm{P}(M_k \mid M_j)$, as well as the usual initial parameters of the Kalman filter for each respective component model M_j.

Variations and extensions on the concepts behind IMM have been developed, specifically to solve problems such as non-linear process or observation models, where linearised or unscented Kalman filters have been used. Further, a smoother, using the mixing probabilities and Gaussian-moment matching, based on IMM, is a leading solution to recursive Bayesian smoothing in Markov-switching model problems (Xiaolong and Pingfang, 2009).

Algorithm 2 Unscented Kalman Filter

procedure $\mathrm{UT}(g(\cdot), m, P; \alpha, \beta, \kappa)$ ▷ .
  $\lambda \leftarrow \alpha^2(n+\kappa) - \kappa$
  $\chi_0 \leftarrow m$
  $w_0 = \lambda/(n+\lambda)$
  for $i \in 1, \ldots, n$ **do**
    $\chi_i,\ \chi_{i+n} \leftarrow \chi_0 \pm \left[\sqrt{(n+\lambda)P}\right]_i$
    $\gamma_i,\ \gamma_{i+n} \leftarrow g(\chi_i),\ g(\chi_{i+n})$
    $w_{i+n} \leftarrow w_i \leftarrow 1/2(n+\lambda)$
  end for
  return $[\chi_i]_{i=0}^{2n},\ [\gamma_i]_{i=0}^{2n},\ [w_i]_{i=0}^{2n}$
end procedure
procedure $\mathrm{UT}^{-1}([\chi_i]_{i=0}^{2n}, [\gamma_i]_{i=0}^{2n}, [w_i]_{i=0}^{2n}; \alpha, \beta, \kappa)$
  $x,\ y \leftarrow \sum_{i=0}^{2n} w_i\chi_i,\ \sum_{i=0}^{2n} w_i\gamma_i$
  $\Sigma \leftarrow (w_0 + 1 - \alpha^2 + \beta)(x-\chi_0)(y-\gamma_0)^{\mathrm{T}} + \sum_{i=1}^{2n} w_i(x-\chi_i)(y-\gamma_i)^{\mathrm{T}}$
  return $x,\ \Sigma$
end procedure
procedure PREDICT($\boldsymbol{x}_{t-1},\ \boldsymbol{P}_{t-1}^{(x)}$)
  $_, \boldsymbol{\chi}_{t|t-1}, \boldsymbol{w} \leftarrow \mathrm{UT}(f(\cdot), \{\boldsymbol{x}_{t-1}, \boldsymbol{P}_{t-1}^{(x)}\})$
  $\boldsymbol{x}_{t|t-1}, \boldsymbol{P}_{t|t-1}^{(x)} \leftarrow \mathrm{UT}^{-1}(\boldsymbol{\chi}_{t|t-1}, \boldsymbol{\chi}_{t|t-1}, \boldsymbol{w})$
  return $\boldsymbol{x}_{t|t-1},\ \boldsymbol{P}_{t|t-1}^{(x)} + \boldsymbol{Q}_t$
end procedure
procedure UPDATE($\boldsymbol{x}_{t|t-1},\ \boldsymbol{P}_{t|t-1}^{(x)},\ \tilde{\boldsymbol{y}}_t$)
  $\boldsymbol{\chi}_t, \boldsymbol{\gamma}_t, \boldsymbol{w} \leftarrow \mathrm{UT}(h(\cdot), \{\boldsymbol{x}_{t|t-1}, \boldsymbol{P}_{t|t-1}^{(x)}\})$
  $_, P_t^{(xy)} \leftarrow \mathrm{UT}^{-1}(\boldsymbol{\chi}_t, \boldsymbol{\gamma}_t, \boldsymbol{w})$
  $\boldsymbol{y}_t, P_t^{(xy)} \leftarrow \mathrm{UT}^{-1}(\boldsymbol{\gamma}_t, \boldsymbol{\gamma}_t, \boldsymbol{w})$
  $\mathbf{K}_t \leftarrow \boldsymbol{P}_t^{(xy)} \left(\boldsymbol{P}_t^{(y)} + \boldsymbol{R}\right)^{-1}$
  $\boldsymbol{x}_{t|t} \leftarrow \boldsymbol{x}_{t|t-1} + \mathbf{K}_t\left(\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_t\right)$
  $\boldsymbol{P}_{t|t}^{(x)} \leftarrow \boldsymbol{P}_{t|t-1}^{(x)} - \mathbf{K}_t(\boldsymbol{P}_t^{(y)} + \boldsymbol{R}_t)\mathbf{K}_t^{\mathrm{T}}$
  return $\boldsymbol{x}_{t|t},\ \boldsymbol{P}_{t|t}^{(x)}$
end procedure
procedure UNSCENTED KALMAN FILTER($\boldsymbol{x}_0,\ \boldsymbol{P}_0^{(x)};\ f, [\boldsymbol{Q}_t]_{t=1}^{T}, h, [\boldsymbol{R}_t]_{t=1}^{T}$)
  $\boldsymbol{x}_{t-1}, \boldsymbol{P}_{t-1}^{(x)} \leftarrow \boldsymbol{x}_0, \boldsymbol{P}_0^{(x)}$
  for $t \in 1, \ldots, T$ **do**
    $\boldsymbol{x}_{t|t-1}, \boldsymbol{P}_{t|t-1}^{(x)} \leftarrow$ PREDICT($\boldsymbol{x}_{t-1},\ u_t,\ \boldsymbol{P}_{t-1}^{(x)}$)
    if measurement $\tilde{\boldsymbol{y}}_t$ available **then**
      $\boldsymbol{x}_{t|t}, \boldsymbol{P}_{t|t}^{(x)} \leftarrow$ UPDATE($\boldsymbol{x}_{t|t-1},\ u_t,\ \boldsymbol{P}_{t|t-1}^{(x)},\ \tilde{\boldsymbol{y}}_t$)
      $\boldsymbol{x}_t, \boldsymbol{P}_t^{(x)} \leftarrow \boldsymbol{x}_{t|t}, \boldsymbol{P}_{t|t}^{(x)}$
    else
      $\boldsymbol{x}_t, \boldsymbol{P}_t^{(x)} \leftarrow \boldsymbol{x}_{t|t-1}, \boldsymbol{P}_{t|t-1}^{(x)}$
    end if
  end for
  return $\left[(\boldsymbol{x}_t, \boldsymbol{P}_t^{(x)})\right]_{t=1}^{T}$
end procedure

4

Fuzzy Logic and Control

4.1 Fuzzy Sets

A fuzzy set differs from a conventional set in that, rather than simply being a list of members contained therein, it is a constituent of two parts: some space over which a member can be taken from, X; and a functional, mapping a member to some degree of truth that it belongs to the given fuzzy set, u (Klir and Yuan, 1995; Zadeh, 1996; Mendel, 2001).

$$\underset{\sim}{A} = (X, u : X \to [0, u_{\max}]) \tag{4.1}$$

For a traditional fuzzy set, $u_{\max} = 1$, but there are exist a class of *intuitionistic* fuzzy sets, whereby an additional functional $v : X \to [0, v_{\max}]$ exists, describing the non-membership of the set. The two functionals are defined such that $u_{\max} + v_{\max} \leq 1$, i.e. $0 \leq u(x) + v(x) \leq 1, \forall x \in X$ (Atanassov, 1986). In this respect, the traditional fuzzy set would define v as the complement of u: $v(x) \equiv 1 - u(x)$.

The definition of a fuzzy set is such that where multiple fuzzy sets are defined for some space, X, a value may belong to multiple sets in varying degrees. In certain extensions on fuzzy set theory, interactions between sets are required, and so operations on fuzzy sets must thus be defined. For crisp (non-fuzzy) sets, operations such as complement, intersection and union may be defined as simple logical deductions, respectively:

$$\begin{aligned} \bar{A} \equiv \neg A &= \{x \notin A \,|\, \forall x\} \\ A \cap B \equiv \neg(\bar{A} \cup \bar{B}) &= \{x \,|\, x \in A \text{ and } x \in B\} \\ A \cup B \equiv \neg(\bar{A} \cap \bar{B}) &= \{x \,|\, x \in A \text{ or } x \in B\} \end{aligned} \tag{4.2}$$

Operators on fuzzy sets, however, cannot rely on these Boolean operations, because all members in the function space can belong to all fuzzy sets to different degrees. The operations of complement, intersection and union are instead defined as functions over the unit interval, $[0, 1]$, representing the range of membership function u. Since the respective result of the operations will be a fuzzy set over the same space, the operations can be defined as compositions of the member functions. Standard fuzzy operations for complement, intersection and union are respectively defined, where fuzzy sets $\underset{\sim}{A} = (X, u_A)$ and $\underset{\sim}{B} = (X, u_B)$:

$$\begin{aligned} \bar{\underset{\sim}{A}} &\equiv (X, u_\neg : u_\neg(x) = 1 - u_A(x), \forall x \in X) \\ \underset{\sim}{A} \cap \underset{\sim}{B} &\equiv (X, u_\cap : u_\cap(x) = \min(u_A(x), u_B(x)), \forall x \in X) \\ \underset{\sim}{A} \cup \underset{\sim}{B} &\equiv (X, u_\cup : u_\cup(x) = \max(u_A(x), u_B(x)), \forall x \in X) \end{aligned} \tag{4.3}$$

A visual representation of the standard fuzzy set operations is shown in Figure 4.1, for two fuzzy sets with triangular membership functions.

While the above-defined operators are arguably the most common implementations of fuzzy set operations, they may be defined over any function that satisfies a set of basic axioms. Notably for set complement, the requirements are that the complement membership function is strictly negative continuous-monotonic, such that $u_\neg(a) < u_\neg(b)$ iff. $u(a) > u(b)$, and that $u_\neg$ is involutory, such that $\neg(u_\neg) \equiv u$. A general interpretation of fuzzy intersection can be defined as some t-norm, and shares all the properties therein, while satisfying the axiom of sub-idempotence: the intersection of a fuzzy set with itself is *at most* itself. The general operations for intersection can thus be defined $u_\cap(x) := \top_\cap(u_A(x), u_B(x))$. Examples of t-norms used include the product, or the Łukasiewicz t-norm. The general union operator shares many of the properties of a t-norm, however the identify for a union operator is 0, as opposed to 1 for a t-norm.

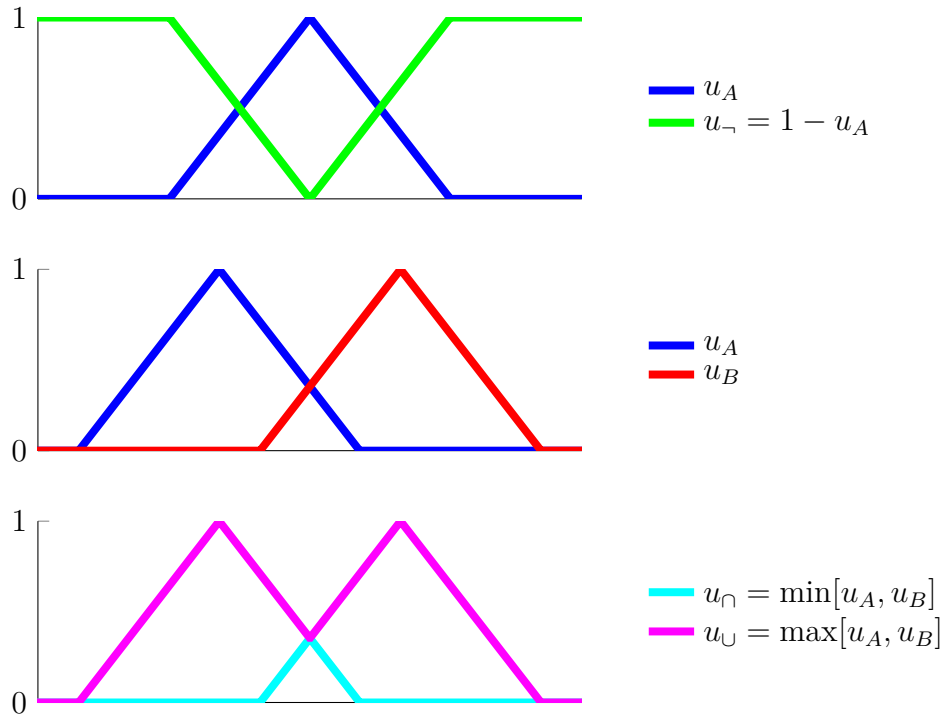


Figure 4.1: An example of standard fuzzy operations, on triangular membership functions u_A and u_B, showing membership functions for complement, $u_\neg$, intersection, $u_\cap$, and union, $u_\cup$

4.2 Fuzzy Inference

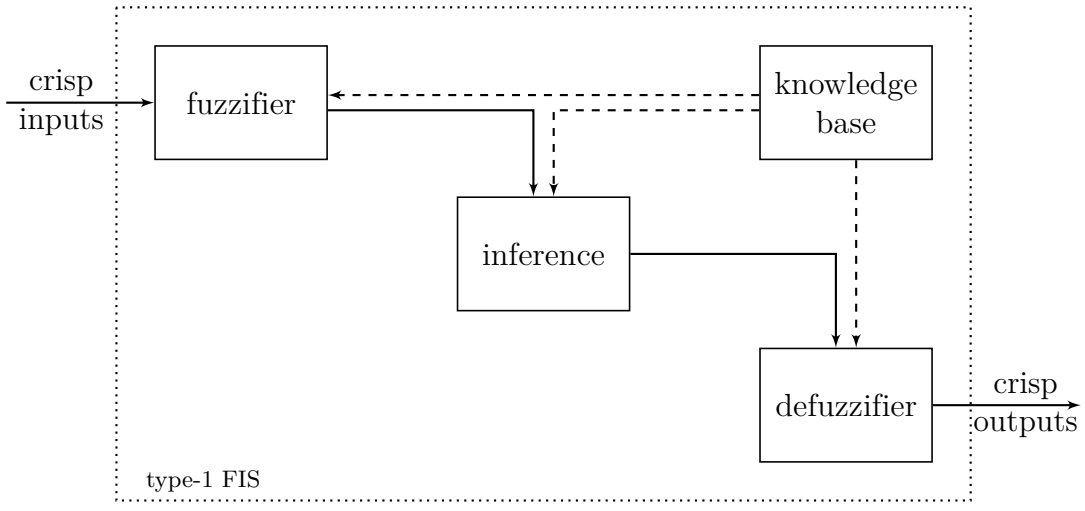


Figure 4.2: Type-1 fuzzy inference system

Fuzzy inference describes the reasoning system by which some input is interpreted using fuzzy logic and operations some fuzzy sets representing *lingustic* values with uncertain distinction. Given a number of fuzzy sets, a fuzzy inference system is predicated on a series of IF–THEN rules that depend on membership to respective sets (Mendel, 2001). A fuzzy rule is typically defined with an antecedent and consequent. The input is *fuzzified* by its membership to the antecedent sets, and the intersection of these memberships for a given rule form a *firing strength*, describing the degree to which a rule is fulfilled. Fuzzy rules are typically denoted as follows:

$$\begin{aligned} \text{IF} \quad & \overbrace{a \text{ is } \underset{\sim}{A} \text{ and } b \text{ is } \underset{\sim}{B} \text{ and } \ldots}^{\text{antecedent}} \\ \text{THEN} \quad & \underbrace{z \text{ is } \underset{\sim}{Z}}_{\text{consequent}} \end{aligned} \tag{4.4}$$

The format of the consequent is dependent on the fuzzy model used. The major categorisations of fuzzy consequents are the Mamdani-Larsen (ML) model, and the Takagi-Sugeno-Kang (TSK) model. The ML considers the consequent as a fuzzy set $\underset{\sim}{Z}$ with its own membership function. In the TSK model, Z is instead some crisp function. The consequent of each rule are aggregated with their corresponding firing strength and some *defuzzification* criteria is used to convert the aggregate into a crisp output.

A visual representation of the fuzzy inference system is shown in Figure 4.2. The system is describes the basic process outlined. Each component part is dependent on the knowledge base of the given problem, which defines the rules and membership functions for the respective fuzzy sets involved in the inference system. Methods for defuzzification are typically dependent on the fuzzy model used, and the application requirements. Examples of such methods include centre-of-area, or bisector-of-area of the consequent fuzzy set in an ML model; or a weighted fuzzy mean for the consequent crisp singleton values in a TSK model.

4.3 Type-2 Fuzzy Logic

4.3.1 Type-2 Fuzzy Sets

Fuzzy logic is designed to address uncertainty in assigning a category or label to a datum. The nature of fuzzy sets in traditional, so-called *type-1*, fuzzy logic uses

crisp functions to describe the membership of fuzzy sets (Mendel and John, 2002; Mendel, 2014). Type-2 fuzzy logic is an extension on this, that can be used to describe both uncertainty in set membership, and the uncertainty in the means by which membership is quantified. The extension of sets over which inference occurs means that in a type-2 fuzzy inference system, the underlying membership criteria can be considered three-dimensional, and in seeking a crisp output, some type reduction must occur. The result of this type reduction should produce some type-1 fuzzy set(s) on which defuzzification can apply. An overview of this is presented in Figure 4.3, demonstrating requirements for extending the type-1 FIS, when compared with Figure 4.2.

A formal definition of a type-2 fuzzy set is defined as follows, where the tuple (X, u) represents the space over which the set is defined and the space of all admissible type-1 membership functions.

$$\underset{\sim}{A} = ((X, J_x), \mu_A : X \times J_x \to [0, \mu_{\max}]) \tag{4.5}$$

The secondary membership function $\mu(x, u)$, where $x \in X$ and $u \in J_x$, characterises the uncertainty surrounding the membership u of x. The area that μ_A covers is described as the footprint of uncertainty (FOU), and effectively blurs the underlying type-1 fuzzy sets. Note that as before, typically $\mu_{\max} = 1$.

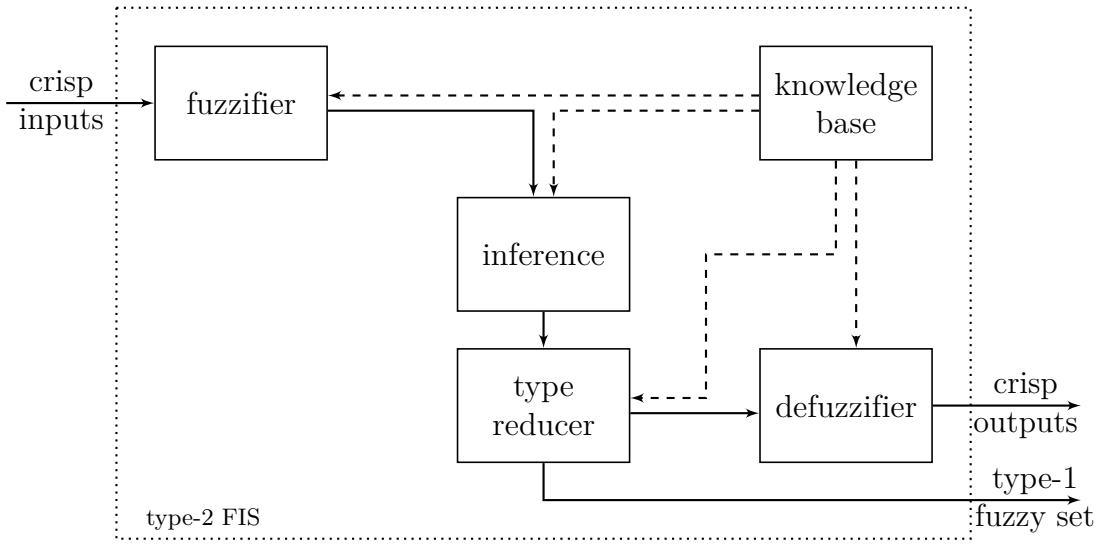


Figure 4.3: Type-2 fuzzy inference system

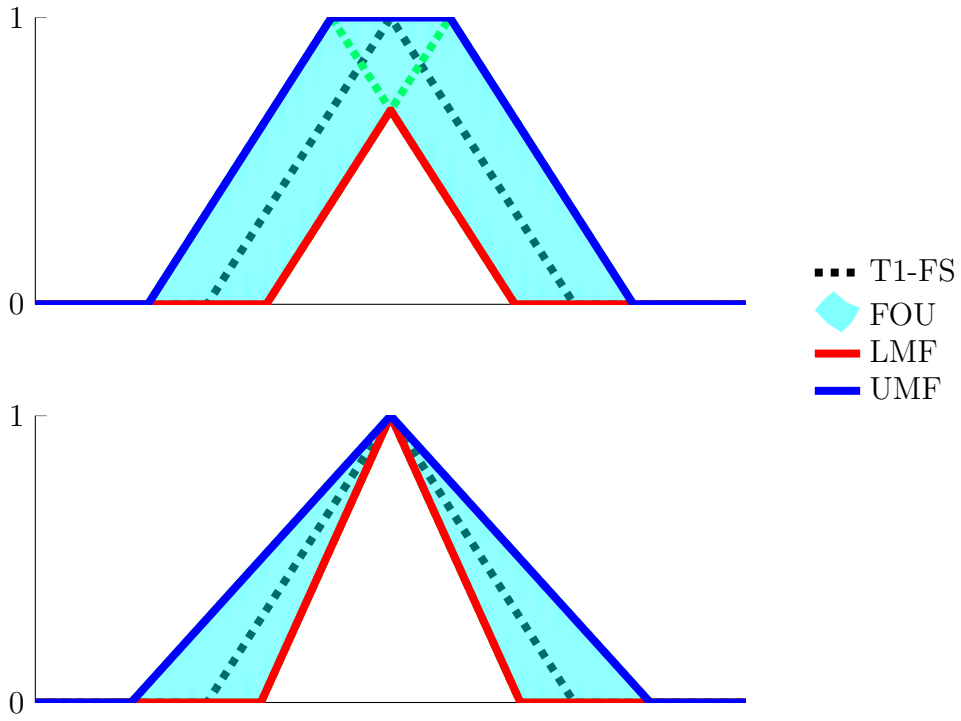


Figure 4.4: Interval type-2 fuzzy sets with uncertain mean (top) and uncertain variance, showing the upper (UMF) and lower (LMF) membership functions, and footprint of uncertainty (FOU) of the IT2 set. Comparatively, a type-1 fuzzy set is underlaid, along with the membership distributions corresponding to the extreme mean values in the membership interval (green, dotted)

4.3.2 Interval Type-2 Fuzzy Logic

While general type-2 fuzzy sets are powerful representations of the uncertainty surrounding the membership criteria, the high dimensional nature of the resulting set leads to increased computational complexity, especially in the continuous case. As such, the most common use of type-2 fuzzy logic is in the use of *interval* type-2 (IT2) fuzzy sets, which bound the uncertainty of the membership with intervals, such that $\mu_A(x,u) = 1, \forall x \in X, \forall u \in J_x$. IT2 fuzzy sets can be associated with the mathamatics of type-1 sets, and provides a simpler, but potentially less powerful, interpretation of type-2 fuzzy logic (Mendel et al., 2006).

An interval is defined as a (crisp) set defined by two extremes such that such that all points between the extremums are contained within the set, i.e. interval defined over space B, $[a,b] = \{c \in B \,|\, a <= c <= b\}$[1]. An IT2 fuzzy set is thus described as an interval of membership functions defined as an upper and lower bound that describe the set J_x:

[1]This definition technically describes a closed set which contains, within itself, its limits

$$\underset{\sim}{A} = \left(X, \left[\underline{\mu}_A : X \to [0,1], \overline{\mu}_A : X \to [0,1]\right]\right) \tag{4.6}$$

As before, the X denotes the space over which a member can be taken from. Note that while $\underset{\sim}{A}$ is a type-2 set, the secondary membership always maps to 1 and so is omitted from the definition. It is also clear here that the membership functions that define the interval J_x are themselves type-1 membership functions. As well as simplifying the definition of the type-2 fuzzy set, this property also allows for the arithmetic of type-1 fuzzy logic to be extended and applied onto the IT2 set. Note that by the nature of interval logic, where $\underline{\mu}(x) = \overline{\mu}(x)$, $\underset{\sim}{A}$ is reduced to a type-1 fuzzy set. The two functions, $\underline{\mu}(x)$ and $\overline{\mu}(x)$ are respectively referred to hereafter as the lower (LMF) and upper (UMF) membership functions of an IT2 fuzzy set.

The key feature of a type-2 fuzzy set is in the effective blurring of the membership of a value. Where a crisp set has binary admissibility criteria and a type-1 fuzzy set is a simple mapping from primary domain into membership space, the type-2 set has a so-called *footprint of uncertainty* (FOU) over the membership space, as defined by J_x. The FOU may be defined as some uncertainty over a variety of parameters, and the LMF and UMF of the proceeding IT2 set are defined as the respective lower and upper bounds of the this FOU.

Consider the two IT2 fuzzy sets in Figure 4.4. In top plot, the footprint results from the case where the mean value of a triangular membership function is uncertain, resulting in thick FOU with a plateau in the upper limit; the green distributions represent the distributions created from the upper and lower estimates of the mean. In the lower plot, the uncertainty arises from uncertain spread, or variance, leading to wide FOU deviating from the mean of the membership functions. In both cases, it can be seen that the UMF is defined as the upper bound of the FOU, and the LMF as the lower bound.

Operations over multiple IT2 fuzzy sets must take into account both membership functions. Since both can be considered as type-1 fuzzy sets, operators defined over such, e.g. in (4.3), can be applied on IT2 sets. Here, the respective set operations on two IT2 fuzzy sets, $\underset{\sim}{A}$ and $\underset{\sim}{B}$, over member space X can be defined as follows:

$$
\begin{aligned}
\bar{\underset{\sim}{A}} &\equiv \left(X, \left[1 - \underline{\mu}_A(x), 1 - \overline{\mu}_A(x)\right], \forall x \in X\right) \\
\underset{\sim}{A} \cap \underset{\sim}{B} &\equiv \left(X, \left[\top_\cap(\underline{\mu}_A(x), \underline{\mu}_B(x)), \top_\cap(\overline{\mu}_A(x), \overline{\mu}_B(x))\right], \forall x \in X\right) \\
\underset{\sim}{A} \cup \underset{\sim}{B} &\equiv \left(X, \left[\bot_\cup(\underline{\mu}_A(x), \underline{\mu}_B(x)), \bot_\cup(\overline{\mu}_A(x), \overline{\mu}_B(x))\right], \forall x \in X\right)
\end{aligned} \tag{4.7}
$$

Here, the $\top_\cap$ represents a t-norm for the intersection operation over two fuzzy sets, typically either the minimum or product. $\bot_\cup$ denotes a t-conorm, most commonly the maximum.

The applications of these operators most commonly lie in inference systems. Consider the following fuzzy ruleset:

$$
\text{Rule } j : \begin{cases} \text{IF} & a(1) \text{ is } \underset{\sim}{A}^j(1) \text{ and } \ldots \text{ and } a(n) \text{ is } \underset{\sim}{A}^j(n) \\ \text{THEN} & z \text{ is } Z^j \end{cases} \tag{4.8}
$$

Here, the antecedent is formed from n IT2 fuzzy sets, where $\underset{\sim}{A}^j(k)$ can be considered the k^{th} fuzzy set corresponding to rule j, with membership defined over the interval $[\underline{\mu}_k^j, \overline{\mu}_k^j]$, and $a(k)$ is the k^{th} value of some vector of inputs, $\mathbf{a} = (a(1), \ldots, a(n))$. As before, the underlying model defines the nature of Z, whereby in an MK model, Z^j are fuzzy sets, and in a TSK model, Z^j are crisp functions.

In the inference system over this ruleset, the firing strength of each rule is defined as an interval $\Omega^j(\mathbf{a}) = [\underline{\omega}^j(\mathbf{a}), \overline{\omega}^j(\mathbf{a})]$. The firing strength for a rule can be taken as the intersection over each set in the antecedent, and so by (4.8), the lower and upper firing strength for a rule are calculated as follows:

$$
\begin{aligned}
\underline{\omega}^j(\mathbf{a}) &= \top_{k=1}^{n} \underline{\mu}_k^j(a(k)) \\
\overline{\omega}^j(\mathbf{a}) &= \top_{k=1}^{n} \overline{\mu}_k^j(a(k))
\end{aligned} \tag{4.9}
$$

Where, $\top_{k=1}^{n} \mu_k$ is the operation of $\top_\cap$ over all μ, i.e. $\top_\cap(\mu_1, \top_\cap(\mu_2, \ldots)))$, exploiting the associative property of t-norms. The resulting firing strengths can be used as is, in the case of a TSK model, or combined with the respective limits of membership interval of fuzzy set $\underset{\sim}{Z}$, in the case of an ML model with

IT2 consequent. The resulting interval is defines the IT2 fuzzy set of the inference system, with the FOU for each rule consequent defined over the interval $[\top_{\cup}(\underline{\omega}^j(\mathbf{a}), \underline{\mu}^j_Z(z)), \top_{\cup}(\overline{\omega}^j(\mathbf{a}), \overline{\mu}^j_Z(z))]$, where $\underline{\mu}_Z = \overline{\mu}_Z = 1, \forall j, z$ in a TSK model.

4.3.3 Type Reduction

The firing strength of a fuzzy rule weights the effects that rule in an inference system. In the case of type-2 fuzzy logic, the firing strength is not a crisp value (Mendel, 2014). Indeed, in the case of IT2 sets, the firing strength over the input is an interval as defined in the previous section. As a result, resolving the output of a T2-FIS requires some type reduction to convert the model from type-2 to type-1 (see Figure 4.3). There are many different type reduction methods, involving centroid and maxima-based calculations over the IT2 fuzzy set. One of the most common set of methods include the set of enhanced Karnik-Mendel (KM) set of methods (Mendel et al., 2006). One such method is the centre-of-sets type-reduction, which results in an interval describing the bounds of a type-1 fuzzy set.

The general method for calculating is to find some lower and upper bounds, $[z_l, z_u]$ such that

$$
\begin{aligned}
z_l &= \min_{L \in [a,b]} \frac{\int_a^L z\overline{\mu}_z(z)\mathrm{d}z + \int_L^b z\underline{\mu}_z(z)\mathrm{d}z}{\int_a^L \overline{\mu}_z(z)\mathrm{d}z + \int_L^b \underline{\mu}_z(z)\mathrm{d}z} \\
z_r &= \max_{R \in [a,b]} \frac{\int_a^R z\underline{\mu}_z(z)\mathrm{d}z + \int_R^b z\overline{\mu}_z(z)\mathrm{d}z}{\int_a^R \underline{\mu}_z(z)\mathrm{d}z + \int_R^b \overline{\mu}_z(z)\mathrm{d}z}
\end{aligned}
\tag{4.10}
$$

The computation of these bounds requires an iterative method to solve the switching points between UMF and LMF as defined by L and R over the support of $z \in Z$. In the discrete case, where the member space Z of the IT2 fuzzy set is defined as a set of points z_i, the integrals may be approximated with summation over the Z. The KM methods seek to calculate the switch points, such that the centroid of the set can be found.

The continuous KM method initialises some estimation of L and R as the normalised integral of the $z\mu_z(z)$, where $\mu_z = (\underline{\mu}_z - \overline{\mu}_z)/2$, i.e. $L = R = \int z\mu_z \mathrm{d}z / \int \mu_z \mathrm{d}z$. The values of L and R are used in the integral sums defined on the right-hand side of (4.10) for the lower and upper intervals respectively. The result of each integral sum, L' and R' respectively are compared with the current estimations, such that if $|L' - L| < \varepsilon$ then $z_l = L'$ and if $|R' - R| < \varepsilon$ then $z_r = R'$. The parameter ε

defines some small error tolerance. If the conditional does not pass, the estimates of switch points L and R are set to L' and R' and the integral is recalculated. This process is iterated for z_l and z_r independently, and once the condition passes (reportedly in 3-4 iterations), the bounds are used to describe the type-reduced set. Defuzzification of the resulting set is then simply the mean of the intervals, $(z_l + z_r)/2$.

In a T2-FLS, where the inference system relies on a MK model over IT2 fuzzy sets, the above calculations are required to output a crisp value. In the case of a TSK model, the support of Z is unary, such that $z = a = b$. Therefore, with the fact that the output is crisp, and thus the membership is 1, it intuitively follows that the integrals reduce down to $z_l = z_r = z$. Therefore, the joint output of an inference system with these conditions over r rules is simply the normalised weighted sum of each rule's consequent:

$$z(\mathbf{a}) = \frac{\sum_{j=1}^{r} \omega^j(\mathbf{a}) z^j}{\sum_{j=1}^{r} \omega^j(\mathbf{a})} \tag{4.11}$$

Where $\omega^j = (\underline{\omega}^j + \overline{\omega}^j)/2$, the type-reduced firing strength of each rule.

4.4 Applications

4.4.1 Fuzzy Clustering

Clustering describes the problem of grouping objects or data points based on particular similarities. There are many approaches to solving this problem, with most methods seeking to minimise some distance function defining each cluster. A notable category of these methods are centroid-based clustering algorithms, defined such that a cluster is defined by all points that are closest to its geometric centre, i.e. its centroid.

The simplest implementation of centroid-based clustering is the k-means clustering. For some k randomised centroids that lie within the range of data, the grouping is initialised by assigning data points to the cluster beloining to the centroid with which it has minimum distance to. The choice of distance function is largely dependent on the type of data, but the most common function used is the L2 (Euclidean) norm. Due the randomised nature of the initialisation, k-means is an iterative method that recalculates the centroid of the clusters using their newly assigned data. The data points are dependently reassigned to the clusters, based on

the same distance relation but with relation to the newly calculated centroid. The process of centroid calculation and data assignment is repeated until the location of each cluster's centroid converges such that it does not change over repeated reassignments.

The k-means method, and others that directly associate data points with a single cluster, are referred to as hard clustering. In hard clustering, each datum has a crisp membership to a single cluster. The top-right plot in Figure 4.5 shows the Voronoi partitioning that results from k-means clustering. The plot serves to demonstrate that k-means can successfully populations where they are distinctly separate, but suffers where there is overlapping. In such cases, there is no way to represent the uncertainty in categorisation when using a hard clustering approach. To overcome this, an extension was developed, in the form of fuzzy c-means.

Fuzzy c-means (FCM) is a method that seeks to associate a membership function for each cluster representative of the population and the uncertainty surrounding associating a datum with it. In FCM, each cluster is a type-1 fuzzy set over the data space, X. For n data points, $X' = \{x_i \in X \mid 1 \leq i \leq n\}$, and c clusters, the fuzzy partition of the data set can be defined by a $c \times n$ real matrix U. The entries of U consist of the membership grades of each datum in X' as evaluated by the membership function of each fuzzy set $\underset{\sim}{C}^j = (X, u^j : X \to [0, 1])$.

In FCM, the entries of partition matrix U must satisfy a set of conditions:

1. $U^j = \{u_j(x_1), u_j(x_2), \ldots, u_j(x_n)\}$ is the membership function of the j^{th} fuzzy set $\underset{\sim}{C}^j$ evaluated over X'
2. $U_i = \{u_1(x_i), u_2(x_i), \ldots, u_c(x_i)\}$ are the membership grades for a datum x_j to each of the c fuzzy sets
3. $\sum_{\forall j} u^j(x_i) = 1,\ 1 \leq i \leq n$, i.e. the sum of all membership grades for a datum is equal to 1
4. $0 < \sum_{\forall i} u^j(x_i) < n,\ 1 \leq j \leq c$, i.e. no fuzzy set $\underset{\sim}{C}^j$ can be empty or contain all of X' entirely

The fuzzy partition that define the final fuzzy sets that describe the result of the clustering can be generated iteratively, much in the same way the final clusters are created in k-means. Based on the defining criteria, an initial partition $U^{(0)}$ is randomly initialised. As with k-means, a value representing the centroid of each cluster must be calculated. For a given cluster, the centroid v^j is calculated as the sum of each datum, weighted by its normalised membership to the cluster.

Algorithm 3 Fuzzy c-means

procedure INITIALISE(c, X) ▷ initialise c clusters over space X
  for $j \in 1, \ldots, c$ **do**
    $\nu^j \sim \mathcal{U}(X)$
  end for
  return $[\nu^j]_{j=1}^c$
end procedure
procedure ASSIGN($\mathbf{x}$, $[\nu^j]_{j=1}^c$) ▷ Assign memberships of data $\mathbf{x}$ to clusters
  for $x_i \in \mathbf{x}$ **do**
    for $j \in 1, \ldots, c$ **do**
      $d^j(x) = \|\nu^j - x_i\|$ ▷ calculate distances from centroids
    end for
    for $j \in 1, \ldots, c$ **do**
      $U_i^j \leftarrow \left(\sum_{k=1}^c \left(\frac{d^j(x)}{d^k(x)} \right)^{\frac{2}{m-1}} \right)^{-1}$
    end for
  end for
  return U
end procedure
procedure RE-WEIGHT(U)
  for $j \in 1, \ldots, c$ **do**
    $\nu^j = \left(\sum_{i=1}^n (U_i^j)^m x_i \right) \left(\sum_{i=1}^n (U_i^j)^m \right)^{-1}$ ▷ $U_i^j \equiv u^j(x_i)$
  end for
  return $[\nu^j]_{j=1}^c$
end procedure
procedure FCM($\mathbf{x}, c$; $m, X, b_{\max}$)
  $[\nu^j]_{j=1}^c \leftarrow$INITIALISE(c, X)
  $U^{(0)} \leftarrow$ASSIGN($\mathbf{x}$, $[\nu^j]_{j=1}^c$)
  while not converged **do**
    $b \leftarrow b + 1$
    $[\nu^j]_{j=1}^c \leftarrow$RE-WEIGHT($U^{(b-1)}$)
    $U^{(b)} \leftarrow$ASSIGN($\mathbf{x}$, $[\nu^j]_{j=1}^c$)
    if $b \geq b_{\max}$ **then**
      break
    end if
  end while
  return $U^{(b)}$, $[\nu^j]_{j=1}^c$
end procedure

$$v^j = \frac{\sum_{i=1}^n u^j(x_i)^m x_i}{\sum_{i=1}^n u^j(x_i)^m} \tag{4.12}$$

Where $x(k)$ is the value of the datum in the k^{th} dimension of X, and m refers to some fixed weighting exponent $m > 1$.

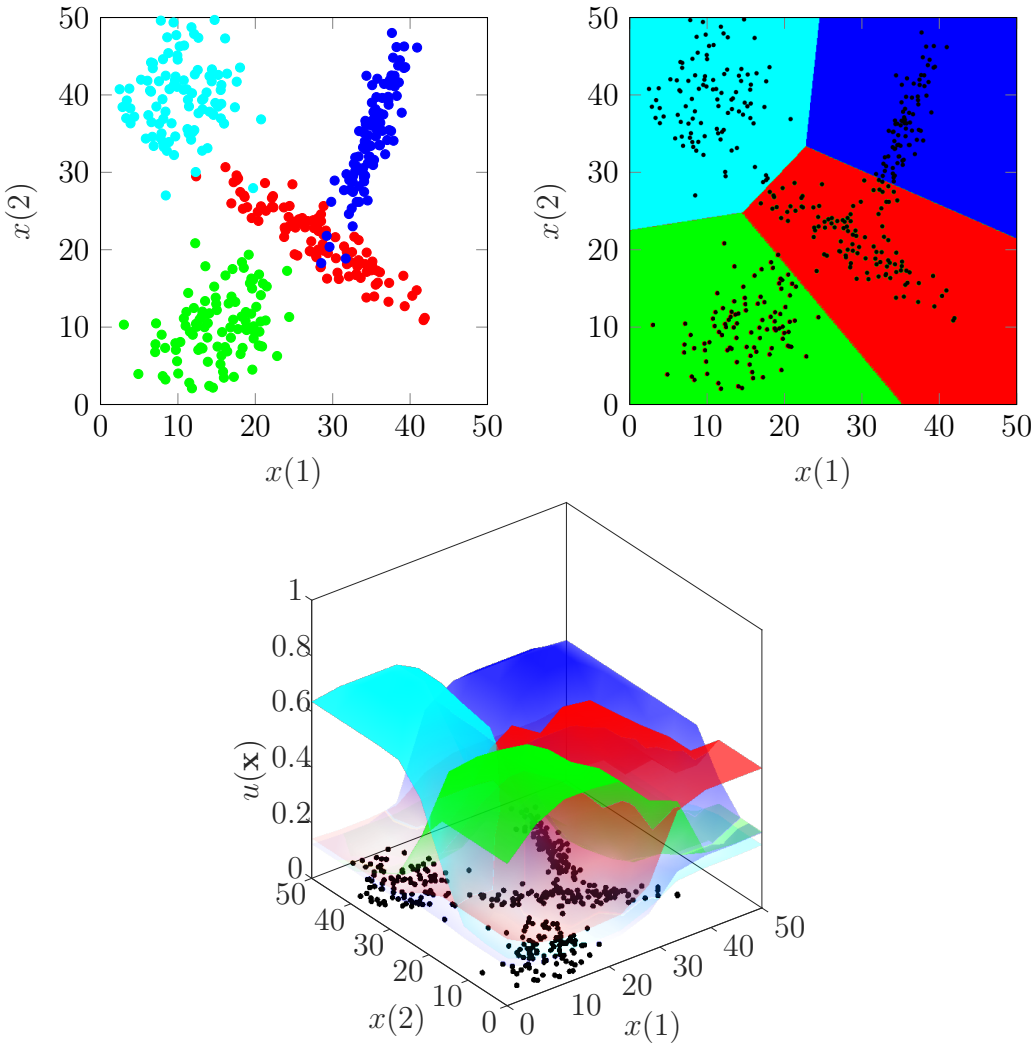


Figure 4.5: Example clustering problem over domain $\mathbf{x} \in [0, 50] \times [0, 50]$, with four partially overlapping populations (top left). The result of k-means hard clustering gives non-overlapping regions for each population (top right), whereas a fuzzy approach like fuzzy c-means returns membership functions for each population giving handling the uncertainty (bottom).

Based on the centroid of a cluster, the membership of each data point may be calculated, using some distance function, $d^j(x) = \|v^j - x\|$. The membership function for a given cluster C^j is thus defined:

$$u^j(x) = \left[\sum_{k=1}^{c} \left(\frac{d^j(x)}{d^k(x)} \right)^{\frac{2}{m-1}} \right]^{-1} \tag{4.13}$$

Since the calculations of centroid v^j and membership function u^j are mutually recursive, the estimations of partition function $U^{(b)}$ can be iterated such that the

calculations of u^j are dependent on the centroids calculated from the previous partitions $U^{(b-1)}$. The algorithm can be iterated until convergence, such that the change in centroid between iterations is minimised, or until some upper threshold of epochs. The resulting partition represents the final clustering membership grades, and each respective cluster be interpreted as a fuzzy rule in some inference system representing the classification of data into corresponding populations. In Figure 4.5, an example of these fuzzy membership grades over a data space is shown using fuzzy c-means and Euclidean distance function. The plot demonstrates, in particular, the uncertainty in boundaries between each population. The membership surfaces can be considered in contrast to the shaded regions of the k-means clustered data, where the hard interfaces between cells lead to misinterpretation of a data points membership to a population due to the absoluteness of the classification. The fuzzy clusters represent soft boundaries that allow for multiple possible interpretations of the data, especially those in the overlapping regions of different populations.

On the weighting exponent, m, the parameter is typically referred to as a control of the fuzziness of the resulting clusters. There are methods for heuristically defining the exponent, often constrained in the interval $[1.5, 2.5]$ Pal and Bezdek (1995), and a theoretical basis behind the choice of exponent is presented in Yu et al. (2004). However, the most commonly reported value is $m = 2$, which typically performs well empirically, and has some physical interpretation Bezdek and Dunn (1975). In the limit, as $m \to \infty$, the only solution is that every cluster shares the geometric centre of the given data as its centroid Yu et al. (2004).

On Similarities to Expectation Maximisation

The nature of fuzzy c-means clustering is similar, in ways, to expectation maximisation (EM): a statistical technique for fitting pdfs to mixture models, i.e. assigning fitting individual distributions to clusters. The nature of expectation maximisation is an iterative process, consisting of so-called E- and M-steps, after appropriate initialisation.

Where k-means clustering is a particular case of EM, giving a point estimate of the underlying distributions in the form of centroids as the expected values, assigning probabilities to 1 to clusters for which the distance from centroid is minimum and 0 otherwise Gupta et al. (2011). As an extension of k-means, FCM could be considered a case of EM. Consider the EM approach for Gaussian mixture models (GMM), in that each *cluster* is represented by a Gaussian pdf with mean and covariance. The E-step in EM describes the assignment of probabilities for each datum to the respective clusters, which is followed by the M-step that updates the

prospective Gaussian parameters based on the weighted sum of the probabilities.

Contrasting the FCM algorithm, as described in Algorithm 3, with EM, it is clear there are similarities in the iterative nature. The assignment corresponds to the E-step, and re-weighting takes the place of an M-step. Indeed, the reweighting of centroids in FCM is equivalent to the updating of the means of the Gaussians in the M-step of the GMM example. The key difference is in the interpretation of cluster shape and the assignment, and interpretation, of probabilities and memberships in the respective methods.

In both k-means and FCM, the geometry of clusters is that they are circular, about some centre-of-mass, by which all points are assigned to a cluster based on some distance. This centroid is the sole parameter, with the *width* of the cluster dictated by the distance between centroids, with borders forming Voronoi cells in the case of k-means. In the extension from k-means to FCM, the adaptation of membership is analogous to probability measures in EM, despite the different representations of uncertainty: the subjective belief of membership of a datum to a given cluster is conceivably the respective likelihood of its membership. In comparison with GMM EM, the uncertainty in clusters is controlled by a predefined weighting exponent, m, and the distance metric used, rather than a parametric representation of spread as described by the Gaussian covariances.

In the past, literature has compared the two approaches in applications, e.g. Prakash and Kumari (2017), or attempted to combine the interpretations Nasser et al. (2006).

4.4.2 Predictive Fuzzy Control

Fuzzy logic has been widely used for control systems, including in the predicting output from a controller, enabling the system to predict changes to measurements (Lu et al., 2001). The TSK model has served as a powerful mechanism for combining linear systems to approximate smooth non-linear systems. With non-linear mechanisms, the parameter uncertainty is typically not handled sufficiently, and so an IT2 fuzzy sets are used (Lam and Seneviratne, 2008). Model predictive control systems are designed to approximate are designed to handle constraints and approximate multivariate and time-varying problems. The problems for which predictive fuzzy control has been applied, include systems of uncertain parameters, or with time-delays (Lu et al., 2015).

