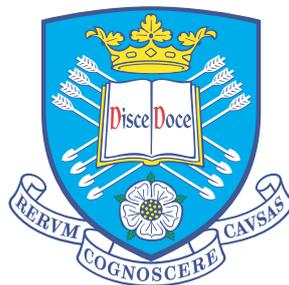


Modelling and Identification of Immune Cell Migration during the Inflammatory Response

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of
Philosophy

by Anastasia Kadochnikova
Department of Automatic Control and Systems Engineering



The University of Sheffield
2019

Abstract

Neutrophils are the white blood cells that play a crucial role in the response of the innate immune system to tissue injuries or infectious threats. Their rapid arrival to the damaged area and timely removal from it define the success of the inflammatory process. Therefore, understanding neutrophil migratory behaviour is essential for the therapeutic regulation of multiple inflammation-mediated diseases. Recent years saw rapid development of various *in vivo* models of inflammation that provide a remarkable insight into the neutrophil function. The main drawback of the *in vivo* microscopy is that it usually focuses on the moving cells and obscures the external environment that drives their migration. To evaluate the effect of a particular treatment strategy on neutrophil behaviour, it is necessary to recover the information about the cell responsiveness and the complex extracellular environment from the limited experimental data. This thesis addresses the presented inference problem by developing a dynamical modelling and estimation framework that quantifies the relationship between an individual migrating cell and the global environment.

The first part of the thesis is concerned with the estimation of the hidden chemical environment that modulates the observed cell migration during the inflammatory response in the injured tail fin of zebrafish larvae. First, a dynamical model of the neutrophil responding to the chemoattractant concentration is developed based on the potential field paradigm of object-environment interaction. This representation serves as a foundation for a hybrid model that is proposed to account for heterogeneous behaviour of an individual cell throughout the migration process. An approximate maximum likelihood estimation framework is derived to estimate the hidden environment and the states of multiple hybrid systems simultaneously. The developed framework is then used to analyse the neutrophil tracking data observed *in vivo* under the assumption that each neutrophil at each time can be in one of three migratory modes: responding to the environment, randomly moving, and stationary. The second part of the thesis examines the process of neutrophil migration at the subcellular scale, focusing on the subcellular mechanism that translates the local environment sensing into the cell shape change. A state space model is formulated based on the hypothesis that links the local protrusions of the cell membrane and the concentration of the intracellular pro-inflammatory signalling protein. The developed model is tested against the local concentration data extracted from the *in vivo* time-lapse images via the classical expectation-maximisation algorithm.

Acknowledgements

I would first and foremost like to thank my supervisor, Prof. Visakan Kadirkamanathan, for providing me with an opportunity to work on an exciting project and for broadening my academic interests. I am most grateful for his expert guidance, his helpful comments, and his patience that I am sure has been put to test many times.

I wish to thank our collaborator, Prof. Stephen A. Renshaw, for making me feel welcome in his lab and providing insightful feedback to even the most incremental results I had to show. I appreciate the support of my second supervisor, Prof. Lyudmila Mihaylova, and thank her for the enthusiasm throughout my application process and the keen interest in my progress during the initial stages of my research. I gratefully acknowledge the ACSE department for funding my PhD project.

I am indebted to Dr Hannah Isles for chaperoning me in the Renshaw lab, repeatedly explaining the basics of cell biology, always finding time to do “just a few more” experiments, and helping me to make sense of my results. I am grateful to Dr Yang Zhang for sharing the tool for cell boundary tracking with me and helping with its application to neutrophils. Her assistance enabled the progress of the final chapter which seemed hopeless for a long time. I appreciate the collaboration of Dr Andrew Hills who introduced me to new programming languages, version control and file management tools that I used to organise my workload. I would also like to thank Dr Olga Isupova for long conversations about the inference methods over numerous tea breaks.

I would not have gotten to this point without constant motivation from Prof. Vladimir T. Bobronnikov – my former supervisor, teacher, and friend. His relentless energy, boundless research interests, and systematic approach to both academic and everyday problems have been a source of my inspiration for the past nine years.

I wish to thank my friends in Belfast for convincing me to pursue a PhD in the UK; my friends in Moscow for their participation and continued support despite the distance; my friends in Sheffield for always finding a way to distract me.

A very special word of thanks goes to my partner, Pablo, whose calm and measured ways gave me comfort at difficult times. Without his patient affection, warm-heartedness, and encouragement this would have been a much more gruelling journey.