Given a system modelled with an IT2 FIS with r rules, with antecedent consisting

of fuzzy sets, and the consequent is a TSK model, a plant rules for predicting the system state x are defined:

$$\text{Rule } j: \begin{cases} \text{IF} & f_1(x(t)) \text{ is } \underset{\sim}{M}^j(1) \text{ and } \ldots \text{ and } f_n(x(t)) \text{ is } \underset{\sim}{M}^j(n) \\ \text{THEN} & x(t+1) = \mathbf{A}^j x(t) + \mathbf{B}^j u(t) \end{cases} \tag{4.14}$$

$\underset{\sim}{M}^j$ is an IT2 fuzzy set for the j^{th} rule that corresonds to functions $\mathbf{f}$. The parameter $u(t)$ represents input control and $\mathbf{A}^j$, and $\mathbf{B}^j$ are some known matrices describing the linear predictive system. The controller, designed to guarantee stability (Lam and Seneviratne, 2008), is as follows, for rule p (independent of j) with antecedent IT2 fuzzy set, $\underset{\sim}{N}^p$.

$$\text{Rule } p: \begin{cases} \text{IF} & g_1(x(t)) \text{ is } \underset{\sim}{N}^p(1) \text{ and } \ldots \text{ and } g_n(x(t)) \text{ is } \underset{\sim}{N}^p(n) \\ \text{THEN} & u_c(t) = \mathbf{C}(t) x_c(t) \end{cases} \tag{4.15}$$

The functions $\mathbf{g}$ are evaluated over the state, and u_c and x_c are the output and input of the controller, corresponding to lossy versions of the control parameter and process state, $u(t) = \nu(t) u_c(t)$ and $x_c(t) = \xi(t) x(t)$, with loss terms described by mutually independent Bernoulli processes, ξ and ν. It follows then that the inferred predictive model is the sum of rule consequent weighted by their firing strength (Lu et al., 2015):

$$\begin{aligned} x(t+1) &= \sum_{j=1}^{r} \hat{\omega}_f^j(x(t)) \left(\mathbf{A}^j x(t) + \mathbf{B}^j u(t) \right) \\ \text{and } u(t) &= \sum_{p=1}^{r} \xi(t) \nu(t) \hat{\omega}_g^p(x(t)) \mathbf{C}^p(t) x(t) \\ \Rightarrow x(t+1) &= \sum_{j=1}^{r} \hat{\omega}_f^j(x(t)) \sum_{p=1}^{r} \hat{\omega}_g^p(x(t)) \left(\mathbf{A}^j + \xi(t) \mathbf{B}^j \mathbf{C}^p(t) \right) x(t) \end{aligned} \tag{4.16}$$

The similarities in structure to the state space model are evident here, in the propagation of the state x by some weight matrix and a control parameters. Here, non-linearity can exist over the fuzzy memberships and in the resulting inference system. The control output is somewhat analogous to the observation model in the SSM, but it is used here as a feedback term that is not present in the HMM structure of the SSM.

5

State of the Art

5.1 Spatio-temporal Modelling

The popularity in state space models stems from their hierarchical structure that accounts for the variability in real-world systems. In particular, SSM models can handle randomness and uncertainty from stochastic temporal processes and measurement errors in their probabilistic nature (Särkkä, 2013). There are a wide range of applications to which state space models are applied, including robotics (Chen, 2012), ecology (Auger-Méthé et al., 2016), sensor control (Auger et al., 2013), and target tracking (Challa et al., 2011). Developments, particularly to Kalman filters, have been made both in terms of their application (Shi et al., 2002; Barut et al., 2005; Shi et al., 2012; Buckland et al., 2004; Patterson et al., 2008; Csilléry et al., 2013; Cao et al., 2014; Molnár et al., 2010) and theoretical extensions to the base linear Kalman filter (Einicke and White, 1999; Wan et al., 2000; Julier and Uhlmann, 2004; Särkkä, 2007; Arasaratnam and Haykin, 2009; Shmaliy, 2011; Ahn and Chan, 2014).

In control and robotics, the Kalman filter and its extensions is widely used for the estimation of parameters in control systems and in predictive control. Examples of this include the estimation of motor speed and flux in sensorless systems (Shi et al., 2002, 2012); and in robot localisation and movement (Huang and Dissanayake, 2007). Indeed, Bayesian estimation and the Kalman filter have been widely applicable in the field of robotics (Roumeliotis and Bekey, 2000; Chen, 2012). Robust solutions to sensor fusion and online learning are needed, and the Kalman filter offers this (Auger et al., 2013).

From the perspective of life sciences, such as ecology and epidemiology, spatio-

temporal modelling can range from tracking and predicting population dynamics and animal movement paths (Patterson et al., 2008; Molnár et al., 2010), to forecasting influenza epidemics (Cao et al., 2014) and tree mortality (Csilléry et al., 2013). In forecasting problems in weather systems, ensemble variants of the Kalman filters are often used (Hunt et al., 2007; Cassola and Burlando, 2012). In recent years, similar approaches have been applied to the inversion of electrical resistivity data, especially in time-lapse monitoring (Nenna et al., 2011; Camporese et al., 2015).

Target tracking has been predominant area of research with which Bayesian filters have been applied (Särkkä, 2013). Some tracking methods deal predominantly with explicit feature detectors (Yilmaz et al., 2006) and deterministically corresponding features between time-slices (Veenman et al., 2001; Shafique and Shah, 2005). Recent research has learnt tracking through recurrent neural networks (Milan et al., 2017). The alternative is the approach of statistical correspondence, whereby instead of creating some fixed pairing between features in adjacent time steps of a discrete-time system, some statistical metric, predominantly a probability distribution, is used to define correspondence over time. In particular, recursive Bayesian estimation, as described in Chapter 3, is used to make predictions with uncertainty over some conditional probability distribution (Gelman et al., 2004; Särkkä, 2013). In terms of theory, there are uses of the Kalman filter in controlling active contours for deformable object tracking (Rathi et al., 2007), and for irregular motion monitoring in 3-D video space using histogram of oriented gradients and the UKF (Tao and Klette, 2012).

For scenarios in which multiple objects must be tracked, there are a wide range of approaches (Luo et al., 2014). The general approach towards this is to represent each tracked object as a set of features, and associate them between time-steps. Examples include graph matching, that joins the set of features in each time-step as some bipartite graph. The cost minimisation function for optimising this graph can be some greedy method (Shu et al., 2012), or using the optimised Hungarian algorithm (Munkres, 1957; Xing et al., 2009; Reilly et al., 2010). Alternative graph-based methods include min-max methods for reducing cost and maximising flow (Butt and Collins, 2013; Wu et al., 2012), and through conditional random fields (Yang and Nevatia, 2012; Milan et al., 2013).

Applications of Bayesian inference towards target tracking include uses of the particle filters and variants of the Kalman filter (Kratz and Nishino, 2010; Ban et al., 2016). The particle filter, and indeed Monte Carlo-based methods, are based

on random sampling and resampling to approximate distributions (Gordon et al., 1993). Unlike the Kalman filter, they are not limited to linear models, and can approximate non-Gaussian, including multi-modal, distributions (Andrieu et al., 2003). In general, particle filters perform well on tracking problems, and with sufficient samples their representation power is high (Breitenstein et al., 2009; Liu et al., 2012; Hu et al., 2012). However, due to the large number of samples required, they suffer from high computational requirements, both in terms of memory and time. Where the Kalman filter's assumptions hold, or nearly hold, it is a near-optimal solution with linear time requirements (Särkkä, 2007).

Since the introduction of the linear Kalman filter (Kalman, 1960), there have been extensions to its capabilities for non-linear approaches. The extended Kalman filter, take both first- and second-order Jacobians to linearise the covariance transformation over non-linear functions (Einicke and White, 1999; Huang and Dissanayake, 2007). Alternative proposals use deterministic samples representing the distribution, including the unscented and cubature Kalman filters (Wan et al., 2000; Julier and Uhlmann, 2004; Särkkä, 2007; Arasaratnam and Haykin, 2009; Ahn and Chan, 2014). These approaches are powerful in their representation of the Gaussian distribution, and the implementation for such methods is relatively simple compared with the EKF, due to the absence of Jacobian calculations. The unscented Kalman filter also consistently outperforms the first-order EKF, and is parallel in its accuracy with the second-order EKF (Ahn and Chan, 2014).

More recent approaches have addressed the non-linearity of real-world processes by using kernel embedding of the state-space model and filtering systems in a reproducing kernel Hilbert space (Song et al., 2009; Lampert, 2015; Gebhardt et al., 2017). While this embedded approach is a powerful technique for predicting and filtering, it typically relies on learning the dynamic process, and has largely been applied to reinforcement learning and control problems. Other approaches to so-called deep state-space models include learning a parametric representation of the underlying system with recurrent neural networks (Krishnan et al., 2015; Karl et al., 2016; Eleftheriadis et al., 2017), or by using multi-layer feed-forward neural networks to parametrise high dimensional measurements to learn a lower dimensional observation model (Haarnoja et al., 2016).

Fuzzy logic and inference systems have been used for time series forecasting and predictive control. In contemporary research, interval type 2 fuzzy sets are predominantly used for a range of applications, including networked control systems (Lu et al., 2015), microgrid control design (Marín et al., 2016), and in regression

problems (Eyoh et al., 2017). Eyoh et al. (2016) also propose an approach for time-series prediction using intuitionistic interval type-2 fuzzy logic, and speculate on the possibility of using the Kalman filter to learn parameters introduced in their algorithm. Fuzzy logic with Kalman filters have been briefly explored in the past (Kobayashi et al., 1995; Chen et al., 1998; Matia et al., 2004). These have been derived under various applications, for example resolving vehicle speed (Kobayashi et al., 1995), or robot localisation (Matia et al., 2004). Each approach has been derived slightly differently, normally tailored to their particular application. For example, Longo et al. (2002) uses the Kalman filter to approximate the state, and uses fuzzy logic for sensor fusion to create an observation model. Matía et al. (2006) approach the problem by replacing the underlying probability assumptions of the state space model with fuzzy memberships, and use fuzzy logic to resolve the filtering problem. In all of these cases, the combinations do not necessarily exploit the features and similarities of probabilistic and fuzzy theories, but rather join up, or switch out, two distinct methods.

There have been various other attempts to both unify probability theory and fuzzy logic (Gan et al., 2005; Viertl, 2008) and to characterise each as a respective subset of the other (Zadeh, 1995). For the last fifteen years or so, there has been conciliation between the two fields and recognition that while they share similarities, probability is not fuzzy logic, and fuzzy logic is not probability (Gan et al., 2005). However, certain aspects of the two have been resolved, for example Gan et al. (2005) and Verma and Hanmandlu (2007) shows that fuzzy logic systems can be equivalent to Gaussian mixture models, demonstrating that a system resolved using fuzzy inference has an analogue in probability theory.

5.2 Data Processing in Near Surface Geophysics

Applying image processing and machine learning techniques is a relatively novel approach to the interpretation and analysis of ERT data. Intrusive investigation, especially drilling, is the most significant and common means of analysing the subsurface properties. For alternative methods, there are often limited baselines available for validation of methods. Approaches to the analysis and interpretation of ERT both in general, and focussed specifically targeting time-lapse imaging are presented in surveys by Loke et al. (2013) and by Singha et al. (2015) respectively. The most significant advances in the field of electrical resistivity tomography are in the inversion process.

The most common difficulty with approaches for analysis of ERT data is the

tendency of the images to exhibit smooth boundaries between differing regions (Loke et al., 2013). Prior research into techniques for boundary resolution in static imaging covers gradient methods, including finding the steepest vertical gradient (Chambers et al., 2012), a Laplacian-based approach (Hsu et al., 2010), and use of the watershed method for edge detection (Elwaseif and Slater, 2010, 2012). Artefacts of the smoothness constraints in the inversion regularisation, as well as true gradational changes in the underlying resistivity profile can mean these methods struggle. However, it was shown that the steepest gradient is typically a high performing estimation of boundaries where the structures are homogeneous with relatively sharp interfaces (Chambers et al., 2013). Mitigation of smoothness constraints on edge detection by incorporating detected boundaries into the inversion process has been attempted, with some success by Elwaseif and Slater (2012) and Zhou et al. (2014). Where subsurface properties exhibit heterogeneity, such as in complex wetland sites, methods for detection have been derived from identifying isosurfaces in the total ERT data (Chambers et al., 2012, 2014a). Studies have also investigated the use of clustering for these methods, including work developed in this project (Ward et al., 2014b), and detailed in Section 6.2. Clustering ERT data has also been used to learn inversion methods, though the approach is relatively simple in its scope (Audebert et al., 2014).

In time-lapse monitoring, electrical resistivity is predominantly used for imaging and modelling groundwater systems (Singha et al., 2015). For example, geophysical monitoring has been used to provide spatial information that complements point monitoring systems to image changes to groundwater conditions, such as in studies by Binley et al. (2002) and Kuras et al. (2009). ERT is particularly effective in resolving groundwater changes due the close relations between electrical resistivity and pore saturation (as discussed in Section 2.3.2). Specific developments in monitoring instrumentation (Ogilvy et al., 2009; Wilkinson et al., 2010; Chambers et al., 2014b; Supper et al., 2014) and time-lapse inversion methods (Hayley et al., 2011; Loke et al., 2014) have increased feasibility and quality of automated monitoring tools such as the automated time-lapse ERT (ALERT) (Kuras et al., 2009; Chambers et al., 2009) and the proactive infrastructure monitoring and evaluation (PRIME) (Chambers et al., 2016) systems. In particular, the PRIME system allows for long-term resistivity monitoring, such as of moisture dynamics of landslides sites (Uhlemann et al., 2017).

In translating time-lapse ERT data to quantitative measures of hydrology and hydrogeology, early methods sought to use direct mappings of resistivity to hydrological parameters directly, such as through empirical relations between laboratory

and field data (Purvance and Andricevic, 2000; Bowling et al., 2006; Linde et al., 2006). The simplicity of the approaches meant they are poor methods of metric conversion. Limitations of the ERT for monitoring, along with the problematic mappings, lead to poor recovery in of meaningful parameters (Singha et al., 2015). One proceeding branch of the state-of-the-art was the direct comparison of consecutive time-slices in resolved ERT images. This is largely done through basic frequentist statistics and temporal correlation analysed through classical temporal and spectral methods, most recently in the work of Johnson et al. (2012) and of Wallin et al. (2013), which both use cross-correlation methods on hydrologic changes. A similar approach to this was researched and developed for quarry dewatering, based on the clustering framework developed in this thesis. This research is briefly discussed in Section 6.3.2 and presented in Chambers et al. (2015b). Problems with methods such as these include the need for the need for both a baseline, and manual intervention for feature selection (Chambers et al., 2015b; Elwaseif and Slater, 2013).

In the case of contaminant transport studies, ERT has been demonstrated a successful tool in the resolution of saline leachate, and dense non-aqueous phase liquids (DNAPLS) (Chambers et al., 2006; Mansoor and Slater, 2007). Studies that directly sample tracer concentration approximate moments to quantify and model the underlying transport processes including advection, dispersion and mass transfer. In time-lapse ERT, these spatio-temporal properties can be approximated from features in ERT and used as surrogates for the tracer. Studies such as those by Singha and Gorelick (2005) and Haarder et al. (2012) approximate first-order moments and moisture content metrics through averaging. Studies of laboratory experiments and high-resolution ERT images have showed that velocities derived from ERT data are strongly correlated with the hydraulic conductivity of the underlying system (Vanderborght et al., 2005; Koestel et al., 2008). Indeed, time-lapse ERT can be used to resolve systems that form advection-diffusion models (Kemna et al., 2002; Müller et al., 2010). These studies are extended upon in this thesis and investigated using an automated feature detection and tracking framework to approximate the first- and second-order centralised moments in a tracer transport system. The full details of this investigation are given in Chapter 7, and published in Ward et al. (2016).

Beyond basic hydrodynamic properties of the subsurface, there is little application of ERT to modelling groundwater parameters (Singha et al., 2015). However, alternative studies on point measurements in transport systems provides insight into potential directions for better utilizing. The work by Xu and Gómez-Hernández

(2016) investigates estimating the initial parameters in a transport system using the ensemble version of the Kalman filter, to resolve an estimation of source location and concentration. The experiments are tested solely on synthetic datasets but draw from previous research that use Bayesian and stochastic procedures to model various properties of the system (Butera et al., 2013; Cupola et al., 2015). The ensemble approach utilises a variant that approximates non-Gaussian noise through a normal-score transform (Zhou et al., 2011). Prior experiments using the ensemble Kalman filter has been limited in scope, and does not address source localisation (Hendricks Franssen and Kinzelbach, 2009; Xu and Gómez-Hernández, 2015a). The lack of visual reference means that initialisation is difficult, requiring ensemble methods.

A table summarising the state of the art with respect to the contributions of this thesis is presented in Table 5.1.

Problem	**State of the Art**
Characterising subsurface features using ERT	Features of interest in both static and temporal ERT models include identifying subsurface boundaries, e.g. geological interfaces (Chambers et al., 2011, 2012, 2013), or tracer moments and structured plumes (Binley et al., 2002; Kuras et al., 2009)
Distinguishing regions in stationary ERT models	Gradient-based approaches are common for boundary detection in ERT (Chambers et al., 2012; Hsu et al., 2010; Elwaseif and Slater, 2010), but using known information for generating isosurfaces has also been used (Chambers et al., 2013); tracer moments are often detected with manual intervention, or taking statistics of the overall model (Chambers et al., 2006; Mansoor and Slater, 2007; Singha and Gorelick, 2005; Haarder et al., 2012)
Detecting change in time-lapse models	Time-lapse ERT advances have largely been limited to developments in instrumentation or inversion methods (Wilkinson et al., 2010; Supper et al., 2014; Hayley et al., 2011; Loke et al., 2014) or specific frameworks with global statistics output for analysis (Kuras et al., 2009; Chambers et al., 2016)
Combining mechanistic models with data	The application of state space models to physical systems has been applied in interpreting localisation in target tracking (Challa et al., 2011) and control problems (Huang and Dissanayake, 2007; Chen, 2012; Auger et al., 2013; Lu et al., 2015)
Handling complex prediction problems; Quantifying the uncertainty in the system and making use of this in predictions; Presenting the uncertainty in the system	Non-linear and multi-process systems have been solved with adaptations to the Kalman filter (Julier and Uhlmann, 2004; Särkkä, 2007, 2013; Ahn and Chan, 2014), as well as Monte Carlo approaches (Liu et al., 2012; Hu et al., 2012), and fuzzy logic control systems Lam and Seneviratne (2008); Lu et al. (2015). Combinations of these approaches, particularly in combining fuzzy and probabilistic representations of uncertainty are less common, although the combination of fuzzy logic and EKF is proposed in (Eyoh et al., 2016) for parameter estimation

Table 5.1: Summary of state of the art

6

Feature Detection in Electrical Resistivity Tomography

6.1 Introduction

The initial focus of the research project sought methods to locate and distinguish features of interest in electrical resistivity images. Initial studies sounded a variety of image processing methods for local feature detection, with the aim of localising notable regions in both static ERT images, and time-series data, for the purposes of change detection.

This research gave light to the fact that many of the assumptions held by local feature extraction algorithms, such as blob and corner detectors, did not hold in the smoothed ERT images. One group of methods, that defined features by their boundaries, is partially suitable to applications within near-surface geophysical imaging with ERT. Indeed, gradient based detectors have previously been used to identify local interfaces between physical media, such as between sand and clay, in static ERT images, to varying degrees of success.

Developments were made in localising features and different regions within ERT images. A density-guided fuzzy clustering method was developed and has proved highly effective in identifying region boundaries and quantifying their uncertainty in the smoothed environment. The boundary identification has been applied to feature localisation and change detection problems in real-world geophysical surveying. Despite the effectiveness of clustering, it is not suitable for identifying single features in time-series data, such as in tracer monitoring. Studies found that the computer vision technique of maximally stable extremal regions can capture

such features.

This chapter highlights the issues with feature detection in ERT, and describe the algorithms developed and adapted for use on ERT data, as well as some of the applications on which they have been used.

6.2 Distribution-Guided Fuzzy c-Means

6.2.1 Feature Detection for ERT

The geoelectrical inverse problem of ERT requires regularisation to be solved, most commonly in the form of a global smoothness constraint. Therefore, it is often difficult to isolate distinct boundaries in the resulting tomographic images. As a result of this, it has been found that in such images, gradient edge detectors do not perform well at identifying a good estimation of the geological interfaces. Alternative methods have included creating a resistivity isosurface based on known resistivity properties at the interface, but this often requires intrusive investigation of the study site.

The gradational transitions between regions in ERT images, as a result of both physical properties and smoothness constraints, means there is significant uncertainty in the specific boundaries obtainable by the images. An extension on fuzzy c-means clustering was developed, that exploits the density distribution of the resistivity in the image cell to pre-estimate resistivity populations corresponding to distinct regions. The populations are modelled with Gaussian kernels and are used to generate a distance metric for each population that is used in each clusters' membership function. The distribution-guided clustering method was developed as a reliable means for analysis of 2-D and 3-D ERT images generated with standard survey and inversion approaches, to delineate mineral volumes and thereby estimate yields (Ward et al., 2014b).

6.2.2 Kernel Density Estimation

Instead of using a pre-defined number of clusters to analyse, as is the case in k-means and standard fuzzy c-means clustering, the developed approach uses the density information of data to be clustered to approximate, and model, the number of distinct populations (Tran, 1990; Sheather and Jones, 1991). Using kernel density estimatation (KDE), it is possible to approximate the density function of the data, giving the probability that a randomly selected cell will have a certain resistivity. The density function of a resistivity image is approximated by the sum

of symmetric weighting functions, scaled logarithmically. The approximation of the pdf is given, as follows, where x_i represents the individual resistivity values of each cell in the ERT model. N is the total number of resistivity cells.

$$\hat{f}_\sigma(x) = \frac{1}{N}\sum_{i=1}^{N} \mathcal{LK}_{10}(x \,|\, x_i, \sigma) \tag{6.1}$$

Here, $\mathcal{LK}_{10}$ is some kernel, applied over the base-10 logarithm of the input data. The control parameter, σ is termed the *bandwidth*, and describes the spread of the kernel.

$$\mathcal{LK}_{10}(x \,|\, m, \sigma) = \mathcal{K}(\log_{10} x \,|\, \log_{10} m, \sigma) \tag{6.2}$$

A list of commonly used kernels for KDE are defined in Table 6.1. The Epanechnikov kernel, as defined, is optimal in reducing the mean-square error, but in practise, the Gaussian kernel is most commonly used. Use of the Gaussian kernel results in an approximated distribution function that is inherently the smoothed sum of Gaussians. The developed clustering algorithm is thus based on the principle that individual populations within the pdf can be approximated as isolated normal distributions that can serve as a guide for the assignment of data to groups, as well as improve uncertainty representations over the standard iterative allocation.

Table 6.1: Common kernel definitions for use with kernel density estimation

Kernel	$\mathcal{K}(x \,\|\, m, \sigma)$
Uniform	$\begin{cases} \frac{1}{2\sigma} & \lvert x-m\rvert \le \sigma \\ 0 & \text{otherwise} \end{cases}$
Triangular	$\begin{cases} \frac{1}{\sigma}\left(1+\frac{\lvert x-m\rvert}{\sigma}\right) & \lvert x-m\rvert \le \sigma \\ 0 & \text{otherwise} \end{cases}$
Gaussian	$\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-m)^2}{2\sigma^2}\right)$
Epanechnikov	$\begin{cases} \frac{3}{4\sigma}\left(1-\frac{(x-m)^2}{\sigma^2}\right) & (x-m)^2 \le \sigma^2 \\ 0 & \text{otherwise} \end{cases}$

Note that choice of bandwidth is dependent on the data used. Bandwidths may be selected manually, depending on the resolution desired, much in the same way bin size is selected in histogram analysis. There are also defined methods for approximating estimating the bandwidth, such as the *rule of thumb* method, which for a density function that is Gaussian, is defined $\sigma = ((4s^5)/(3N))^{0.2}$, where s is the sample standard deviation of the all data points the kernels are applied to (Sheather and Jones, 1991; Botev et al., 2010). However, where the distribution is not normally distributed, e.g. when the pdf is multi-modal, the resulting bandwidth can lead to the KDE approximating suffering from over-smoothing.

6.2.3 Distribution-Weighted Distance Metric

Given the estimate density function, $\hat{f}_\sigma$, the number of clusters can be defined as the number of distinct populations. Each modal point, $\bar{x}^j$ in the distribution, corresponds to the mean value of a population, in turn corresponding to the centroid of a cluster. Given c-modal distribution, $\bar{x}^j$ corresponds to the j^{th} local maxima in $\hat{f}_\sigma$. A Gaussian can be fitted by then estimating the standard deviation of each population distribution. To obtain this, the points of $\hat{f}$, where its second derivative is equal to 0, that are closest to $\bar{x}^j$ are taken as a *support point*, $\dot{x}^j$. The variance of the j^{th} population distribution, $(\sigma^j)^2$ is thus approximated by an inverse Gaussian centred on $\hat{x}^j$, applied to $\dot{x}^j$. This is formally defined:

$$\begin{aligned} \dot{x}^j &= \underset{\forall \dot{x}}{\operatorname{argmin}} \left(|\bar{x}^j - \dot{x}| \left| \frac{\mathrm{d}^2}{\mathrm{d}x^2} \hat{f} \right|_{x=\dot{x}} = 0,\ \hat{f}(\bar{x}^j) > \hat{f}(\dot{x}) \right) \\ \left(\sigma^j\right)^2 &\approx \frac{(\bar{x}^j - \dot{x}^j)^2}{2 \ln \left(\hat{f}(\bar{x}^j) / \hat{f}(\dot{x}^j) \right)} \end{aligned} \tag{6.3}$$

With the mean and variance of an individual cluster population, it can be approximated by some distribution function. This function, effectively modelling the probability a datum belongs to a cluster, given its resistivity value, can be incorporated into a distance metric used to allocate points to a cluster. For c fuzzy clusters, each with a population mean and variance, the fuzzy membership can be defined as in FCM, in (4.13), as follows, with $d^j(x)$ denoting the distribution weighted distance of a datum x to the j^{th} cluster centroid.

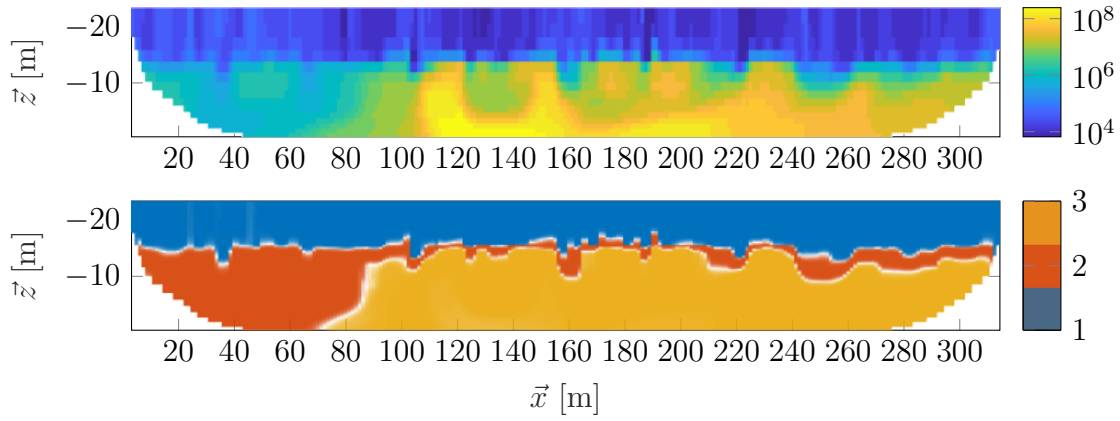


Figure 6.1: Static electrical resistivity model and corresponding clustering. Intensity of clusters corresponds to certainty of clustering, proportional to each resistivity cell's membership to its primary allocated cluster

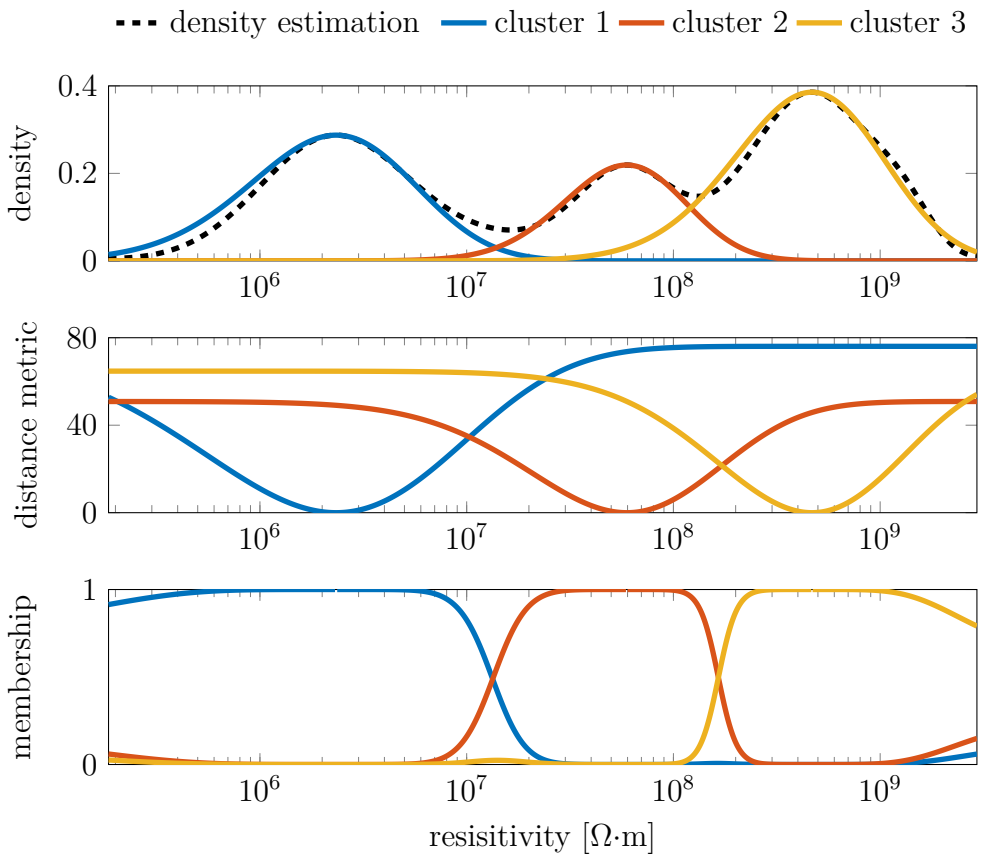


Figure 6.2: Density of resistivity in ERT model in Figure 6.1 calculated by KDE with bandwidth $\sigma = 0.25$. Fitted Gaussians are overlaid and used to calculated the corresponding GFCM distance metric and membership functions for each cluster

$$
\begin{aligned}
d^j(x) &= \left|x - \bar{x}^j\right| \cdot \left(1 - \exp\left(-\frac{\left(x - \bar{x}^j\right)^2}{2\left(\sigma^j\right)^2}\right)\right) \\
u^j(x) &= \left[\sum_{p=1}^{c}\left(\frac{d^j(x)}{d^p(x)}\right)^2\right]^{-1}
\end{aligned}
\tag{6.4}
$$

In Figure 6.1, an example ERT image is clustered using the density-guided fuzzy c-means (GFCM). The density of the resistivity distribution of the ERT cells is approximated by KDE in Figure 6.2. Three populations Gaussians are fitted to the density estimation, and these are used to form the respective distance metrics and membership functions for each cluster's fuzzy set. Evaluating the each membership function over the resistivity cells of the ERT image returns a fuzzy subset with membership of each cell to a cluster. The clustering can be defuzzified by allocating a cell to the cluster to which it has maximum membership. In Figure 6.1, the defuzzified *crisp* clustering is shown, but here the intensity of each cell is scaled by its membership to its allocated cluster.

A comparison of membership functions as generated by k-means, iterative fuzzy c-means and GFCM is shown in Figure 6.3, along with the corresponding clustering by each method. Both k-means and FCM were initialised with random centroids, and converged in 18 and 15 iterations respectively. In terms of runtime, k-means clustered the image in 0.02029 seconds, and FCM in 0.02560, which is a factor of 5-6 times that of GFCM's 0.00414 seconds. Indeed, the density-based information can produce a representative set of membership functions to cluster the data. In Figure 6.3, the membership functions are steeper than that of the iterated FCM, and the corresponding clustering, plotted with membership-scaled intensity, shows similar crisp clustering, with less quantified uncertainty about the bounds.

6.3 GFCM in Near-Surface Geophysics Applications

6.3.1 Subsurface Interface Detection

6.3.1.1 Study Site: Willington

One of the datasets in this study is taken from a site in the Great Ouse valley, near Willington, a village approximately 4 km easy of Bedford, UK. The Great Ouse is part of the Wash fluvial network, and preserves records of late Quaternary uplift, changes in climate, and Palaeolitic human activity. The geology of the site

composes river terrace sand and gravel, overlying a bedrock formation of Oxford Clay, a brownish grey mudstone. The formation crops out at both the southeast and northwest of the survey area, and the thickness is approximately 20 m, in part due to exposure by extractive activities in the valley.

The study site is situated on terrace deposits that overly the bedrock. The deposits have been part of sand and gravel operations in the area, and at the time of study, the topsoil was stripped and banked, exposing alluvium at the surface. A number of borehole logs are available to aid interpretation and calibration of the geophysical modelling. The deposits were unsaturated at the time of study, due to unrelated dewatering in mineral works immediately south of the survey site.

For ERT data collection and inversion, to generate the image data, the survey was carried out over an area of 93 m × 93 m, using 16 survey lines that were

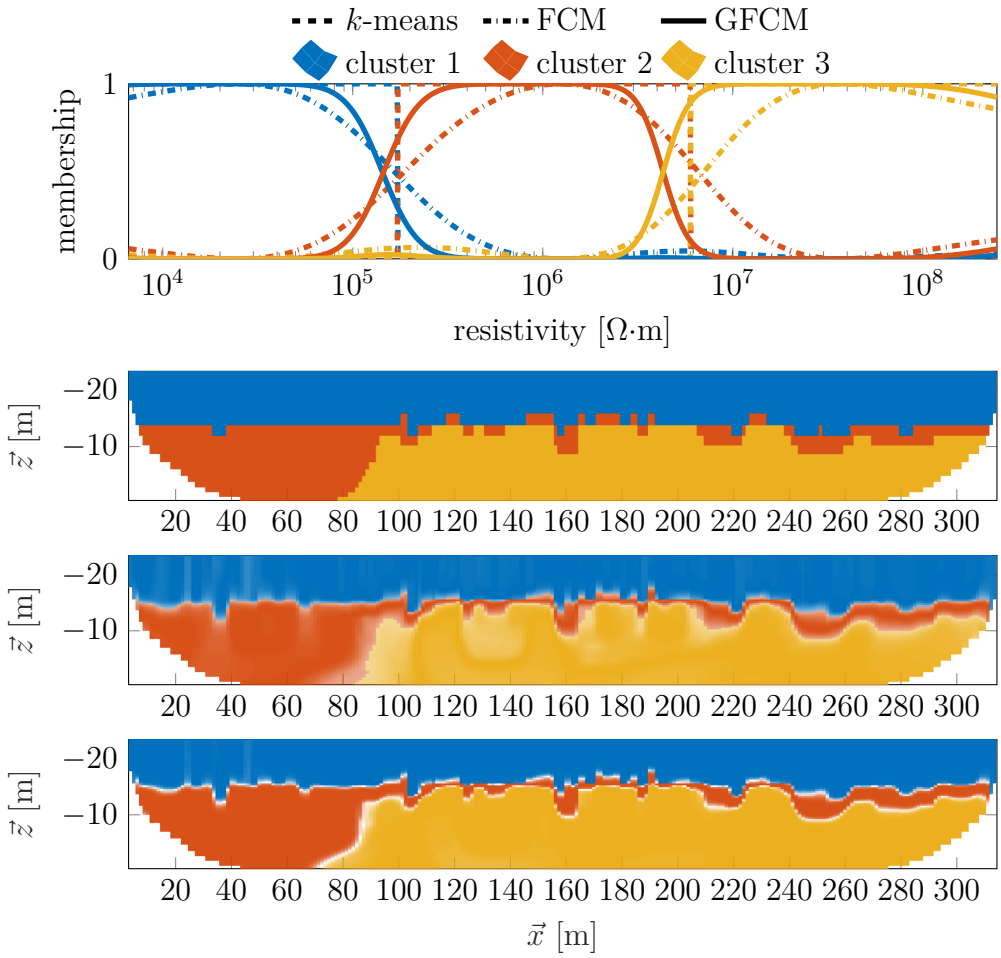


Figure 6.3: Membership functions (top) and clustering by k-means, iterative fuzzy c-means and distribution-guided fuzzy c-means (ordered respectively) on the ERT image in Figure 6.1

Algorithm 4 Distribution Guided Fuzzy Clustering Method

procedure INITIALISE($[x_i]_{i=1}^n; \mathbf{x}, \sigma$) ▷ $\mathbf{x}$ is KDE mesh, σ is the bandwidth
  $\hat{f}_\sigma(\mathbf{x}) \leftarrow \frac{1}{n}\sum_{i=1}^n \mathcal{LK}_{10}(\mathbf{x}|x_i, \sigma)$
  $c \leftarrow$ number of local maxima (peaks) in $\hat{f}(\mathbf{x})$
  for $j \in 1, \ldots, c$ **do**
    $\bar{x}^j \leftarrow j^{\text{th}}$ peak of $\hat{f}(\mathbf{x})$
    $\dot{x}^j \leftarrow \text{argmin}_{\forall \dot{x} \in \mathbf{x}} \left(|\bar{x}^j - \dot{x}| \left| \frac{\mathrm{d}^2 \hat{f}}{\mathrm{d}x^2} \right|_{\dot{x}} = 0, \hat{f}(\bar{x}^j) > \hat{f}(\dot{x}) \right)$
    $(\sigma^j)^2 \leftarrow (\bar{x}^j - \dot{x}^j)^2/2 \left(\ln(\hat{f}(\bar{x}^j)) - \ln(\hat{f}(\dot{x}^j)) \right)$
  end for
  return $[\bar{x}^j, (\sigma^j)^2]_{j=1}^c$
end procedure
procedure ASSIGN($[x_i]_{i=1}^n, [\bar{x}^j, (\sigma^j)^2]_{j=1}^c$)
  for $x_i \in \mathbf{x}$ **do**
    for $j \in 1, \ldots, c$ **do**
      $d^j(x_i) = |\bar{x}^j - x_i| \left(1 - \exp\left(-(x - \bar{x}^j)^2/2(\sigma^j)^2\right)\right)$
    end for
    for $j \in 1, \ldots, c$ **do**
      $U_i^j \leftarrow \left(\sum_{k=1}^c \left(\frac{d^j(x)}{d^k(x)} \right)^2 \right)^{-1}$
    end for
  end for
  return U
end procedure
procedure DGFCM($[x_i]_{i=1}^n; \mathbf{x}, \sigma$) ▷ x_i is data, $\mathbf{x}, \sigma$ are KDE parameters
  $[\bar{x}^j, (\sigma^j)^2]_{j=1}^c \leftarrow$INITIALISE($[x_i]_{i=1}^n, \mathbf{x}, \sigma$)
  $U \leftarrow$ASSIGN($\mathbf{x}, [\bar{x}^j, (\sigma^j)^2]_{j=1}^c$)
  return $U, [\bar{x}^j, (\sigma^j)^2]_{j=1}^c$
end procedure

positioned at intervals 6 m apart, along both the x- and y-directions. A dipole-dipole configuration, with dipole lengths $a = 3$ m and 6 m and seperations of $1a$ to $8a$. Both normal and reciprocal measurement sets were collected and the line, with orthogonal survey lines set up to minimise single-line bias. The dipole-dipole arrays produce relatively high level of resolution, with the difference between normal and reciprocal measures used to assess and correct random and systematic errors in the survey. The field data was inverted using a 3-D regularised least-squares optimisation, and the forward problem was solved with finite-element method. The resulting ERT models was produced using an L2-norm constraint, and is made up of 10,571 resistivity model cells, with dimension $31 \times 31 \times 11$ ($x \times y \times z$). The resistivity image can be seen in Figure 6.4.

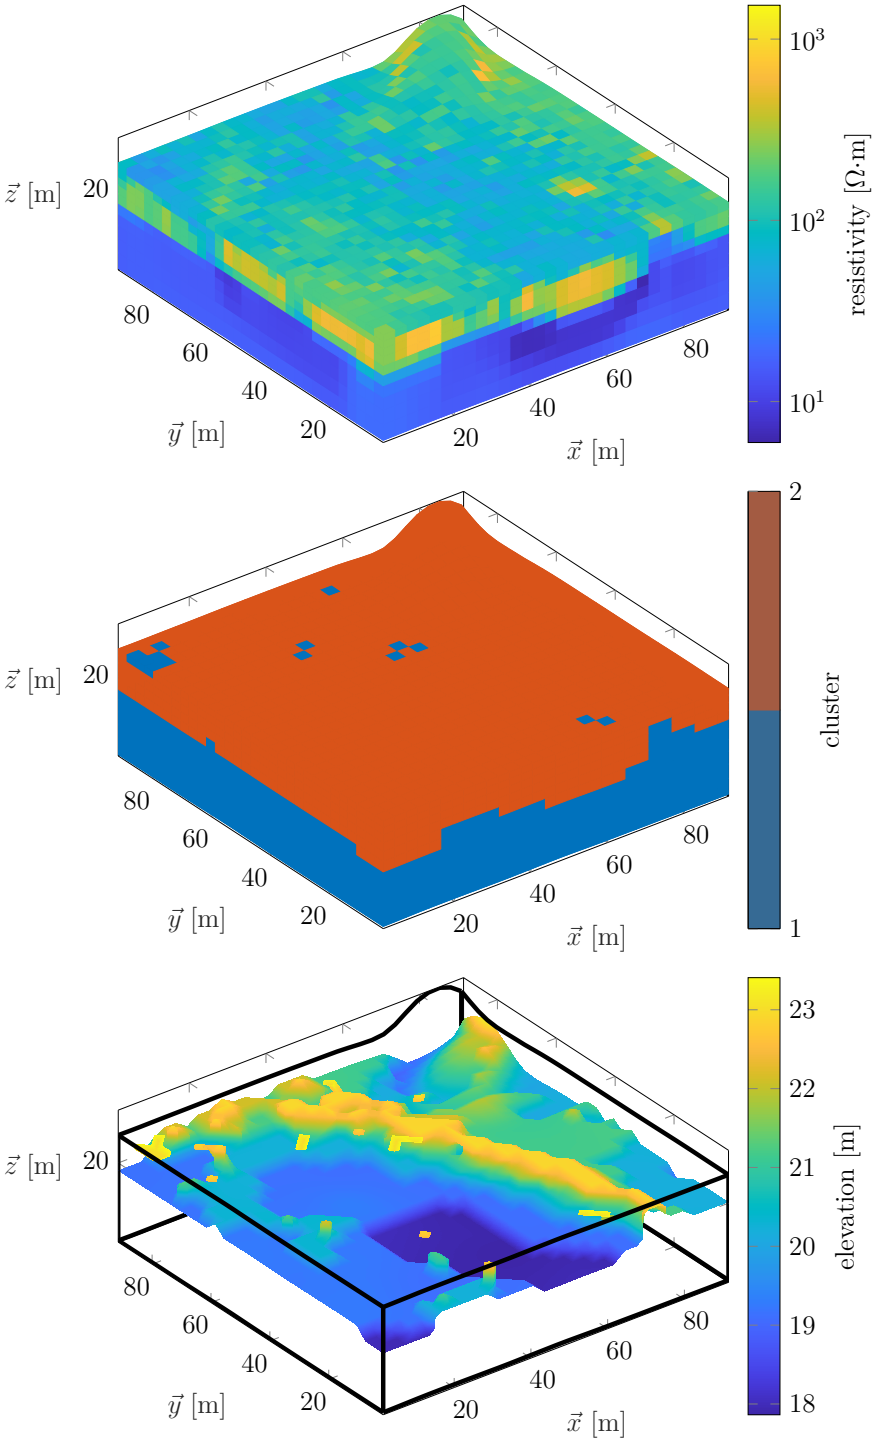


Figure 6.4: ERT model of the Willington site (top) with 2-clustering obtained by GFCM, with bandwidth 0.25. The interface between the two clusters is taken to be the bedrock surface, and an elevation map is shown (bottom)

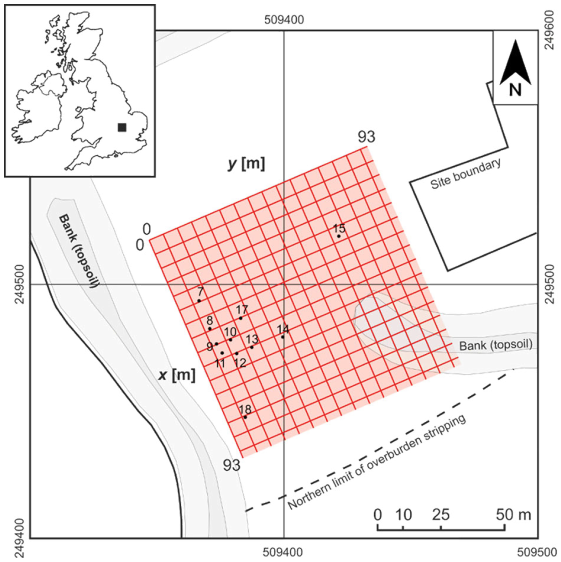


Figure 6.5: 3-D ERT survey design for Willington site, showing survey areas (red shading), lines (red lines) and borehole positions (black dots)

6.3.1.2 Bedrock Detection

The GFCM clustering method was applied to the Willington site. The bandwidth parameter, controlling smoothness of the density estimation, was selected as $\sigma = 0.25$, which produced two distinct populations in the resistivity model. The interface between the main bodies of like-resistivity represented by the clusters is considered to approximate the bedrock surface. In Figure 6.4, models of the clusters and interface between them are plotted. The elevation surface shows the topology of the bedrock surface, and presents a visual representation for use in surveying.

In terms of the effectiveness of the clustering's estimation, the elevation of the bedrock at a series of boreholes is compared with that of the approximated surface. In Figure 6.5, the borehole locations are identified in the survey design, and the respective measured and approximated bedrock elevation is given in Table 6.2 for each borehole in the survey area.

6.3.2 Groundwater Drawdown and Rebound Monitoring

6.3.2.1 Study Site: Rock Common

The ERT survey in this study was developed with the aim to monitor groundwater dynamics in an active quarry, and to assess the quality of clustering with automated

Table 6.2: Detected bedrock elevation by GFCM at boreholes in Willington survey area, compared with ground truth. Error is given as the absolute difference between actual and approximated bedrock depths, in metres

Borehole	7	8	9	10	11	12	13	14	15	17	18
$\vec{x}$	29.5	41.2	49.3	48.5	51.9	54.4	54.5	55.6	27.8	42.2	78.8
$\vec{y}$	8.6	8.1	8.4	13.9	8.8	14.0	20.4	33.2	69.2	21.0	7.2
Bedrock	21.02	20.76	19.95	19.93	19.83	19.68	20.77	20.73	20.86	19.87	21.76
GFCM	20.31	18.66	18.68	18.67	18.68	19.65	19.79	20.00	19.65	18.66	20.80
error	0.709	2.096	1.273	1.264	1.154	0.033	0.985	0.731	1.215	1.209	0.960

time-lapse ERT. The quarry site is under the effect of well-point dewatering, whereby groundwater is pumped from the quarry floor. The ERT survey was performed to detect a controlled lapse in the quarry pumping, to measure subsurface water levels. The clustering was applied to provide detailed information on the groundwater rebound and the drawdown that followed the pump being switched back on. The resistivity measurements and clustering information were compared to direct observations of the groundwater levels at certain times, by use of monitoring boreholes with piezometers.

The quarry at Rock Common is approximately 11 km north of Worthing, in West Sussex. The activity in the quarry has been ongoing since the 1920s, and the quarry depth, at the time of the ERT survey, was approximately 40 m at its deepest point. The setting of the survey is given in Figure 6.6. The dewatering of the

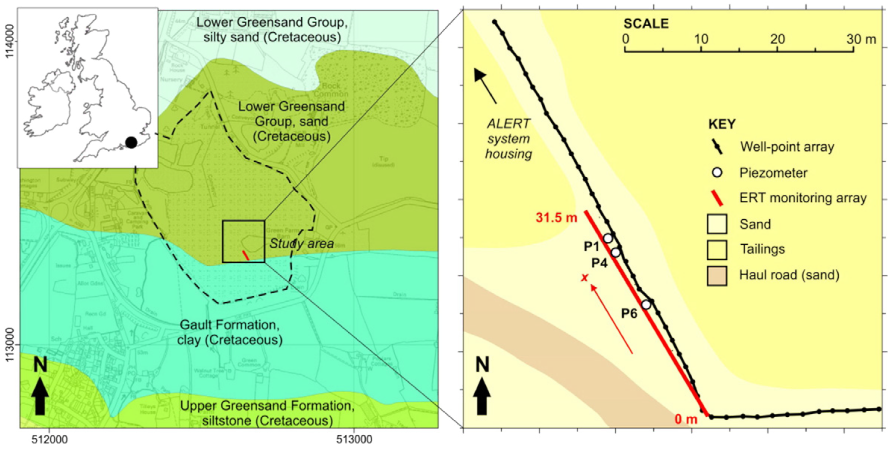


Figure 6.6: Setting of the Rock Common survey, with geological map with Ordnance Survey base layer (© Crown copyright and database rights 2015 Ordnance Survey 100021290), showing quarry boundary (dashed black line) and the study site location, with site location at national scale inset (left). Schematic field layout showing monitoring array, piezometers and well-point system (right)

quarry is carried out using well-point arrays connected to pumps installed across the quarry floor. The well-point system is a linked array of shallow boreholes from which the groundwater is pumped, positioned at 2 m intervals, and extending to a depth of 5 m below ground level. The well-point array within the quarry floor is shown in Figure 6.6. Under normal conditions, the water level is maintained to between 0.5 m and 2 m below the lowest point of the quarry floor, by continuous pumping. The hydrological regime is dynamic, and subject to flash flooding during heavy rain, especially during winter. The groundwater levels are also sensitive to the effects of pumping, and changes made to it: reduction in pumping has led to rebound in the order of tens of centimetres within a few hours.

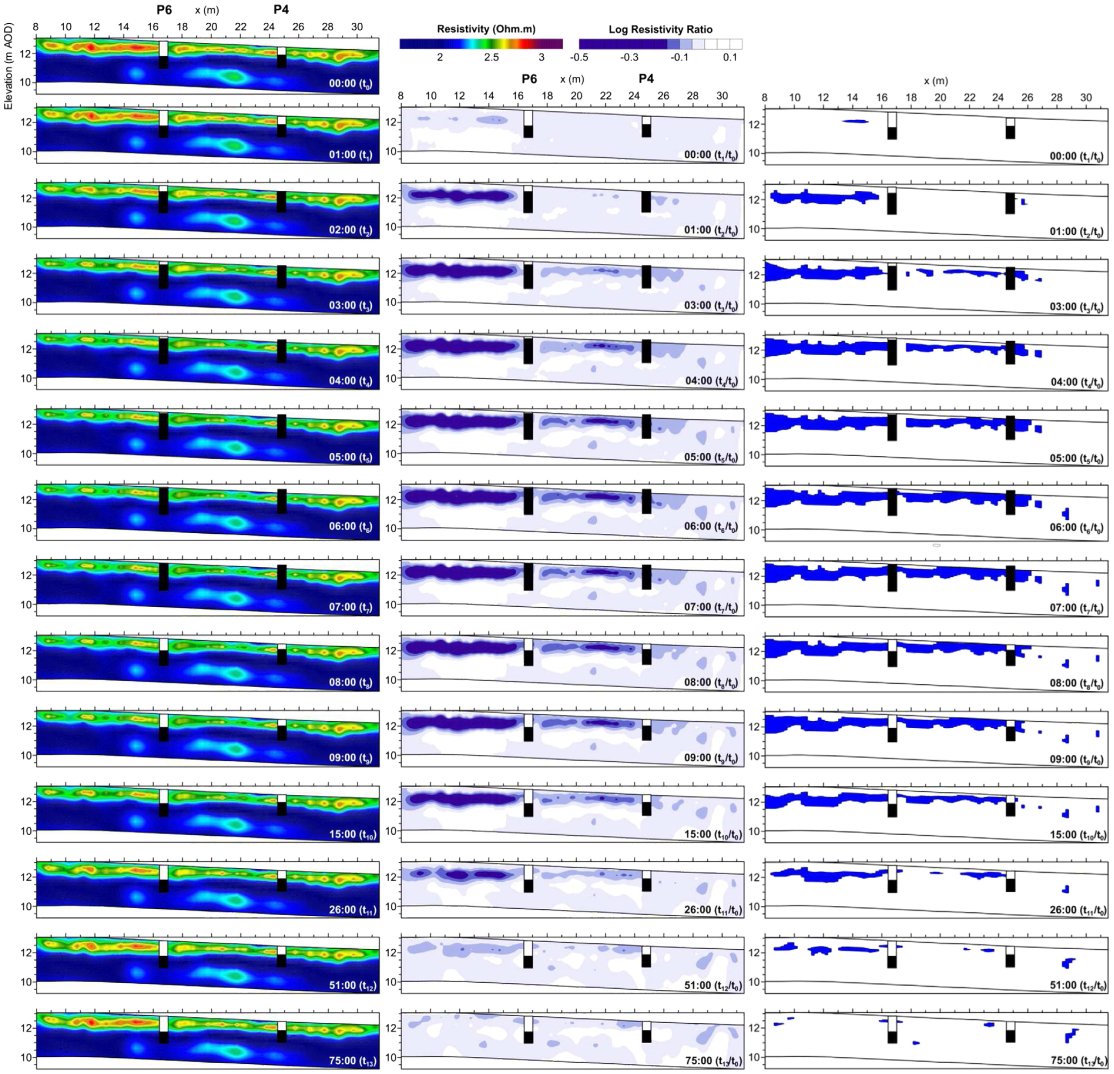


Figure 6.7: Resistivity monitoring for pumping interruption between t_0 and t_13 (left), and log resistivity ratio plots showing change relative to base line t_0, and corresponding principle change cluster (right). Water level records for peizometer 4 and 6 are shown, and time is given in hours and minutes (hh:mm)

The ERT survey line was installed along the well-point line in the deepest part of the quarry. The location was partly selected such that the survey would not be disturbed by ongoing quarry activity, but the proximity to the well-point extraction allowed for observation of groundwater drawdown and rebound that occur as a direct result of changes in the pumping regime. The monitoring was taken routinely over 12 months, with ERT measurements taken at a frequency of once per day. However, the clustering was applied during a period of increased frequency, in a controlled scenario. The pumping was switched off for a period of approximately six hours, during the month of May.

6.3.2.2 Results

During the pump switch-off, the groundwater levels rebounded by up to −0.9 m. This was measured by the piezometers, and indicated in the ERT data. In Figure 6.7, the time-lapse resistivity, and log resistivity ratio sections shown the increasing moisture content, in the form of decreasing resistivity, which corresponds with the water levels corresponding to piezometers 4 and 6. After the pumps are switch on a approximately t_7, the decrease in water levels is followed by a gradual decrease in resistivity. The cluster analysis of the data is taken as the major cluster corresponding to the largest decrease in log resistivity, and this principle change cluster used as a quantifier for saturated model cells at each time step.

The number of cells in the principle change cluster plotted against time is shown in

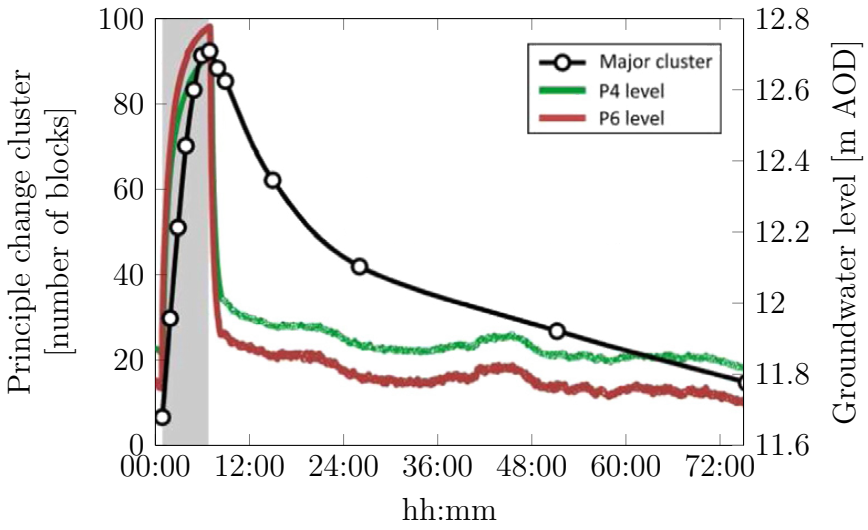


Figure 6.8: Principle change cluster area versus the water level records for piezometers 4 and 6 during the pumping experiment. The shaded region indicates the period during which pumping was reduced

Figure 6.8, overlaid by the groundwater level as obtained by the two piezometers. The general patterns of moisture change between the ERT and corresponding cluster quantification is largely similar to the groundwater levels measured by the peizometers at boreholes 4 and 6. The resistivity images exhibit a lag in both the rebound and drawdown phase. The discrepancies between the measured peizometers and the clustered information can be explained by the lag between the pressure response, as observed in the boreholes, compared with the slower changes in sediment saturation as measured in the resistivity images.

The GFCM method provides a means for automatically detecting and quantifying spatial changes within ERT data, which in turn can effectively image changes in moisture and groundwater content in the near subsurface. In this example, GFCM successfully isolated and identified the regions in which the water levels rise, as validated by the reference boreholes. There was some user intervention in the selection of the principle change cluster, but in general, the technique provides a near-automatic image analysis tool.

6.3.3 Discussion

The results in Table 6.2 and in Figure 6.4 show that GFCM is an effective means for approximating interfaces in near-surface geology, and that the method can serve as an aid for visual interpretation in deposit surveying. Further, the quantitative estimation of bedrock depths show that the GCFM can provide a reasonable approximation of localisation. Further, there is little user intervention required for analysis of the data, presenting the GFCM as a quick analysis tool for preliminary investigation of a site. Limitations fall largely along the lines of those in both resistivity imaging and edge detections on then. Generating ERT models involves an exponential decrease in the resolution of the data as a proportion to the distance from the measurement electrodes. In addition, interfaces with a gradational nature are produced by the smoothness-constrained inversion, and traditional edge detectors can struggle here. However, general studies demonstrate that the GFCM is indeed a powerful tool for detecting region boundares in ERT data, with a quantifiable measure of certainty.

In terms of change detection in time-lapse monitoring, the distribution-guided clustering allowed for effective quantification. The distance metrics were applied over the entire time-series, and populations were found independent of time constraints. This allowed for clusters to be approximated, and the area principle change to be localised, in the context of the entire time-series. The results, as shown in Figure 6.8 serve to demonstrate the effectiveness of both ERT and the

proposed GFCM method for feature analysis, with real-world capabilities.

The downsides to the clustering method, especially in the time-lapse method, is that the clustering method labels all voxels to a particular category, and cannot automatically give any information of each respective cluster. For time-lapse data, the issue in identifying a feature based on user intervention may be undesirable, especially in automatic processing and monitoring.

6.4 Maximally Stable Volumes in ERT

6.4.1 Algorithm

The maximally stable volumes (MSV) algorithm is a generalisation of the maximally stable extremal regions (MSER) method, used to find connected regions that are considered "stable" (Matas et al., 2002). An input volume is binarised at multiple successive thresholds and a tree data structure is created using the spatially connected components. The MSER algorithm 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 the automatic selection of ROIs, and has been found to be especially robust in the presence of noise (Yu et al., 2013).

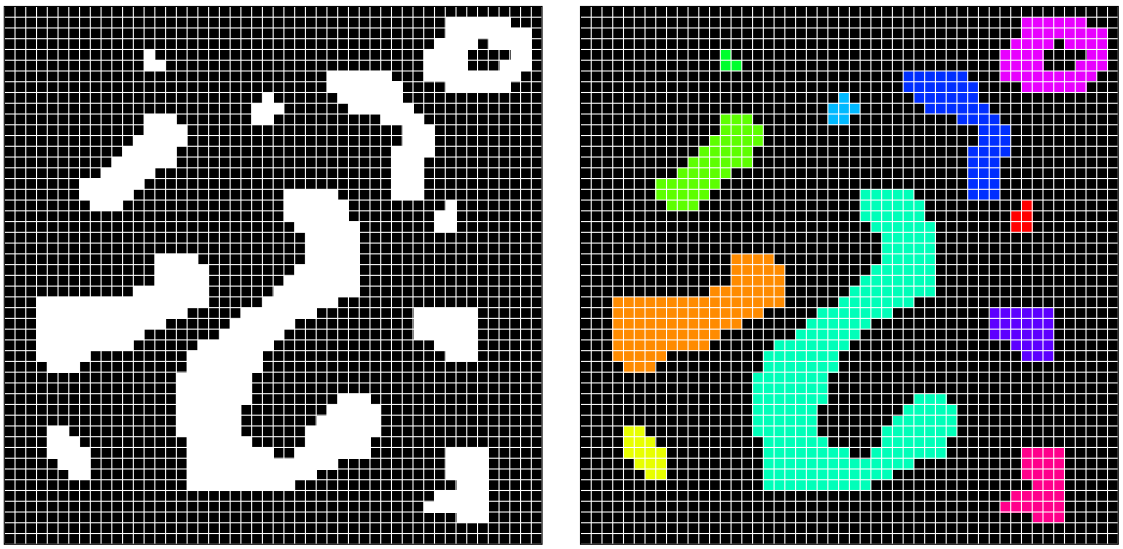


Figure 6.9: A binary image with indendent connected components individually coloured (right)

Algorithm 5 Maximally Stable Volumes

procedure MSV($\mathbf{V}; \Delta_{1:n}$) ▷ $\mathbf{V}$ is volume over space X
  $C_0 \leftarrow \{X\}$ ▷ or $\{\emptyset\}$ if $\Delta_{1:n}$ is decending
  root$(T) \leftarrow (0, \mathbf{C}_0)$ ▷ T is a tree with depth $n+1$
  for $i \in 1, \ldots, n$ **do**
    for all $\mathbf{x} \in X$ **do**
      $\mathbf{B}(\mathbf{x}) \leftarrow (\mathbf{V}(\mathbf{s}) \geq \Delta_i)$
    end for
    $C_i \leftarrow$ set of connected components in $\mathbf{B}$
    for all $\mathbf{c} \in C_{i-1}$ **do**
      children$(T, (i-1, \mathbf{c})) \leftarrow \{(i, \mathbf{c}') \,|\, \mathbf{c}' \in C_i, \mathbf{c}' \subseteq \mathbf{c}\}$
    end for
    for all $(i, \mathbf{c}) \in T$ **do**
      if $i = 0$ **or** children$(T, (i, \mathbf{c})) = \emptyset$ **then**
        $\Psi(i, \mathbf{c}) \leftarrow 0$
      else ▷ Note $|\cdot|$ here represents cardinality
        $M_{\text{parent}} \leftarrow |\text{parent}(T, (i, \mathbf{c}))|$
        $M_{\text{child}} \leftarrow \max\{|\mathbf{c}'| \,|\, (_, \mathbf{c}') \in \text{children}(T, (i, \mathbf{c}))\}$
        $\Psi(i, \mathbf{c}) \leftarrow |\mathbf{c}|^{-1}(M_{\text{parent}} - M_{\text{child}})$
      end if
    end for
  end for
  return $\{\mathbf{c} \,|\, (i, \mathbf{c}) \in T, \Psi(i, \mathbf{c}) \geq \Psi(\text{parent}(T, (i, \mathbf{c})))$
    and $\Psi(i, \mathbf{c} \geq \Psi(i+1\mathbf{c}'), \forall (i+1, \mathbf{c}') \in \text{children}(T, (i, \mathbf{c}))\}$
end procedure

The algorithm of MSV is set up such that a set of thresholds, $\mathbf{\Delta}$, is specified, that spans the range of unique values in the image data, discretised over some interval. Each threshold is used to binarise the input image, based on whether the value is below or above the threshold. An individual connected component is identified as a group of pixels in the binary image that are connected to each other, but spatially disconnected from all other connected components. An example of connected components in a binary image is shown in Figure 6.9, with the binary image on the left and connected components individually coloured on the right. The algorithm for calculating connected components simply involves finding an unlabelled voxel with value 1 and using the flood-fill algorithm to label all the voxels in the connected region containing that voxel. The process is repeated for every remaining unlabelled 1-value voxels in the model. In calculating connectivity, a voxel is adjacent both along principal coordinate axes, and diagonally, i.e. 8-connected for 2-D images, and 26-connected for 3-D.

To construct a connected component tree, each component forms a tree node at its threshold level. The node is a child of the component in which it is strictly

contained at the previous threshold level. By the nature of the thresholding and connectivity criterion, there are no overlapping regions. 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 at threshold Δ_i, $(i, \mathbf{c})$ has a corresponding stability, $\Psi(i, \mathbf{c})$. The root node (with no parent) and leaf nodes (with no children) have 0 stability. The measure for stability for all other nodes in the tree T is defined:

$$\begin{aligned} \Psi(i, \mathbf{c}) &= |\mathbf{c}|^{-1} \left(|\mathbf{c}^+| - \max_{\forall c^-} \left\{ |\mathbf{c}^-| \right\} \right) \\ \text{where} & \\ (_, \mathbf{c}^+) &\leftarrow \operatorname*{parent}_T(i, \mathbf{c}) \\ (_, \mathbf{c}^-) &\in \operatorname*{children}_T(i, \mathbf{c}) \end{aligned} \tag{6.5}$$

Traversing the tree depth-first to calculate local maxima in the stability returns the set of nodes that correspond to *maximally stable volumes*. Some branches in the tree will have monotonically decreasing stability, and will result in no MSVs, while others may have multiple maxima and result in nested output MSVs. Post-processing constraints can be included to remove MSVs that are either too large or too small, proportional to the input image. Stability can also be calculated based on the sizes of connected components some fixed interval (≥ 1) threshold levels above or below the current component.

6.4.2 Feature Extraction

Figure 6.10 shows a comparison of distribution-guided fuzzy c-means algorithm with the MSV applied to a 2-D ERT image from a tracer monitoring image. 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, whereby every point is given an arbitrary label, including the "background" regions.

MSVs are well suited for the purpose of detecting tracers in ERT images, due to both their successes at identifying distinct regions in a ackground of noise, and also the presence of nested volumes. In Figure 6.10, the nested volumes mark steps in stability in the gradual shift in resistivity between the tracer plume and the background media. The boundaries of these nested volumes give discrete contours describing the contain plume. Detected volumes are useful to extract

local imformation in an ERT time series. The information can be used for tracking purposes, and analysed to form distinct measurements.

6.5 Discussion

This chapter has introduced and summarised feature detectors for use in ERT imaging for different types of features. Applying clustering solutions to static ERT models has shown that this is an effective approach for identifying regions of distinct resistivity, that correspond with physical features such as bedrock and soil

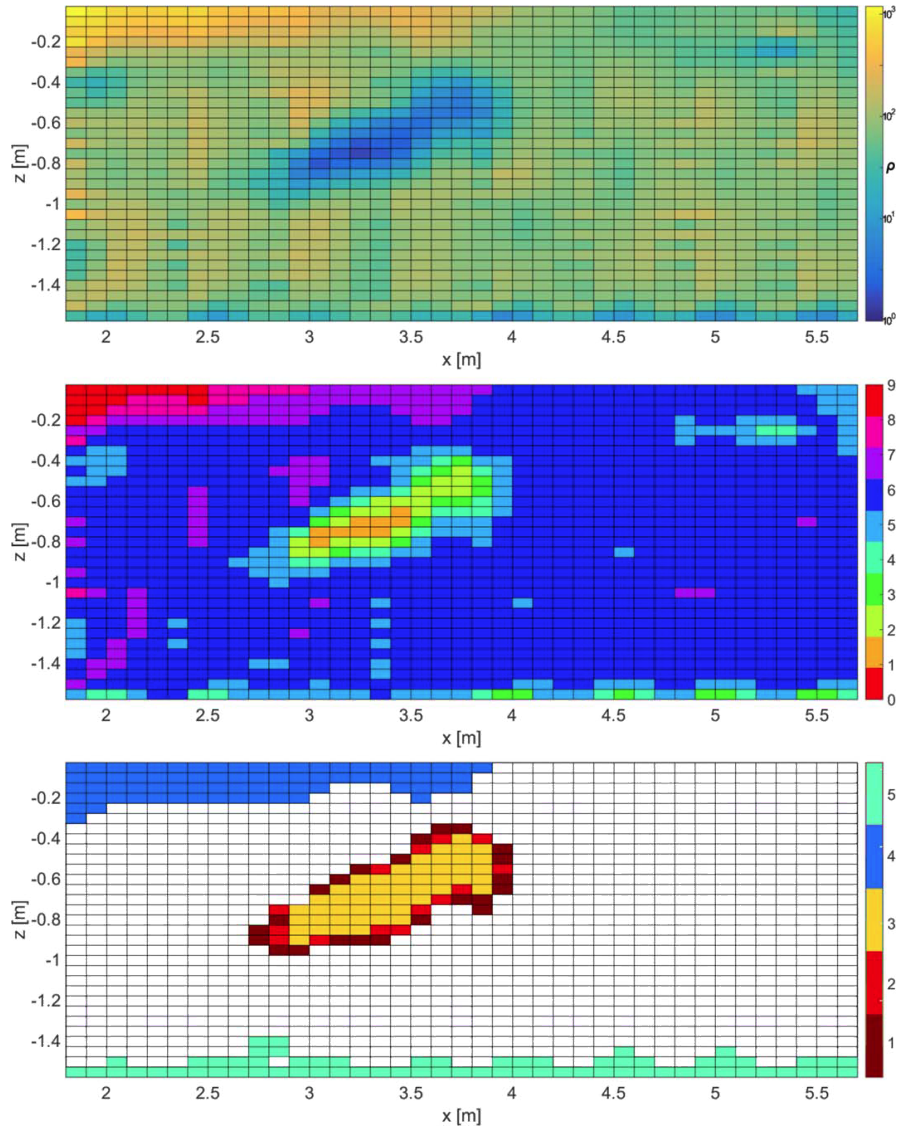


Figure 6.10: 2-D resistivity model with feature segmentation results from distribution-guided fuzzy clustering (middle) and maximally stable volumes (bottom). Individually coloured regions represent distinctly seperate detected regions, which, in the MSV case, represent nesting

deposits. Furthermore, cluster boundaries can be used to localise interfaces in the subsurface, giving complete spatial metrics such as bedrock depth. The method was applied to case studies at real-world sites and was shown to be effective comparing to known results. The development of the distribution-guided fuzzy clustering method was demonstrated as being effective and efficient for the task at hand, and the resulting boundary uncertainty is tighter than the basic fuzzy c-means, as it relies on density information rather than random initialisation. Additionally, there is no requirement to set the number of clusters, as this parameter is inferred from the density information.

Extending the clustering approach to a time-lapse ERT model, namely the quarry dewatering experiment, highlighted both the benefits and pitfalls of such an approach to spatio-temporal analysis. While the clustering approach was successful at identifying the regions of high moisture content indicative of the groundwater rebound, that was consistent with physical interpretation and direct measurement, it required manual intervention to select the corresponding cluster representing the high moisture change. This motivated using a different approach to feature detection that would distinguish specific features-of-interest from the background resistivity model. The maximally stable volumes method was adapted for and applied to ERT data, highlighting its capabilities versus clustering, where the background is not considered a *feature*. A summary of the contributions in this chapter is given in Table 6.3

Problem	**Contribution**
Characterising subsurface features using ERT	Clustering methods were applied to identify distinct regions and determine boundaries in subsurface geology, as well as detection of change in moisture content in the quarry dewatering experiment
Distinguishing regions in stationary ERT models	The developed clustering method was effective in identifying a bedrock surface based on resistivty measurements that was consistent with borehole logs. Similarly MSV was shown to identify features in static instances of a time-lapse ERT experiment
Detecting change in time-lapse models	The quarry dewatering experiment made use of clustering resistivity change as a result of increasing moisture content that was consistent with physical and observed models
Quantifying the uncertainty in the system and making use of this in predictions	The GFCM method used density information to initialise clusters and used a fuzzy interpretation of membership based on distance weighted by population variance
Presenting the uncertainty in the system	The uncertainty available in boundary identification presented by the GFCM can be used to inform on the reliability of the estimate, with greater reliance on the data than randomly-initialised FCM

Table 6.3: Contributions in methods for feature detection in ERT models

7

Linear Filtering for Tracer Moment Monitoring

7.1 Problem Definition

Modelling physical processes in hydrodynamic systems is a widely studied field, with complex analysis and models used to approximate the underlying hydrogeophysical and hydrogeological properties of an area. Analysis and modelling can occur at a micro, porous level, to large scale systems describing the groundwater movement and distribution over whole regions. Analysing areas with ERT gives a reasonable spatial resolution in the order of centimetres to metres. The resistivity properties of saturated media, and ionic concentrates dissolved within, means that time-lapse ERT can provide an excellent tool for monitoring hydrogeophysical processes.

While coupling measured information with hydrodynamic processes has been a widely studied area, there is limited interpretation and analysis of ERT using automated computational tools (see Chapter 6), and fewer using this analysis as measurements to model subsurface processes. Given a system, such as a contaminant plume advecting and dispersing along some subsurface flow field, automatic ERT analysis can be used as an observation of its motion. The observations can be used to correct and interpret system process models, for estimation of properties and parameters in the observed system.

7.1.1 Plume Tracking

When monitoring tracers, it is possible to visualise and distinguish the regions of contrasting resistivity in ERT imaging. Using a method, such as maximally

stable volumes (MSV), such regions can be automatically identified in 2-D and 3-D ERT images, as discussed in in Section 6.4. A tracer tracking framework has been developed incorporating MSVs as observations for a set of Kalman filters. 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 plume, but also to stationary regions, and random or systematic artefacts in the model. Additionally, it is likely that nested volumes occur, given the gradational nature of boundaries in ERT, and the diffusion of the tracer. Using a multi-track approach avoids the need to identify a single prominent region in each time-slice. Rather, minimising a cost function between predicted observations, $\boldsymbol{y}_{t|t-1}$, in each track and the measured observations allows the assignment of the detected regions to the independent Kalman filters.

The set of measurements are formed of the MSVs in the ERT image at time t, $\tilde{\boldsymbol{y}}_t = \mathrm{MSV}[V_\rho(t, \vec{\mathbf{x}})]$. The cost function used to associate predictions with observations is based on a variant of the Hungarian assignment method (Munkres, 1957; Reilly et al., 2010). For some given matrix of distances between all combinations of predictions and observations, the method will find the optimal combination of assignments by minimising the sum of distances. The distance measure for a given measurement and a Kalman filter track's predicted observation is defined

$$d(\boldsymbol{y}_t, \tilde{\boldsymbol{y}}_t) = (\boldsymbol{y}_t - \tilde{\boldsymbol{y}}_t)^{\mathrm{T}} \boldsymbol{P}_t^{(y)} (\boldsymbol{y}_t - \tilde{\boldsymbol{y}}_t) \tag{7.1}$$

Costs for non-assignment for both tracks and observed measurements are also incorporated into the assignment algorithm. If a predicted or measured observation has no associated distance lower than the non-assignment cost, it will not be assigned to a measurement or Kalman filter track respectively. Where there is no measurement corresponding to Kalman filter track, the prediction is propagated to the next time-slice with no correction step. Where no Kalman filter is associated with a measurement, that measurement is used to initialise a new Kalman filter, thereby starting a new track. If a track has no measurements assigned for some predetermined number of consecutive time-steps, the region it is tracking is

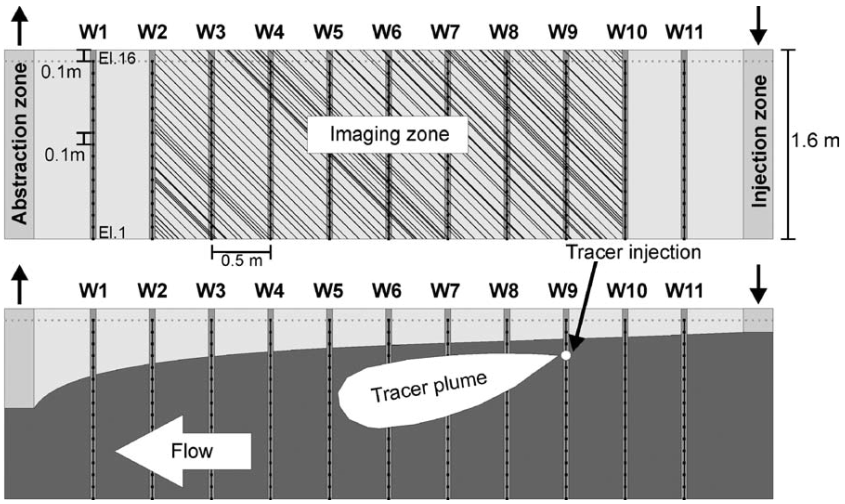


Figure 7.1: Cross-sectional view of 2-D pit experiment setup, and demonstration of flow field and tracer injection

considered *lost*. Lost tracks will consist of detected regions that were either simple noisy artefacts in the detection, or tracked regions that have left the imaging volume. Once a track is lost, it will no longer be updated or used in the assignment function.

The basic outline of the TTF algorithm for a time-lapse ERT model, $V_\rho : \{1, \ldots, T\} \times \mathbb{R}^d \to \mathbb{R}$, is as follows. For the first time-slice, MSVs are identified. From each stable volume, obtain a measurement vector and initialise a Kalman filter with it. These Kalman filters will then form the initial set of tracks. For consecutive time-steps, the algorithm will continue by first propagating each track's state through a process model, i.e. Kalman prediction, and identifying the MSVs from the current ERT time-slice and converted to observation domain for the filters. The assignment method will be used for the predicted observations and the MSV measurements. As described above, the tracks will be updated and new tracks created. The process is 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 tracked regions. Post-processing can be applied to remove extraneous tracks and detected noise, either manually or automatically, based on the number of time-steps a region was tracked, the expected velocity of the track, or the magnitude of the error covariance.

7.1.2 Pit Experiment (2-D time-lapse)

Algorithm 6 Tracer Tracking Framework

procedure INITTRACK($\boldsymbol{y}$) ▷ Initialise a new track based on a measurement
  $\boldsymbol{x} \leftarrow \mathbf{H}^{-1}\boldsymbol{y}$
  $\boldsymbol{P}^{(x)} \leftarrow \mathbf{F}\mathbf{I}\mathbf{F}^{\mathrm{T}} + \boldsymbol{Q}$
  $\kappa(0) \leftarrow (\boldsymbol{x}, \boldsymbol{P}^{(x)})$
  return κ
end procedure
procedure TTF($V_\rho(X, 1:T)$; ε, M, N) ▷ ε is non-assignment cost and M, N control lost tracks
  $\mathcal{T} \leftarrow \emptyset$ ▷ Set of active tracks
  $\mathcal{L} \leftarrow \emptyset$ ▷ Set of lost tracks
  for $t \in 1, \ldots, T$ **do**
    $\boldsymbol{\nu} \leftarrow V_\rho(X, t)$
    for $\kappa \in \mathcal{T}$ **do**
      $\kappa(t) \leftarrow$PREDICT($\kappa(t-1)$) ▷ $\kappa(t)$ has a state $(\boldsymbol{x}, \boldsymbol{P}^{(x)})$
      age(κ) $\leftarrow$ age(κ) + 1
      invisible(κ) $\leftarrow$ invisible(κ) + 1
    end for
    $C_\nu \leftarrow$MSV($\boldsymbol{\nu}$)
    $\delta(t, \mathbf{x}) \leftarrow V_\rho(t, \mathbf{x})^{-1} - V_\rho(0, \mathbf{x})^{-1}$
    $Y_\nu \leftarrow \{[\bar{\mathbf{x}}, \bar{\rho}]^{\mathrm{T}} \,|\, \mathbf{c} \in C_\nu, \mathbf{x} \in \mathbf{c},$
      $\bar{\mathbf{x}} \leftarrow \sum \mathbf{x}\delta(t, \mathbf{x}) / \sum \delta(t, \mathbf{x}), \bar{\rho} \leftarrow \sum \boldsymbol{\nu}(\mathbf{x})/|\mathbf{x}|\}$
    for all $\tilde{\boldsymbol{y}} \in Y_\nu$ **do**
      $\kappa^* \leftarrow \mathrm{argmin}_{\kappa \in \mathcal{T}}\{\mathrm{cost}(\tilde{\boldsymbol{y}}, \kappa(t))\}$
      if cost($\tilde{\boldsymbol{y}}, \kappa(t)$) $< \varepsilon$ **then**
        $\kappa^*(t) \leftarrow$UPDATE($\kappa^*, \tilde{\boldsymbol{y}}$)
        invisible(κ^*) $\leftarrow 0$
      else
        $\kappa' \leftarrow$INITTRACK($\tilde{\mathbf{y}}$)
        $\mathcal{T} \leftarrow \mathcal{T} \cup \{\kappa'\}$
        age(κ') $\leftarrow 1$
        invisible(κ') $\leftarrow 0$
      end if
    end for
    $\mathcal{L} \leftarrow \mathcal{L} \cup \{\kappa \,|\, \kappa \in \mathcal{T}, \mathrm{age}(\kappa) > M$ **and** invisible(κ) $> N\}$
    $\mathcal{T} \leftarrow \mathcal{T} \backslash \mathcal{L}$
  end for
  return $\mathcal{T}$, $\mathcal{L}$
end procedure

To develop and test the plume tracking framework, the experiments were performed using two time-lapse cross-borehole ERT tracer experiments, with one using 2-D cells, and the other using 3-D.

The 2-D time-lapse experiment consists of a saline tracer monitored using time-lapse ERT in a reduced field scale hydraulic test cell. The experimental area was

an unlined elongated trench, 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/day) compared to the rate of ERT image acquisition. This was controlled to avoid image blurring on the horizontal scale of the model cells (0.1 m).

The cell was instrumented with 11 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 dipole-dipole measurements between adjacent arrays, with a full dataset being measured every 40 minutes over a 24 hour period, 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 hours 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/day. A diagram of the experimental setup is shown in Figure 7.1. Full technical details of the experiment were published in (Kuras et al., 2009). The full inverted dataset is shown in Figure 7.2.

7.1.3 Tank Experiment (3-D time lapse)

The 3-D experiment setup is similar to the pit experiment, however it was carried out in a laboratory tank. The tank, with volume 1.30 m × 0.80 m × 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 1.7 m/day. The saline tracer had a concentration of 20 g/l, and had a density-driven, vertical Darcy velocity of 0.4 m/day. A volume of 1.3 l of the tracer was injected over 0.5 hours at a depth of 0.2 m and a distance of 0.1 m upstream of the ERT imaging volume.

The ERT volume has dimensions of 0.80 m × 0.40 m × 0.75 m, and is defined by 8 vertical electrode arrays, situated in tank as referenced by A1-A8 in Figure 7.3. Each electrode array comprised of 16 electrodes at 0.05 m spacing between depths of 0.1 m and 0.85 m. Cross-borehole dipole-dipole measurements were carried out between pairs of arrays shown by the numbered dashed lines in Figure 7.3 . Full datasets were acquired every hour over a 16 hour period and were inverted using a

3-D version of the pit experiment constraints, with L2 model smoothness and an L1 measure of data misfit.

7.2 Tracking Centre-of-Mass

7.2.1 Model Setup

The simplest way to model the tracer plume is to track its centre-of-mass, i.e. its first-order central moment. Given some region in the t^{th} time-slice of an ERT model, $\vec{X} \subset \mathbb{R}^d$, the centre-of-mass can be calculated using a weighted mean, such

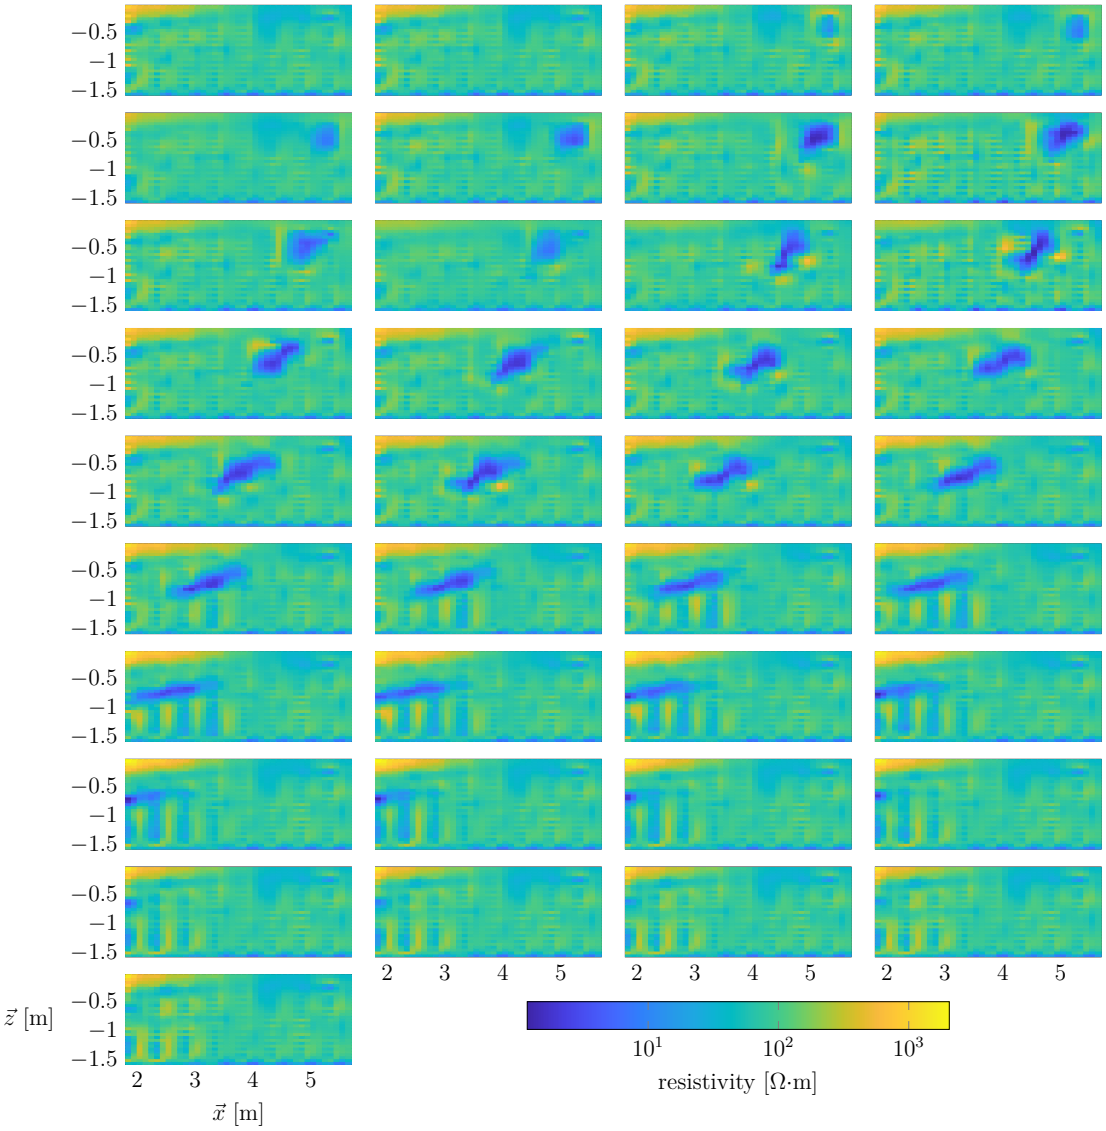


Figure 7.2: Time-lapse resistivity model for 2-D pit experiment over 37 time-steps showing the low resistivity tracer plume advection (ordered chronologically left to right, top to bottom)

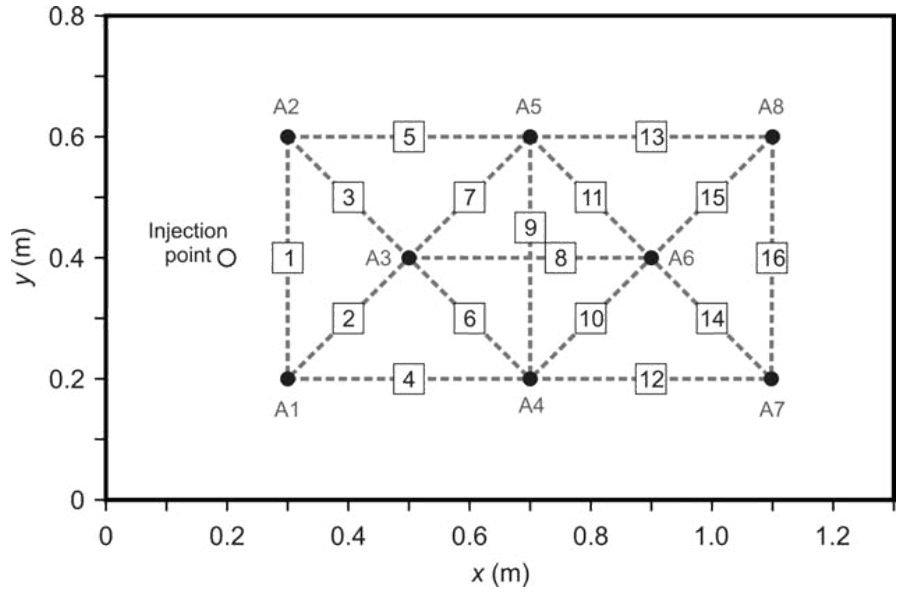


Figure 7.3: Arrangement of vertical electrodes in the 3-D laboratory tank experiment, showing measurement array pairs and tracer injection point

that each spatial cell in an MSV is weighted by its conductivity change from some baseline, $t = 0$.

$$\begin{aligned} \bar{\mathbf{x}} &= \frac{\sum_{\forall \vec{\mathbf{x}} \in \vec{X}} \vec{\mathbf{x}} \cdot \delta_c(t, \vec{\mathbf{x}})}{\sum_{\forall \vec{\mathbf{x}} \in \vec{X}} \delta_c(t, \vec{\mathbf{x}})} \\ \delta_c(t, \vec{\mathbf{x}}) &= V_\rho(t, \vec{\mathbf{x}})^{-1} - V_\rho(0, \vec{\mathbf{x}})^{-1} \end{aligned} \tag{7.2}$$

Here, δ_c denotes the change in conductivity. The centre-of-mass is the vector of Cartesian coordinates in 2-D or 3-D, defined as $\bar{\mathbf{x}} = (\vec{x}, \vec{y}, \vec{z})$ (ignoring the $\vec{z}$ coordinate in the 2-D case). The measured resistivity of the detected tracer region is taken simply as the mean of the resistivity of each cell in the detected region.

$$\bar{\rho} = \frac{1}{|\vec{X}|} \sum_{\forall \vec{\mathbf{x}} \in \vec{X}} V_\rho(t, \vec{\mathbf{x}}) \tag{7.3}$$

In both the pit and the tank experiment, the systems are assumed to be subject to a constant horizontal velocity field and density-driven sinking. The basic velocity model is used as the process that drives the state over time. This is a linear system, and the linear-discrete Kalman filter can be used. The state vector is thus formed of the tracer plume centre-of-mass and mean resistivity, as well as the velocity

along the principle axes. Table 7.1 and Table 7.2 show the setup of the filters used in the tracks for the 2-D pit experiment and the 3-D tank experiment.