This thesis is dedicated to my mother, Elena, and to her husband, Eduard. My heart is filled with gratitude to my mother for her love, kindness, and wisdom. Her unwavering belief in my success gave me confidence to pursue the most interesting path in life that has led me to where I am now. I wish I had the words to express how much Ed’s support meant to me. I will always admire the fortitude and optimism with which he encouraged me to keep fighting my battles even when he was losing his.

Contents

List of Figures	xi
List of Tables	xvi
List of Algorithms	xvii
Nomenclature	xix
1. INTRODUCTION	1
1.1 Background	1
1.2 Systems perspective on neutrophil migration	3
1.3 Aims and objectives	4
1.4 Thesis overview	5
1.5 Disseminated results	7
2. MATHEMATICAL MODELLING OF CELL MIGRATION	9
2.1 Background on inflammation	9
2.1.1 Neutrophils in the inflammatory response	10
2.1.2 The mechanism of chemotaxis	12
2.1.3 Role of in vivo experiments in chemotaxis research	13
2.2 Mathematical models of chemotaxis	16
2.2.1 Population level	17
2.2.2 Individual cell level	18
2.2.3 Subcellular level	21
2.2.4 Cell morphodynamics	24
2.3 Systems approach to the modelling of biological processes	26
2.3.1 State space representation	27

2.3.2	Preliminaries	28
2.4	Parameter estimation	30
2.4.1	Bayesian estimation	32
2.4.2	Maximum Likelihood estimation	33
2.4.3	Maximum a posteriori estimation	35
2.4.4	Least Squares estimation	36
2.4.5	Parameter inference from the incomplete data	37
2.4.6	The expected log likelihood for a state space model	40
2.5	State estimation	41
2.5.1	Optimal state estimation	41
2.5.2	Filtering in linear Gaussian case	43
2.5.3	Smoothing in linear Gaussian case	45
2.5.4	Suboptimal state estimation	48
2.5.5	State estimation of jump Markov systems	50
3.	INFERENCE OF THE HIDDEN CHEMOATTRACTANT FIELD FROM OBSERVED CELL MIGRATION PATTERNS	55
3.1	Background	55
3.2	Methods	57
3.2.1	Data acquisition	57
3.2.2	Model of neutrophil dynamics	60
3.2.3	Chemoattractant field decomposition	61
3.3	Chemoattractant concentration inference	64
3.3.1	The log likelihood function	64
3.3.2	The forecasting step	65
3.3.3	The maximisation step	67
3.3.4	The expectation step	67
3.3.5	The estimation algorithm	70
3.4	Results	71
3.4.1	Simulation examples	72
3.4.2	Estimation results: cell velocities	77
3.4.3	Estimation results: normal injury	80

3.4.4	Estimation results: mild and severe injury	81
3.5	Discussion	83
3.6	Summary	84
4.	AN APPROXIMATE MAXIMUM LIKELIHOOD FRAMEWORK FOR ESTIMATING THE ENVIRONMENT DRIVING OBJECTS WITH HYBRID DYNAMICS	87
4.1	Background	88
4.2	The problem statement	90
4.3	The estimation framework	93
4.3.1	The likelihood function	94
4.3.2	The maximisation step	96
4.3.3	The expectation step	96
4.3.4	The approximate algorithm	98
4.4	Simulation examples	99
4.4.1	Experiment I: uniformly distributed starting points	100
4.4.2	Experiment II: single starting point	105
4.4.3	Experiments III and IV: starting points on the line	107
4.4.4	Experiment V: dynamical model mismatch	108
4.4.5	Experiment VI: field model mismatch	110
4.5	Discussion	112
4.6	Summary	117
5.	CHARACTERISATION OF THE ENVIRONMENT INFLUENCE ON NEUTROPHILS DURING STAGES OF INFLAMMATION	119
5.1	Background	119
5.2	Methods	122
5.2.1	Data acquisition	122
5.2.2	Describing heterogeneous cell behaviour	122
5.2.3	Environment inference with a hybrid model of neutrophil dy- namics	126
5.3	Results	128
5.3.1	Chemoattractant concentration inference during neutrophil re- cruitment stage	129