7.2.2 Results

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 contrast difference with the background. MSVs are removed automatically if they were outside the range of tolerated sizes which, in both the 2-D and 3-D

Table 7.1: 2-D Pit Model Setup

Process	$\boldsymbol{x}_t$	**F**	$\boldsymbol{Q}$
	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ \bar{\rho} \\ v_{\vec{x}} \\ v_{\vec{y}} \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$\text{diag}\left(\begin{bmatrix} 0.2 \\ 0.2 \\ 0.01 \\ 0.001 \\ 0.001 \end{bmatrix}\right)$
Observation	**H**		$\boldsymbol{R}$
	$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$		$\text{diag}\left(\begin{bmatrix} 0.1 \\ 0.01 \\ 0.02 \end{bmatrix}\right)$

Table 7.2: 3-D Tank Model Setup

Process	$\boldsymbol{x}_t$	**F**	$\boldsymbol{Q}$
	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ \vec{z} \\ \bar{\rho} \\ v_{\vec{x}} \\ v_{\vec{y}} \\ v_{\vec{z}} \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$0.01\mathbf{I}_5$
Observation	**H**		$\boldsymbol{R}$
	$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$		$\text{diag}\left(\begin{bmatrix} 0.5 \\ 0.1 \\ 0.2 \\ 0.2 \end{bmatrix}\right)$

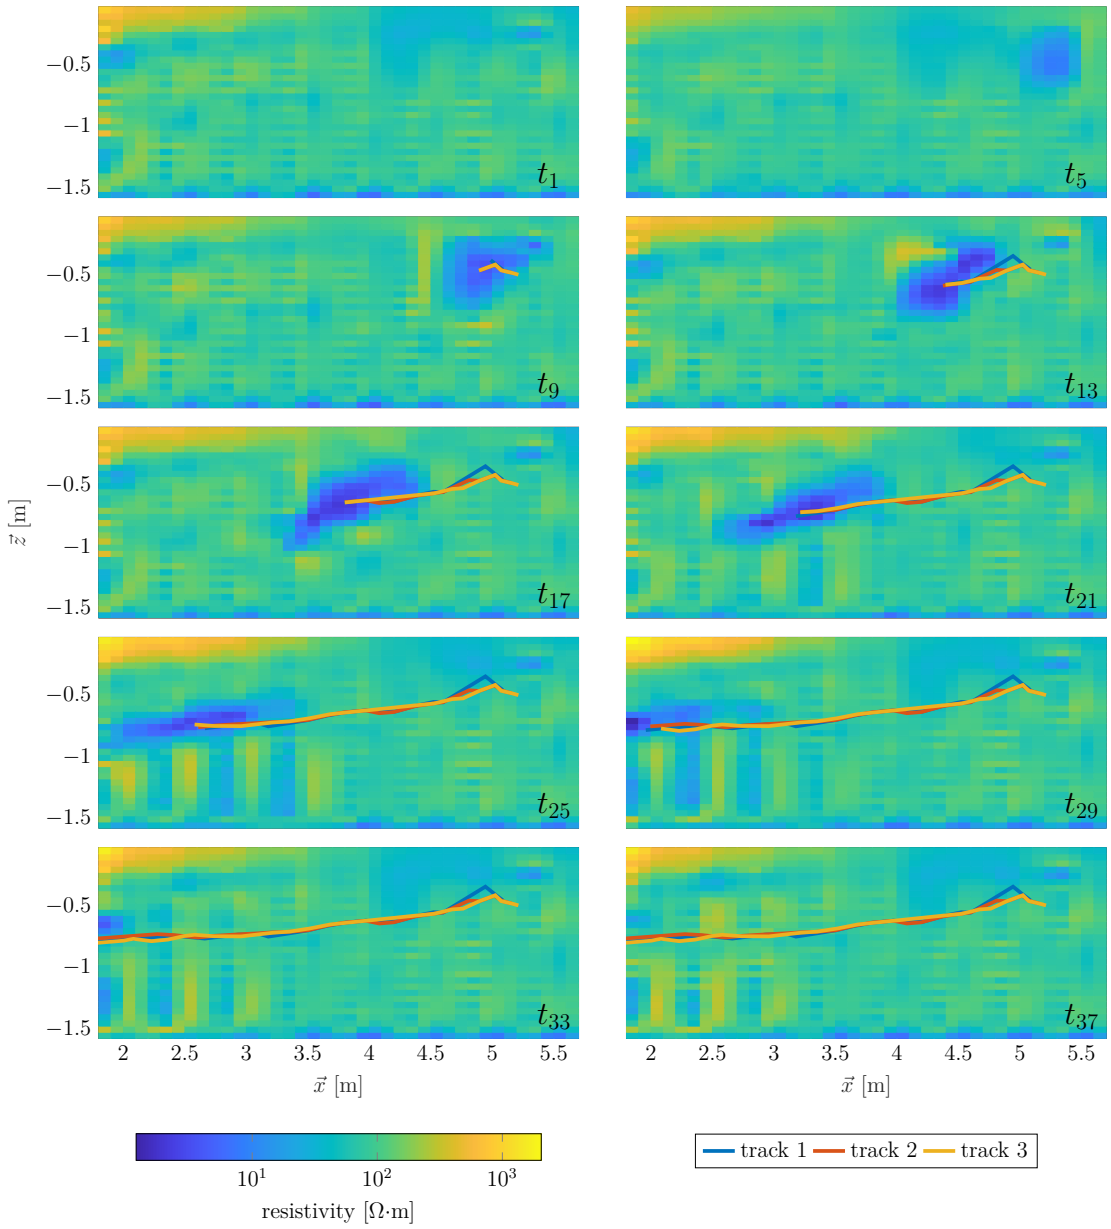


Figure 7.4: Time-slices of the 2-D time-lapse ERT model of the pit tracer experiment, with Kalman filter tracks showing the tracked position of the centre-of-mass of the tracer plume up to the current time-slice

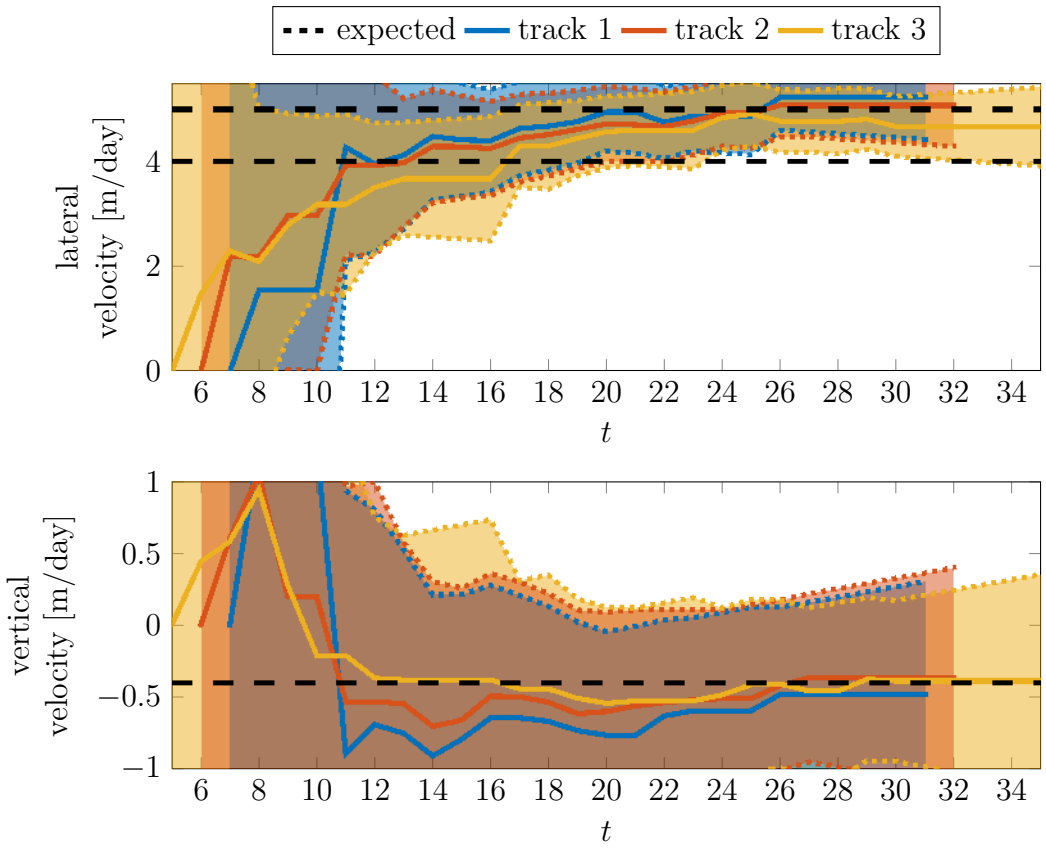


Figure 7.5: Plots showing horizontal (top) and vertical (bottom) velocity state estimates (m/day) for the three tracks shown in Figure 7.4. Calculated expected velocities are marked with dashed black lines: 4-5 and -0.4 m/day respectively. Shaded regions represent the estimation error to ±1 standard deviation

---

experiments, as fixed between 2.5% and 25% of the total volume of the image, inclusive.

The results at the end of the tracking in the pit experiment are displayed at intervals of 4 time-slices in Figure 7.4. At the end of the 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/day 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 were 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 7.4 are represented by the connecting solid lines. The three tracks representing the tracer path progress show the assignment and detection of the tracer plume: the blue region of low resistivity, ∼1-30 Ω·m.

The tracked state estimates of horizontal and vertical velocity of these tracks are plotted against time in Figure 7.5. The top plot shows the estimated velocity of the tracked region reaching approximately 4-5 m/day, consistent with the controlled flow field of the experiment. Toward the end of the tracking period, the measured velocity is slightly higher. Despite this, the error region representing ± 1 standard deviation bounds the upper velocity limit. The density-driven sinking estimated by the vertical state velocity is also plotted, and this decreases to a value that is consistent with the calculated Darcy velocity of the system, at -0.4 m/day. The pattern is present in each of the tracer tracks and shows the reliability of the tracking model.

The results of the MSV algorithm for each time-step in the 3-D tank experiment are shown in Figure 7.6. 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 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 artefacts in the ERT data, such as those resolved in t_1.

The tracks shown in Figure 7.6 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, $\|(v_{\vec{x}}, v_{\vec{y}})\|$, and vertical velocity $v_{\vec{z}}$ of these tracks are plotted in Figure 7.7. Each path shown corresponds to the tracer plume which, at later time-steps, is represented by a set of nested MSVs, and show that both horizontal and vertical estimated velocities match the expected values. 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 begin to decrease in the final four time steps. This corresponds to the tracer plume moving outside the imaging region, which occurs over multiple time steps. Detections of the tracer in this plume are thus incomplete and the detected centre-of-mass begins to affect the Kalman state for these tracks. The phenomenon does not affect all tracks, as one is no longer dictated by detection and is driven only by the predictions based on the prior state. 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.

7.3 Moment Tracking

7.3.1 Model Setup

Extending the representation of the tracked object in a tracer experiment, it is possible to incorporate more information form the detected features to better model the system observed. For tracer plume tracking, it is possible to exploit some of

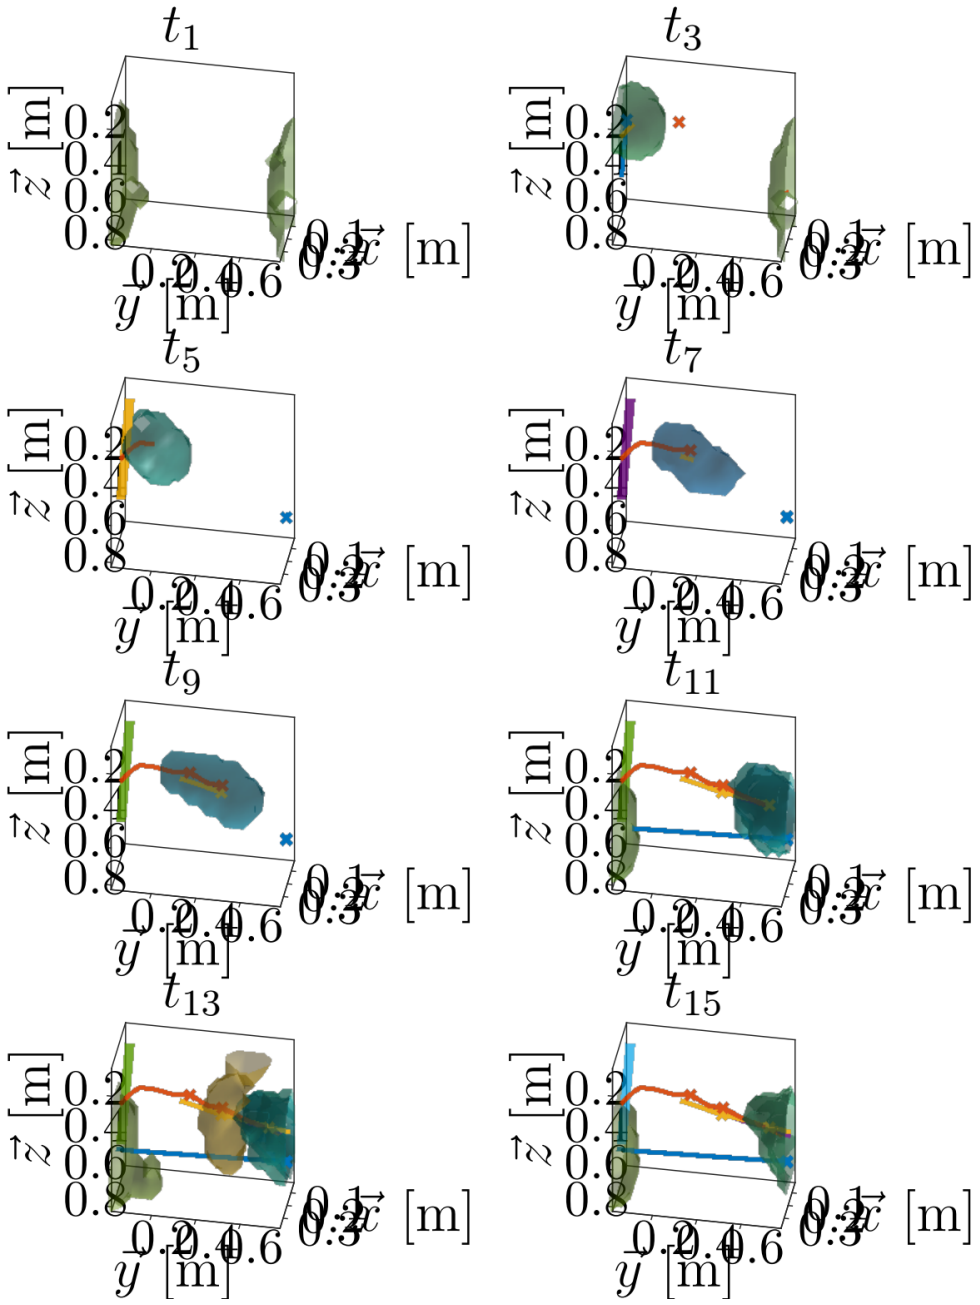


Figure 7.6: Time-slices of 3-D time-lapse ERT model of tank experiment, showing detected MSVs and Kalman filter tracks of the tracked positions up to the current time-slice. MSV surfaces are shaded to represent log mean resistivity

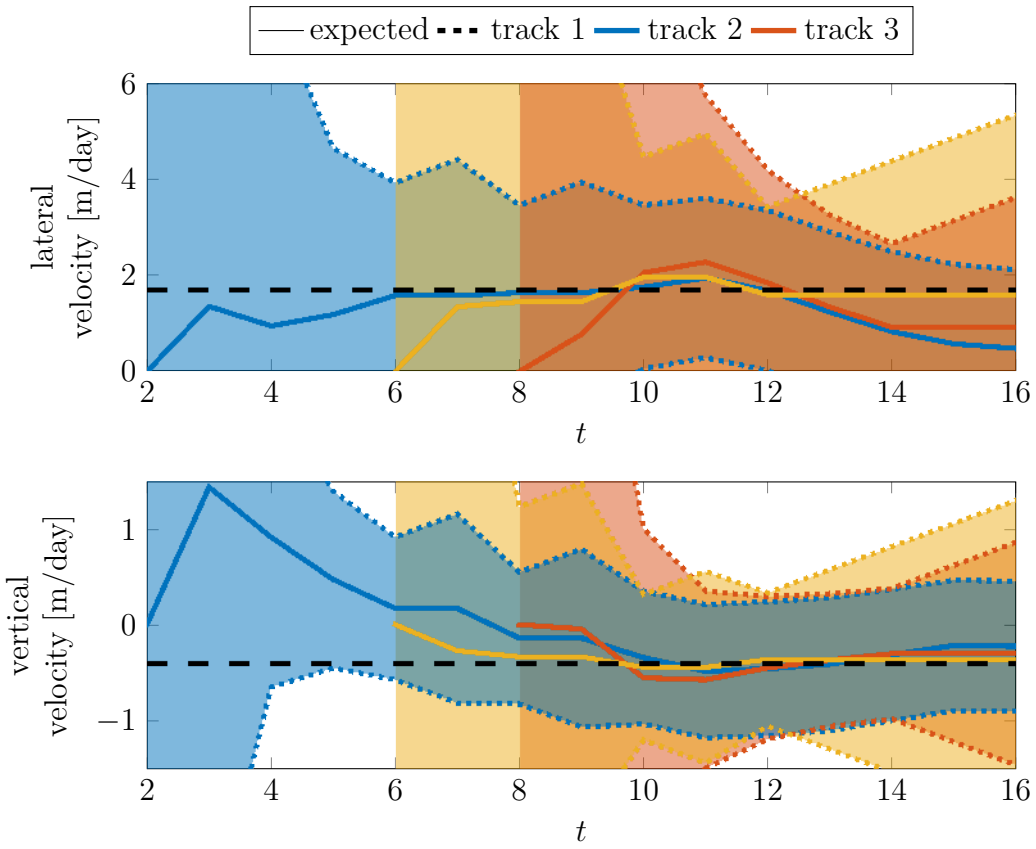


Figure 7.7: Plots showing the lateral (top) and vertical velocity state estimates (m/day) for the tracks filtered in the 3-D tank experiment. The shaded bars repreresent ±1.0 s.d. and dashed lines indicated expected velocities of 1.8 m/day and 0.4 m/day respectively

---

the simple physics behind the hydrodynamics involved. A reasonable assumption, based on the known information, is that the tracer will experience advection along the flow field at a fixed velocity. The velocity parameters are presumed unknown in the observation model, and estimated with the process model, as defined in Table 7.1 and Table 7.2. An additional expectation is that the tracer plume will experience dispersion, both chemically and mechanically.

Estimating tracer dispersion as a linear system can be done by approximating the plume's spread, i.e. second-order moment, with a bounding ellipsoid, represented by a transformed semi-axis. A semi-axis magnitude $\vec{\mathbf{s}}$ describes scaling factors for orthonormal unit axes, that correspond to the principle coordinate axes, $\vec{\mathbf{i}} = (1, 0, 0)$, $\vec{\mathbf{j}} = (0, 1, 0)$, and $\vec{\mathbf{k}} = (0, 0, 1)$, and θ describes the angle of orientation about the $\vec{\mathbf{k}}$-axis, i.e. perpendicular to the expected flow field. Using the parameters described, a new observation can be defined as the combination of centre-of-mass and orientation of a tracer plume in an ERT model: $\boldsymbol{y} = (\vec{\mathbf{x}}, \bar{\rho}, \vec{\mathbf{s}}, \theta)$.

For a basic linear representation of advection-dispersion, the parameters for velocity and rate of dispersion must be estimated within the dynamic state model. Including parameters for velocity, as in the centre-of-mass model, and dispersion coefficients along the component axes representing the tracer's second-order moment, gives a model setup as defined in Table 7.3. In this model, both the velocity and rate of dispersion are considered approximately constant, and the estimated orientation will remain the same, as with the mean resistivity, due to constant effects of the flow field and gravity. As such, the dynamic system is be defined as with process and observation models for the the first-order and second-order moments are the

Table 7.3: 2-D second-order moment tracking model setup

Process	$\boldsymbol{x}_t$	**F**
	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ \bar{\rho} \\ v_{\vec{x}} \\ v_{\vec{y}} \\ \vec{s}_i \\ \vec{s}_j \\ \vec{\theta} \\ d_i \\ d_j \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$
Observation		**H**
		$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}$
Noise Covariance	$\boldsymbol{Q}$	$\boldsymbol{R}$
	$\text{diag}\left(\begin{bmatrix} 0.1 \\ 0.01 \\ 0.01 \\ 0.01 \\ 0.0001 \\ 0.1 \\ 0.1 \\ 0.1 \\ 0.001 \\ 0.001 \end{bmatrix}\right)$	$\text{diag}\left(\begin{bmatrix} 1 \\ 0.2 \\ 0.2 \\ 0.1\pi \\ 0.1 \\ 0.1 \end{bmatrix}\right)$

same.

7.3.2 Results

The extended model, accounting for both velocity and dispersion properties give a greater level of detail for prediction and estimation of the system state. A subset of the tracking results on the 2-D pit tracer experiment are shown in Figure 7.8, which visualises the second-order moment semi-axes as ellipsoids, showing detected features and tracked moments: the spread of the ellipsoid represents the spread of the tracer plume, and the centre of the ellipsoid corresponds to the central first-order moment, i.e. the centre-of-mass.

From the tracked ellipsoids, it can be seen that the constraining model of constant linear dispersion restricts deviations in the detections caused by temporally independent maximally stable volumes. In the early time-slices, there are misdetections of low resistivity artefacts in t_1 and t_5. These detections are not associated with tracks as the tracer plume is as it progresses through the imaging plane. Additional regions are tracked, such as the highly resistive regions in the top left corner of the imaging plane, and the resistivity artefacts in the bottom of the image plane that correspond to the bases of the electrode boreholes every 0.5 m. These stationary regions tend to be tracked because, although they do not correspond with the tracer, they share the assumed properties of the process model, namely constant velocity and rate of dispersion, albeit constantly zero.

In comparison with the tracks resulting from only centre-of-mass tracking, in Figure 7.4, the extension of the representation serves to demonstrate an effective tracking regime. Visual confirmation of the results shows equivalent ability of tracking first order moments, although with more false tracks. In addition, the semi-axis representation of the tracer plume is successfully tracked, and in some cases corresponds better with the expectation of the tracer plume that the detected features.

7.4 Discussion

The tracer tracking framework developed serves as an effective means of automatically identifying tracer movement with limited user intervention in isolating specific tracks. The results take into account inaccuracies in detection methods and variance from the expected model. The results have clearly shown the capabilities of tracking a moving object, demonstrating simple post processing steps that can

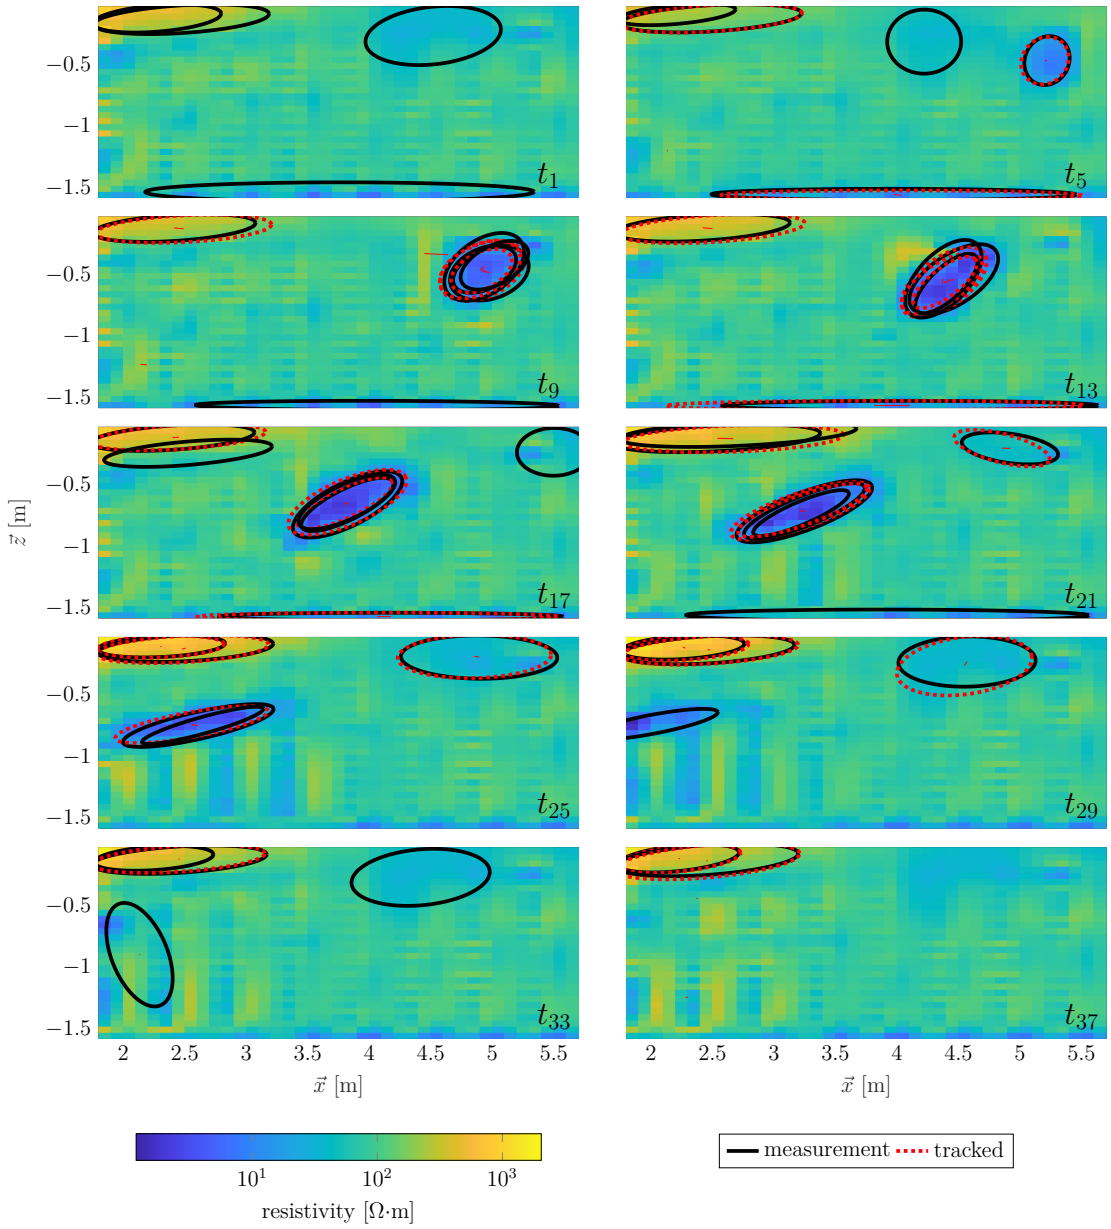


Figure 7.8: Time-slices of the 2-D time-lapse ERT model of the pit tracer experiment, with detected and filtered ellipsoids representing the automatically tracked second-order moments of tracer at the current time-step

be both automated and manually controlled. In this case, short or stationary tracks can be easily identified for the given state transition parameters.

The detection method can effectively identify and isolate MSVs, and the capabilities for nested volumes allow for multiple tracks to be assigned to the same tracer. This approach can be useful for observing different levels of concentration in the tracer, for example, one track following the movement of the high concentration region and another tracking the larger region as the tracer disperses into the outlying environment. This is especially the case in the in the semi-axis representation of the tracer spread, as discrete iso-levels of the tracer's concentration can be tracked as independent ellipsoids.

Using the Bayesian approach of the Kalman filter has been effective in dealing with the uncertainties involved with ERT imaging. Taking into account potential errors in both the detection method and expected trajectory allows an appropriate level of response to changes in the system. The Kalman gain, in particular, serves as an optimal, or near-optimal, weighting of the estimated versus measured observations of the system, based on the process model and estimate error covariance of the tracked state. It is evident that the tracer tracking framework described can provide useful functionality in the detection and tracking of subsurface hydrogeophysical processes, both in assisting and automating monitoring.

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 using ERT data, there are greater assumptions required for the certainty of detection. Due to the inversion process, it is always the case that image resolution varies throughout the imaging space. A degree of uncertainty is present that is non-uniform in the imaging survey, something that is not explicitly dealt with in the tracking framework. In addition to this, the detection method relies on local connectivity between resistivity cells to calculate connected components. The assumption here is that the cells are sized and distributed uniformly. Uniform model cells are present in the ERT experiments used here, but this is not a necessity in ERT imaging and is not always a valid assumption. The current implementation of MSV relies on the structured grids for calculating connectivity.

The estimation of the tracer moments are also based on a pseudo-measure of concentration that weights cells based on their observed change from some baseline. Where there is no convenient baseline, for example 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 con-

ductivity would additionally allow a focussed 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 artefacts that may be removed with the use of a baseline (Ward et al., 2016).

The methods are also limited by the need for feature detection, and capturing the boundaries of individual features is non-trivial, especially where they are gradational. In addition, feature detectors can approximate multiple objects that do not correspond with regions of interest. The artefacts in the image must be filtered out, hence the need for track assignment. However, a preferable approach may model the entire imaging plane as a single observation model. For example the tracer experiment may be modelled using the advective dispersion equation, so as to model the concentration over a given spatial domain. Archie's law may be used to relate the resistivity to concentration, and serve as the observation model, mapping binned concentration to resistivity cells.

Problem	**Contribution**
Detecting change in time-lapse models	The development of the tracer tracking framework has shown clear efficacy for predicting moments in tracer plumes; and the identification of system parameters using ERT
Combining mechanistic models with data	Combining a basic mechanistic model, first in the form of transport only, and then incorporating diffusion, with ERT has allowed the underlying hydrogeological system to be modelled with minimal direct observation
Quantifying the uncertainty in the system and making use of this in predictions	The use of Bayesian methods allow for the capture of uncertainty from the detection method, namely mis-identifiers, and lost tracks, has given a robustness to the system that is applicable in both 2-D and 3-D experiment data
Presenting the uncertainty in the system	The parameters obtained as part of the unobserved state model are presented as probability distributions, parameterised by the mean and covariance of a normal distribution. This has been used to quantify confidence in the results.

Table 7.4: Contributions of the tracer tracking framework

8

Fuzzy Rules for Filtering in Multi-Process Systems

8.1 Motivation

Systems under the influence of uncertain dynamic processes can pose a distinct challenge for predictive estimators, especially in the case where multiple non-linear processes influence the system state to varying degree. In a wide range of application-domain problems, including sensing data, maneuvering target tracking, and hydrodynamic modelling, there are complex system processes that occur simultaneously, or consecutively at unknown intervals. In the general case, consider this problem of a state space model with multiple processes, as introduced in Section 3.6

$$x_t^j = f^j(t, x_{t-1}, u_t^j; \xi_t^j) \tag{8.1}$$

The state is propagated over time by a set of r process models f^j, each with some control parameters u_t^j and unknown noise ξ_t^j. As reported previously, there are existing methods that offer suboptimal filtering of the SSM, one such is the interacting multiple model. While the IMM is a powerful tool, it relies on certain assumptions, and has limitations. The key assumption in using IMM is that there is some known, or estimable model of switching probabilities that describe the probability that a given state process will switch to another (see Figure 3.5). In many real world cases, this is an unknown and potentially highly variable parameter, and the choice of these probabilities can drastically affect the result of

the filter. In the case of IMM, a state estimate for each individual process model is maintained, and they are joined at each time-step based on the switching model and measurement. The problem can lead to issues of scalability whereby multiple states for each time-slice, and large switching models, can lead to poor results.

As an alternative approach to handling the multi-process SSM problem, a solution is proposed that embeds a fuzzy inference system into each time-step of a recursive Bayesian inference system. Based on predictions of the state, and corresponding probability distribution functions of the estimated observation for each corresponding process model, a set of membership functions can be defined to approximate the true state. Here, each process can be considered a *linguistic* value in the overall SSM. For example, consider a basic motion model of an object on a plane. Possible linguistic models can include moving forward, and turning (left/right). If each value can be considered to occur at any point, seemingly independently, to varying degree, i.e. at a given time, the tracked object may turn left by an unknown degree, then a fuzzy inference system can be used to weight each possible model based on some membership functions, and output a crisp state. The use of fuzzy logic also seeks to simplify the modelling process, and handle the uncertainty that exists in switching, without the need for prior probabilities to be determined or estimated.

8.2 Type-1 Fuzzy Kalman Filter

8.2.1 Formulation

Consider the (non-linear) Markov-switching SSM, with r independent component processes, defined in (3.28). Without loss of generality, define each component process model f^j such that it maps the system state, $\boldsymbol{x}_t$ to the domain of a component state $\boldsymbol{x}^j_{t|t-1}$, the *a priori* estimate of j^{th} component state. Consider also a single observed measurement (vector) at time t, $\tilde{\boldsymbol{y}}_t$ and define each observation model h^j such that it maps the component state estimate $\boldsymbol{x}^j_t$ into the observation domain. The proposed filter seeks to estimate the true (unified) process state x_t, by first calculating the *a priori* estimates of each component state, x^j_t, such that

$$\begin{aligned} p\Big(x^j_t \,|\, x_{t-1}\Big) &= \mathcal{N}\Big(f^j(t, \boldsymbol{x}_{t-1}, u^j_t), \boldsymbol{Q}^j_t\Big) \\ p\Big(y_t \,|\, x^j_t\Big) &= \mathcal{N}\Big(h^j(t, \boldsymbol{x}^j_t, u^j_t), \boldsymbol{R}_t\Big) \end{aligned} \tag{8.2}$$

Here, $p(y_t \,|\, x^j_t)$ represents the conditional observation pdf on given only the effect of the j^{th} process at time t. The expected value of the conditional distribution is

denoted $\boldsymbol{y}_t^j = \mathbb{E}[y_t \,|\, x_t^j]$. The process control for the j^{th} component is denoted u_t^j and the noise term for each process model is assumed to be a zero-mean normal distribution, $\xi_t^j \sim \mathcal{N}(\mathbf{0}, \boldsymbol{Q}_t^j)$.

The predictive distributions of each component process can be independently estimated by means of a Kalman prediction method, considering each is given an estimate of the previous state, for component model j, the *a priori* estimate and corresponding error covariance are defined:

$$\begin{aligned} \boldsymbol{x}_{t|t-1}^j &= f^j(t, \boldsymbol{x}_{t-1}, u_t^j) \\ \boldsymbol{P}_{t|t-1}^{j(x)} &= \operatorname{cov}\left[\boldsymbol{x}_{t|t-1}^j - \tilde{\boldsymbol{x}}_{t-1}, \boldsymbol{x}_{t|t-1}^j - \tilde{\boldsymbol{x}}_{t-1}\right] \end{aligned} \tag{8.3}$$

8.2.2 Linear Models

The most simple case for prediction is where all underlying process and observation models are linear, in which case there exist for each component SSM a process and observation matrix $\mathbf{F}^j$ and $\mathbf{H}^j$, respectively. Here, $\mathbf{F}^j$ is an $n_j \times n_x$ matrix, where $n_j := |\boldsymbol{x}^j|$ and $n_x := |\boldsymbol{x}|$, the dimensionality of the component and unified states respectively; and $\mathbf{H}^j$ is a $n_y \times n_j$ matrix, where $n_y := |\boldsymbol{y}|$. The prediction step for linear models is thus

$$\begin{aligned} \boldsymbol{x}_{t|t-1}^j &= \mathbf{F}^j \boldsymbol{x}_{t-1} + \mathbf{U}_t^f u_t^j + \xi_t^j \\ \boldsymbol{y}_{t|t-1}^j &= \mathbf{H}^j \boldsymbol{x}_{t|t-1}^j + \mathbf{U}_t^h u_t^j + \nu_t^j \end{aligned} \tag{8.4}$$

8.2.3 Estimate Error Covariances

The corresponding component calculations for process, observation and joint covariances are derived as follows. Consider that the covariance can be defined as the matrix derived from the expectated product of the deviation of each variate from its expected value: $\operatorname{cov}[\boldsymbol{a}, \boldsymbol{b}] \equiv \mathbb{E}\left[(\boldsymbol{a} - \mathbb{E}[\boldsymbol{a}])\,(\boldsymbol{b} - \mathbb{E}[\boldsymbol{b}])^{\mathrm{T}}\right]$. Then, from (8.3), the predictive estimation error covariance can be defined as such.

$$
\begin{aligned}
\boldsymbol{P}_{t|t-1}^{j(x)} &= \mathbb{E}\Big[\Big(\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big) - \mathbb{E}\Big[\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big]\Big) \\
&\qquad \Big(\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big) - \mathbb{E}\Big[\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big]\Big)^{\mathrm{T}}\Big] \\
&= \mathbb{E}\Big[\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big)\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big)^{T}\Big] \\
&= \mathbb{E}\Big[\Big(\mathbf{F}^{j}\boldsymbol{x}_t + \mathbf{U}_t^{f}u_t^{j} + \xi_t^{j} - \tilde{\boldsymbol{x}}_t\Big) \\
&\qquad \Big(\mathbf{F}^{j}\boldsymbol{x}_t + \mathbf{U}_t^{f}u_t^{j} + \xi_t^{j} - \tilde{\boldsymbol{x}}_t\Big)^{\mathrm{T}}\Big] \\
&= \mathbb{E}\Big[\Big(\mathbf{F}^{j}\boldsymbol{x}_t + \mathbf{U}_t^{f}u_t^{j} - \tilde{\boldsymbol{x}}_t\Big) \\
&\qquad \Big(\mathbf{F}^{j}\boldsymbol{x}_t + \mathbf{U}_t^{f}u_t^{j} - \tilde{\boldsymbol{x}}_t\Big)^{\mathrm{T}}\Big] + \mathbb{E}\Big[\xi_t^{j}\xi_t^{j\mathrm{T}}\Big] \\
&\approx \mathbb{E}\Big[\Big(\mathbf{F}^{j}\boldsymbol{x}_{t-1} - \mathbf{F}^{j}\tilde{\boldsymbol{x}}_{t-1}\Big)\Big(\mathbf{F}^{j}\boldsymbol{x}_{t-1} - \mathbf{F}^{j}\tilde{\boldsymbol{x}}_{t-1}\Big)^{T}\Big] + \mathbb{E}\Big[\xi_t^{j}\xi_t^{j\mathrm{T}}\Big] \\
&= \mathbb{E}\Big[\Big(\mathbf{F}^{j}\boldsymbol{x}_{t-1} - \mathbf{F}^{j}\tilde{\boldsymbol{x}}_{t-1}\Big)\Big(\mathbf{F}^{j}\boldsymbol{x}_{t-1} - \mathbf{F}^{j}\tilde{\boldsymbol{x}}_{t-1}\Big)^{T}\Big] + \boldsymbol{Q}_t^{j} \\
&= \mathbf{F}^{j}\mathbb{E}\Big[\big(\boldsymbol{x}_{t-1} - \tilde{\boldsymbol{x}}_{t-1}\big)\big(\boldsymbol{x}_{t-1} - \tilde{\boldsymbol{x}}_{t-1}\big)^{T}\Big]\mathbf{F}^{j\mathrm{T}} + \boldsymbol{Q}_t^{j} \\
&= \mathbf{F}^{j}\boldsymbol{P}_{t-1}^{(x)}\mathbf{F}^{j\mathrm{T}} + \boldsymbol{Q}_t^{j}
\end{aligned}
\tag{8.5}
$$

Where the filter seeks to minimise the error in approximation, the $\mathbb{E}\Big[\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big] = \mathbf{0}$. Since the derivation is based on the assumption that the component process is representative of the current state, such that $\tilde{\boldsymbol{x}}_t \approx f^{j}(t, \tilde{\boldsymbol{x}}_{t-1}, u_t^{j}) = \mathbf{F}^{j}\tilde{\boldsymbol{x}}_{t-1} + \mathbf{U}_t^{f}u_t^{j}$. Where the number of component processes, $r = 1$, this is an equivalence relation and trivially the component estimate error covariance is equivalent to that of the state error covariance in the linear-discrete Kalman filter. It follows that observation and joint estimate covariances for each component process can be derived in a similar way, again assuming for an optimal filter, the expected difference between estimation and true state is zero.

$$
\begin{aligned}
\boldsymbol{P}_t^{j(xy)} &= \mathbb{E}\Big[\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big)\Big(\boldsymbol{y}_{t|t-1}^{j} - \tilde{\boldsymbol{y}}_t\Big)^{\mathrm{T}}\Big] \\
&\approx \mathbb{E}\Big[\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big)\Big(\mathbf{H}^{j}\boldsymbol{x}_{t|t-1}^{j} - \mathbf{H}_t^{j}\tilde{\boldsymbol{x}}_t\Big)^{\mathrm{T}}\Big] \\
&= \mathbb{E}\Big[\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big)\Big(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t\Big)^{\mathrm{T}}\Big]\mathbf{H}_t^{j\mathrm{T}} \\
&= \boldsymbol{P}_{t|t-1}^{j(x)}\mathbf{H}_t^{j\mathrm{T}}
\end{aligned}
\tag{8.6}
$$

$$
\begin{aligned}
\boldsymbol{P}_t^{j(y)} &= \mathbb{E}\Big[\Big(\boldsymbol{y}_{t|t-1}^j - \tilde{\boldsymbol{y}}_t\Big)\Big(\boldsymbol{y}_{t|t-1}^j - \tilde{\boldsymbol{y}}_t\Big)^{\mathrm{T}}\Big] \\
&\approx \mathbb{E}\Big[\Big(\mathbf{H}^j\boldsymbol{x}_{t|t-1}^j + \nu_t^j - \mathbf{H}_t^j\tilde{\boldsymbol{x}}_t\Big)\Big(\mathbf{H}^j\boldsymbol{x}_{t|t-1}^j + \nu_t^j - \mathbf{H}_t^j\tilde{\boldsymbol{x}}_t\Big)^{\mathrm{T}}\Big] \\
&= \mathbb{E}\Big[\Big(\mathbf{H}^j\boldsymbol{x}_{t|t-1}^j - \mathbf{H}_t^j\tilde{\boldsymbol{x}}_t\Big)\Big(\mathbf{H}^j\boldsymbol{x}_{t|t-1}^j - \mathbf{H}_t^j\tilde{\boldsymbol{x}}_t\Big)^{\mathrm{T}}\Big] + \mathbb{E}\Big[\nu_t^j\nu_t^{j\mathrm{T}}\Big] \\
&= \mathbf{H}_t^j\mathbb{E}\Big[\Big(\boldsymbol{x}_{t|t-1}^j - \tilde{\boldsymbol{x}}_t\Big)\Big(\boldsymbol{x}_{t|t-1}^j - \tilde{\boldsymbol{x}}_t\Big)^{\mathrm{T}}\Big]\mathbf{H}_t^{j\mathrm{T}} + \boldsymbol{R}_t^j \\
&= \mathbf{H}_t^j\boldsymbol{P}_{t|t-1}^{j(x)}\mathbf{H}_t^{j\mathrm{T}} + \boldsymbol{R}_t^j \\
&= \mathbf{H}_t^j\boldsymbol{P}_t^{j(xy)} + \boldsymbol{R}_t^j
\end{aligned}
\tag{8.7}
$$

As with the state estimation, the component observation process is considered approximately representative of the true observation value, such that $\tilde{\boldsymbol{y}}_t \approx h^j(t, \tilde{\boldsymbol{x}}_t, u_t^j) = \mathbf{H}_t^j\tilde{\boldsymbol{x}}_t + \mathbf{U}_t^h u_t^j$. It also follows that when $r = 1$, the approximation holds as equivalency and the above estimation error covariances are their corresponding optimal covariance matrix counterparts in the linear-discrete Kalman filter, as defined in (3.11).

8.2.4 Component Gain

The estimation for each component state seeks to minimise the error between the estimated state, which can be attained as the a minimum mean-square error (MSE) using the estimation error covariance:

$$
\boldsymbol{x}_{t|t}^j = \underset{\boldsymbol{x}}{\operatorname{argmin}}\,\mathbb{E}\Big[(\boldsymbol{x} - \tilde{\boldsymbol{x}}_t)^{\mathrm{T}}(\boldsymbol{x} - \tilde{\boldsymbol{x}}_t)\Big] \tag{8.8}
$$

It can be noted that the mean-square error is equivalent to the trace of the estimate error covariance matrix, where the trace is defined as the sum of all diagonal elements: $\mathrm{tr}[\mathbf{A}] = \sum_{\forall j} A_{jj}$.

Consider the following correction term for component state estimate $\boldsymbol{x}_t^j$, with matrix gain K_t^j that minimises the MSE based on the residual between measured and estimated observations:

$$
\boldsymbol{x}_{t|t}^j = \boldsymbol{x}_{t|t-1}^j + \mathbf{K}_t^j\Big(\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_{t|t-1}^j\Big) \tag{8.9}
$$

The corresponding error covariance matrix $\boldsymbol{P}_{t|t}^{j(x)}$ can be defined substituting in this correction term.

$$
\begin{aligned}
\boldsymbol{P}_{t|t}^{j(x)} &= \operatorname{cov}\left[\boldsymbol{x}_{t|t}^{j} - \tilde{\boldsymbol{x}}_t, \boldsymbol{x}_{t|t}^{j} - \tilde{\boldsymbol{x}}_t\right] \\
&= \mathbb{E}\left[\left(\boldsymbol{x}_{t|t}^{j} - \tilde{\boldsymbol{x}}_t\right)\left(\boldsymbol{x}_{t|t}^{j} - \tilde{\boldsymbol{x}}_t\right)^{\mathrm{T}}\right] \\
&= \mathbb{E}\Big[\left(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t + \mathbf{K}_t^{j}\left(\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_{t|t-1}^{j}\right)\right) \\
&\qquad \left(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t + \mathbf{K}_t^{j}\left(\tilde{\boldsymbol{y}}_t - \boldsymbol{y}_{t|t-1}^{j}\right)\right)^{\mathrm{T}}\Big] \\
&\approx \mathbb{E}\Big[\left(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t + \mathbf{K}_t^{j}\left(\mathbf{H}^{j}\tilde{\boldsymbol{x}}_t + \nu_t - \mathbf{H}^{j}\boldsymbol{x}_{t|t-1}^{j}\right)\right) \\
&\qquad \left(\boldsymbol{x}_{t|t-1}^{j} - \tilde{\boldsymbol{x}}_t + \mathbf{K}_t^{j}\left(\mathbf{H}^{j}\tilde{\boldsymbol{x}}_t + \nu_t - \mathbf{H}^{j}\boldsymbol{x}_{t|t-1}^{j}\right)\right)^{\mathrm{T}}\Big] \\
&= \mathbb{E}\Big[\left(\left(\mathbf{I} - \mathbf{K}_t^{j}\mathbf{H}^{j}\right)\left(\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_{t|t-1}^{j}\right) + \mathbf{K}_t^{j}\nu_t^{j}\right) \\
&\qquad \left(\left(\mathbf{I} - \mathbf{K}_t^{j}\mathbf{H}^{j}\right)\left(\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_{t|t-1}^{j}\right) + \mathbf{K}_t^{j}\nu_t^{j}\right)^{\mathrm{T}}\Big] \\
&= \left(\mathbf{I} - \mathbf{K}_t^{j}\mathbf{H}^{j}\right)\mathbb{E}\left[\left(\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_{t|t-1}^{j}\right)\left(\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_{t|t-1}^{j}\right)^{\mathrm{T}}\right]\left(\mathbf{I} - \mathbf{K}_t^{j}\mathbf{H}^{j}\right)^{\mathrm{T}} \\
&\quad + \mathbf{K}_t^{j}\mathbb{E}\left[\nu_t^{j}\nu_j^{j\mathrm{T}}\right]\mathbf{K}_t^{j\mathrm{T}} \\
&= \left(\mathbf{I} - \mathbf{K}_t^{j}\mathbf{H}^{j}\right)\boldsymbol{P}_{t|t-1}^{j(x)}\left(\mathbf{I} - \mathbf{K}_t^{j}\mathbf{H}^{j}\right)^{\mathrm{T}} + \mathbf{K}_t^{j}\boldsymbol{R}_t^{j}\mathbf{K}_t^{j\mathrm{T}}
\end{aligned}
\tag{8.10}
$$

The diagonals of this matrix represent the square error of each element in the component state, and thus the trace is equal the mean square error of the state estimate. While expanding out the above formulation of $\boldsymbol{P}_{t|t}^{j(x)}$, it follows that:

$$
\begin{aligned}
\operatorname{tr}\left[\boldsymbol{P}_{t|t}^{j(x)}\right] &= \operatorname{tr}\Big[\boldsymbol{P}_{t|t-1}^{j(x)} - \boldsymbol{P}_{t|t-1}^{j(x)}\mathbf{H}^{j\mathrm{T}}\mathbf{K}_t^{j\mathrm{T}} - \mathbf{K}_t^{j}\mathbf{H}^{j}\boldsymbol{P}_{t|t-1}^{j(x)} \\
&\qquad + \mathbf{K}_t^{j}\mathbf{H}^{j}\boldsymbol{P}_{t|t-1}^{j(x)}\mathbf{H}^{j\mathrm{T}}\mathbf{K}_t^{j} + \mathbf{K}_t^{j}\mathbf{R}^{j}\mathbf{K}_t^{j\mathrm{T}}\Big] \\
&= \operatorname{tr}\left[\boldsymbol{P}_{t|t-1}^{j(x)}\right] - 2\operatorname{tr}\left[\mathbf{K}_t^{j}\mathbf{H}^{j}\boldsymbol{P}_{t|t-1}^{j(x)}\right] \\
&\qquad + \operatorname{tr}\left[\mathbf{K}_t^{j}\left(\mathbf{H}^{j}\boldsymbol{P}_{t|t-1}^{j(x)}\mathbf{H}^{j\mathrm{T}} + \boldsymbol{R}^{j}\right)\mathbf{K}_t^{j\mathrm{T}}\right] \\
&= \operatorname{tr}\left[\boldsymbol{P}_{t|t-1}^{j(x)}\right] - 2\operatorname{tr}\left[\mathbf{K}_t^{j}\mathbf{H}^{j}\boldsymbol{P}_{t|t-1}^{j(x)}\right] + \operatorname{tr}\left[\mathbf{K}_t^{j}\boldsymbol{P}_t^{j(y)}\mathbf{K}_t^{j\mathrm{T}}\right]
\end{aligned}
\tag{8.11}
$$

The above simplication considers that the trace of a matrix transpose is equal to the trace of the matrix, and considering the symmetry of a covariance matrix, $\boldsymbol{P}_{t|t-1}^{j(x)}\mathbf{H}^{j\mathrm{T}}\mathbf{K}_t^{j\mathrm{T}} \equiv \left(\mathbf{K}_t^{j}\mathbf{H}^{j}\boldsymbol{P}_{t|t-1}^{j(x)}\right)^{\mathrm{T}}$. To minimise the error with respect to gain $\mathbf{K}_t^{j}$, the MSE can be differentiated, and the result set to 0 and solved for $\mathbf{K}_t^{j}$.

$$\begin{aligned}\frac{\mathrm{d}\mathrm{tr}\left[\boldsymbol{P}_{t|t}^{j(x)}\right]}{\mathrm{d}\mathbf{K}_t^j} &= -2\boldsymbol{P}_{t|t-1}^{j(x)}\mathbf{H}^{j\mathrm{T}} + 2\mathbf{K}_t^j\boldsymbol{P}_t^{j(y)} \\ \Rightarrow \mathbf{K}_t^j &= \mathbf{P}_{t|t-1}^{j(x)}\mathbf{H}^{j\mathrm{T}}\left(\boldsymbol{P}_t^{j(y)}\right)^{-1} \\ &= \mathbf{P}_t^{j(xy)}\left(\boldsymbol{P}_t^{j(y)}\right)^{-1}\end{aligned} \tag{8.12}$$

Substituting the calculation of gain $\mathbf{K}_t^j$ back into (8.10) and applying simplification similar to that in (8.11) yields the following correction term for $\boldsymbol{P}_{t|t}^{j(x)}$.

$$\boldsymbol{P}_{t|t}^{j(x)} \approx \boldsymbol{P}_{t|t-1}^{j(x)} - \mathbf{K}_t^j\boldsymbol{P}_{t|t-1}^{j(y)}\mathbf{K}_t^{j\mathrm{T}} \tag{8.13}$$

As with the calculations of all previous error covariance matrices, the assumption for optimality is that the j^{th} component state estimate is truly representative of the true state, in addition to the respective normality assumptions for state and observation models and noise. Note that where $r = 1$, and all other assumptions hold, K_t^j is the optimal gain for correcting the state estimate.

8.2.5 Type-1 Fuzzy Inference

The state of each component SSM is defined in its own domain, but is dependent on some surjective process model $f^j : \mathbb{R}^{n_x} \times U^f \to \mathbb{R}^{n_j}$ from the unified state domain to that of the component state, where U^f defines the control space for the process model. During the combination of component SSMs, in estimating the unified state, there must be some functional defined for each component, $g^j : \mathbb{R}^{n_j} \times U^g \to \mathbb{R}^{n_x}$, that maps from component to unified state domain. Like h^j, g^j need not be injective or surjective, with U^g denoting the control space for the mapping. In the linear-discrete case, it is assumed that the mapping is some system of linear equations $g^j(\boldsymbol{x}_t, u_t) := \mathbf{G}^j\boldsymbol{x}_t + \mathbf{U}_t^g u_t$.

Given a multi-process system with r component SSMs, a (type-1) fuzzy inference system can be defined as a set of rules representing each independent process. Consider the measured observation as the independent input variable for the FIS, such that the antecedent is formed of fuzzy membership functions based on the estimated observations for each respective component process:

$$\text{Rule } j: \begin{cases} \text{IF} & \tilde{\boldsymbol{y}}_t \text{ is } \underset{\sim}{Y}^j \\ \text{THEN} & \boldsymbol{x}_{t|t} \text{ is } g^j\left(\boldsymbol{x}_{t|t}^j, u_t^j\right) \end{cases} \tag{8.14}$$

The function g^j represents the mapping from j^{th} component state domain to unified state domain, and the underlying fuzzy consequent is a Takagi-Sugeno-Kang model of order n_y. The antecedent may be formed of a single fuzzy set defined over the entire observation domain, or as independent fuzzy sets each defined over the space of a single element variable in the observation vector.

$$\text{Rule } j: \begin{cases} \text{IF} & \tilde{y}_t(1) \text{ is } \underset{\sim}{Y}^j(1) \text{ and } \ldots \text{ and } \tilde{y}_t(n_y) \text{ is } \underset{\sim}{Y}^j(n_y) \\ \text{THEN} & \boldsymbol{x}_{t|t} \text{ is } g^j\left(\boldsymbol{x}_{t|t}^j, u_t^j\right) \end{cases} \tag{8.15}$$

Here, $\tilde{y}_t(k)$ and $\underset{\sim}{Y}^j(k)$ denote the k^{th} element variable in the measured observation $\tilde{\boldsymbol{y}}_t$, and the fuzzy set generated based on the corresponding element in the estimated observation for component model j.

Given an estimated observation for the j^{th} SSM, $\boldsymbol{y}_{t|t-1}^j = \mathbb{E}\left[\tilde{y}_t \,|\, \boldsymbol{x}_{t|t-1}\right]$, a fuzzy set can be defined over the observation domain with membership function derived from the pdf of the estimation.

$$\underset{\sim}{Y}^j = \left(\mathbb{R}^{n_y}, \mu_Y^j : \mu_Y^j(\tilde{\boldsymbol{y}}_t) = \mathcal{G}\left(\tilde{\boldsymbol{y}}_t \,|\, \boldsymbol{y}_{t|t-1}^j, \boldsymbol{P}_t^{j(y)}\right)\right) \tag{8.16}$$

Here, the membership function is a kernel associated with the observation estimate. An appropriate choice of kernel could be an non-normalised Gaussian, which is the natural option when all normality assumptions of the underlying processes hold. Alternatively, a triangular kernel may be used to generate a triangular membership, when the element variables of $\tilde{\boldsymbol{y}}$ are independent, with spread dependent on the estimate error covariance. Regardless of the choice of kernel, the uncertainty of the observation estimation, including associated noise, is captured in the observation error covariance, and as such this should be represented in the generated fuzzy membership function. The function definitions for example membership kernels are given in Table 8.1.

The weighting effect of each component SSM on the final unified state estimate can be determined by the firing strength of each rule. The firing strength for rule j^{th}, $\omega^j(\tilde{\boldsymbol{y}}_t)$, is defined as the t-normal of each fuzzy set in the consequence

$$\omega^j(\tilde{\boldsymbol{y}}_t) = \top_{k=1}^{n} \mu_k^j(\tilde{y}_t(k)) \tag{8.17}$$

Where n is the number of fuzzy sets in the antecedent, $n = 1$ in (8.14), and $n = n_y$ in (8.15). $\mu_k^j(\tilde{y}_t(k))$ is the membership grade of the k^{th} element variable in $\tilde{y}_t$ to the k^{th} fuzzy set $\tilde{Y}^j(k)$ in the antecedent of the FIS. When $n = 1$, $\mu_k^j(\tilde{y}_t(k)) \equiv \mu_Y^j(\tilde{\boldsymbol{y}}_t)$.

The final defuzzification of the FIS, given measured observation $\tilde{\boldsymbol{y}}_t$, is thus the sum of the consequent of each component rule, weighted by its normalised firing strength. This crisp output corresponds to the unified state estimate $\boldsymbol{x}_{t|t}$.

$$\boldsymbol{x}_{t|t} = \sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t) g^j\left(\boldsymbol{x}_{t|t}^j, u_t^j\right) \tag{8.18}$$

Where $\hat{\omega}^j(\tilde{\boldsymbol{y}}_t) = \omega^j(\tilde{\boldsymbol{y}}_t) / \sum_{p=1}^{r} \omega^p(\tilde{\boldsymbol{y}}_t)$ is the normalised firing strength for Rule j. The unified estimate error covariation is not as simple to derive, even for linear mapping functionals g^j.

Table 8.1: Kernel definitions for generative fuzzy membership functions in the fuzzy Kalman filter, scaled such that $\mathcal{G}(\boldsymbol{x} = \boldsymbol{m}) = 1$

Kernel (univariate)	$\mathcal{G}\left(x \,\middle|\, m, \sigma^2\right)$
Triangular	$\begin{cases} 1 + \frac{x-m}{\sigma\sqrt{6}} & m - \sigma\sqrt{6} \leq x < m \\ 1 + \frac{m-x}{\sigma\sqrt{6}} & m \geq x \geq m + \sigma\sqrt{6} \\ 0 & \text{otherwise} \end{cases}$
Gaussian	$\exp\left(-\frac{(x-m)^2}{2\sigma^2}\right)$
Kernel (multivariate)	$\mathcal{G}\left(\boldsymbol{x} \,\middle|\, \boldsymbol{m}, \boldsymbol{\Sigma}\right)$
Gaussian	$\exp\left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{m})^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{x}-\boldsymbol{m})\right)$

$$
\begin{aligned}
\boldsymbol{P}_{t|t}^{(x)} &= \mathbb{E}\left[\left(\boldsymbol{x}_{t|t} - \tilde{\boldsymbol{x}}_t\right)\left(\boldsymbol{x}_{t|t} - \tilde{\boldsymbol{x}}_t\right)^{\mathrm{T}}\right] \\
&= \mathbb{E}\left[\left(\sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\left(\mathbf{G}^j \boldsymbol{x}_{t|t}^j + \mathbf{U}_t^g u_t^j\right) - \tilde{\boldsymbol{x}}_t\right)\right. \\
&\qquad \left.\left(\sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\left(\mathbf{G}^j \boldsymbol{x}_{t|t}^j + \mathbf{U}_t^g u_t^j\right) - \tilde{\boldsymbol{x}}_t\right)^{\mathrm{T}}\right] \\
&\approx \mathbb{E}\left[\left(\sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\mathbf{G}^j \boldsymbol{x}_{t|t}^j - \sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\mathbf{G}^j \mathbf{F}^j \tilde{\boldsymbol{x}}_{t-1}\right)\right. \\
&\qquad \left.\left(\sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\mathbf{G}^j \boldsymbol{x}_{t|t}^j - \sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\mathbf{G}^j \mathbf{F}^j \tilde{\boldsymbol{x}}_{t-1}\right)^{\mathrm{T}}\right] \\
&= \mathbb{E}\left[\left(\sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\mathbf{G}^j\left(\boldsymbol{x}_{t|t}^j - \mathbf{F}^j \tilde{\boldsymbol{x}}_{t-1}\right)\right)\right. \\
&\qquad \left.\left(\sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\mathbf{G}^j\left(\boldsymbol{x}_{t|t}^j - \mathbf{F}^j \tilde{\boldsymbol{x}}_{t-1}\right)\right)^{\mathrm{T}}\right] \\
&= \sum_{j=1}^{r}\sum_{p=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\hat{\omega}^p(\tilde{\boldsymbol{y}}_t)\mathbf{G}^j \mathbb{E}\left[\left(\boldsymbol{x}_{t|t}^j - \mathbf{F}^j \tilde{\boldsymbol{x}}_{t-1}\right)\left(\boldsymbol{x}_{t|t}^p - \mathbf{F}^p \tilde{\boldsymbol{x}}_{t-1}\right)^{\mathrm{T}}\right] \mathbf{G}^{p\mathrm{T}}
\end{aligned}
\tag{8.19}
$$

It can be seen here that the non-diagonal elements in the sum of $P_{t|t}^{(x)}$ are unknown, and these elements refer to the cross-correlation between component rules. If the individual component processes are considered independent, then only the leading diagonal of the resulting expectation matrix, such that $p = j$, are required. Therefore, the unified covariance for a set of linear models, can be approximated as follows:

$$
\boldsymbol{P}_{t|t}^{(x)} \approx \sum_{j=1}^{r}\left(\hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\right)^2 \mathbf{G}^j \boldsymbol{P}_{t|t}^{j(x)} \mathbf{G}^{j\mathrm{T}} \tag{8.20}
$$

As a recursive filter over a time-series, the updated unified state estimate, $\boldsymbol{x}_{t|t}$, forms the input for the component process models over the next time step. The overall control system for the proposed fuzzy Kalman filter is illustrated in Figure 8.1, which shows the combination of recursive prediction correction with the generative fuzzy inference system that forms the correction and unification of the component processes in a multi-process SSM. Here, measured observation vector, $\tilde{\boldsymbol{y}}_t$ represents the crisp input of the FIS, with the unified state estimate at time t, $\boldsymbol{x}_{t|t}$ forming the crisp output.

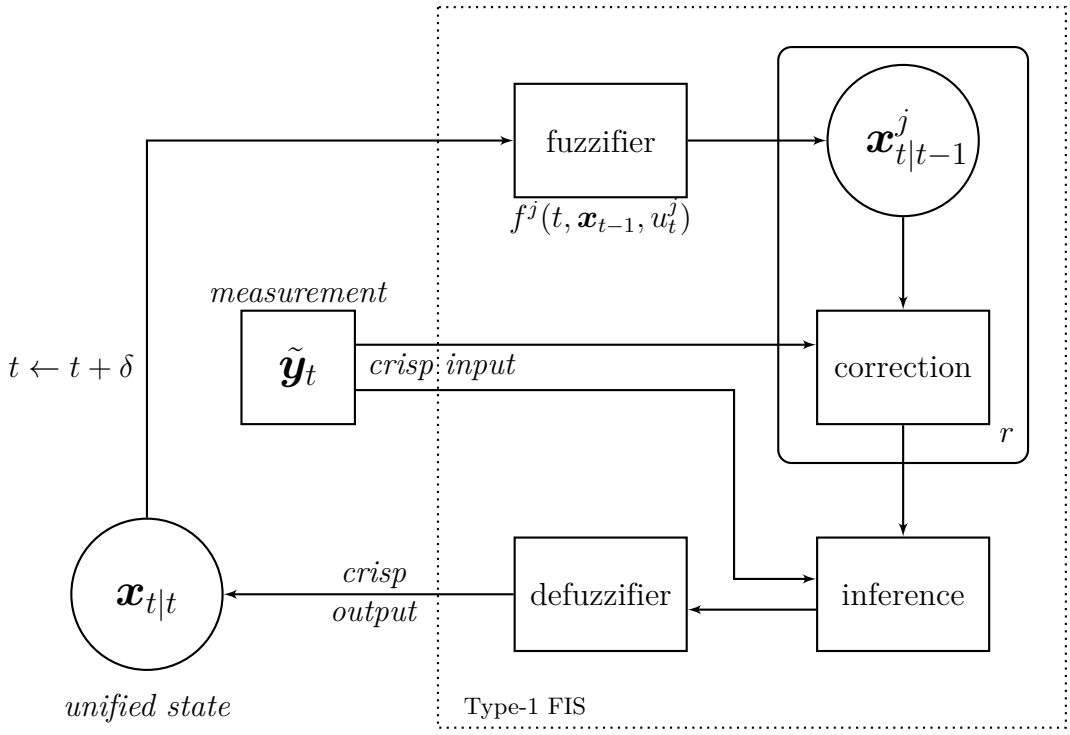


Figure 8.1: Type-1 fuzzy Kalman filter process, showing fuzzy inference and correction cycle over increments of t, with regular measurements

8.3 Illustrative Example

As an illustration of the proposed type-1 fuzzy Kalman filter, T1-FKF, consider a multi-process system of some tracked object moving on a coordinate plane $\vec{x} \times \vec{y}$. The object is observed by its positional coordinate at a given time $\tilde{\boldsymbol{y}}_t = (\vec{x}_t, \vec{y}_t)^{\mathrm{T}}$. The object is moving in three possible directions, to different degrees, in either a *forward* direction ($+\vec{x}$), or it is moving *left* or *right* ($+\vec{y}$ and $-\vec{y}$) respectively. Each direction is modelled as a component SSM, as defined in Table 8.2, where velocity control parameters d_x and $\pm d_y$ were set as 0.5 and ± 0.25 respectively. The unified state is defined as the coordinate positions and the velocity values in each direction, $\tilde{\boldsymbol{x}}_t = (\vec{x}_t, \vec{y}_t, v_{\vec{x}}, v^+_{\vec{y}}, v^-_{\vec{y}})$, where $v^+_{\vec{y}}$ is the velocity in the *left* direction, and $v^-_{\vec{y}}$ is the velocity in the *right direction*. The switching between models is set such that the *forward* model is constant, with *left* and *right* alternating along a mirrored period-2 sinusoid. A linear-discrete Kalman filter (LKF) model is defined in Table 8.3 for comparison.

The true state modelled by the described system is illustrated in Figure 8.2, along

with the estimation by linear-discrete Kalman filter and the proposed T1-FKF. The firing strength of each component process rule is also plotted showing that the embedded fuzzy inference system was generally accurate in approximating the most dominant process, even with noisy measurements. A sample ruleset is plotted in terms of the respective membership functions of each set on each element variable of $\tilde{\boldsymbol{y}}_t$ for $t = 50$, in Figure 8.3 As observed in Figure 8.2, the true

Table 8.2: Model setup for type-1 fuzzy Kalman filter example system

Process	$\boldsymbol{x}_t^j$	$\mathbf{F}^j$	$\boldsymbol{Q}^j$
forward	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{x}} \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & d_x & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$	$0.1 \cdot \mathbf{I}_3$
left	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{y}}^+ \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & d_y & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$	$0.1 \cdot \mathbf{I}_3$
right	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{y}}^- \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & -d_y \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$0.1 \cdot \mathbf{I}_3$
Observation	$\mathbf{H}^j$		$\boldsymbol{R}^j$
forward/ left / right	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$		$0.25 \cdot \mathbf{I}_2$
Mapping	$\mathbf{G}^j$	$\mathbf{U}^g$	u_t^j
forward	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$\boldsymbol{x}_{t-1}$
left	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$\boldsymbol{x}_{t-1}$
right	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$	$\boldsymbol{x}_{t-1}$
FIS	t-norm ($\top_\cap$)		Kernel ($\mathcal{G}$)
	minimum		Gaussian (univariate)

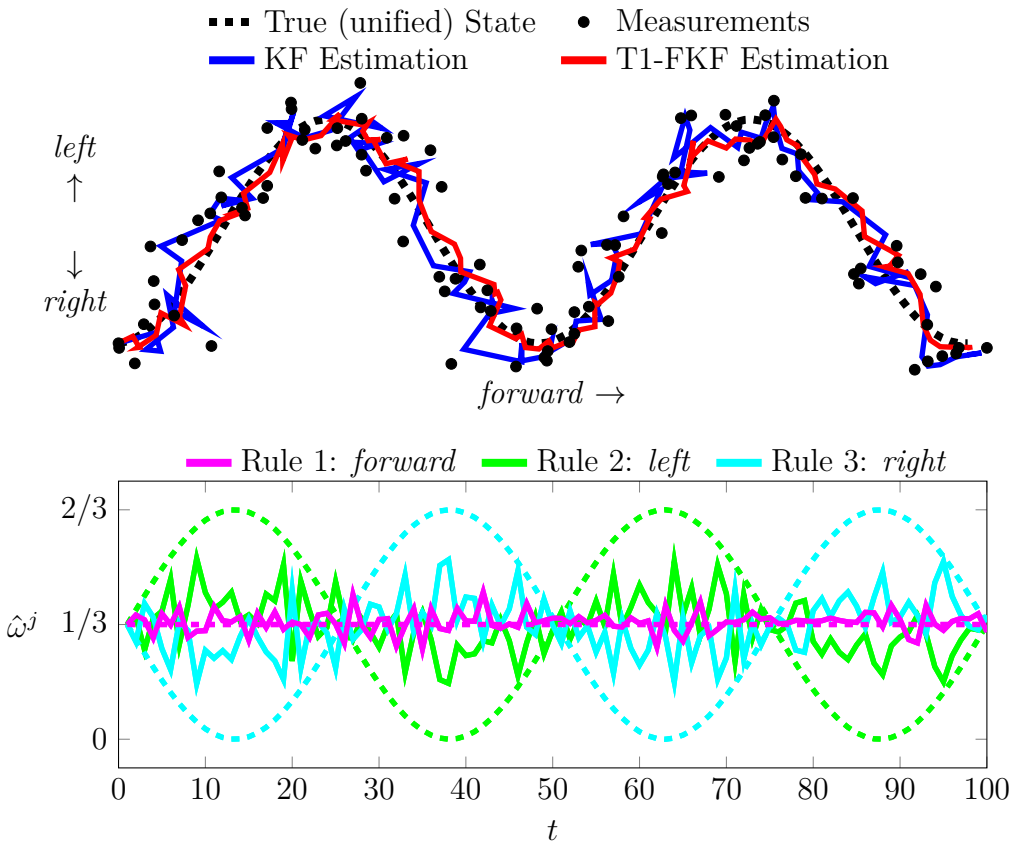


Figure 8.2: Multi-process system with three component processes (*forward*, *left*, *right*) with sinusoidal switching, estimated with linear-discrete Kalman filter and type-1 fuzzy Kalman filter (top), with firing strengths of each component rule plotted against true system switching (dashed)

state processes are in equilibrium. This is successfully modelled, with calculated firing strengths found to be $\hat{\boldsymbol{\omega}} = \{0.3520, 0.3163, 0.3316\}$ when using the minimum

Table 8.3: Model setup for linear-discrete Kalman filter for comparison against T1-FKF illustrative example

Process	$\boldsymbol{x}_t$	**F**	$\boldsymbol{Q}$
	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{x}} \\ v_{\vec{y}}^{+} \\ v_{\vec{y}}^{-} \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & d_x & 0 & 0 \\ 0 & 1 & 0 & d_y & -d_y \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$0.1 \cdot \mathbf{I}_5$
Observation	**H**		$\boldsymbol{R}$
	$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix}$		$0.25 \cdot \mathbf{I}_2$

Table 8.4: Root mean-square error for estimation of element variables in the unified state for example multi-process system by linear-discrete Kalman filter and type-1 fuzzy Kalman filter

RMSE	$\tilde{x}(1)$	$\tilde{x}(2)$	$\tilde{x}(3)$	$\tilde{x}(4)$	$\tilde{x}(5)$
LKF	0.1801	0.1962	0.8195	0.2439	0.2439
T1-FKF	0.1071	0.1723	0.0284	0.0106	0.0173

t-norm, and $\hat{\boldsymbol{\omega}} = \{0.3666, 0.3101, 0.3233\}$ with the product t-norm. Also shown in Figure 8.3 is that the membership functions for both the *left* and *right* for the first element variable of $\boldsymbol{y}_t$. This is true always, since the process models in the respective component SSMs simply propagate $y(1) = \vec{x}$ by the identity, such that the predicted value will always be the same.

It is visibly clear that the T1-FKF can handle the multi-process system with better approximation of the position state. The root mean-squared error (RMSE), shown in Table 8.4 for each parameter is consistently lower for the T1-FKF estimate. This is apparent in Figure 8.4, which plots the estimated velocity parameters for each respective component model against the true state. The true state for each of the parameters is constant, as each is a fixed velocity model. While the correction phase of the LKF is powerful, when estimating the underlying parameters, it suffers, due to the inability to seperate the effects of each variable. The pattern of the firing strengths is partly observable in the estimates of parameters of $v_{\vec{y}}^{+}$ and $v_{\vec{y}}^{-}$, but the reliance on a single linear process model means the system is poorly represented, and therefore the estimations are low quality. In contrast, the T1-FKF is successful in approximating these parameters in what is effectively a non-linear SSM, when considering only the unified state, due to its ability to estimate the separate effects of each component model and weight them according to measured data.

8.4 Unscented Fuzzy Kalman Filter

8.4.1 Formulation

In many real-world scenarios, from target tracking to parameter estimation problems, the embedded models that drive the system are often not linear. It is true, also, that many scenarios the underlying probability distributions are only be approximately Gaussian. Given a multi-process SSM, any of the component process f^j or observation h^j models may be non-linear, and indeed if the component process is non-linear, the respective domain mapping g^j may also be. As in the

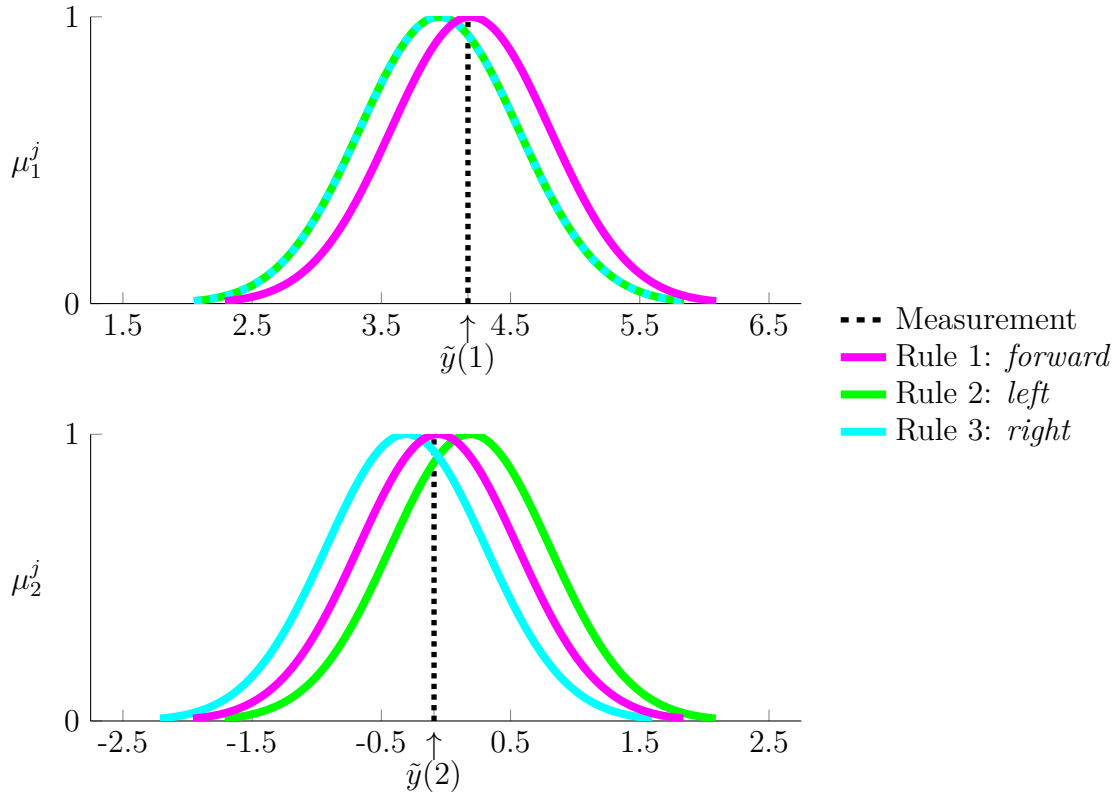


Figure 8.3: Membership functions for each rule at $t = 50$, where the true case should result in a firing strength of $\hat{\omega}_1 = \hat{\omega}_2 = \hat{\omega}_3 = \frac{1}{3}$. The crisp input is highlighted by the dashed line. Calculated firing strengths were $\hat{\boldsymbol{\omega}} = \{0.3520, 0.3163, 0.3316\}$ for the minimum t-norm, and $\hat{\boldsymbol{\omega}} = \{0.3666, 0.3101, 0.3233\}$ for the product t-norm

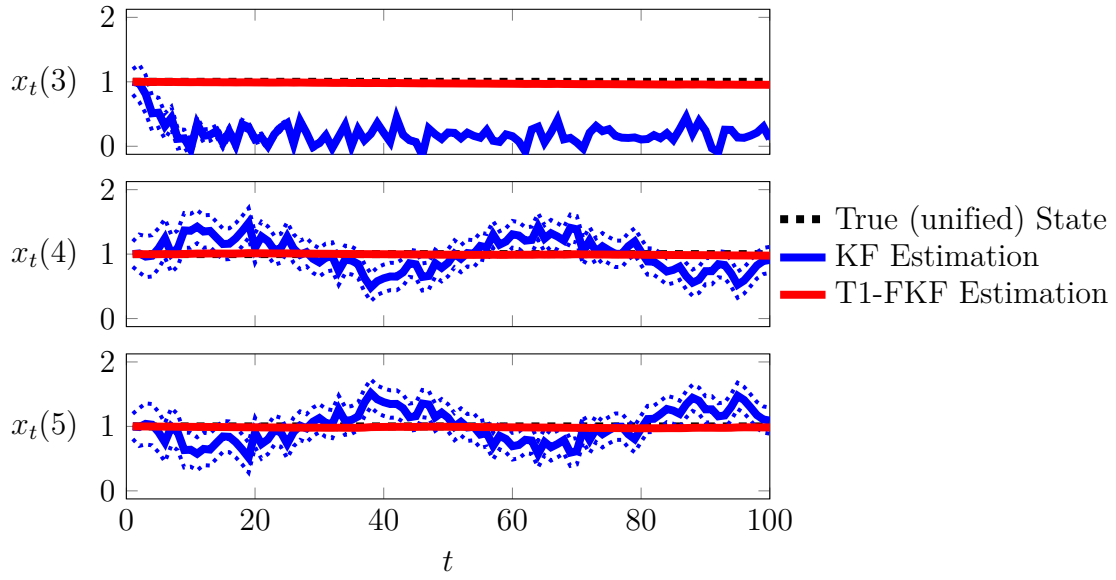


Figure 8.4: Estimates by linear-discrete Kalman filter and type-1 fuzzy Kalman filter for state velocity parameters for each component motion model, *forward* ($x_t(3)$), *left* ($x_t(4)$), and *right* ($x_t(5)$). The dotted lines represents the 95% confidence interval for the estimates

Algorithm 7 Unscented Type-1 Fuzzy Kalman Filter

procedure UT1-FKF($\boldsymbol{x}_0$, $\boldsymbol{P}_0^{(x)}$; $[f^j]_{j=1}^r$, $[\boldsymbol{Q}_t]_{t=1}^T$, $[h^j]_{j=1}^r$, $[\boldsymbol{R}_t]_{t=1}^T$, $[g^j]_{j=1}^r$)
  $\boldsymbol{x}_{t-1}, \boldsymbol{P}_{t-1}^{(x)} \leftarrow \boldsymbol{x}_0, \boldsymbol{P}_0^{(x)}$
  for $t \in 1, \ldots, T$ **do**
    for $j \in 1, \ldots, r$ **do**
      $\boldsymbol{x}_{t|t-1}^j, \boldsymbol{P}_{t|t-1}^{j(x)} \leftarrow$ **with** f^j: PREDICT($\boldsymbol{x}_{t-1}$, u_t, $\boldsymbol{P}_{t-1}^{(x)}$)
      $\boldsymbol{x}_{t|t}^j, \boldsymbol{P}_{t|t}^{j(x)} \leftarrow$ **with** h^j: UPDATE($\boldsymbol{x}_{t|t-1}^j, \boldsymbol{P}_{t|t-1}^{(x)}, \tilde{\boldsymbol{y}}_t$)
      $\boldsymbol{y}_t^j, \boldsymbol{P}_t^{j(y)} \leftarrow$ **with** h^j: OBSERVE($\boldsymbol{x}_{t|t-1}^j, \boldsymbol{P}_{t|t-1}^{(x)}$)
      $\omega^j \leftarrow \min\{\mathcal{G}(\tilde{y}_t(k)|y_t^j(k), P_t^{j(y)}(k,k)) \,|\, k = 1, \ldots, n_y\}$
      $_, \boldsymbol{\gamma}^j, \boldsymbol{w} \leftarrow \mathrm{UT}(g^j(\cdot), \boldsymbol{x}_{t|t}^j, \boldsymbol{P}_{t|t}^{j(x)})$
      $\boldsymbol{x}_g^j, \boldsymbol{P}_g^j \leftarrow \mathrm{UT}^{-1}(\boldsymbol{\gamma}^j, \boldsymbol{\gamma}^j, \boldsymbol{w})$
    end for
    $\boldsymbol{x}_{t|t} \leftarrow \sum_{j=1}^r \omega^j \boldsymbol{x}_g^j$
    $\boldsymbol{P}_{t|t}^{(x)} \leftarrow \sum_{j=1}^r (\omega^j)^2 \boldsymbol{P}_g^j$
  end for
  return $\left[(\boldsymbol{x}_t, \boldsymbol{P}_t^{(x)})\right]_{t=1}^T$
end procedure

single-process filtering problem, it is possible to represent the distributions using deterministically sampled *sigma points*, which can be individually propagated through the respected transformation by means of the *unscented transform* (see Section 3.4.3).

As with the unscented Kalman filter, the (unified) state estimate in the multi-process SSM may be represented by its sigma points. To recap, the prior state estimate sigma points $\chi_{i;t-1}$ can be defined as in (3.18).

$$\begin{aligned} \chi_{0;t-1} &= \boldsymbol{x}_{t-1} \\ \chi_{i;t-1} &= \begin{cases} \chi_{0;t-1} + \left[\sqrt{(n_x + \lambda)\boldsymbol{P}_{t-1}^{(x)}}\right]_i & i = 1, \ldots, n_x \\ \chi_{0;t-1} - \left[\sqrt{(n_x + \lambda)\boldsymbol{P}_{t-1}^{(x)}}\right]_{i-n_x} & i = n_x + 1, \ldots, 2n_x \end{cases} \end{aligned} \tag{8.21}$$

Given a non-linear component process model f^j, the sigma points can be propagated through and an recombined to calculate the *a priori* state estimate and its error covariance.

$$
\begin{aligned}
\chi^{j}_{i;t|t-1} &= f^{j}(t, \chi_{i;t-1}, u^{j}_{t}) \\
\boldsymbol{x}^{j}_{t|t-1} &= \mathbb{E}\Big[f^{j}(t, \tilde{\boldsymbol{x}}_{t-1}, u^{j}_{t})\Big] \\
&\approx \sum_{i=0}^{2n_x} w^{(x)}_{i} \chi^{j}_{i;t|t-1} \\
\boldsymbol{P}^{j(x)}_{t|t-1} &= \mathbb{E}\Big[\Big(\boldsymbol{x}^{j}_{t|t-1} - f^{j}(t, \tilde{\boldsymbol{x}}_{t-1}, u^{j}_{t})\Big) \\
&\qquad \Big(\boldsymbol{x}^{j}_{t|t-1} - f^{j}(t, \tilde{\boldsymbol{x}}_{t-1}, u^{j}_{t})\Big)^{\mathrm{T}}\Big] \\
&\approx \sum_{i=0}^{2n_x} w^{(p)}_{i} \left(\boldsymbol{x}^{j}_{t|t-1} - \chi^{j}_{t|t-1}\right) \left(\boldsymbol{x}^{j}_{t|t-1} - \chi^{j}_{t|t-1}\right)^{\mathrm{T}} + \boldsymbol{Q}^{j}_{t}
\end{aligned}
\tag{8.22}
$$

Here, $w^{(x)}_{i}$ and $w^{(p)}_{i}$ represent the respective expectation and covariance recombination weight for the i^{th} sigma point, defined as in (3.17). Control parameters, α, β and κ are defined based on the problem, with $\beta = 2$ chosen when underlying distributions can be considered exactly Gaussian. α and κ are spread parameters, with the sigma points generated with dependent parameter $\lambda = \alpha^2(n_x + \kappa) - n_x$, and n_x is the dimension of $\boldsymbol{x}^{j}$.

The processed sigma points for the *a priori* state estimate for each component process can continue to be propagated through the steps of UKF prediction. Sigma points represented the distribution of the component observation pdf can be found by propagating the *a priori* sigma points, $\chi^{j}_{i;t|t-1}$ through the respective observation functions.

$$
\begin{aligned}
\gamma^{j}_{i;t|t-1} &= h^{j}(t, \chi^{j}_{i;t|t-1}, u^{j}_{t}) \\
\boldsymbol{y}^{j}_{t|t-1} &= \mathbb{E}\Big[h^{j}\Big(t, f^{j}(t, \tilde{\boldsymbol{x}}_{t-1}, u_{t}), u_{t}\Big)\Big] \\
&\approx \sum_{i=0}^{2n_x} w^{(x)}_{i} \gamma^{j}_{i;t|t-1}
\end{aligned}
\tag{8.23}
$$

Indeed, the observation estimate error covariance and the joint state and observation error covariance can also be approximated through weighted sums of sigma points.

$$
\begin{aligned}
\boldsymbol{P}_t^{j(xy)} &= \mathbb{E}\Big[\Big(\boldsymbol{x}_{t|t-1}^j - f^j(t, \tilde{\boldsymbol{x}}_{t-1}, u_t^j)\Big) \\
&\qquad \Big(\boldsymbol{y}_{t|t-1}^j - h^j\Big(t, f^j(t, \tilde{\boldsymbol{x}}_{t-1}, u_t), u_t\Big)\Big)^{\mathrm{T}}\Big] \\
&\approx \sum_{i=0}^{2n_x} w_i^{(p)} \left(\boldsymbol{x}_{t|t-1}^j - \chi_{t|t-1}^j\right) \left(\boldsymbol{y}_{t|t-1}^j - \gamma_{t|t-1}^j\right)^{\mathrm{T}} \\
\boldsymbol{P}_t^{j(y)} &= \mathbb{E}\Big[\Big(\boldsymbol{y}_{t|t-1}^j - h^j\Big(t, f^j(t, \tilde{\boldsymbol{x}}_{t-1}, u_t), u_t\Big)\Big) \\
&\qquad \Big(\boldsymbol{y}_{t|t-1}^j - h^j\Big(t, f^j(t, \tilde{\boldsymbol{x}}_{t-1}, u_t), u_t\Big)\Big)^{\mathrm{T}}\Big] \\
&\approx \sum_{i=0}^{2n_x} w_i^{(p)} \left(\boldsymbol{y}_{t|t-1}^j - \gamma_{t|t-1}^j\right) \left(\boldsymbol{y}_{t|t-1}^j - \gamma_{t|t-1}^j\right)^{\mathrm{T}} + \boldsymbol{R}_t^j
\end{aligned}
\tag{8.24}
$$

The correction terms for the UKF in component SSMs is the same as with that defined in the linear model-based T1-FKF, specifically (8.9) and (8.13), for the *a posteriori* component estimates $\boldsymbol{x}_{t|t}^j$ and their estimate error covariance, $\boldsymbol{P}_{t|t}^{j(x)}$, respectively. The correction gain $\boldsymbol{K}_t^j$ is suboptimal, but can be approximated as in the T1-FKF.

$$
\mathbf{K}_t^j = \boldsymbol{P}_t^{j(xy)} \left(\boldsymbol{P}_t^{j(y)}\right)^{-1} \tag{8.25}
$$

8.4.2 Unscented Fuzzy Inference

In a non-linear multi-process SSM, where at least one component process model is non-linear, it follows that the mapping from component state domain to unified state domain, $g^j(\boldsymbol{x}_t^j, u_t^j)$ may also be non-linear. Given the estimate of each component process state and observation of a multi-process SSM, in the form of sigma points and weights, the mapped state estimate may be approximated by the unscented transform using a set of corrected sigma points.

$$
\chi_{i;t|t}^j = \chi_{i;t|t-1}^j + \boldsymbol{K}_t^j\left(\tilde{\boldsymbol{y}}_t - \gamma_{i;t|t-1}^j\right) \tag{8.26}
$$

Note that this calculation of corrected sigma points is equivalent to reinitialising a set of sigma points based on corrected component state estimate $\boldsymbol{x}_{t|t}^j$. Given the sigma points, the underlying inference system in the unscented type-1 fuzzy Kalman filter (UT1-FKF) can be adapted from (8.15) and incorporated with the unscented transform over domain mapping g^j:

$$\text{Rule } j: \begin{cases} \text{IF} & \tilde{y}_t(1) \text{ is } \underset{\sim}{Y}^j(1) \text{ and } \ldots \text{ and } \tilde{y}_t(n_y) \text{ is } \underset{\sim}{Y}^j(n_y) \\ \text{THEN} & \boldsymbol{x}_{t|t} \text{ is } \sum_{i=0}^{2n_x} w_i^{(x)} g^j\left(\chi_{i;t|t}^j, u_t^j\right) \end{cases} \tag{8.27}$$

Generating the component fuzzy sets for each element variable in the same manner as with the linear T1-FKF, the unified state is calculated through defuzzification of the rule consequents.

$$\boldsymbol{x}_{t|t} = \sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t) \sum_{i=0}^{2n_x} w_i^{(x)} g^j\left(\chi_{i;t|t}^j, u_t^j\right) \tag{8.28}$$

Here, $\hat{\omega}^j(\tilde{\boldsymbol{y}})$ denote the normalised firing strength of each component rule, as defined in (8.17). Here, $w_i^{(x)}$, denote the UT weight of each sigma point. Using this approximate of the unified state, it is also possible to approximate the unified state estimate error covariance.

$$\begin{aligned} \boldsymbol{P}_{t|t}^{(x)} &= \mathbb{E}\left[\left(\boldsymbol{x}_{t|t} - \tilde{\boldsymbol{x}}_t\right)\left(\boldsymbol{x}_{t|t} - \tilde{\boldsymbol{x}}_t\right)^{\mathrm{T}}\right] \\ &= \mathbb{E}\left[\left(\boldsymbol{x}_{t|t} - \sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t) g^j\left(f^j(t, \tilde{\boldsymbol{x}}_{t-1}, u_t), u_t\right)\right)\right. \\ &\qquad \left.\left(\boldsymbol{x}_{t|t} - \sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t) g^j\left(f^j(t, \tilde{\boldsymbol{x}}_{t-1}, u_t), u_t\right)\right)^{\mathrm{T}}\right] \\ &\approx \sum_{j=1}^{r} \left(\hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\right)^2 \sum_{i=0}^{2n_x} w_i^{(p)} \left(\boldsymbol{x}_{t|t} - g^j\left(\chi_{i;t|t}^j, u_t^j\right)\right)\left(\boldsymbol{x}_{t|t} - g^j\left(\chi_{i;t|t}^j, u_t^j\right)\right)^{\mathrm{T}} \end{aligned} \tag{8.29}$$

8.5 Interval Type-2 Fuzzy Kalman Filter

8.5.1 Capturing the Uncertainty of State Estimation

In the proposed (type-1) fuzzy Kalman filter, the uncertainty surrounding which component process is in effect has been handled solely by the posterior distribution over which the observation estimate is defined. In an optimal filter, the calculation of the observation posterior distribution is representative of the underlying system processes, and therefore distribution about $\boldsymbol{y}_{t|t-1}$ captures the uncertainty of the

system in its entirety. However, given the underlying assumption that the true state, and therefore its observable value, is affected by multiple possible processes, each to a different degree, the estimated observation for a given component process, and therefore its fuzzy membership, may not capture the uncertainty in the predictive distribution of each component set. Therefore, using a type-1 fuzzy set to denote the antecedent of each rule means the generative fuzzy inferfence system does not take into account the uncertainty of the state estimation.

The notion of a type-2 fuzzy set that may represent both the membership over the observation domain, i.e. the current type-1 definition, and the membership of each proposed observation to the total distribution over which the state may lie, is a complicated combination of two continuous (and infinite) pdfs. The proposed solution is to represent the uncertainty of the state estimate as an interval, and therefore define the fuzzy sets of the antecedents as interval type-2 fuzzy sets. The interval should capture a certain degree of credibility of the estimate such that it is near representative of the whole distribution of state estimates.

8.5.2 Credible Region

The credible region of a posterior distribution describes the interval over which the population parameters are considered likely to fall, to a given degree of credibility, or confidence. Common means of identifying the credible intervals include the regions of highest density, such that some percentage of the distribution density is contained within; and the regions over some quantile interval such that some percentage of possible credible parameter values are contained. Given that with Kalman and Kalman-derivative filters, including the fuzzy filter proposed, the underlying distributions of posterior and predictive distributions are assumed to be a unimodal and normally distributed, it follows that the region of highest density to a degree of credibility is equivalent to the quantile interval of the same degree. It is also assumed that the distributions are symmetric, and as such the interval on either side of the mean is equidistant from the mean. In these conditions, the credible region is analagous to the confidence interval in frequentist statistics. The confidence interval denotes the confidence that the true parameter falls within the interval to a given degree of confidence.

The credible region for the approximation of a univariate, symmetric Gaussian probability density function over x_t can thus be calculated

$$\mathrm{CR}(x_t) = \mathbb{E}[x_t] \pm z^* \mathrm{Var}[x_t]^{\frac{1}{2}} \tag{8.30}$$

The critical value, z^*, controls the level of confidence that the true value is contained within the interval. For 95% confidence, $z^* = 1.96$.

Approximating the credible region over the state estimate by capturing the % confidence interval about the mean, means that some upper and lower bounded estimates of the observation can be approximated. Assuming a non-linear system, and the unscented transform used to calculate the *a priori* state estimate of each component rule, the interval over which the observation distribution may lie can be approximated using the transformed sigma points offset by the confidence interval to calculate the respective upper and lower state sigma points:

$$\begin{aligned} \underline{\chi}^j_{i;t|t-1} &= \chi^j_{i;t|t-1} - z^* \mathrm{diag}\left(\sqrt{\boldsymbol{P}^{j(x)}_{t|t-1}}\right) \\ \overline{\chi}^j_{i;t|t-1} &= \chi^j_{i;t|t-1} + z^* \mathrm{diag}\left(\sqrt{\boldsymbol{P}^{j(x)}_{t|t-1}}\right) \end{aligned} \tag{8.31}$$

The approximations of prediction sigma points and estimate error are as defined in (8.22). The interval of each individual state estimate variable is calculated independently, such that only the diagonal elements of $\sqrt{\boldsymbol{P}^{j(x)}_{t|t-1}}$ are used, representing the standard deviation of each element variable in the state. Some numerically stable matrix square root is used, such as Cholesky decomposition, as in the generation of sigma points for the UT.

Given the sigma interval bounds for the state estimate, an interval distribution for the observation can be generated by weighting the sigma points of each bound.

$$\begin{aligned} \underline{\gamma}^j_{i;t|t-1} &= h\big(t, \underline{\chi}^j_{i;t|t-1}, u_t\big) \\ \underline{\boldsymbol{y}}^j_{t|t-1} &= \sum_{i=0}^{2n_x} w^{(x)}_i \underline{\gamma}^j_{i;t|t-1} \\ \underline{\boldsymbol{P}}^{j(y)}_t &= \sum_{i=0}^{2n_x} w^{(p)}_i \left(\underline{\boldsymbol{y}}^j_{t|t-1} - \underline{\gamma}^j_{i;t|t-1}\right)\left(\underline{\boldsymbol{y}}^j_{t|t-1} - \underline{\gamma}^j_{i;t|t-1}\right)^{\mathrm{T}} + \boldsymbol{R}^j_t \end{aligned} \tag{8.32}$$

$$
\begin{aligned}
\overline{\gamma}^{j}_{i;t|t-1} &= h\Big(t, \overline{\chi}^{j}_{i;t|t-1}, u_t\Big) \\
\overline{\boldsymbol{y}}^{j}_{t|t-1} &= \sum_{i=0}^{2n_x} w^{(x)} \overline{\gamma}^{j}_{i;t|t-1} \\
\overline{\boldsymbol{P}}^{j(y)}_{t} &= \sum_{i=0}^{2n_x} w^{(p)}_{i} \left(\overline{\boldsymbol{y}}^{j}_{t|t-1} - \overline{\gamma}^{j}_{i;t|t-1}\right) \left(\overline{\boldsymbol{y}}^{j}_{t|t-1} - \overline{\gamma}^{j}_{i;t|t-1}\right)^{\mathrm{T}} + \boldsymbol{R}^{j}_{t}
\end{aligned}
\tag{8.33}
$$

The upper and lower intervals can be for the observation estimate can now be used to form an interval type-2 fuzzy set, with membership function uncertain mean and standard deviation. Each membership element variable is defined over its own interval, so it follows that the generative FIS described in (8.27) may defined where the antecedent sets $\underset{\sim}{Y}^{j}(k)$ are IT2 fuzzy sets.

8.5.3 Membership Functions

The observation estimate defined over an interval can be represented by the respective (Gaussian) distributions at the bounds of the interval. For each element variable in the measurement vector, an IT2 fuzzy membership function can be defined over the upper and lower distributions. Consider the membership function defined for a generated type-1 set in (8.16), the kernel function may be applied to the boundary estimates to create the upper and lower limits for constructing the IT2 membership function.

For each element variable, sort the boundaries estimates and assign them such that $\overline{y}^{j}_{t|t-1}(k) \geq \underline{y}^{j}_{t|t-1}(k)$ where now $\overline{y}^{j}_{t|t-1}(k)$ denotes the maximum of the two element variable estimates, and $\underline{y}^{j}_{t|t-1}(k)$ the minimum. The corresponding variance for the upper and lower element variable bounds are then denoted $\overline{P}^{j(y)}_{t}(k,k)$ and $\underline{P}^{j(y)}_{t}(k,k)$ respectively. Where the mean values of the estimate are equal, the bounds should be ordered such that $\overline{P}^{j(y)}_{t}(k,k) \geq \underline{P}^{j(y)}_{t}(k,k)$. Note that if both the expected observation estimate and their covariances are equal, the fuzzy set reduces to a type-1 set, with a single membership function.

The kernel function applied over the bounds is defined, for notational simplicity, as follows:

$$
\begin{aligned}
\underline{\mathcal{G}}^{j}_{k;t}(\tilde{y}_t(k)) &= \mathcal{G}\Big(\tilde{y}_t(k) \,|\, \underline{y}^{j}_{t|t-1}(k), \underline{P}^{j(y)}_{t}(k,k)\Big) \\
\overline{\mathcal{G}}^{j}_{k;t}(\tilde{y}_t(k)) &= \mathcal{G}\Big(\tilde{y}_t(k) \,|\, \overline{y}^{j}_{t|t-1}(k), \overline{P}^{j(y)}_{t}(k,k)\Big)
\end{aligned}
\tag{8.34}
$$

As with the T1-FKF, the kernel function can be selected based on the problem at hand, for example triangular or Gaussian. Given that the underlying distributions are assumed to be normally distributed, consider, for example, the Gaussian kernel applied over the two boundary estimates of the observation interval. The upper and lower membership functions of the fuzzy set, $\overline{\mu}_t^j$ and $\underline{\mu}_t^j$ respectively are defined as the maximum and minimum of the footprint of uncertainty over the domain of the fuzzy set. In the case of Gaussian kernels, this is defined as the respective lower and upper boundary functions, each corresponding to both the upper and lower membership functions depending on some switching roots.

The switching roots are defined where the two boundary functions intersect, i.e. where $\underline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) = \overline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k))$. Equating the two functions returns a quadratic, that needs to be solved for $\tilde{y} = r$, such that $ar^2 + br + c = 0$, where r denotes the points of intersection and

$$\begin{aligned} a &= \frac{1}{2\underline{P}_t^{(y)}(k,k)} - \frac{1}{2\overline{P}_t^{(y)}(k,k)} \\ b &= \frac{\overline{y}_{t|t-1}^j(k)}{\overline{P}_t^{(y)}(k,k)} - \frac{\underline{y}_{t|t-1}^j(k)}{\underline{P}_t^{(y)}(k,k)} \\ c &= \frac{\left(\underline{y}_{t|t-1}^j(k)\right)^2}{\underline{P}_t^{(y)}(k,k)} - \frac{\left(\overline{y}_{t|t-1}^j(k)\right)^2}{\overline{P}_t^{(y)}(k,k)} \end{aligned} \tag{8.35}$$

Where the expected values are uncertain, but variances are equal, there will be single intersection and the function to solve becomes linear (as $a = 0$). The root will exist between the boundary expected values, $\underline{r} = r$. However, where the variance and means both differ, there will be two roots, some *interior* root lying between the expected values, and one lying either side of the upper or lower bound of expected values. In this case, the *interior* root is denoted by $\underline{r}$ and the exterior by $\overline{r}$. These values are referred to as switching roots as in the definition of upper and lower membership, the usage of boundary kernels in the respective definitions switches between upper and lower at these points. For example, the IT2 membership function over a Gaussian over uncertain mean and variance is defined:

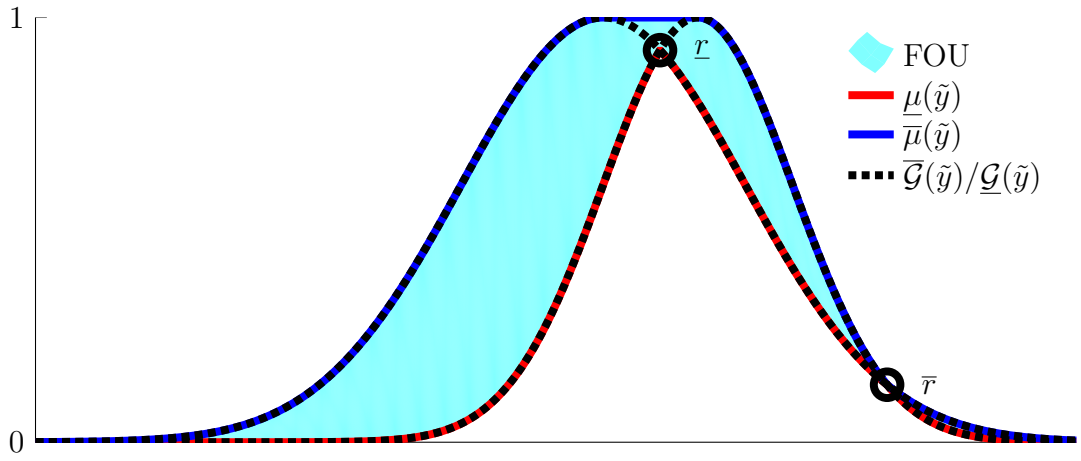


Figure 8.5: Sample generated Gaussian membership functions for IT2 fuzzy set, with upper and lower membership functions generated by (8.36) using Gaussian kernels on the propagated confidence interval. Interior and exterior switching roots are highlighted with black discs

$$\overline{\mu}_k^j(\tilde{y}_t(k)) = \begin{cases} \overline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \overline{r}_{k;t}^j < \underline{y}_{t|t-1}^j(k) \;\&\; \tilde{y}_t(k) < \overline{r}_{k;t}^j \\ \underline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \tilde{y}_t(k) < \underline{y}_{t|t-1}^j(k) \\ 1 & \underline{y}_{t|t-1}^j(k) \leq \tilde{y}_t(k) \leq \overline{y}_{t|t-1}^j(k) \\ \overline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \tilde{y}_t(k) < \overline{y}_{t|t-1}^j(k) \\ \underline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \overline{r}_{k;t}^j > \overline{y}_{t|t-1}^j(k) \;\&\; \tilde{y}_t(k) > \overline{r}_{k;t}^j \end{cases} \tag{8.36}$$

$$\underline{\mu}_k^j(\tilde{y}_t(k)) = \begin{cases} \underline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \overline{r}_{k;t}^j < \underline{y}_{t|t-1}^j(k) \;\&\; \tilde{y}_t(k) < \overline{r}_{k;t}^j \\ \overline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \tilde{y}_t(k) \leq \underline{r}_{k;t}^j \\ \underline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \tilde{y}_t(k) > \underline{r}_{k;t}^j \\ \overline{\mathcal{G}}_{k;t}^j(\tilde{y}_t(k)) & \overline{r}_{k;t}^j > \overline{y}_{t|t-1}^j(k) \;\&\; \tilde{y}_t(k) > \overline{r}_{k;t}^j \end{cases}$$

The defined upper and lower membership functions are visualised and plotted in Figure 8.5, where the switching roots are annotated to indicate where the respective functions are defined by upper and lower boundary kernels $\overline{\mathcal{G}}\left(\tilde{y}_t(k)\right)$ and $\underline{\mathcal{G}}\left(\tilde{y}_t(k)\right)$. Note that were $\overline{r}$ is undefined, in the case of a single root, the conditions on it in (8.36) are ignored. It is also clear that when the mean and variance of both bounds are equal, i.e. a type-1 MF, both roots are undefined, as $a = b = c = 0$.

In the case of a triangular distribution, it is trivial to derive the upper and lower membership functions, but requires a large number of conditionals and so is not

defined here. An example of a triangular IT2 fuzzy set and its membership function was discussed in Section 4.3.2, and visualised in Figure 4.4.

8.5.4 Inference and Type-Reduction

As each fuzzy set of the component rule antecedents defines an interval, with some upper and lower membership, it follows that the firing strength of each rule is so too defined as an interval, $\Omega^j(\tilde{\boldsymbol{y}}_t) = [\underline{\omega}^j(\tilde{\boldsymbol{y}}_t), \overline{\omega}^j(\tilde{\boldsymbol{y}}_t)]$. The firing strength of the rules is defined as a joining t-norm, $\top_\cap$, over the respective upper and lower membership functions of the fuzzy sets of a rule, $\underset{\sim}{Y}^j(k) = (\mathbb{R}, [\underline{\mu}^j_k, \overline{\mu}^j_k])$.

$$\begin{aligned} \overline{\omega}^j(\tilde{\boldsymbol{y}}_t) &= \top_{k=1}^{n_y} \overline{\mu}^j_k(\tilde{y}_t(k)) \\ \underline{\omega}^j(\tilde{\boldsymbol{y}}_t) &= \top_{k=1}^{n_y} \underline{\mu}^j_k(\tilde{y}_t(k)) \end{aligned} \tag{8.37}$$

The respective firing strengths can be used to approximate an upper and lower bound for the correction term in the filtering process. Given the upper and lower firing strength, the defuzzifed *a posteriori* state estimate can be calculated, as in the UT1-FKF, by the sum of each rule consequent. In the (unscented) IT2 fuzzy Kalman filter (UIT2-FKF), the firing strength is defined as the normalised type reduction of $\Omega^j(\tilde{\boldsymbol{y}}_t)$:

$$\begin{aligned} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t) &= \left(\overline{\omega}^j(\tilde{\boldsymbol{y}}_t) - \underline{\omega}^j(\tilde{\boldsymbol{y}})\right) \left(2\sum_{p=0}^{r} \frac{\overline{\omega}^p(\tilde{\boldsymbol{y}}_t) - \underline{\omega}^p(\tilde{\boldsymbol{y}}_t)}{2}\right)^{-1} \\ \boldsymbol{x}_{t|t} &= \sum_{j=1}^{r} \hat{\omega}^j(\tilde{\boldsymbol{y}}_t) \sum_{i=0}^{2n_x} w_i^{(x)} g^j\left(\chi^j_{i;t|t}, u^j_t\right) \\ \boldsymbol{P}^{(x)}_{t|t} &= \sum_{j=1}^{r} \left(\hat{\omega}^j(\tilde{\boldsymbol{y}}_t)\right)^2 \sum_{i=0}^{2n_x} w_i^{(p)} \left(\boldsymbol{x}_{t|t} - g^j\left(\chi^j_{i;t|t}, u^j_t\right)\right)\left(\boldsymbol{x}_{t|t} - g^j\left(\chi^j_{i;t|t}, u^j_t\right)\right)^{\mathrm{T}} \end{aligned} \tag{8.38}$$

Note that since the generative FIS in the UIT2-FKF has a TSK model consequent, the output is crisp, and the type reduction operation reduces to a simple mean operation, as discussed in Section 4.3.3.

Problem	Contribution
Handling complex prediction problems	The development of the fuzzy Kalman filter and its derivatives allows for estimation and prediction in systems that can be described by a number of complex, possibly non-linear models
Quantifying the uncertainty in the system and making use of this in predictions	The FKF combines the uncertainty in its prediction through the means of probabilistic representation with assumed density filtering, with fuzzy representation of component processes representing the system using both type-1 and type-2 membership functions
Presenting the uncertainty in the system	The uncertainty of the predictions can be presented both in terms of the probabilistic representation of state measurements, which are parameterised by a mean and covariance, and the firing strengths represent the decision-making during the modelling

Table 8.5: Contributions of the development of the fuzzy Kalman filters

8.6 Discussion

The general conclusions that can be drawn from the method introduced in this chapter are that using fuzzy inference to supplement the Bayesian filtering. Indeed, the multiple model approach to solving complex systems is an effective means in resolving estimations in problems with uncertainty both in the observation and process models, but in the effects of different processes enacting on the system's state. It is notable that the fuzzy approach infers the weighting of processes without need of a prior set of switching probabilities which, in systems with a large number of processes, r, comprises an $r \times r$ matrix. Simultaneously, multiple tracks must be simultaneously maintained, alongside an expected estimate of the unified state, meaning that computational requirements are multiplied by order r. These contributions are summarised in Table 8.5

Under particular conditions, the setup can be defined as an ensemble method, defined in its correction term as the weighted sum of ensembles, each propagated individually, i.e. ensemble averaging.

$$\boldsymbol{x}_{t|t} = \sum_{j=1}^{r} \omega^j \boldsymbol{x}_{t|t}^j \tag{8.39}$$

The unified state estimation here is the combination of each component state estimation. In the fuzzy approach described in this chapter, the component state is mapped into the relevant domain and represent an ensemble estimate, and the

weight is the firing strength of each component rule. The format of this method is equivalent to that defined in (8.18) and elsewhere for the respective unscented and type-2 derivations. Where an ensemble may be generated over a set of preconditions, the FKF is more inline with the bucket-of-models approach, in that it takes a set of models, each representing a component process in the switching SSM. Instead of selecting only the best model, the fuzzy approach provides a means to weight each model.

Considering the underlying Gaussian assumptions, in type-1 case, using the Gaussian kernel for the generative fuzzy membership of each model, the inference system takes up the form of a Gaussian mixture model. Consider the j^{th} generative fuzzy rule under the product t-norm and the Gaussian kernel. The firing strength of the rule is defined as follows, using (8.16) and (8.17), the definitions of a rule set and corresponding firing strength.

$$\begin{aligned} \omega^j &\propto \prod_{k=1}^{n} \mathcal{N}\left(y^j_{t|t-1}(k), P^{j(y)}_t(k,k)\right) \\ &\approx \mathcal{N}\left(\boldsymbol{y}^j_{t|t-1}, \boldsymbol{P}^{j(y)}_t\right) \end{aligned} \tag{8.40}$$

From a probabilistic perspective, the firing strength of each component rule is now a normally distributed joint probability distribution over the whole observation space, to be evaluated on measurement $\tilde{y}_t$. Given that the firing strength can be considered Gaussian, the defuzzification criteria, as defined by (8.18), assuming here that both component estimate and error covariance have been propagated by g^j, or that g^j is an identity.

$$\begin{aligned} p(\tilde{x}_{t|t}) &\propto \sum_{j=1}^{r} \mathcal{N}\left(\boldsymbol{y}^j_{t|t-1}, \boldsymbol{P}^{j(y)}_t\right) \mathcal{N}\left(\boldsymbol{x}^j_{t|t}, \boldsymbol{P}^{j(x)}_{t|t}\right) \\ &= \sum_{j=1}^{r} \mathcal{N}\left(\begin{bmatrix} \boldsymbol{x}^j_{t|t} \\ \boldsymbol{y}^j_{t|t-1} \end{bmatrix}, \begin{bmatrix} \boldsymbol{P}^{j(x)}_{t|t} & \\ & \boldsymbol{P}^{j(y)}_t \end{bmatrix} \right) \end{aligned} \tag{8.41}$$

Thus, estimation of the state can be characterised as a mixture model over the joint distributions of state and observation estimates, with the expectation corresponding to the updated state estimate and predicted observation. Note that the distribution defined in (8.41) does not take into account the weightings over the each rule, and in practise each distribution can be evaluated seperately, using the state estimate,

$\boldsymbol{x}_{t|t}^{j}$ and evaluating the measurements over ω^{j}, as described in the derivation of the fuzzy approach.

9

Experimental Evaluation

9.1 Method Comparisons and Metrics

This chapter details some controlled experiments to demonstrate the successes of the proposed method against similar filtering methods. First the methods are tested against a synthetic problem common in evaluating filtering methods, especially those targetted at approximating the state Markov-switching system. Target tracking, as discussed briefly in Section 5.1, can be modelled in 2-D and 3-D space, and modelled as a series of individual manoeuvres. Such systems are common in applications such as radar tracking, where combined with noisy bearings, the uncertainty of a given manoeuvre can mean traditional single processes struggle to model or adapt to specific changes in the object's movement.

For these experiments, the proposed UT1- and UIT2-FKF is compared against non-linear implementations of the interacting multiple model, with component filtering equations adapted from linearised extended Kalman filter (EIMM) and, like the unscented FKFs, with the unscented Kalman filter (UIMM). Comparison against single process models compares against approximations by the UKF and by sequential importance resampling (SIR, a.k.a. the bootstrap particle filter). While the IMM methods will be initialised with the same parameters as the both UT1- and UIT2-FKF, in terms of component processes and noise terms, the UKF and SIR filter will only approximate a single velocity model.

For the synthetic object tracking methods, the true state parameters will be compared against the estimates of each algorithm. The root mean-squared error (RMSE) of each component parameter will be used to demonstrate the accuracy of state estimation, and the mean distance error (MSD) of the positional component

will be used to compare the estimation of observable parameters indicate the effects of correcting noisy measurement by the respective filters.

Calculations for these metrics are defined, as follows, where T is the total number of time steps in the filtering, and where $\vec{\mathbf{x}}$ and $\tilde{\vec{\mathbf{x}}}$ are the Cartesian coordinate elements of the state estimate and corresponding true state, $(\vec{x}, \vec{y}, \vec{z})$.

$$\begin{aligned} \mathrm{RMSE}(\boldsymbol{x}_{1:T}) &= \left(T^{-1} \sum_{t=1}^{T} \left(\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_{t|t} \right)^2 \right)^{\frac{1}{2}} \\ \mathrm{MDE}(\boldsymbol{x}_{1:T}) &= T^{-1} \sum_{t=1}^{T} \| \tilde{\vec{\mathbf{x}}}_t - \vec{\mathbf{x}}_{t|t} \| \end{aligned} \tag{9.1}$$

The final controlled experiment uses a non-synthetic system, and serves to demonstrate the effects of using the FKF to separate the physical processes in a system as component state processes. The example is a falling bouncing ball, with simple state dynamics, viewed using a fixed-position camera. The measurement vector forms a bounding box, and there is no ground truth for the hidden state variables. For approximated parameters in the ERT experiment, they were visually compared with calculated values where available.

9.2 Coordinated Turning Models

9.2.1 Planar System

The basic motion model of a tracked object in Cartesian space can be described by the linear, discretised Weiner velocity model. Assuming constant velocity, i.e. movement subject to no acceleration or deceleration, over each interval in time, the objects position is simply incremented by the respective directional velocity estimates scaled by some parameter controlling the measurement time interval. To model the possible change in direction of the tracked object, a coordinate turn model is considered in parallel with the velocity model. This forms a multi-process SSM where the object can be considered moving *forward* or performing a *coordinated turn* by some angle θ.

In this experiment, the tracked object's movement is fixed to a plane $(\vec{x}, \vec{y})$. As such, the object may only move across this plane, and the velocity model is considered 2-dimensional $(v_{\vec{x}}, v_{\vec{y}})$. In addition, the turning of the object can only occur about a single axis, perpendicular to the plane, so there is only a single turning angle

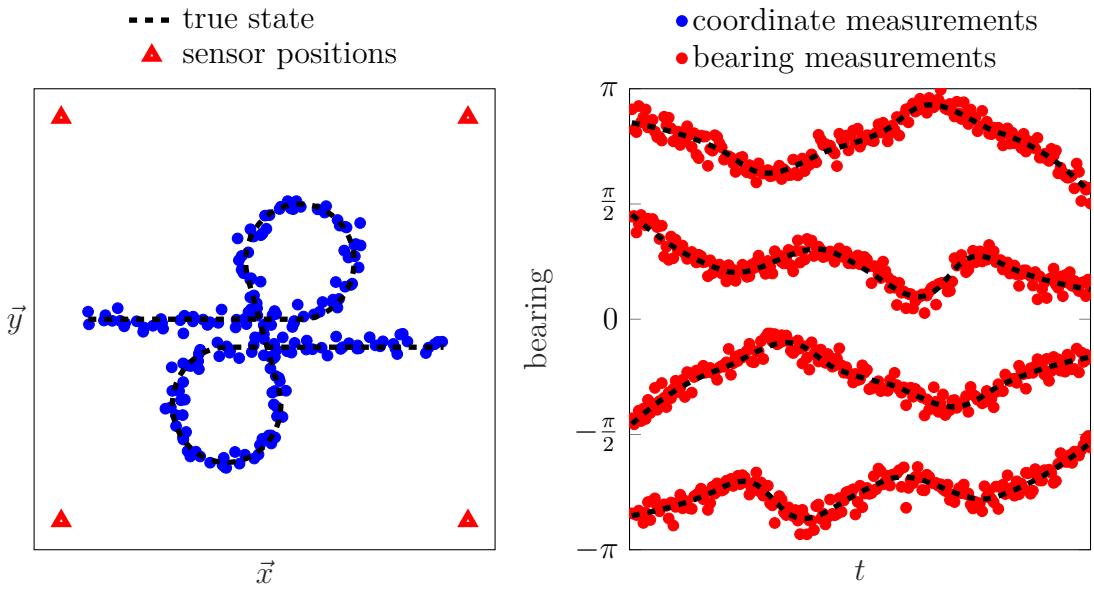


Figure 9.1: State observation and noisy measurements for the planar coordinated turn experiment using the coordinate and bearings-only observation models. The sensor locations used to calculate bearings are also plotted. Bearings are given in radians

to model. For the planar velocity and turning model, the unified system state is defined:

$$\boldsymbol{x} = \begin{bmatrix} \vec{x} & \vec{y} & v_{\vec{x}} & v_{\vec{y}} & \theta \end{bmatrix}^{\mathrm{T}} \tag{9.2}$$

To compare robustness of the algorithms, the filtering experiments were run using two separate observation models. The first, a simple linear coordinate model, takes the coordinate components of the state estimate with some additive Gaussian noise parameter. The second simulates radar tracking, by observing the object's position as a set of angular bearings about some placed sensors on the model plane. In the bearings-only tracking experiment, four sensors were placed around the plane such that the entire trajectory of the object is contained within. Noise was additively added each individual bearing angle, such that the measurement bearings no longer have a single intersection indicating the position of a tracked object. This method of observation modelling gives a greater uncertainty than a single noisy point, as in the coordinate-only model. Example measurements for the two measurement criteria are shown in Figure 9.1. Full details of the system setup is described in Table 9.1. Note that the domain mappings g^j are simply linear mappings of respective component state element variables to the unified state,

with the previous unified state estimate of non-component variables propagated as is. In future model setup tables, where this is trivially the case, the mapping functions are omitted.

Experiments were run over 100 Monte Carlo simulations, each with freshly generated noisy measurements, with noise sampled randomly from a Gaussian distribution with covariance $\mathbf{R}[= \mathbf{R}^j,\ \forall j]$. Each experiment set of measurement was

Table 9.1: Model setup for coordinated turn filtering on a plane with two component models, and two possible measurements: coordinates and angular bearings about some n_s sensors, s^k, where $s^k_{\vec{x}}$ are $s^k_{\vec{y}}$ are the coordinates of the k^{th} sensor position. Control parameter d_t denotes the discrete time interval between states, and ϵ is some very small ($< 10^{-15}$) value to force $\mathbf{Q}^j$ to be positive-definite

Process	$\boldsymbol{x}_t^j$	f^j	$\boldsymbol{Q}^j$
forward	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{x}} \\ v_{\vec{y}} \end{bmatrix}$	$\begin{matrix} \vec{x} + d_t v_{\vec{x}} \\ \vec{y} + d_t v_{\vec{y}} \\ v_{\vec{x}} \\ v_{\vec{y}} \end{matrix}$	$\begin{bmatrix} \frac{d_t^3}{3} & 0 & \frac{d_t^2}{2} & 0 \\ 0 & \frac{d_t^3}{3} & 0 & \frac{d_t^2}{2} \\ \frac{d_t^2}{2} & 0 & d_t & 0 \\ 0 & \frac{d_t^2}{2} & 0 & d_t \end{bmatrix}$
coordinated turn	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{x}} \\ v_{\vec{y}} \\ \theta \end{bmatrix}$	$\begin{matrix} \vec{x} + \theta^{-1}\left(v_{\vec{x}}\sin(\theta d_t) + v_{\vec{y}}\left(\cos(\theta d_t) - 1\right)\right) \\ \vec{y} + \theta^{-1}\left(v_{\vec{x}}\left(1 - \cos(\theta d_t)\right) + v_{\vec{y}}\sin(\theta d_t)\right) \\ v_{\vec{x}}\cos(\theta d_t) - v_{\vec{y}}\sin(\theta d_t) \\ v_{\vec{x}}\sin(\theta d_t) + v_{\vec{y}}\cos(\theta d_t) \\ \theta \end{matrix}$	$\begin{bmatrix} \epsilon \cdot \mathbf{I}_4 & \\ & 1 \end{bmatrix}$

Observation	h^j	$\boldsymbol{R}^j$
coordinate	$\begin{matrix} \vec{x} \\ \vec{y} \end{matrix}$	$0.5 \cdot \mathbf{I}_2$
bearings only	$\begin{matrix} \arctan\left(\dfrac{\vec{y} - s^k_{\vec{y}}}{\vec{x} - s^k_{\vec{x}}}\right), \\ k\ =\ 1, \ldots, n_s \end{matrix}$	$0.1 \cdot \mathbf{I}_{n_s}$

Mapping	$\mathbf{G}^j$	$\mathbf{U}^g$	u_t^j
forward	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$\boldsymbol{x}_{t-1}$
coordinated turn	$\mathbf{I}_5$	*n/a*	*n/a*

FIS	t-norm ($\top_\cap$)	Kernel ($\mathcal{G}$)
	product	Gaussian CR

filtered with each of the filtering methods, and RMSE and MDE of the Cartesian and bearings-only experiments are given in Table 9.2 and Table 9.3 respectively. The results consistently show the superior performance of multiple model filters over the single filter UKF. It is clear that the estimates as calculated by the UIMM over the EIMM is consistent with the general idea that the unscented transform outperforms linearisation, at least with first-order Jacobians.

The errors in approximating the observable parameters, i.e. the $(\vec{x}, \vec{y})$ coordinates,

Table 9.2: Mean (± 1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for six non-linear filtering methods on planar coordinated turning experiment, using Cartesian coordinates as the observation model

RMSE	UKF	SIR	EIMM	UIMM	UT1-FKF	UIT2-FKF
$\vec{x}$	0.4767 ± 0.0079	0.1173 ± 0.0078	0.2223 ± 0.0078	0.1655 ± 0.0090	0.1043 ± 0.0066	0.1049 ± 0.0066
$\vec{y}$	0.4725 ± 0.0084	0.1105 ± 0.0076	0.2219 ± 0.0071	0.1587 ± 0.0078	0.0942 ± 0.0071	0.0944 ± 0.0070
$v_{\vec{x}}$	0.6590 ± 0.0032	0.4265 ± 0.0065	0.4091 ± 0.0067	0.3342 ± 0.0076	0.3935 ± 0.0031	0.4001 ± 0.0031
$v_{\vec{y}}$	0.6395 ± 0.0036	0.3928 ± 0.0072	0.3929 ± 0.0069	0.2933 ± 0.0076	0.3113 ± 0.0040	0.3163 ± 0.0039
θ	–	–	0.5227 ± 0.0043	0.5129 ± 0.0025	0.5115 ± 0.0071	0.5021 ± 0.0069
MDE	0.5611 ± 0.0077	0.1370 ± 0.0068	0.2198 ± 0.0072	0.1631 ± 0.0062	0.1328 ± 0.0063	0.1335 ± 0.0063

Table 9.3: Mean (± 1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for six non-linear filtering methods on planar coordinated turning experiment, using bearings from four sensors as the observation model

RMSE	UKF	SIR	EIMM	UIMM	UT1-FKF	UIT2-FKF
$\vec{x}$	0.7692 ± 0.0300	0.3388 ± 0.0285	0.6096 ± 0.0357	0.5170 ± 0.0783	0.2631 ± 0.0289	0.2692 ± 0.0285
$\vec{y}$	0.6859 ± 0.0238	0.2763 ± 0.0221	0.4271 ± 0.0228	0.4513 ± 0.1367	0.2126 ± 0.0222	0.2175 ± 0.0222
$v_{\vec{x}}$	0.7324 ± 0.0074	0.5813 ± 0.0166	0.7038 ± 0.0207	0.6040 ± 0.0450	0.5066 ± 0.0109	0.5214 ± 0.0107
$v_{\vec{y}}$	0.7056 ± 0.0068	0.4780 ± 0.0130	0.5721 ± 0.0160	0.5479 ± 0.0360	0.3934 ± 0.0096	0.4054 ± 0.0091
θ	–	–	0.6668 ± 0.0075	0.6459 ± 0.0227	0.5812 ± 0.0211	0.5684 ± 0.0211
MDE	0.8791 ± 0.0251	0.3847 ± 0.0241	0.6168 ± 0.0301	0.5269 ± 0.0865	0.3130 ± 0.0233	0.3211 ± 0.0234

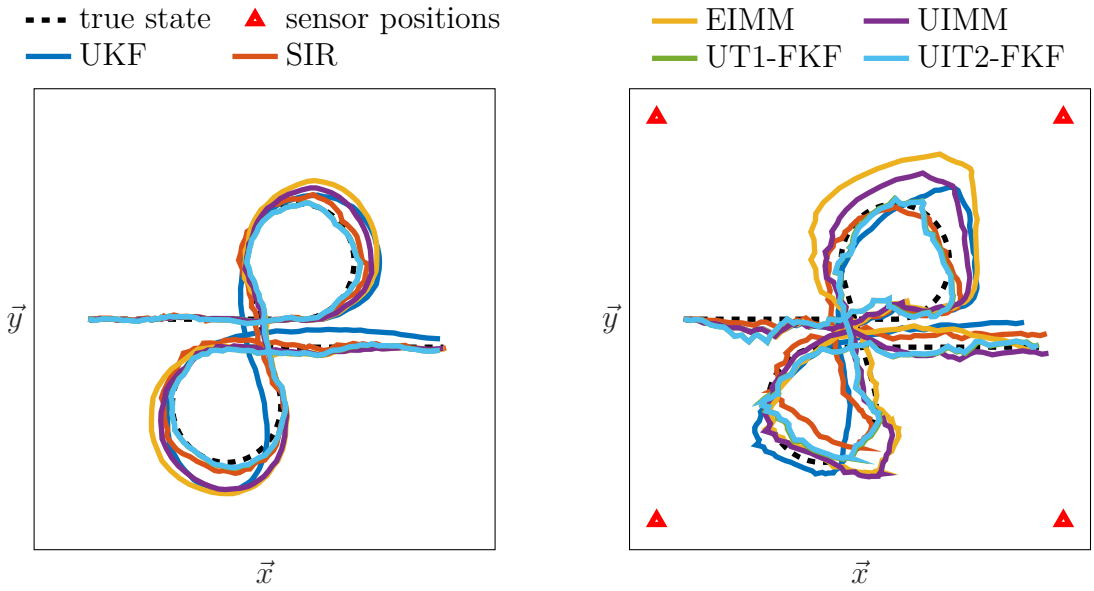


Figure 9.2: Estimates of the $(\vec{x}, \vec{y})$ components of the object tracking experiment by six methods for non-linear filtering, using two different observation models. The left plot shows the filtered estimates based on Cartesian coordinate measurements; the right plot shows the estimates based on bearing-only measurements, based on four sensors

especially highlight the performance of the fuzzy filters proposed. Indeed, when visualised in Figure 9.2, the position estimated by both the UT1- and UIT2-FKF are located consistently along the path of the true state, notably where the path deviates with sudden change in direction. Even when the observation model is non-linear, in the bearings-only experiment, the fuzzy filters can successfully and consistently estimate the current state. It's evident from the error results that SIR particle filter is also effective at state estimate, especially on the observable parameters, despite being set up with only the linear-discrete velocity model.
Indeed, the SIR filter with a high volume of particles, is expected to be a strong performer, due to its use of Monte Carlo sampling. However, it is less effective than the fuzzy Kalman filters, in estimating all parameters, and is less accurate than UIMM in approximating state parameters that are not directly observable. Boxplots of the parameter estimate RMSE over the Monte Carlo simulations for are presented in Figure 9.3 and Figure 9.4 for the two respective observation models. The accuracy of the parameter estimation shows that the three unscented transform-based multi-model methods are best, especially in the Cartesian coordinate observation model. In this instance, the performance of UIMM and the FKF methods are similar, but the FKF methods are better at approximating the observable parameters. This is likely because while switching in the UIMM is controlled by predefined probabilities,

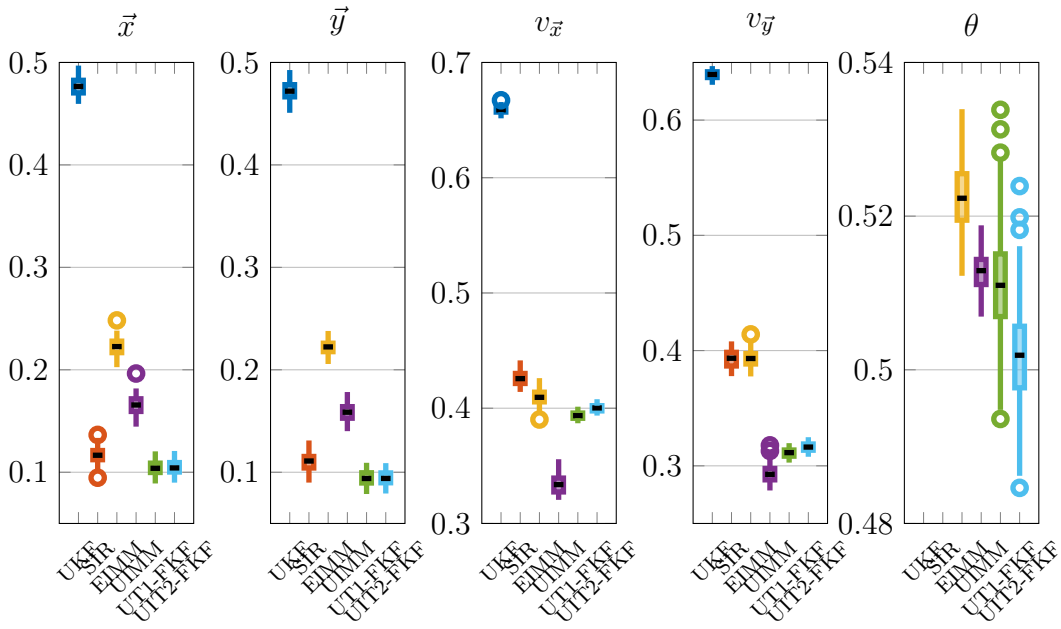


Figure 9.3: Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the planar tracking experiment using the Cartesian coordinate observation model. Turning angle, θ, was not estimated by UKF and SIR. Individual markers indicate outliers that fall outside 1.5±IQR

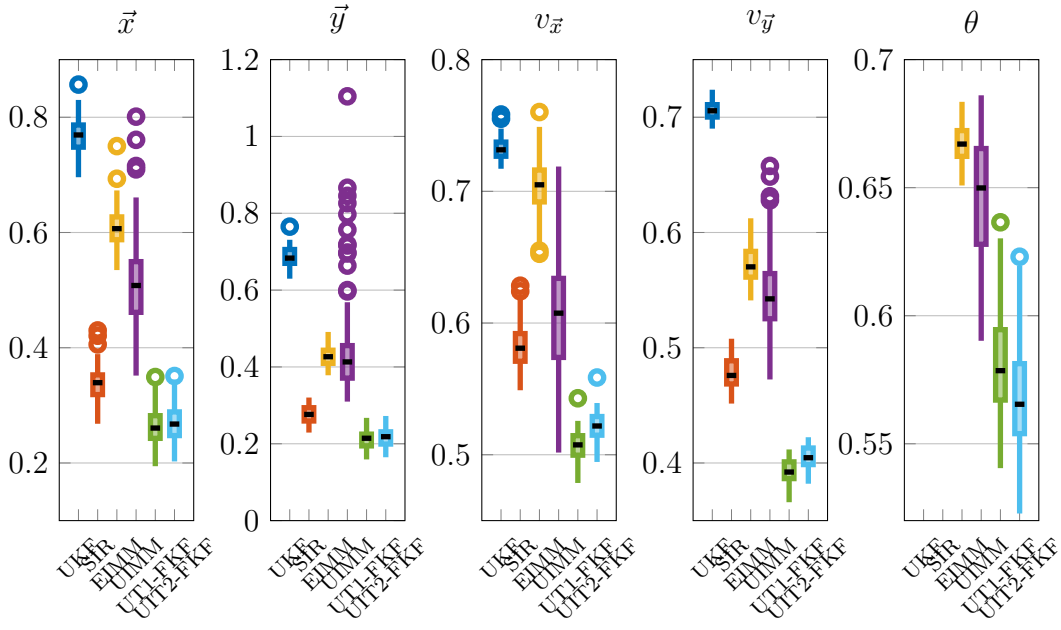


Figure 9.4: Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the planar tracking experiment using the bearings-only observation model

the firing strength of the FKF methods are driven solely by observation likelihood.

In the non-linear bearings-only observation model, the accuracy of all methods deteriorates, however the effect of this occurs to a greater degree by the UIMM, especially in comparison to the FKF estimates. The long tails and wide IQR of the UIMM RMSE values indicates that it is not robust to noise in the non-linear case. In contrast, the detrimental effects of increased uncertainty in measurements, due to noise and non-linearity, is not as high in the FKF methods. Indeed, the consistent RMSE, with low standard deviation, indicate a high degree of robustness to the observation-based fuzzy approach.

9.2.2 3-D System

Extending on the previous experiment, an object is now tracked in 3-D Cartesian space. The state now requires an additional coordinate and respective velocity parameter. To describe motion in coordinated turning 3-D, alongside the horizontal turning angle, θ, a pitch angle, φ describes the angle about the vector along the $\vec{x} - \vec{y}$ plane that is perpendicular to direction of the objects trajectory. The state for the coordinate turning in 3-D is defined:

$$\boldsymbol{x} = \begin{bmatrix} \vec{x} & \vec{y} & \vec{z} & v_{\vec{x}} & v_{\vec{y}} & v_{\vec{z}} & \theta & \varphi \end{bmatrix}^{\mathrm{T}} \tag{9.3}$$

The component processes influencing the model in this experiment are defined as the linear-discrete velocity model, describing the *forward* motion of the tracked object; a coordinated turn about the $\vec{z}$-axis, named the *yaw* model hereafter, although it is independent of velocity $v_{\vec{z}}$; and the coordinate turn of the object's *pitch*, describing the upward or downward turning of the object. As with the previous example, two observation models are used to model the accuracy of the filtering algorithms being compared. The first is simply position in Cartesian coordinates $(\vec{x}, \vec{y}, \vec{z})$, with the observation being a linear mapping of the corresponding state variables. The second model is again a bearings-only model. Angular bearings of four sensors are again modelled, however given the 3-D nature of the system, the angular position is formed of two angles, the elevation and the azimuth with respect to the sensor. Each sensor is also modelled on a plane with constant $\vec{z}$ position, below the modelled trajectory. Full details of the model setup for component processes and observation models are described in Table 9.4. Examples of noisy measurements for the two models are plotted in Figure 9.5.

The filtering algorithms were tested over 100 Monte Carlo simulations, with noisy measurements for each observation model randomly generated in each simulation. The unscented transform methods were compared here, with EIMM omitted from experiments as it is considered a worse-performing analogue to UIMM. In addition, the SIR was used, again with 5000 randomly sampled particles.

Table 9.4: Model setup for object tracking on in 3-D Cartesian space with three component models, and two possible measurements: coordinates and angular bearings about some n_s sensors, s^k, where $s^k_{\vec{x}}$, $s^k_{\vec{y}}$ and $s^k_{\vec{z}}$ are the coordinates of the k^{th} sensor position. Control parameter d_t denotes the discrete time interval between states, and ϵ is some very small ($< 10^{-15}$) value to force $\mathbf{Q}^j$ to be positive-definite

Process	$\boldsymbol{x}^j_t$	f^j	$\boldsymbol{Q}^j$
forward	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ \vec{z} \\ v_{\vec{x}} \\ v_{\vec{y}} \\ v_{\vec{z}} \end{bmatrix}$	$\begin{matrix} \vec{x} + d_t v_{\vec{x}} \\ \vec{y} + d_t v_{\vec{y}} \\ \vec{z} + d_t v_{\vec{z}} \\ v_{\vec{x}} \\ v_{\vec{y}} \\ v_{\vec{z}} \end{matrix}$	$\begin{bmatrix} \frac{d_t^3}{3}\mathbf{I}_3 & \frac{d_t^2}{2}\mathbf{I}_3 \\ \frac{d_t^2}{2}\mathbf{I}_3 & d_t\mathbf{I}_3 \end{bmatrix}$
coordinated turn (yaw)	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{x}} \\ v_{\vec{y}} \\ \theta \end{bmatrix}$	$\begin{matrix} \vec{x} + \theta^{-1}\left(v_{\vec{x}}\sin(\theta d_t) + v_{\vec{y}}\left(\cos(\theta d_t) - 1\right)\right) \\ \vec{y} + \theta^{-1}\left(v_{\vec{x}}\left(1 - \cos(\theta d_t)\right) + v_{\vec{y}}\sin(\theta d_t)\right) \\ v_{\vec{x}}\cos(\theta d_t) - v_{\vec{y}}\sin(\theta d_t) \\ v_{\vec{x}}\sin(\theta d_t) + v_{\vec{y}}\cos(\theta d_t) \\ \theta \end{matrix}$	$\begin{bmatrix} \epsilon \cdot \mathbf{I}_4 & \\ & 1 \end{bmatrix}$
coordinated turn (pitch)	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ \vec{z} \\ v_{\vec{x}} \\ v_{\vec{y}} \\ v_{\vec{z}} \\ \varphi \end{bmatrix}$	$\begin{matrix} \vec{x} + \varphi^{-1}\left(v_{\vec{x}}\sin(\varphi d_t) + v_{\vec{z}}\left(\cos(\varphi d_t) - 1\right)\right) \\ \vec{y} + \varphi^{-1}\left(v_{\vec{x}}\left(1 - \cos(\varphi d_t)\right) + v_{\vec{y}}\sin(\varphi d_t)\right) \\ \vec{z} + \varphi^{-1}\left(\left\|(v_{\vec{x}}, v_{\vec{y}})\right\|\left(1 - \cos(\varphi d_t)\right) + v_{\vec{z}}\sin(\varphi d_t)\right) \\ v_{\vec{x}}\cos(\varphi d_t) - v_{\vec{z}}\sin(\varphi d_t) \\ v_{\vec{y}}\cos(\varphi d_t) - v_{\vec{z}}\sin(\varphi d_t) \\ \left\|(v_{\vec{x}}, v_{\vec{y}})\right\|\sin(\varphi d_t) + v_{\vec{z}}\cos(\varphi d_t) \\ \varphi \end{matrix}$	$\begin{bmatrix} \epsilon \cdot \mathbf{I}_6 & \\ & 1 \end{bmatrix}$
Observation	h^j	$\boldsymbol{R}^j$	
coordinate	$\begin{matrix} \vec{x} \\ \vec{y} \\ \vec{z} \end{matrix}$	$0.5 \cdot \mathbf{I}_3$	
bearings only	$\begin{matrix} \arccos\left(\frac{\vec{z} - s^k_{\vec{z}}}{\left\|(\vec{x},\vec{y},\vec{z}) - \mathbf{s}^k\right\|}\right), \\ \arctan\left(\frac{\vec{y} - s^k_{\vec{y}}}{\vec{x} - s^k_{\vec{x}}}\right) \\ k \;=\; 1, \ldots, n_s \end{matrix}$	$0.1 \cdot \mathbf{I}_{2n_s}$	
FIS	t-norm ($\top_\cap$)	Kernel ($\mathcal{G}$)	
	product	Gaussian CR	

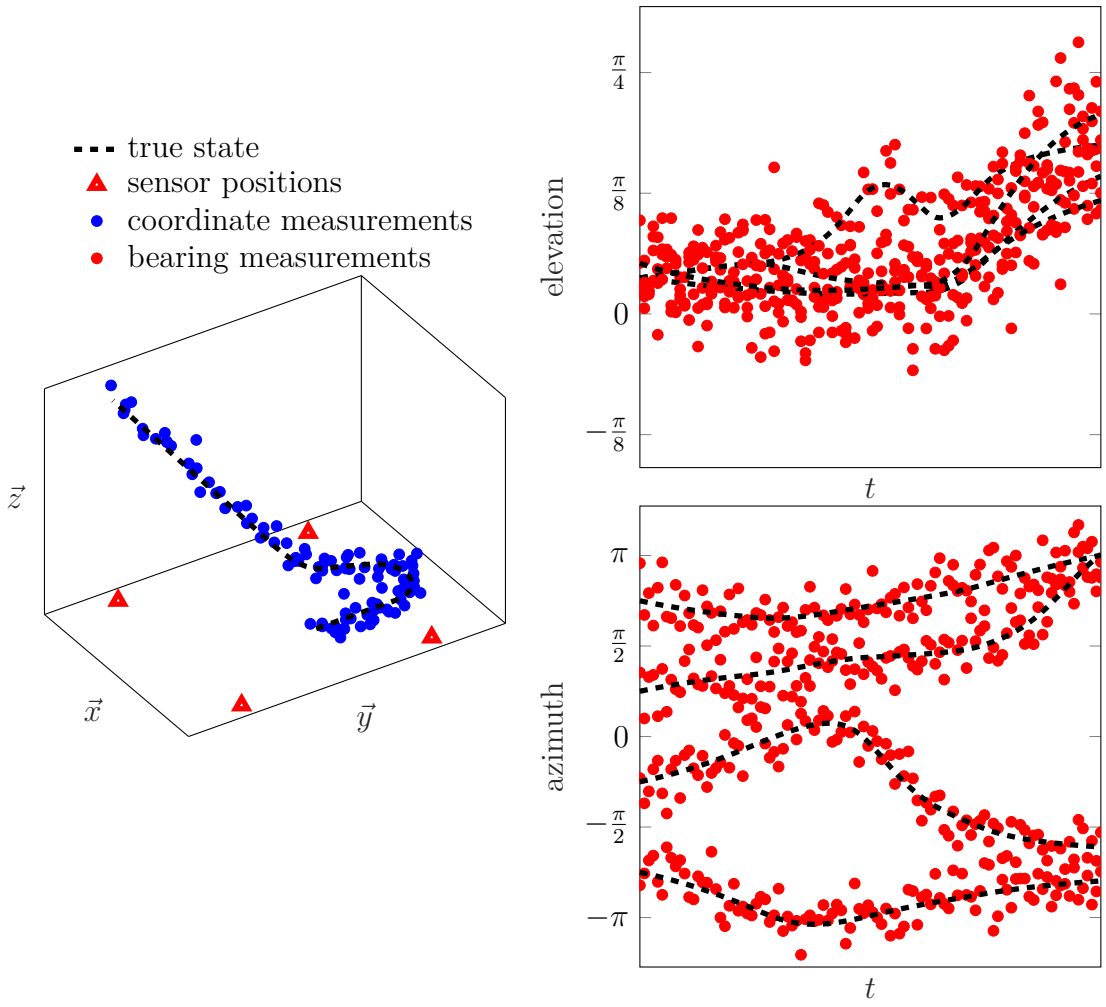


Figure 9.5: State observation and noisy measurements for the 3-D coordinated turn experiment using the coordinate and bearings-only observation models. The sensor locations used to calculate bearings are also plotted. Elevation and azimuth are given in radians

---

The error metrics for each filtering algorithm are shown in Table 9.5 and Table 9.6 for the coordinate and bearings-only observation models respectively. As in the previous example, the non-linear measurement criteria leads to a deterioration in the accuracy of the estimates. Notable, in this experiment, is the poor performance of the SIR particle filter according to the RMSE metric: the increased dimensionality of both the state and observation domains may be a factor in this, as the number of particles is not increased for this simulation.

Boxplots of the RMSE of are plotted in Figure 9.6 and Figure 9.7 for the two observation criteria, omitting the SIR, comparing the performance of the state parameter estimation by each of the unscented transform based methods. The

Table 9.5: Mean (±1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for five non-linear filtering methods on 3-D coordinated turning experiment, using Cartesian coordinates as the observation model

RMSE	UKF	SIR	UIMM	UT1-FKF	UIT2-FKF
$\vec{x}$	0.4364 ± 0.0472	5.5244 ± 0.0690	0.3806 ± 0.0505	0.3579 ± 0.0620	0.3587 ± 0.0620
$\vec{y}$	0.4372 ± 0.0451	4.5325 ± 0.0797	0.3747 ± 0.0475	0.3500 ± 0.0579	0.3509 ± 0.0577
$\vec{z}$	0.4330 ± 0.0420	1.8307 ± 0.0657	0.3541 ± 0.0392	0.3258 ± 0.0542	0.3236 ± 0.0550
$v_{\vec{x}}$	0.5144 ± 0.0500	1.1677 ± 0.0329	0.4065 ± 0.0485	0.3815 ± 0.0583	0.3839 ± 0.0583
$v_{\vec{y}}$	0.5211 ± 0.0444	1.5364 ± 0.0311	0.4089 ± 0.0334	0.3628 ± 0.0458	0.3665 ± 0.0457
$v_{\vec{z}}$	0.5062 ± 0.0471	0.4564 ± 0.0433	0.4781 ± 0.0469	0.5158 ± 0.0641	0.5227 ± 0.0626
θ	–	–	0.3421 ± 0.0562	0.4358 ± 0.0713	0.4260 ± 0.0683
φ	–	–	0.3473 ± 0.0396	0.4312 ± 0.0610	0.4346 ± 0.0619
MDE	0.6950 ± 0.0436	0.6682 ± 0.0442	0.5906 ± 0.0454	0.5537 ± 0.0586	0.5536 ± 0.0586

Table 9.6: Mean (±1 standard deviation) RMSE and MDE over 100 Monte Carlo simulations for five non-linear filtering methods on 3-D coordinated turning experiment, using bearings from four sensors as the observation model

RMSE	UKF	SIR	UIMM	UT1-FKF	UIT2-FKF
$\vec{x}$	1.4876 ± 0.1326	5.0235 ± 0.2148	0.8171 ± 0.1439	0.7859 ± 0.1742	0.7760 ± 0.1711
$\vec{y}$	0.8180 ± 0.1807	4.5223 ± 0.1711	0.6621 ± 0.1022	0.6895 ± 0.1622	0.6855 ± 0.1590
$\vec{z}$	0.9974 ± 0.1933	3.2411 ± 0.2622	0.7133 ± 0.1700	0.6366 ± 0.1771	0.6310 ± 0.1743
$v_{\vec{x}}$	0.5156 ± 0.0395	0.5886 ± 0.0451	0.4828 ± 0.0715	0.4639 ± 0.1010	0.4593 ± 0.0970
$v_{\vec{y}}$	0.4088 ± 0.0283	1.0732 ± 0.0462	0.5575 ± 0.0655	0.5239 ± 0.0910	0.5178 ± 0.0869
$v_{\vec{z}}$	0.3686 ± 0.0544	1.0275 ± 0.0651	0.8737 ± 0.1821	0.7978 ± 0.1928	0.7953 ± 0.1947
θ	–	–	0.3160 ± 0.0575	0.4149 ± 0.0877	0.4197 ± 0.0889
φ	–	–	0.3375 ± 0.0430	0.4026 ± 0.1095	0.4039 ± 0.1077
MDE	1.6758 ± 0.1447	1.7116 ± 0.1196	1.1264 ± 0.1360	1.1086 ± 0.1887	1.0990 ± 0.1862

single-process UKF here is the worst performer, as is to be expected. The median of both the UT1- and UIT2-FKF filter across parameter estimates are better than that of the UIMM filter, but the general patterns in the distribution of errors over experiments is similar between the methods, suggesting similar performance. The consistency of the results and success in estimating states in higher dimensions is further indication of the robustness of the FKF approach.

9.3 Separating Physical Processes

The previous examples have been based on simulated models. In this section, the system is setup to model a ball falling and bouncing, as observed with a static camera over 60 frames. The ball is detected in a frame using background subtraction and morphological operations to remove noise. The ball is identified as the single largest connected object in the resulting scene. The observation is

Table 9.7: Model setup for tracking the falling ball with three seperate models describing seperate process effects on the ball's state.

Process	$\boldsymbol{x}_t^j$	f^j	$\boldsymbol{Q}^j$
lateral	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{x}} \\ \vec{r} \end{bmatrix}$	$\begin{matrix} \vec{x} + d_t v_{\vec{x}} \\ \vec{y} \\ v_{\vec{x}} \\ \vec{r} \end{matrix}$	$\begin{bmatrix} 0.1 & 0 & 0.05 & 0 \\ 0 & 0.1 & 0 & 0 \\ 0.05 & 0 & \epsilon & 0 \\ 0 & 0 & 0 & 0.1 \end{bmatrix}$
vertical	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{y}} \\ \vec{r} \end{bmatrix}$	$\begin{matrix} \vec{x} \\ \vec{y} + d_t v_{\vec{y}} \\ v_{\vec{y}} \\ \vec{r} \end{matrix}$	$\begin{bmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.1 & 0.05 & 0 \\ 0 & 0.05 & 0.1 & 0 \\ 0 & 0 & 0 & 0.1 \end{bmatrix}$
accelerated falling	$\begin{bmatrix} \vec{x} \\ \vec{y} \\ v_{\vec{y}} \\ a_{\vec{y}} \\ \vec{r} \end{bmatrix}$	$\begin{matrix} \vec{x} \\ \vec{y} + v_{\vec{y}} - a_{\vec{y}} \\ v_{\vec{y}} - a_{\vec{y}} \\ a_{\vec{y}} \\ \vec{r} \end{matrix}$	$\begin{bmatrix} 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0.1 & 0 & 0 & 0 \\ 0 & 0 & 0.2 & 0.1 & 0 \\ 0 & 0 & 0.1 & 0.2 & 0 \\ 0 & 0 & 0 & 0 & 0.1 \end{bmatrix}$
Observation	h^j		$\boldsymbol{R}^j$
bounding box	$\begin{matrix} \vec{x} - \vec{r} \\ \vec{y} - \vec{r} \\ 2\vec{r} \\ 2\vec{r} \end{matrix}$		$0.5 \cdot \mathbf{I}_4$
FIS	t-norm ($\top_\cap$)		Kernel ($\mathcal{G}$)
	product		Gaussian CR

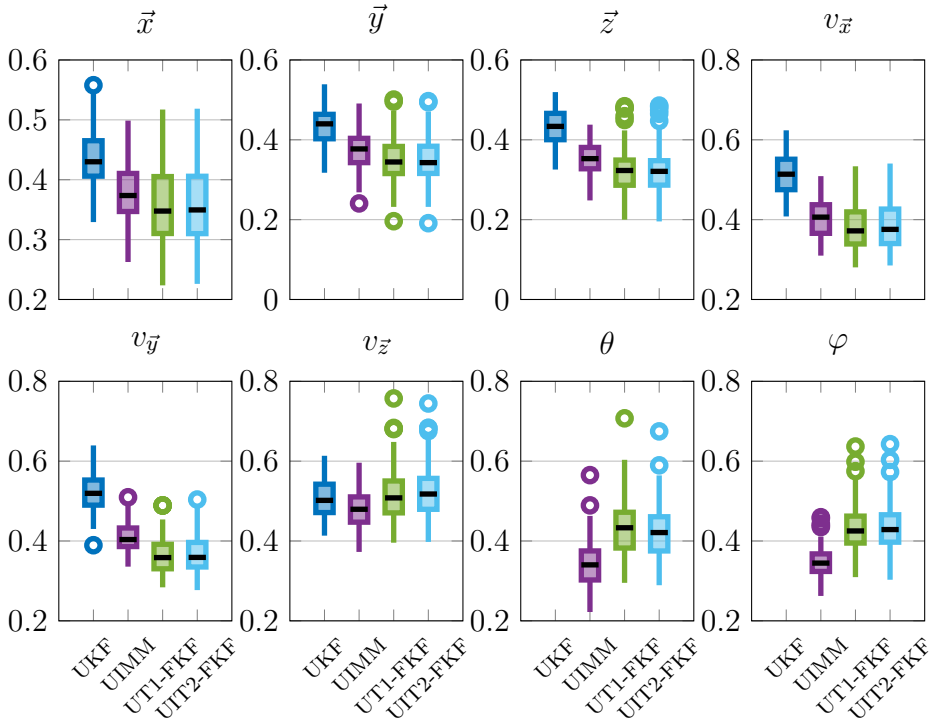


Figure 9.6: Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the 3-D tracking experiment using the Cartesian coordinate observation model

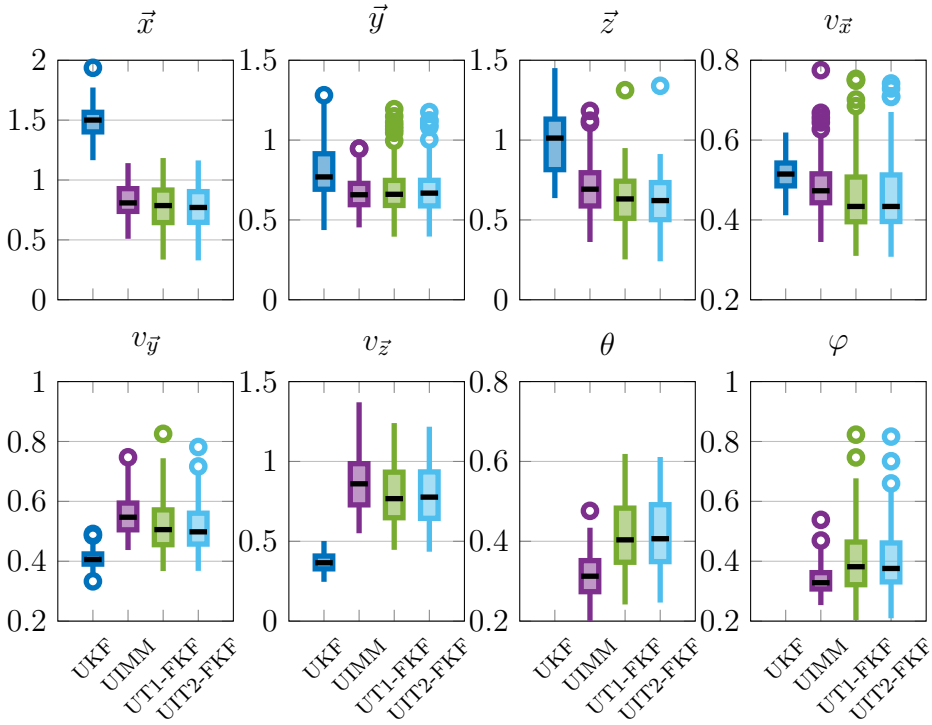


Figure 9.7: Boxplots of the RMSE for each state variable over 100 Monte Carlo simulations of the 3-D tracking experiment using the bearings-only observation model

formed of a bounding box, denoting the position of the top corner in $(\vec{x}, \vec{y})$ space, and the width and height of the box, which is assumed to be square, and equal to two-times the approximated radius of the ball, $2\vec{r}$. This experiment serves to highlight the capabilities of the proposed fuzzy filter to a non-synthetic system, and also to demonstrate how separating underlying physical processes and filtering in parallel using a multi-process SSM can help model the system under the effect of changes to respective parameters.

The falling ball experiences minimal horizontal forces, so lateral $v_{\vec{x}}$ is expected to deviate to only a small degree, if at all, over the whole observed system. The vertical velocity, however, describes the balls motion as it falls, as well as bounces. This is affected by accelerated falling, subject to gravity, and any elasticity properties of the ball when it bounces. When separating these three processes into separate states, the effect of each are considered independent and joined by the FKF based on the observed data. The unified state parameters are described by:

$$\boldsymbol{x} = \begin{bmatrix} \vec{x} & \vec{y} & v_{\vec{x}} & v_{\vec{y}} & a_{\vec{y}} & \vec{r} \end{bmatrix}^{\mathrm{T}} \tag{9.4}$$

The setup for each component SSM process is given in Table 9.7, which also defines the linear observation function describing a bounding box capturing the ball's position in the captured frame. Here, the UIT2-FKF is implemented, and compared with the estimate as found by LKF, which is implemented using the linear-discrete Weiner velocity process. The filtered observations are plotted over each frame and presented in Figure 9.8. The state is initialised in the top-middle of the image, with zero velocity and acceleration, approximately the position the ball is first dropped. This can be observed in the first six frames, where the ball has not entered the view of the camera.

It is seen here that both methods struggle to filter the position of the ball in the early frames, where the ball enters the frame. This is in part due to the initialisation parameters, but it is clear that the UIT2-FKF, with it's modelled acceleration process, has a good estimation of the ball's location before the bounce in frame 15. The individual processes also give the UIT2-FKF increase ability to handle the sudden change in direction at the point of the bounce and at the top of the arc of each bounce.

In Figure 9.9, the filtered vertical position, $\vec{y}$ is plotted for each frame. Comparing the LKF estimate with the fuzzy filter, the plot shows better response to changes

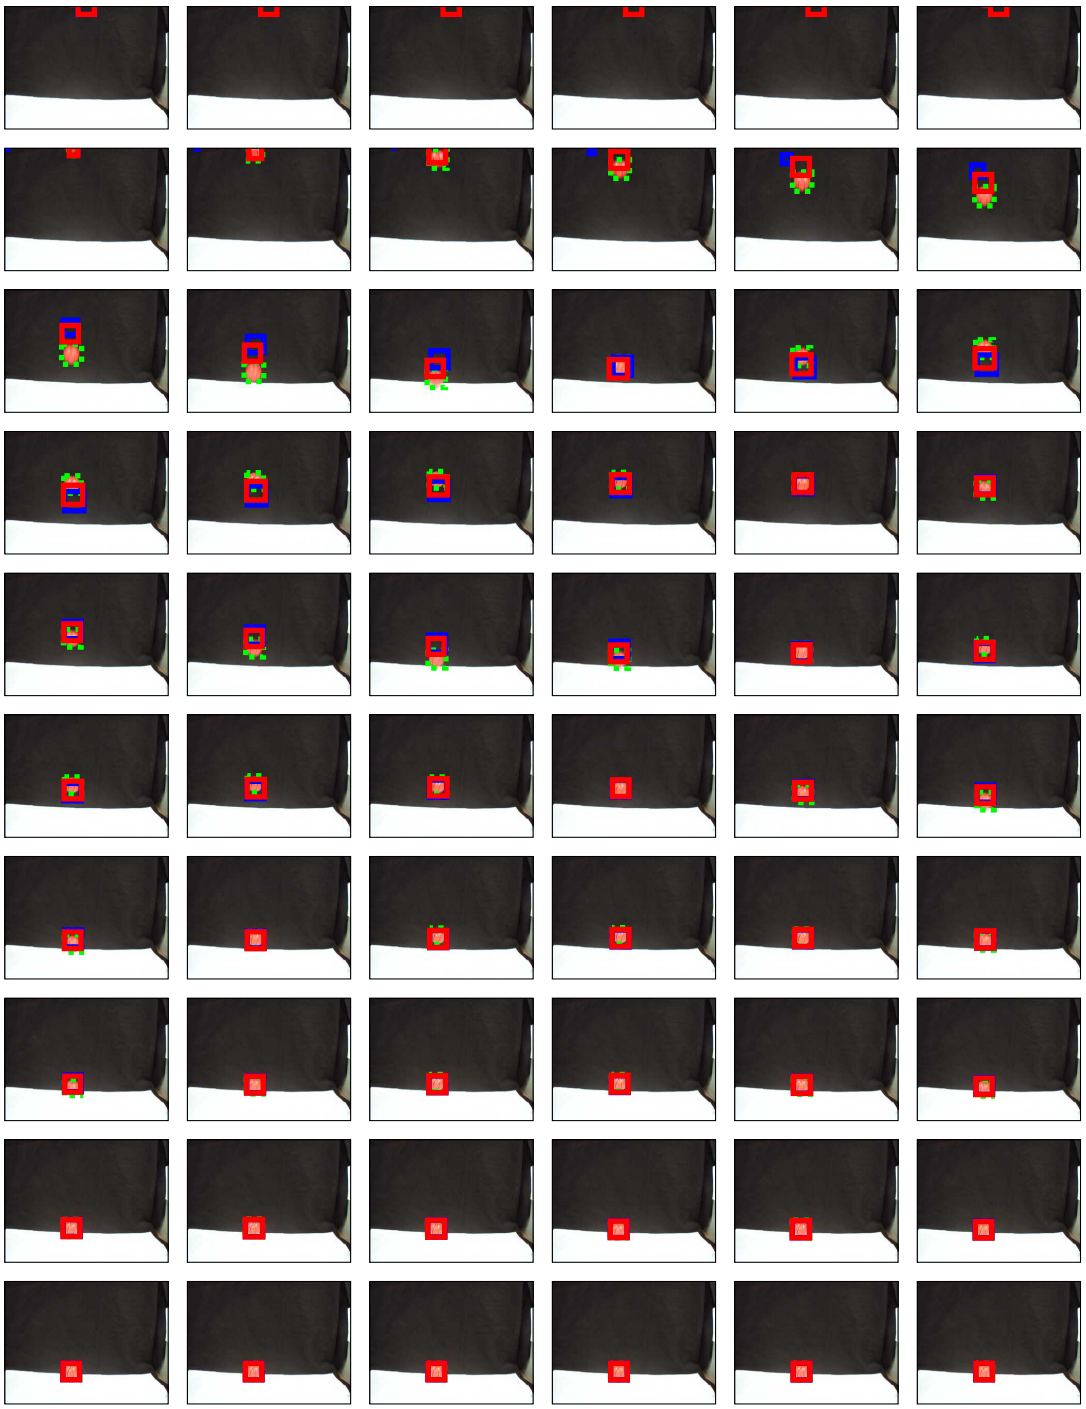


Figure 9.8: 60 frames of a video showing a ball falling and dropping, overlaid with bounding boxes denoting the filtered estimates of the ball's position at the current frame by LKF (blue) and UIT2-FKF (red). The green dashed bounding boxs are the measurements provided to the filters

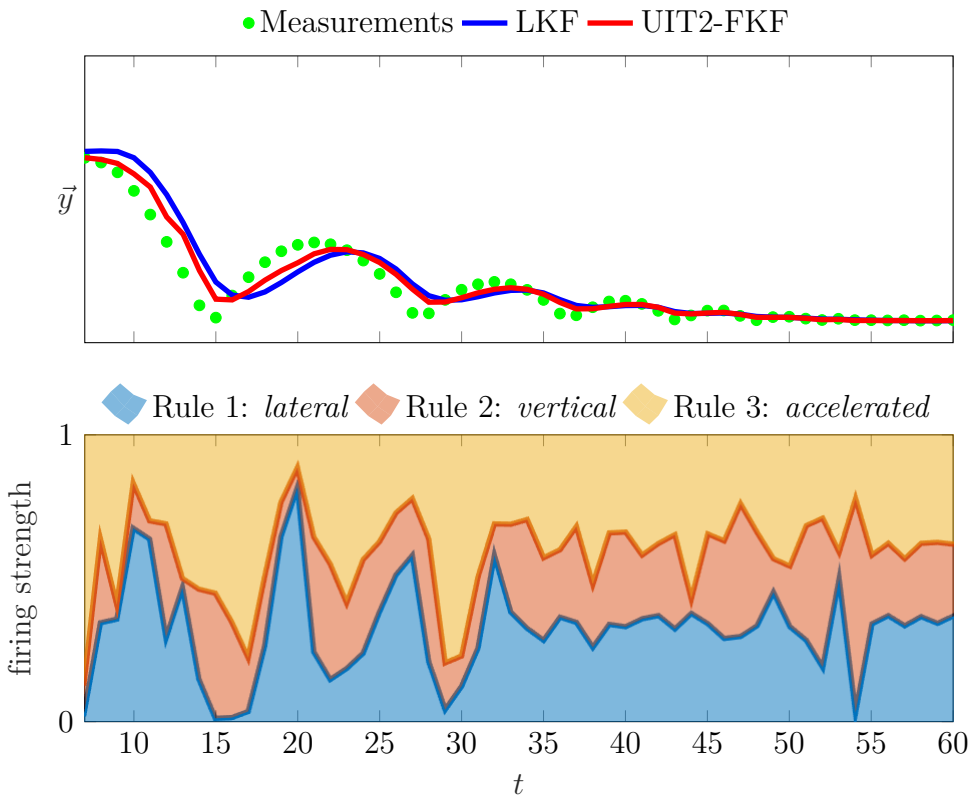


Figure 9.9: Estimates of the vertical position of the ball, $\vec{y}$ at frame t, filtered by the LKL and UIT2-FKF (top) and an area plot of corresponding firing strength of each component rule per frame

in the ball position by the UIT2-FKF. In terms of identifying individual processes under effect, the firing strength of each rule for a given frame is also plotted. Here, the patterns in the firing strength, especially that of the accelerated model, follow approximately that of the ball's position. Where the extrema occur in the plot of $\vec{y}$, corresponding to the points at which the ball is stationary before a bounce, or at the top of its arc, there is a marked reduction in the firing of both the velocity model and the accelerated model. Contrasting with the stationary periods, the moments directly following are characterised by a dominant firing of the accelerated model, and to a lesser extent the vertical velocity model. This is reduced as the bounces in the ball become less pronounced, until the near-stationary ball is filtered such that each process is equally in effect, albeit with 0 lateral and vertical velocities and acceleration.

9.4 Electrical Resistivity Tomography

9.4.1 Process Models

Referring back to the problem described for the 2-D case in Section 7.1, it can be difficult to model a system with state-space model using detected features in such a noisy environment. The additional problems surrounding boundary resolution and multi-target tracking, even in scenarios with known singular features. An alternative perspective for the ERT data is to treat the imaging cells as a single observation vector, whereby the ERT image, $V_\rho(\vec{\mathbf{x}}, t)$, is mapped to a column vector denoting the measured observations, $\tilde{\boldsymbol{y}}_t$. Here, the state space models are defined with large observation dimension. The corresponding process must then be described by a system that can resolve an observable value at each given resistivity cell.

The approach followed in this study is to model the motion of a saline tracer in a velocity field with the advection-dispersion equation. Here, the processes of advection, dispersion and density-driven sinking must be considered, and where it is reasonable to resolve the two mechanical transport models in a single model, i.e. the ADE, the density-driven sinking is a separately driven process independent of the velocity field. Deconstructing the processes into seperate SSMs and forming a switching Markov model, we can define the system processes linguistically corresponding to their respective forces, for example, *advection*, and *sinking*.

From the perspective of purpose, there are several problems that may be needed to solve with a set of data observed either as a measure of concentration, or resistivity. The first, as discussed by Xu and Gómez-Hernández (2016), is the identification of source location and initial conditions of a system. The second is similar to the problem addressed in Chapter 7, whereby the underlying system parameters of velocity and dispersion are estimated. In both cases, with sufficient models, the estimated parameters can be used to approximate the system over time.

In the absence of specialist modelling software, the underlying models were resolved by random walk particle models. In the case of resolving resistivity cells, a space-based random walk method was used (see Section 2.4.4.1), binning the particles at each time step in bins corresponding to the spatial resolution of the resistivity cell. For point-based concentration methods, a time-domain random walk was used. In both instances, due to the updating of parameters, each method was reinitialised at every time-step. A consequence of this is that the complexity of each method exponentially scales with time.

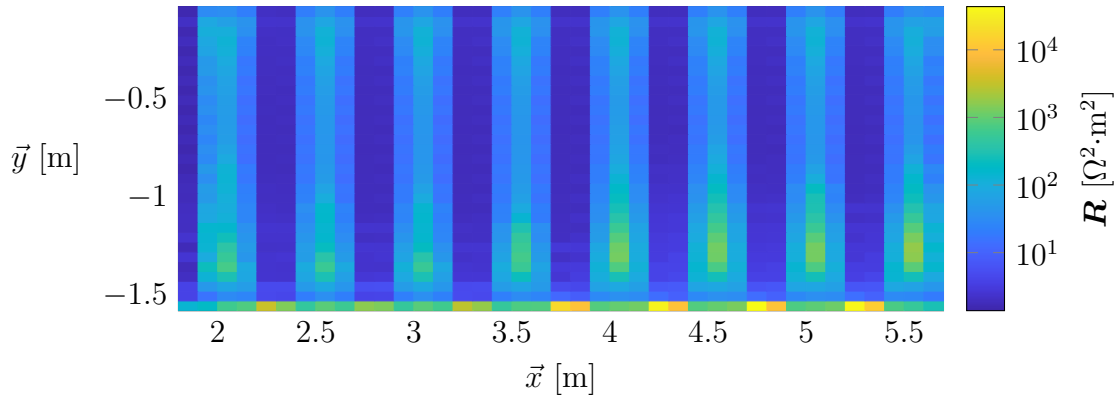


Figure 9.10: Measurement error covariance for each resistivity cell, taken from model resolution inversion output. The uncertainty is inversely proportional to the distance of each cell from measurement electrodes

9.4.2 Injection Types

In the system described by the pit test cell experiment, the tracer was injected over a period of hours. When modelling systems such as this, typically the initial conditions are preset as either instantaneous slug, or continuous. Where the injection is not continuous, i.e. there is an end to the injection period within the modelled scenario, there cannot by an instantaneous injection in a real-world scenario. As such, purely theoretical models describing the underlying system do not suitably model the initial conditions. For non-linear systems, small changes in the initial conditions may have drastic effects on the output of process models. In the case of the pit tracer test, the unscented type-1 fuzzy Kalman filter can be used to alleviate this issue by simultaneously modelling a system with both initial conditions. Consider the measurement setup as defined in Section 7.1.2 with 37 ERT time-slices imaged over 40×32 equal-sized cells. To simulate the resistivity the model, a space-based random walk particle solver can be used, with each particle's spatial location binned in a rectilinear grid equivalent to the resistivity cell structure. Weighting particles by their relative mass, a measure of concentration can be obtained in g/l. Assuming a linear relationship between the concentration, and the underlying fluid's electrical conductivity, Archie's law (Section 2.3.2) can be used to map the concentration to resistivity.

$$\rho = \kappa_s \left(\kappa_c + \kappa_m \theta_s\right)^{-1} \tag{9.5}$$

Here, the linear relationship between electrical conductivity κ (S/m) and solute concentration θ_s (g/l) is defined by the linear relationship $\kappa = \kappa_c + \kappa_m \theta_s$. The linear parameters were calculated as $\kappa_c = 0.07$ and $\kappa_m = 0.1465$, obtained by lab experiments (Kuras et al., 2009). This is all divided by the base conductivity of the system, and the reciprocal is found to obtain a measure of electrical resistivity (Ωm).

In terms of identifying uncertainty in the measurement domain, the output of the inversion process returns a measure of resolution for each cell. As discussed in Section 2.2, the model resolution is approximately inversely proportional to the distance of a cell from a measurement electrode. Assuming this resolution is independent between each cell, a diagonal matrix can be generated with the inverse resolution of each cell, weighted by some fixed scalar. The matrix is thus used to form the measurement noise covariance, $\boldsymbol{R}_t$ for all time-steps (as the cells are jointly inverted over time). In the experiments, the matrix is scaled by 0.1, though this value is chosen experimentally to be scaled appropriately to the change in log resistivity. The covariance value for each cell is plotted in its $(\vec{x}, \vec{z})$ position in Figure 9.10, showing clearly the high uncertainty along the bottom layer in the spaces between electrode boreholes.

The approximation of the initial conditions, as estimated using the UT1-FKF with two component SSMs can be viewed by the resulting firing strength distribution of each condition in the model. Figure 9.11 shows how the algorithm determines the effect of each model. It's clear that after the first 3 or 4 time-steps, the firing strength is largely confident in the slug model. This shift corresponds to the end of injection period after about 1.5 h. There remains some uncertainty in the system, with firing strength in some time-steps being almost 50/50. This is likely due to unresolved uncertainty, that is greatest towards the end of the model where, despite improved approximation of parameters, is affected by the reduced quality of observations as the tracer moves out of the imaging (and simulated) volume. While the initial injection period was set at (4.5, -0.4) m, in line with the true injection, the velocity field in the horizontal and vertical directions were approximated. The initial settings for the horizontal flow field were assigned randomly between 4 and 5 m/day, and the vertical velocity was initialised at 0 m/day.

It's clear from the results that the method is reasonably effective at switching to the correct model during the filtering. The velocity estimations at each time-step in the filter are shown in Figure 9.12. The approximation of both velocities is relatively accurate. In the horizontal velocity, the velocity estimate is contained

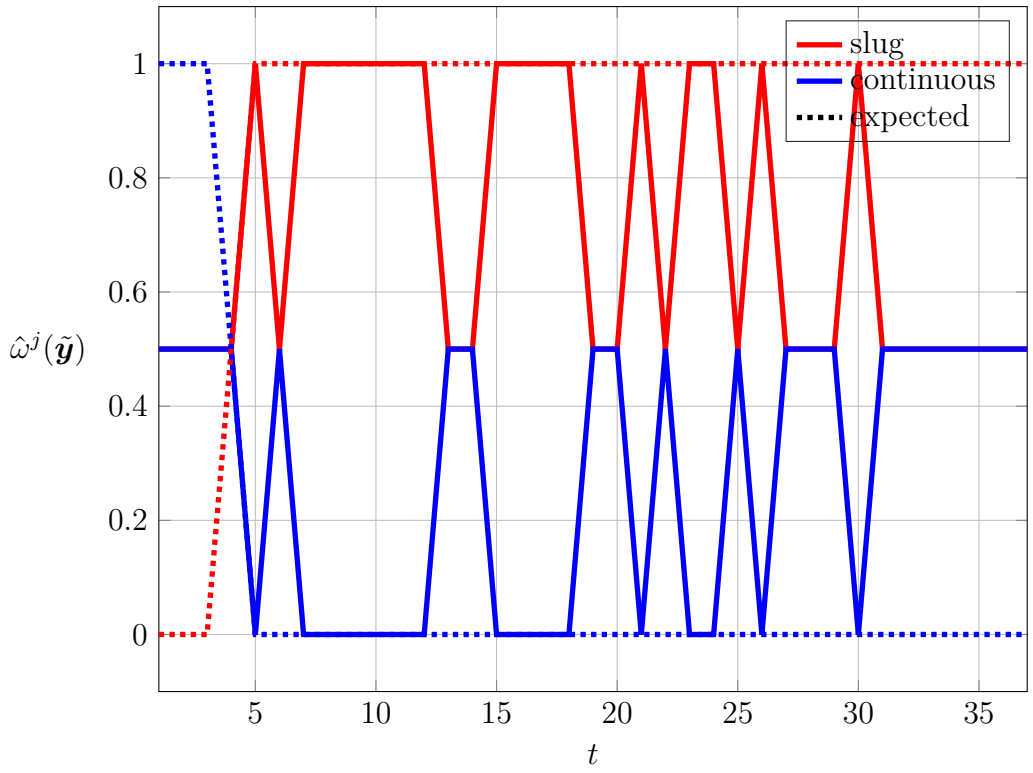


Figure 9.11: Firing strength for component processes of the continuous and instantaneous injection models. The dashed lines represent the expected firing strengths

within the expected range, with convergence to an estimate ≈4.6 m/day. Here, the confidence interval shrinks, perhaps displaying some over-confidence. For the vertical velocity, corresponding to the tracer sinking, however, the estimated velocity field is approximated reasonably well despite not being explicitly modelled as a separate process. The filtered estimation converges to approximately −0.38 m/day after 8-9 time steps, but as the filtering progress, the confidence in the estimation increases. This may, in part, be due to both the sinking implicitly modelled as a constant velocity field, where there may be some gravity-driven acceleration. Additionally, it uncertainty may be due to resolution degradation in the lower cells of the ERT model.

9.4.3 Density-Driven Sinking

A switching model for approximating the horizontal velocity field and the density driven sinking separately was used to approximate the tracer motion in the system. Using the random walk model in (2.14) to model particles from the known injection source, acted upon by the horizontal velocity field and gravity-accelerated

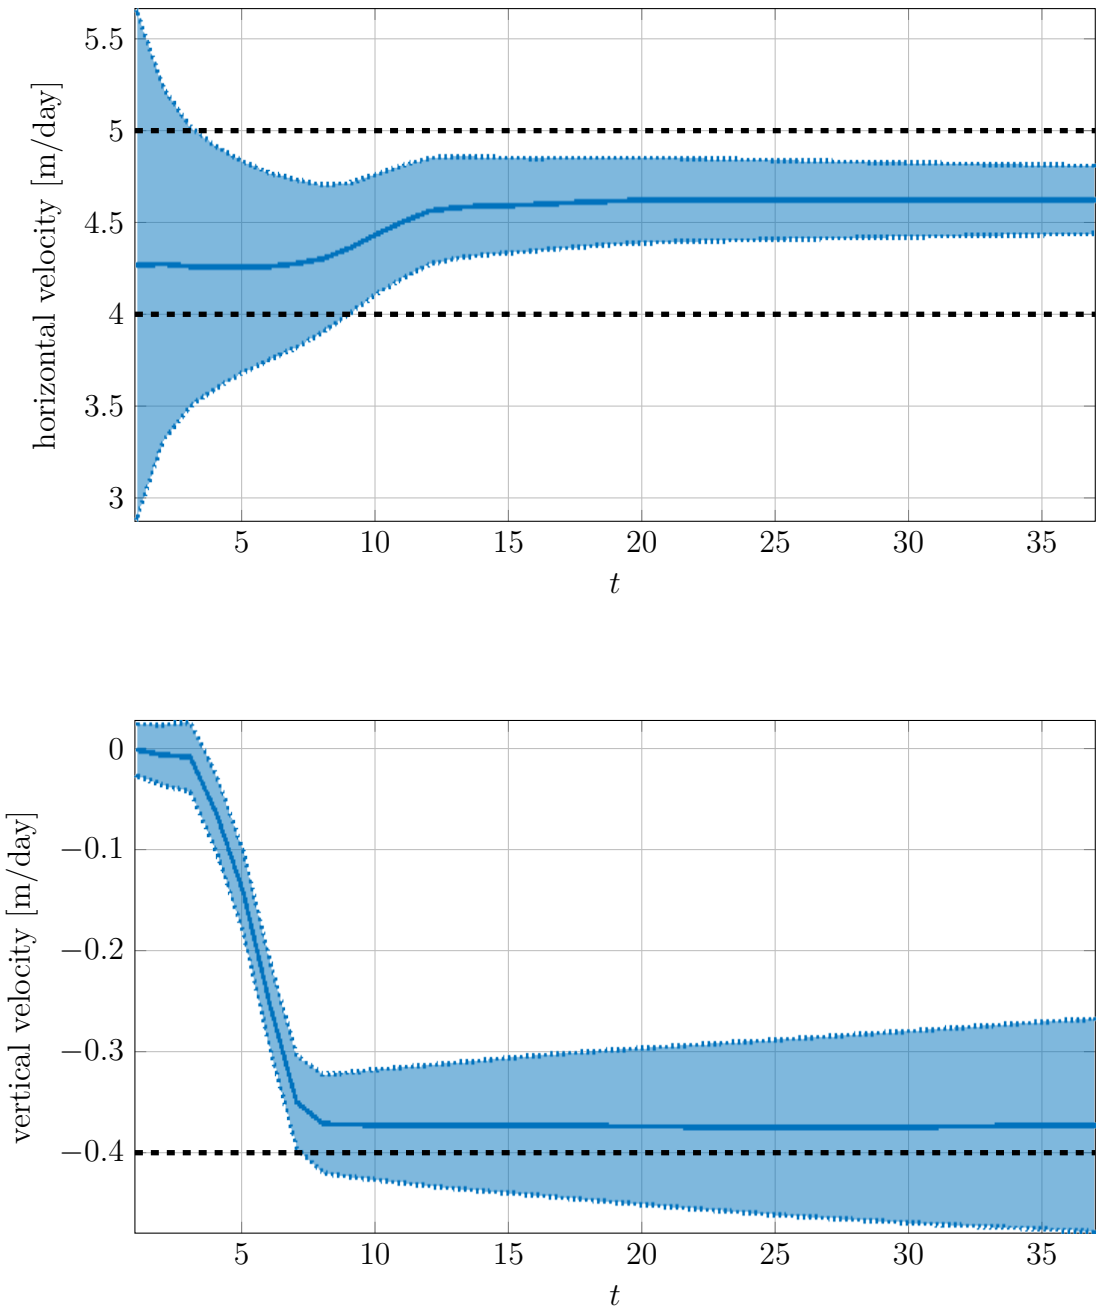


Figure 9.12: Velocity estimates for mixed initial condition UT1-FKF applied to the pit test cell ERT experiment. Shaded regions represent 95% confidence intervals, and the black dashed lines denote calculated expected values

density-driven sinking as two seperate processes. Using the ERT cells as measured observations, as in Section 9.4.2, with resolution terms guiding observation noise, the UT1-FKF was used to model the centre of mass of the tracer. The filtered centre-of-mass of the tracer is shown in Figure 9.13.

Here it shows predominance by the vertical model, with almost no deviation in

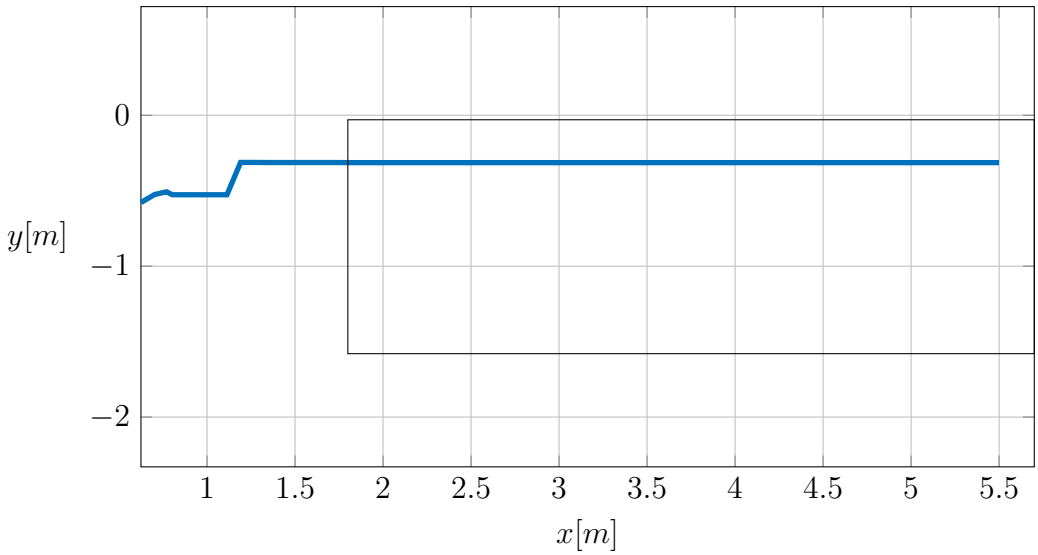


Figure 9.13: Filtered centre-of-mass of the tracer plume using separable physical processes, horizontal velocity and density-driven sinking. The embedded black rectangle denotes the ERT imaging plane

---

the vertical position for the main duration of the filtering. Towards the end, the detected tracer begins to sink, though this is out of the imaging plane. The results from the filtering here are not as high performing as previous experiments, notably from the tracer tracking framework in Chapter 7 which exploits feature detectors and simple linear models. However, it has shown that there is some learning ability for using the switching model. However, the representation used is limited in resolving the ERT. This might be due to limitations of the fuzzy inference engine but, given the positive results over previous results, it is more likely due to ill-constructed model and poorly tuned parameters. This serves to highlight some of the limitations in the proposed model that exist in the Kalman filter and the unscented Kalman filter are not necessarily resolved by the fuzzy inference system, and highlights the potential for further study.

10

Discussion

10.1 Contributions

The key contributions of this thesis fall into two categories. The first is the contributions to the field of geophysical monitoring, specifically in the automated monitoring for electrical resistivity monitoring. The initial project sought to identify techniques that could be used to estimate quantifiable properties of the shallow subsurface in 2-D and 3-D static and time-lapse ERT models. In Chapter 6, the detection of distinct features in ERT data was studied. The developed clustering methodology takes the distinguished regions and identifies boundaries between them. The initial version of this methodology was published in Geophysical Journal International (Ward et al., 2014b), applied to two sites in the UK for means of quantifying subsurface structure and isolating the bedrock layer interface as an elevation map. The clustering approach was found to be high performing, both in terms of comparison with other methods for bedrock detection, and also against empirical measurements of the bedrock depth. The uncertainty captured by the fuzzy clusters allowed for intervals indicating estimation uncertainty to be formulated that captured entirely the true bedrock interface. This study demonstrated the effectiveness of both ERT for static site investigation, and of the clustering algorithm. The proposed GFCM was further used in the study of a wetland at a site of special scientific interest, published in Water Resources Research (Chambers et al., 2014b), further demonstrating its use for automated analysis.

The extension of the developed clustering method was applied to a time-lapse problem, in the monitoring groundwater drawdown and rebound in the quarry dewatering and published in Engineering Geology (Chambers et al., 2015b). The

clustering method worked well at isolating regions of significant change that corresponded with the changes of water level and moisture content in the quarry floor. However this work highlighted limitations in the methodology, specifically that while the clustering was good at isolating separate regions, and therefore the boundaries between them, isolating distinct features, or regions-of-interest, required some manual intervention. These limitations served as motivation for appropriate change detection and tracking methods for time-lapse monitoring. The maximally stable volumes adapted for ERT were shown to be effective at identifying regions, and in combination with the linear-discrete Kalman filter and multi-target tracking methods, a framework for filtering and tracking contaminant and tracer plumes was developed. The tracer tracking framework served as an effective tool for both visualising and automating tracer monitoring, but also proved useful for estimating system parameters. This work was published in Water Resources Research (Ward et al., 2016).

The contributions of this thesis to geoelectrical monitoring have been relatively substantial, in that it has furthered the case for electrical resistivity tomography for remote, non-invasive monitoring. The research developed over Chapter 6 and Chapter 7 has been the subject of, or a significant part of, a number of talks and presentations at a range of workshops and conferences, including the European Geosciences Union AGM (Ward et al., 2014a), the American Geophysical Union Fall Meeting (Ward et al., 2015a), the International Workshop on Geoelectrical Monitoring (Chambers et al., 2015a; Ward et al., 2015b); and the Near Surface Geoscience European Meeting of Environmental and Engineering Geophysics (Uhlemann et al., 2016).

The limitations in the representation provided by the tracer tracking framework served as motivations for the fuzzy Kalman filter adaptations. This work, formalised in Chapter 8, introduces a means to model Markov-switching state-space models by incorporating fuzzy inference over the switching of component processes. The derivation shows that the component processes minimises mean-squared error in each component process, and uses fuzzy sets to infer between them. An illustrative example showed how a non-linear system could be approximated through a set of linear-discrete processes by effectively offloading the inference of non-linearity to the fuzzy inference system. The proposed methodology was further developed to handle non-linear observation models, by utilising the unscented transform. In addition, the characterisation of uncertainty was extended by deriving the switching model as an interval type-2 fuzzy inference system. Experiments were performed on synthetic and real-world systems, and the empirical outcome demonstrated

high performance by the proposed methodology. Results further indicated that the capabilities of the type-2 version was no better, nor worse, than the type-1 filter, highlighting issues about the approximation of uncertainty. It was demonstrated that under particular conditions in the set-up of the linear type-1 fuzzy Kalman filter, that the inference system was analogous to a Gaussian mixture model. Overall, the fuzzy Kalman filters derived in this thesis present a novel means for handling switching models, that are theoretically rigorous and require fewer assumptions than similar methods, with equal or better performance on the tested examples. The type-1 FKF and its unscented extension were submitted for review in IEEE Transactions on Signal Processing (Ward et al., submitted to).

10.2 Future Work

Continuing the work developed in this thesis, the findings of machine learning applications to geoelectrical monitoring may be extended in scope to a variety of scenarios. Already, the flexibility of the clustering methods in static and time-lapse models of both 2-D and 3-D resistivity tomography models has been demonstrated in a number of domain publications. The tracer tracking framework outlined the effectiveness of Kalman filters and feature tracking for ERT monitoring. With sufficiently representative process models incorporated into the framework, the technique may be applied to other monitoring problems, and serve as a part-of-the-package analysis tool for remote monitoring systems such as PRIME (Chambers et al., 2016).

The development of the FKF and its extensions allows for the improvement of and better exploitation of the uncertainty inherent to the modelling problem. Given domain knowledge, the FKF may be applied to a variety of problems in subsurface modelling, both with indirect measurements, such as in the case of ERT, or with direct sensing, such as in the work by Xu and Gómez-Hernández (2015a). For example, in source detection for tracer tracking, the initial value problem may be solved for multiple possible process models, as demonstrated briefly in Section 9.4.2. Bayesian filtering in similar systems has also been used for hydraulic map and facies detection (Le Borgne et al., 2008; Xu and Gómez-Hernández, 2015b), and in these areas the UT1-FKF may be applicable for both estimation of hydraulic parameters, and for filtering of the state process given noisy measurements (Bianchi and Pedretti, 2017).

In discussion of the FKF, it was shown how, for Gaussian membership, the inference over component processes formed a mixture of Gaussians, with the

firing strength itself drawing from a Gaussian. The proceeding derivation of the mixture model joint distribution has the respective correlation components of the covariance matrix set to zero. However in the filter, an estimate of the cross-covariance is maintained in the form of $\boldsymbol{P}_t^{j(xy)}$, as used in the definition of gain $\mathbf{K}_t^j$. The extension of the derived fuzzy Kalman filter to use type-2 fuzzy logic seeks to incorporate the uncertainty in the state into the membership functions. The limitations of propagating an entire distribution across a non-linear function mean that approximations in type-2 fuzzy logic suffer in their simplicity. It's observable in the case of the interval type-2 FKF, where the results do not show any significant performance increase against the unscented type-1 implementation. Part of this reason is likely because under linear or near-linear transitions, the IT2 representation would collapse into a T1 fuzzy set by type reduction if the propagated confidence bounds are approximately symmetrical about the expected value. Indeed, the limited representation of intervals may be extended by increasing either the confidence value, z^*, for 99% or 99.9%, etc. confidences, or indeed representing the uncertainty in the distribution by an α-cut type-2 set. Such a representation would only involve finite-discrete approximations of the state uncertainty by propagating perturbations of the state by evaluations of $P_{t|t-1}^{j(x)}$, but the process is limited in its requirements for discretisation.

An ideal, if not optimal, approximation of the unified state by this derived fuzzy Kalman filter under the conditions that leads to a Gaussian mixture model, would by proportional to the joint distributions of each SSM. Given the corresponding correlation variance, the estimated unified state would be proportional to the mixture of joint distributions for each process.

$$\begin{aligned} p(x_t \,|\, y_t) &\propto \sum_{j=1}^{r} p\Big(x_t^j, y_{t|t-1}^j \,|\, x_{t-1}, y_t\Big) \\ &= \sum_{j=1}^{r} \mathcal{N}\left(\begin{bmatrix} \boldsymbol{x}_{t|t}^j \\ \boldsymbol{y}_{t|t-1}^j \end{bmatrix}, \begin{bmatrix} \boldsymbol{P}_{t|t}^{j(x)} & \boldsymbol{P}_t^{j(xy)} \\ \boldsymbol{P}_t^{j(xy)\mathrm{T}} & \boldsymbol{P}_t^{j(y)} \end{bmatrix} \right) \end{aligned} \tag{10.1}$$

How this cross-covariance might be assimilated into the mixture model is an potential area for future study. The inference system modelling the mixture model would need to optimise the parameters to maximise the likelihood. This might be done through some iterative method, like expectation maximisation, where the number of components is predefined based on the switching model assumptions (Moon, 1996), or by presupposing the weights over some iterative filter, similar to the Gaussian sum particle filter proposed by Kotecha and Djuric (2003). Extending

the fuzzy inference system to utilise concepts from general type-2 fuzzy logic that may fully represent the uncertainty over the component process models: this is an idea that may benefit from further investigation.

10.3 Conclusions

This thesis sought to answer the three hypotheses outlined, on robust identification of features-of-interest; their use in the estimation and prediction of system dynamics; and how uncertainty in such systems can be exploited to both make meaningful predictions and represent the certainty around these estimations. The importance of knowing the degree to which a model prediction is confident in its estimation lies both in the reasonable weighting of observations and of prior expectations. Knowing how confident the expectation of a system and how reliable an observation of it is allows for reasoned inference with reliable confidence in the posterior prediction of a future state. Furthermore, the quantified uncertainty as a feature of the output in any prediction is a useful feature for the application of tools that make use of these models, as they can better inform decisions that have real-world implications.

Through experiments focussed on the localising region boundaries in static resistivity model, the need for quantitative metrics of uncertainty gave way to the exploitation of fuzzy logic by means of the fuzzy clustering. The choice of clustering detracted from the desired automation of the analysis, and here a representation of probability density served as an initialiser and distance metric. The methodology proved effective for the application, with favourable performance compared with state-of-the-art methods, with the added benefit of a confidence measure in the output.

Moving beyond static analysis to time-lapse modelling, limitations in the clustering approach lead to the utilisation of recursive Bayesian filtering in the form of the Kalman filter. The tracer tracking framework exploited the uncertainty in the observations, here defined by a set of connected components per time-slice, when correcting the prior assumptions. The distance metric in the association of distinct objects was the Mahalanobis distance, minimising the overall distances of each measurement with the filter observation likelihood. The use of uncertainty here comes from both the probabilistic measure of distance, and that the tracks conforming to the expected process gave more reliable estimates, both allowing for easy automated post-processing, and in giving an expectation of the state parameters and the degree of confidence of the estimate.

Extending the dimensionality of the observation model to remove the necessity for feature detection served as motivation for adapting the tracking framework and its underlying components. The limitations of the need for tracking multiple features in a known single-feature system, coupled with the modelling challenges of the pit test cell experiment, led to the need for a filter that could factor in multiple assumptions and switch between them automatically. Posing the problem as a Markov switching gave rise to the need for a robust and data-driven method that did not rely on additional assumptions, such as a matrix of switching probabilities. Taking inspiration from the combination of fuzzy logic and probabilistic modelling in the distribution-guided fuzzy clustering method, a fuzzy inference system placed over the switching of the state-space models. This removed the need to explicitly calculate and maintain switching probabilities, and also reduced computational and memory requirements in the maintenance of parallel Kalman filters, a feature that distinguishes the methodology from both the tracer tracking framework, and the IMM filters.

Given limited scope of the ERT experiments in validating the FKF, the problem of target tracking in the context of radar-like systems was used for testing the performance of the proposed filter. Synthetic models over multiple observation types gave a sufficiently difficult task to test both the type-1 and type-2 versions of the FKF on. The algorithms implemented on a switching assumption consistently outperformed the single SSM filters, both using discrete-Gaussian moment filtering in the unscented Kalman filter, and by Monte Carlo approximation using sequential importance resampling. Of the switching model-based approaches, the proposed FKF performed better than, or no worse than, the unscented interacting multiple model, with both consistently outperforming the extended IMM. The synthetic experiments demonstrated the effectiveness of fuzzy inference over the switching model, compared with pre-defined probabilities. While the switching matrices were not optimised, the given values were set based on previously published set-up for the same experiment (Hartikainen et al., 2011).

The formulation of the type-1 fuzzy Kalman filter means that from a mean squared error perspective, the approximation of state expectation and error covariance is optimal, and near-optimal for the UT1-FKF, in its weighting by $\mathbf{K}^j$, the gain corresponding to the j^{th} component process. Given that the switching model encompasses multiple processes with unknown effect at a given time-step, the methodology is suboptimal in the inference across the process switching. This is true both for the IMM and the FKF. Where the interval type-2 fuzzy inference system is used, with confidence intervals defined over the prior estimation of the state,

the experimental results show significant difference in capability with the type-1 version. As discussed previously, this is likely due to the approximation methods being insufficient to capture more meaningful information of the uncertainty, both in the nature of the discrete interval, and in part due to the Bayesian nature of the likelihood of the observation estimation. While a general type-2 fuzzy logic approach may increase representation of the uncertainty, it may follow that the uncertainty of the state is sufficiently propagated through the observation model, and therefore represented in the fuzzy membership definition of the observation estimation likelihood.

In the application of the FKF to ERT, there is clear scope for continued study, both for use of geoelectrical monitoring, but also in general hydrogeological modelling. In resolving the initial condition problem with ERT measurements, the experiment shows that the switching model and the proposed FKF are suitable for modelling these situations. Here there is scope for wider modelling, with more groundwater models guided by expert input. Likewise, the observation model is not limited to geoelectrical images, but may be resolve using direct sampling, resolved through point models such as time-domain random walk, or proprietary purpose-built software such as MODFLOW (Harbaugh et al., 2000), the potential capabilities of the proposed filtering methodology is wide-reaching. Experiments of this manner were beyond the scope of testing for this thesis due to lack of appropriate software or high-level groundwater models available. However it can be concluded that the methods developed in this thesis are both contributory in their theory and also in to the field of geoelectrical monitoring for a wide range of purposes.

References

K. W. Ahn and K.-S. Chan. Approximate conditional least squares estimation of a nonlinear state-space model via an unscented Kalman filter. *Computational Statistics & Data Analysis*, 69:243–254, Jan. 2014.

C. Andrieu, N. de Freitas, A. Doucet, and M. I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50(1-2):5–43, 2003.

I. Arasaratnam and S. Haykin. Cubature Kalman filters. *IEEE Transactions on automatic control*, 54(6):1254–1269, 2009.

K. T. Atanassov. Intuitionistic fuzzy sets. *Fuzzy sets and Systems*, 20(1):87–96, 1986.

M. Audebert, R. Clément, N. Touze-Foltz, T. Günther, S. Moreau, and C. Duquennoi. Time-lapse ERT interpretation methodology for leachate injection monitoring based on multiple inversions and a clustering strategy (MICS). *Journal of Applied Geophysics*, 2014.

F. Auger, M. Hilairet, J. M. Guerrero, E. Monmasson, T. Orlowska-Kowalska, and S. Katsura. Industrial applications of the Kalman filter: A review. *IEEE Transactions on Industrial Electronics*, 60(12):5458–5471, 2013.

M. Auger-Méthé, C. Field, C. M. Albertsen, A. E. Derocher, M. A. Lewis, I. D. Jonsen, and J. M. Flemming. State-space models dirty little secrets: even simple linear Gaussian models can have estimation problems. *Scientific reports*, 6:26677, 2016.

Y. Ban, S. Ba, X. Alameda-Pineda, and R. Horaud. Tracking multiple persons based on a variational bayesian model. In *European Conference on Computer Vision*, pages 52–67. Springer, 2016.

M. Barut, S. Bogosyan, and M. Gokasan. An EKF-based estimator for the speed sensorless vector control of induction motors. *Electric power components and systems*, 33(7):727–744, 2005.

J. Bear. *Dynamics of fluids in porous media*. Courier Corporation, 1972.

J. C. Bezdek and J. C. Dunn. Optimal fuzzy partitions: A heuristic for estimating the parameters in a mixture of normal distributions. *IEEE Transactions on Computers*, 100(8):835–838, 1975.

M. Bianchi and D. Pedretti. Geological entropy and solute transport in heterogeneous porous media. *Water Resources Research*, 53(6):4691–4708, 2017.

A. Binley, G. Cassiani, R. Middleton, and P. Winship. Vadose zone flow model parameterisation using cross-borehole radar and resistivity imaging. *Journal of Hydrology*, 267(3):147–159, 2002.

J. Bodin. From analytical solutions of solute transport equations to multidimensional time-domain random walk (TDRW) algorithms. *Water Resources Research*, 51(3):1860–1871, 2015.

Z. I. Botev, J. F. Grotowski, D. P. Kroese, et al. Kernel density estimation via diffusion. *The Annals of Statistics*, 38(5):2916–2957, 2010.

J. C. Bowling, C. Zheng, A. B. Rodriguez, and D. L. Harry. Geophysical constraints on contaminant transport modeling in a heterogeneous fluvial aquifer. *Journal of contaminant hydrology*, 85(1):72–88, 2006.

M. D. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. Van Gool. Robust tracking-by-detection using a detector confidence particle filter. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 1515–1522. IEEE, 2009.

S. Buckland, K. Newman, L. Thomas, and N. Koesters. State-space models for the dynamics of wild animal populations. *Ecological modelling*, 171(1):157–175, 2004.

I. Butera, M. G. Tanda, and A. Zanini. Simultaneous identification of the pollutant release history and the source location in groundwater by means of a geostatistical approach. *Stochastic Environmental Research and Risk Assessment*, 27(5):1269–1280, 2013.

A. A. Butt and R. T. Collins. Multi-target tracking by lagrangian relaxation to min-cost network flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1846–1853, 2013.

M. Camporese, G. Cassiani, R. Deiana, P. Salandin, and A. Binley. Coupled and uncoupled hydrogeophysical inversions using ensemble Kalman filter assimilation of ERT-monitored tracer test data. *Water Resources Research*, 2015.

P.-H. Cao, X. Wang, S.-S. Fang, X.-W. Cheng, K.-P. Chan, X.-L. Wang, X. Lu, C.-L. Wu, X.-J. Tang, R.-L. Zhang, et al. Forecasting influenza epidemics from multi-stream surveillance data in a subtropical city of China. *PloS one*, 9(3): e92945, 2014.

F. Cassola and M. Burlando. Wind speed and wind energy forecast through Kalman filtering of numerical weather prediction model output. *Applied Energy*, 99:154 – 166, 2012.

S. Challa, M. R. Morelande, D. Mušicki, and R. J. Evans. *Fundamentals of object tracking*. Cambridge University Press, 2011.

J. Chambers, P. Meldrum, D. Gunn, P. Wilkinson, O. Kuras, A. Weller, and R. Ogilvy. Hydrogeophysical monitoring of landslide processes using automated time-lapse electrical resistivity tomography (alert). In *Near Surface 2009-15th EAGE European Meeting of Environmental and Engineering Geophysics*, 2009.

J. Chambers, P. Meldrum, D. Gunn, P. Wilkinson, S. Uhlemann, O. Kuras, and R. Swift. Proactive infrastructure monitoring and evaluation (prime): a new electrical resistivity tomography system for remotely monitoring the internal condition of geotechnical infrastructure assets. In *3rd International Workshop on Geoelectrical Monitoring (GELMON)*. 3rd International Workshop on Geoelectrical Monitoring (GELMON), 2015a.

J. E. Chambers, O. Kuras, P. I. Meldrum, R. D. Ogilvy, and J. Hollands. Electrical resistivity tomography applied to geologic, hydrogeologic, and engineering investigations at a former waste-disposal site. *Geophysics*, 71(6):B231–B239, 2006.

J. E. Chambers, P. B. Wilkinson, and P. I. Meldrum. Investigation and monitoring of aggregate deposits using geo-electrical resistivity imaging (IMAGER). Technical report, British Geological Survey, Nottingham, 2011.

J. E. Chambers, P. B. Wilkinson, D. Wardrop, A. Hameed, I. Hill, C. Jeffrey, M. H. Loke, P. I. Meldrum, O. Kuras, M. Cave, and D. A. Gunn. Bedrock detection beneath river terrace deposits using three-dimensional electrical resistivity tomography. *Geomorphology*, 177-178:17–25, Dec. 2012.

J. E. Chambers, P. I. Meldrum, P. B. Wilkinson, B. Matthews, O. Kuras, D. A. Gunn, S. S. Uhlemann, and L. S. Oxby. Automated electrical resistivity tomography monitoring of quarry dewatering. In *Near Surface*, number September 2013, 2013.

J. E. Chambers, D. A. Gunn, P. B. Wilkinson, P. I. Meldrum, E. P. Haslam, S. Holyoake, M. Kirkham, O. Kuras, A. Merritt, and J. Wragg. 4D electrical resistivity tomography monitoring of soil moisture dynamics in an operational railway embankment. *Near Surface Geophysics*, (12):61–72, Feb. 2014a.

J. E. Chambers, P. B. Wilkinson, S. S. Uhlemann, J. P. R. Sorensen, C. Roberts, A. J. Newell, W. O. C. Ward, A. M. Binley, P. J. Williams, D. C. Gooddy, G. Old, and L. Bai. Derivation of lowland riparian wetland deposit architecture using geophysical image analysis and interface detection. *Water Resources Research*, 50:5886–5905, 2014b.

J. E. Chambers, P. I. Meldrum, P. B. Wilkinson, W. O. C. Ward, C. Jackson, B. Matthews, P. Joel, O. Kuras, L. Bai, S. Uhlemann, and D. Gunn. Spatial monitoring of groundwater drawdown and rebound associated with quarry dewatering using automated time-lapse electrical resistivity tomography and distribution guided clustering. *Engineering Geology*, 193:412–420, 2015b.

J. E. Chambers, P. I. Meldrum, P. B. Wilkinson, D. Gunn, S. Uhlemann, O. Kuras, R. Swift, C. Inauen, and S. Butler. Remote condition assessment of geotechnical assets using a new low-power ERT monitoring system. In *Near Surface Geoscience 2016-22nd European Meeting of Environmental and Engineering Geophysics*, 2016.

G. Chen, Q. Xie, and L. S. Shieh. Fuzzy Kalman filtering. *Information Sciences*, 109(1-4):197–209, 1998.

S. Chen. Kalman filter for robot vision: a survey. *IEEE Transactions on Industrial Electronics*, 59(11):4409–4420, 2012.

K. Csilléry, M. Seignobosc, V. Lafond, G. Kunstler, and B. Courbaud. Estimating long-term tree mortality rate time series by combining data from periodic inventories and harvest reports in a Bayesian state-space model. *Forest ecology and management*, 292:64–74, 2013.

F. Cupola, M. G. Tanda, and A. Zanini. Laboratory sandbox validation of pollutant source location methods. *Stochastic environmental research and risk assessment*, 29(1):169–182, 2015.

P. J. Denning, D. E. Comer, D. Gries, M. C. Mulder, A. Tucker, A. J. Turner, and P. R. Young. Computing as a discipline. *Computer*, 22(2):63–70, 1989.

G. A. Einicke and L. B. White. Robust extended Kalman filtering. *IEEE Transactions on Signal Processing*, 47(9):2596–2599, 1999.

S. Eleftheriadis, T. F. Nicholson, M. P. Deisenroth, and J. Hensman. Identification of Gaussian process state space models. *arXiv preprint arXiv:1705.10888*, 2017.

M. Elwaseif and L. D. Slater. Quantifying tomb geometries in resistivity images using watershed algorithms. *Journal of Archaeological Science*, 37(7):1424–1436, July 2010.

M. Elwaseif and L. D. Slater. Improved resistivity imaging of targets with sharp boundaries using an iterative disconnect procedure. *Journal of Environmental & Engineering Geophysics*, 17(2):89–101, 2012.

M. Elwaseif and L. D. Slater. Reconstruction of discrete resistivity targets using coupled artificial neural networks and watershed algorithms. In *Near Surface Geophysics*, volume 11, pages 517–530, 2013.

I. Eyoh, R. John, and G. De Maere. Interval type-2 intuitionistic fuzzy logic system for non-linear system prediction. In *Systems, Man, and Cybernetics (SMC), 2016 IEEE International Conference on*, pages 001063–001068. IEEE, 2016.

I. Eyoh, R. John, and G. De Maere. Interval type-2 intuitionistic fuzzy logic for regression problems. *IEEE Transactions on Fuzzy Systems*, 2017.

C. W. Fetter. *Applied Hydrogeology*. Pearson Education, Edinburgh, 2001.

M.-T. Gan, M. Hanmandlu, and A. H. Tan. From a Gaussian mixture model to additive fuzzy systems. *IEEE Transactions on Fuzzy Systems*, 13(3):303–316, 2005.

G. H. Gebhardt, A. G. Kupcsik, G. Neumann, et al. The kernel Kalman rule-efficient nonparametric inference with recursive least squares. In *AAAI*, pages 3754–3760, 2017.

A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian data analysis*. Chapman & Hall/CRC, Boca Raton, Florida, second edition, 2004.

M. Gharibi and L. R. Bentley. Resolution of 3-D electrical resistivity images from inversions of 2-D orthogonal lines. *Journal of Environmental & Engineering Geophysics*, 10(4):339–349, 2005.

N. Gordon, D. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F - Radar and Signal Processing*, 140(2):107–113, 1993.

M. R. Gupta, Y. Chen, et al. Theory and use of the em algorithm. *Foundations and Trends® in Signal Processing*, 4(3):223–296, 2011.

E. B. Haarder, A. Binley, M. C. Looms, J. Doetsch, L. Nielsen, and K. H. Jensen. Comparing plume characteristics inferred from cross-borehole geophysical data. *Vadose Zone Journal*, 11(4), 2012.

T. Haarnoja, A. Ajay, S. Levine, and P. Abbeel. Backprop kf: Learning discriminative deterministic state estimators. In *Advances in Neural Information Processing Systems*, pages 4376–4384, 2016.

J. D. Hamilton. Specification testing in markov-switching time-series models. *Journal of econometrics*, 70(1):127–157, 1996.

A. W. Harbaugh, E. R. Banta, M. C. Hill, and M. G. McDonald. Modflow-2000, the u. s. geological survey modular ground-water model-user guide to modularization concepts and the ground-water flow process. *Open-file Report. U. S. Geological Survey*, (92):134, 2000.

J. Hartikainen, A. Solin, and S. Särkkä. EKF/UKF toolbox for MATLAB v1.3, 2011. Software url:http://becs.aalto.fi/en/research/bayes/ekfukf [Accessed 22/11/17].

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

H. J. Hendricks Franssen and W. Kinzelbach. Ensemble Kalman filtering versus sequential self-calibration for inverse modelling of dynamic groundwater flow systems. *Journal of Hydrology*, 365(3):261–274, 2009.

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

W. Hu, X. Li, W. Luo, X. Zhang, S. Maybank, and Z. Zhang. Single and multiple object tracking using log-euclidean riemannian subspace and block-division appearance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(12):2420–2440, 2012.

S. Huang and G. Dissanayake. Convergence and consistency analysis for extended Kalman filter based SLAM. *IEEE Transactions on robotics*, 23(5):1036–1049, 2007.

B. R. Hunt, E. J. Kostelich, and I. Szunyogh. Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D: Nonlinear Phenomena*, 230(1):112–126, 2007.

T. C. Johnson, L. D. Slater, D. Ntarlagiannis, F. D. Day-Lewis, and M. Elwaseif. Monitoring groundwater-surface water interaction using time-series and time-frequency analysis of transient three-dimensional electrical resistivity changes. *Water Resources Research*, 48(7), 2012.

S. Julier and J. Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422, Mar. 2004.

R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME - Journal of Basic Engineering*, 82(D):35–45, 1960.

M. Karl, M. Soelch, J. Bayer, and P. van der Smagt. Deep variational Bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.

A. Kemna, J. Vanderborght, B. Kulessa, and H. Vereecken. Imaging and characterisation of subsurface solute transport using electrical resistivity tomography (ERT) and equivalent transport models. *Journal of Hydrology*, 267:125–146, 2002.

G. Klir and B. Yuan. *Fuzzy sets and fuzzy logic*, volume 4. Prentice hall New Jersey, 1995.

K. Kobayashi, K. C. Cheok, and K. Watanabe. Fuzzy logic rule-based Kalman filter for estimating true speed of a ground vehicle. *Intelligent Automation & Soft Computing*, 1(2):179–190, 1995.

J. Koestel, A. Kemna, M. Javaux, A. Binley, and H. Vereecken. Quantitative imaging of solute transport in an unsaturated and undisturbed soil monolith with 3-D ERT and TDR. *Water Resources Research*, 44(12), 2008.

J. H. Kotecha and P. M. Djuric. Gaussian sum particle filtering. *IEEE Transactions on signal processing*, 51(10):2602–2612, 2003.

L. Kratz and K. Nishino. Tracking with local spatio-temporal motion patterns in extremely crowded scenes. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 693–700. IEEE, 2010.

R. G. Krishnan, U. Shalit, and D. Sontag. Deep Kalman filters. *arXiv preprint arXiv:1511.05121*, 2015.

O. Kuras, J. D. Pritchard, P. I. Meldrum, J. E. Chambers, P. B. Wilkinson, R. D. Ogilvy, and G. P. Wealthall. Monitoring hydraulic processes with automated time-lapse electrical resistivity tomography (ALERT). *Comptes Rendus - Geoscience*, 341(10-11):868–885, Oct. 2009.

H.-K. Lam and L. D. Seneviratne. Stability analysis of interval type-2 fuzzy-model-based control systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 38(3):617–628, 2008.

C. H. Lampert. Predicting the future behavior of a time-varying probability distribution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 942–950, 2015.

T. Le Borgne, M. Dentz, and J. Carrera. Spatial markov processes for modeling lagrangian particle dynamics in heterogeneous porous media. *Physical Review E*, 78(2):026308, 2008.

N. Linde, J. Chen, M. B. Kowalsky, and S. Hubbard. Hydrogeophysical parameter estimation approaches for field scale characterization. *Applied Hydrogeophysics*, pages 9–44, 2006.

J. S. Liu and R. Chen. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93(443):1032–1044, Feb. 1998.

Y. Liu, H. Li, and Y. Q. Chen. Automatic tracking of a large number of moving targets in 3-D. In *European Conference on Computer Vision*, pages 730–742. Springer, 2012.

M. Loke, J. Chambers, D. Rucker, O. Kuras, and P. Wilkinson. Recent developments in the direct-current geoelectrical imaging method. *Journal of Applied Geophysics*, 95:135–156, 2013.

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

M. H. Loke. Tutorial: 2-d and 3-d electrical imaging surveys. Technical Report October, GeoTomo, 2012.

D. Longo, G. Muscato, and V. Sacco. Localization using fuzzy and Kalman filtering data fusion. In *Proceedings of the 5th International Conference on Climbing and Walking Robots, Paris, France*, pages 263–270, 2002.

J. Lu, G. Chen, and H. Ying. Predictive fuzzy pid control: theory, design and simulation. *Information Sciences*, 137(1):157–187, 2001.

Q. Lu, P. Shi, H.-K. Lam, and Y. Zhao. Interval type-2 fuzzy model predictive control of nonlinear networked control systems. *IEEE Transactions on Fuzzy systems*, 23(6):2317–2328, 2015.

W. Luo, J. Xing, X. Zhang, X. Zhao, and T.-K. Kim. Multiple object tracking: A literature review. *arXiv preprint arXiv:1409.7618*, 2014.

N. Mansoor and L. Slater. Aquatic electrical resistivity imaging of shallow-water wetlands. *Geophysics*, 72(5):F211–F221, 2007.

L. G. Marín, F. Valencia, and D. Sáez. Prediction interval based on type-2 fuzzy systems for wind power generation and loads in microgrid control design. In *Fuzzy Systems (FUZZ-IEEE), 2016 IEEE International Conference on*, pages 328–335. IEEE, 2016.

J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. In *British Machine Vision Conference*, volume 22, pages 761–767, Sept. 2002.

F. Matia, A. Jimenez, D. Rodriguez-Losada, and B. M. Al-Hadithi. A novel fuzzy Kalman filter for mobile robots localization. In *Proc. of IPMU*, 2004.

F. Matía, A. Jiménez, B. M. Al-Hadithi, D. Rodríguez-Losada, and R. Galán. The fuzzy Kalman filter: State estimation using possibilistic techniques. *Fuzzy Sets and Systems*, 157(16):2145–2170, 2006.

E. Mazor, A. Averbuch, Y. Bar-Shalom, and J. Dayan. Interacting multiple model methods in target tracking: a survey. *IEEE Transactions on aerospace and electronic systems*, 34(1):103–123, 1998.

J. M. Mendel. *Uncertain rule-based fuzzy logic systems: introduction and new directions*. Prentice Hall PTR Upper Saddle River, 2001.

J. M. Mendel. General type-2 fuzzy logic systems made simple: a tutorial. *IEEE Transactions on Fuzzy Systems*, 22(5):1162–1182, 2014.

J. M. Mendel and R. B. John. Type-2 fuzzy sets made simple. *IEEE Transactions on fuzzy systems*, 10(2):117–127, 2002.

J. M. Mendel, R. I. John, and F. Liu. Interval type-2 fuzzy logic systems made simple. *IEEE transactions on fuzzy systems*, 14(6):808–821, 2006.

W. Menke. *Geophysical Data Analysis: Discrete Inverse Theory*. Academic Press, San Diego, California, 92 edition, 1984.

A. Milan, K. Schindler, and S. Roth. Detection-and trajectory-level exclusion in multiple object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3682–3689, 2013.

A. Milan, S. H. Rezatofighi, A. R. Dick, I. D. Reid, and K. Schindler. Online multi-target tracking using recurrent neural networks. In *AAAI*, pages 4225–4232, 2017.

J. Milsom and A. Eriksen. *Field geophysics*. John Wiley & Sons, fourth edition, 2011.

P. K. Molnár, A. E. Derocher, G. W. Thiemann, and M. A. Lewis. Predicting survival, reproduction and abundance of polar bears under climate change. *Biological Conservation*, 143(7):1612–1622, 2010.

T. K. Moon. The expectation-maximization algorithm. *IEEE Signal processing magazine*, 13(6):47–60, 1996.

K. Müller, J. Vanderborght, A. Englert, A. Kemna, J. A. Huisman, J. Rings, and H. Vereecken. Imaging and characterization of solute transport during two tracer tests in a shallow aquifer using electrical resistivity tomography and multilevel groundwater samplers. *Water Resources Research*, 46(3), 2010.

J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society of Industrial and Applied Mathematics*, 5(1):32–38, March 1957.

S. Nasser, R. Alkhaldi, and G. Vert. A modified fuzzy k-means clustering using expectation maximization. In *Fuzzy Systems, 2006 IEEE International Conference on*, pages 231–235. IEEE, 2006.

V. Nenna, A. Pidlisecky, and R. Knight. Application of an extended Kalman filter approach to inversion of time-lapse electrical resistivity imaging data for monitoring recharge. *Water Resources Research*, 47(10), 2011.

R. Ogilvy, P. Meldrum, O. Kuras, P. Wilkinson, J. Chambers, M. Sen, A. Pulido-Bosch, J. Gisbert, S. Jorreto, I. Frances, et al. Automated monitoring of coastal aquifers with electrical resistivity tomography. *Near Surface Geophysics*, 7(5-6): 367–375, 2009.

N. R. Pal and J. C. Bezdek. On cluster validity for the fuzzy c-means model. *IEEE Transactions on Fuzzy systems*, 3(3):370–379, 1995.

T. A. Patterson, L. Thomas, C. Wilcox, O. Ovaskainen, and J. Matthiopoulos. State–space models of individual animal movement. *Trends in ecology & evolution*, 23(2):87–94, 2008.

R. M. Prakash and R. S. S. Kumari. Spatial fuzzy c means and expectation maximization algorithms with bias correction for segmentation of mr brain images. *Journal of medical systems*, 41(1):15, 2017.

D. T. Purvance and R. Andricevic. On the electrical-hydraulic conductivity correlation in aquifers. *Water Resources Research*, 36(10):2905–2913, 2000.

Y. Rathi, N. Vaswani, A. Tannenbaum, and A. Yezzi. Tracking deforming objects using particle filtering for geometric active contours. *IEEE transactions on pattern analysis and machine intelligence*, 29(8):1470–5, Aug. 2007.

H. E. Rauch, F. Tung, C. T. Striebel, et al. Maximum likelihood estimates of linear dynamic systems. *AIAA journal*, 3(8):1445–1450, 1965.

V. Reilly, H. Idrees, and M. Shah. Detection and tracking of large number of targets in wide area surveillance. *Computer Vision–ECCV 2010*, pages 186–199, 2010.

M. Roth and F. Gustafsson. An efficient implementation of the second order extended Kalman filter. In *Information Fusion (FUSION), 2011 Proceedings of the 14th International Conference on*, pages 1–6. IEEE, 2011.

S. I. Roumeliotis and G. A. Bekey. Bayesian estimation and Kalman filtering: A unified framework for mobile robot localization. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, volume 3, pages 2985–2992. IEEE, 2000.

D. Rucker. Investigating motion blur and temporal aliasing from time-lapse electrical resistivity. *Journal of Applied Geophysics*, 111:1–13, Dec. 2014.

S. Särkkä. On unscented Kalman filtering for state estimation of continuous-time nonlinear systems. *Automatic Control, IEEE Transactions on*, 52(9):1631–1641, 2007.

S. Särkkä. Unscented Rauch–Tung–Striebel smoother. *IEEE Transactions on Automatic Control*, 53(3):845–849, 2008.

S. Särkkä. *Bayesian filtering and smoothing*, volume 3. Cambridge University Press, 2013.

K. Shafique and M. Shah. A noniterative greedy algorithm for multiframe point correspondence. *IEEE transactions on pattern analysis and machine intelligence*, 27(1):51–65, Jan. 2005.

S. J. Sheather and M. C. Jones. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 683–690, 1991.

K. Shi, T. Chan, Y. Wong, and S. Ho. Speed estimation of an induction motor drive using an optimized extended Kalman filter. *IEEE Transactions on Industrial Electronics*, 49(1):124–133, 2002.

Y. Shi, K. Sun, L. Huang, and Y. Li. Online identification of permanent magnet flux based on extended Kalman filter for IPMSM drive with position sensorless control. *IEEE Transactions on Industrial Electronics*, 59(11):4169–4178, 2012.

Y. S. Shmaliy. An iterative Kalman-like algorithm ignoring noise and initial conditions. *IEEE Transactions on Signal Processing*, 59(6):2465–2473, 2011.

G. Shu, A. Dehghan, O. Oreifej, E. Hand, and M. Shah. Part-based multiple-person tracking with partial occlusion handling. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1815–1821. IEEE, 2012.

K. Singha and S. M. Gorelick. Saline tracer visualized with three-dimensional electrical resistivity tomography: Field-scale spatial moment analysis. *Water Resources Research*, 41(5), 2005.

K. Singha, F. D. Day-Lewis, T. Johnson, and L. D. Slater. Advances in interpretation of subsurface processes with time-lapse electrical imaging. *Hydrological Processes*, 29(6):1549–1576, 2015.

L. Song, J. Huang, A. Smola, and K. Fukumizu. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 961–968. ACM, 2009.

R. Supper, D. Ottowitz, B. Jochum, J.-H. Kim, A. Römer, I. Baron, S. Pfeiler, M. Lovisolo, S. Gruber, and F. Vecchiotti. Geoelectrical monitoring: an innovative method to supplement landslide surveillance and early warning. *Near Surface Geophysics*, 12(1):133–150, 2014.

J. Tao and R. Klette. Tracking of 2-D or 3-D irregular movement by a family of unscented Kalman filters. *Journal of information and communication convergence engineering*, 10(3):307–314, 2012.

M. Tedre and E. Sutinen. Three traditions of computing: What educators should know. *Computer Science Education*, 18(3):153–170, 2008.

L. T. Tran. Kernel density estimation under dependence. *Statistics & probability letters*, 10(3):193–201, 1990.

P. Tsourlos. *Modelling, interpretation and inversion of multielectrode resistivity survey data.* PhD thesis, University of York, 1995.

T. Tuytelaars and K. Mikolajczyk. Local invariant feature detectors: A survey. *Foundations and Trends in Computer Graphics and Vision*, 3(3):177–280, 2007.

S. Uhlemann, J. Chambers, P. Wilkinson, M. Kirkham, and M. Beamish. Archaeological and resources assessment of sand and gravel deposits using 3-D ERT and automated interface detection. In *Near Surface Geoscience 2016-22nd European Meeting of Environmental and Engineering Geophysics*, 2016.

S. Uhlemann, J. Chambers, P. Wilkinson, H. Maurer, A. Merritt, P. Meldrum, O. Kuras, D. Gunn, A. Smith, and T. Dijkstra. Four-dimensional imaging of moisture dynamics during landslide reactivation. *Journal of Geophysical Research: Earth Surface*, 122(1):398–418, 2017.

J. Vanderborght, A. Kemna, H. Hardelauf, and H. Vereecken. Potential of electrical resistivity tomography to infer aquifer transport characteristics from tracer studies: A synthetic case study. *Water Resources Research*, 41(6), 2005.

C. Veenman, M. Reinders, and E. Backer. Resolving motion correspondence for densely moving points. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(1):54–72, 2001.

N. K. Verma and M. Hanmandlu. From a Gaussian mixture model to nonadditive fuzzy systems. *IEEE Transactions on Fuzzy Systems*, 15(5):809–827, 2007.

R. Viertl. Fuzzy bayesian inference. In *SMPS*, pages 10–15. Springer, 2008.

E. L. Wallin, T. C. Johnson, W. J. Greenwood, and J. M. Zachara. Imaging high stage river-water intrusion into a contaminated aquifer along a major river corridor using 2-d time-lapse surface electrical resistivity tomography. *Water Resources Research*, 49(3):1693–1708, 2013.

E. Wan, R. Van Der Merwe, et al. The unscented Kalman filter for nonlinear estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pages 153–158. IEEE, 2000.

W. O. C. Ward, P. B. Wilkinson, J. E. Chambers, and L. Bai. Comparative investigation of guided fuzzy clustering and mean shift clustering for edge detection in electrical resistivity tomography images of mineral deposits. In *EGU AGM Geophysical Research Abstracts*, volume 16, 2014a.

W. O. C. Ward, P. B. Wilkinson, J. E. Chambers, L. S. Oxby, and L. Bai. Distribution-based fuzzy clustering of electrical resistivity tomography images for interface detection. *Geophysical Journal International*, 197(1):310–321, Feb. 2014b.

W. O. C. Ward, P. B. Wilkinson, J. Chambers, and L. Bai. Predicting and tracking spatiotemporal moments in electrical resistivity tomography. In *AGU Fall Meeting Abstracts*, 2015a.

W. O. C. Ward, P. B. Wilkinson, J. Chambers, and L. Bai. Tracer moment tracking and forecasting in time-lapse electrical resistivity tomography. In *3rd International Workshop on Geoelectrical Monitoring (GELMON)*. 3rd International Workshop on Geoelectrical Monitoring (GELMON), 2015b.

W. O. C. Ward, P. B. Wilkinson, J. E. Chambers, H. Nilsson, O. Kuras, and L. Bai. Tracking tracer motion in a 4-D electrical resistivity tomography experiment. *Water Resources Research*, 52(5):4078–4094, 2016.

W. O. C. Ward, H. Nilsson, and L. Bai. A type-1 fuzzy inference system for the Kalman filter in switching models. *IEEE Transactions on Signal Processing*, submitted to.

P. B. Wilkinson, P. I. Meldrum, O. Kuras, J. E. Chambers, S. Holyoake, and R. D. Ogilvy. High-resolution electrical resistivity tomography monitoring of a tracer test in a confined aquifer. *Journal of Applied Geophysics*, 70(4):268–276, Apr. 2010.

Z. Wu, A. Thangali, S. Sclaroff, and M. Betke. Coupling detection and data association for multiple object tracking. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1948–1955. IEEE, 2012.

D. Xiaolong and Z. Pingfang. A new interacting multiple model algorithm based on the unscented particle filter. In *Information Assurance and Security, 2009. IAS'09. Fifth International Conference on*, volume 1, pages 419–422. IEEE, 2009.

J. Xing, H. Ai, and S. Lao. Multi-object tracking through occlusions by local tracklets filtering and global tracklets association with detection responses. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1200–1207. IEEE, 2009.

M. Xu and Y. Eckstein. Use of weighted least-squares method in evaluation of the relationship between dispersivity and field scale. *Ground Water*, 33(6):905–908, 1995.

T. Xu and J. J. Gómez-Hernández. Probability fields revisited in the context of ensemble Kalman filtering. *Journal of Hydrology*, 531:40–52, 2015a.

T. Xu and J. J. Gómez-Hernández. Inverse sequential simulation: Performance and implementation details. *Advances in Water Resources*, 86:311–326, 2015b.

T. Xu and J. J. Gómez-Hernández. Joint identification of contaminant source location, initial release time, and initial solute concentration in an aquifer via ensemble Kalman filtering. *Water Resources Research*, 52(8):6587–6595, 2016.

B. Yang and R. Nevatia. An online learned crf model for multi-target tracking. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2034–2041. IEEE, 2012.

A. Yilmaz, O. Javed, and M. Shah. Object tracking: A survey. *ACM Computing Surveys*, 38(4):13–es, Dec. 2006.

J. Yu, Q. Cheng, and H. Huang. Analysis of the weighting exponent in the fcm. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 34(1):634–639, 2004.

T.-H. Yu, O. J. Woodford, and R. Cipolla. A performance evaluation of volumetric 3-D interest point detectors. *International Journal of Computer Vision*, 102 (1-3):180–197, Sept. 2013.

L. A. Zadeh. Discussion: Probability theory and fuzzy logic are complementary rather than competitive. *Technometrics*, 37(3):271–276, 1995.

L. A. Zadeh. Fuzzy sets. In *Fuzzy Sets, Fuzzy Logic, And Fuzzy Systems: Selected Papers by Lotfi A Zadeh*, pages 394–432. World Scientific, 1996.

H. Zhou, J. J. Gomez-Hernandez, H.-J. H. Franssen, and L. Li. An approach to handling non-Gaussianity of parameters and state variables in ensemble Kalman filtering. *Advances in Water Resources*, 34(7):844–864, 2011.

J. Zhou, A. Revil, M. Karaoulis, D. Hale, J. Doetsch, and S. Cuttler. Image-guided inversion of electrical resistivity data. *Geophysical Journal International*, 197(1): 292–309, 2014.