5.3.2	Chemorepellent concentration inference during inflammation resolution stage	139
5.4	Discussion	141
5.4.1	Estimation framework	143
5.4.2	Recruitment stage	145
5.4.3	Resolution stage	146
5.5	Summary	147
6.	STATISTICAL ANALYSIS OF NEUTROPHIL MORPHODYNAMICS REGULATED BY THE SUBCELLULAR SIGNALLING	149
6.1	Background	149
6.2	Methods	152
6.2.1	Data acquisition	152
6.2.2	Model of cell boundary evolution	152
6.2.3	Cell boundary tracking	154
6.2.4	Extracting PIP ₃ concentration data from the time-lapse images	155
6.2.5	Estimating parameters of the boundary evolution model . . .	156
6.3	Results	158
6.3.1	Image processing results	158
6.3.2	Correlation analysis results	160
6.3.3	Example of model identification from <i>in vivo</i> data	163
6.3.4	Example of model identification from <i>in vitro</i> data	166
6.4	Discussion	169
6.5	Summary	171
7.	CONCLUSION	175
7.1	Summary of contributions	175
7.2	Future Work	178
	Appendices	183
A.	DERIVATION OF THE MAXIMISATION STEP FOR THE EM ALGORITHM	185
B.	PROPERTIES OF THE EXPECTATION	189
C.	GEOMETRY OF THE CELL BOUNDARY	191
	Bibliography	195

List of Figures

2.1	A polarised neutrophil sensing the environment.	13
2.2	Time-lapse <i>in vivo</i> images of the PH-Akt-GFP neutrophil: frame-by-frame morphological changes of the cell sensing the environment. . .	14
2.3	Time-lapse <i>in vivo</i> images for zebrafish larva.	15
2.4	Models of subcellular signalling in environment sensing and polarisation.	23
2.5	A directed graph of a hidden Markov model.	31
2.6	Observation sets utilised in filtering and smoothing algorithms. . . .	48
2.7	A directed acyclic graph of a generic left-continuous hybrid model. .	50
2.8	Computational complexity of the exact JMS state estimation.	52
2.9	Multiple model algorithms for a jumping system with two modes. . .	53
3.1	An example of tracking data obtained from <i>in vivo</i> observation of the injured zebrafish larva.	57
3.2	Tracking results for various wound types.	59
3.3	Decomposition of the environment with a grid of basis functions. . .	62
3.4	The potential field modelled with a 4×4 grid of tensor product cubic B-splines.	72
3.5	Field estimation results obtained from 100 MC simulations.	73
3.6	State estimation results obtained from the ERTS smoother for a simulated track.	76
3.7	State estimation results obtained from the URTS smoother for a simulated track.	76
3.8	RMSEs of the state estimates obtained from EKF and ERTS with varying noise covariances.	78

3.9	RMSEs of the state estimates obtained from UKF and URTS with varying noise covariances.	78
3.10	Comparison of state estimation results obtained from two non-linear algorithms	79
3.11	Histograms of cell velocity estimates obtained via UKF-URTS.	80
3.12	Inferred chemoattractant fields for the normal injury data. All colour bars are normalised to scale 0-600 a.u..	81
3.13	Inferred chemoattractant fields for various wound types.	82
4.1	Directed acyclic graph of the considered jump Markov system.	92
4.2	Examples of migration patterns generated by the interaction of hybrid systems with the modelled potential field.	101
4.3	Estimation results for Experiment I.	102
4.4	Example confidence regions for the parameter estimates obtained from Monte-Carlo simulations.	104
4.5	Parameter estimate convergence for 100 simulations in log scale.	105
4.6	Estimation results for Experiment II.	106
4.7	Estimation results for Experiments III and IV.	108
4.8	Estimation results w for Experiment V.	110
4.9	Estimation results for Experiments VI.	112
5.1	Tracking results from the inflammation resolution stage.	123
5.2	Examples of neutrophils switching between three proposed behavioural modes.	125
5.3	A schematic overview of the estimation framework utilised for the joint estimation of the global chemoattractant concentration field and the individual cell behaviour.	127
5.4	Examples of migratory mode probabilities estimated via IMM-URTS algorithm.	131
5.5	Comparison of cell velocity histograms for actively migrating and randomly moving neutrophils during the recruitment stage of the inflammatory response.	132
5.6	The chemoattractant concentration field estimated from the normal injury data via the framework with IMM algorithm. The colour bars are normalised to scale 0-500 a.u..	133

5.7	Estimation results obtained from the EM algorithm with mode history pruning.	134
5.8	Computed mode transition frequencies for the normal injury datasets.	134
5.9	Migratory mode probabilities estimated via IMM-URTS algorithm from the tail mild injury data.	135
5.10	The chemoattractant concentration field estimated from the mild injury data. The colour bars are normalised to scale 0-200 a.u..	136
5.11	Computed mode transition frequencies for the mild injury datasets.	136
5.12	Migratory mode probabilities estimated via IMM-URTS algorithm from the tail fin nick injury data.	137
5.13	The chemoattractant concentration field estimated from the normal injury data. The colour bars are normalised to scale 0-250 a.u..	137
5.14	Computed mode transition frequencies for the tail fin nick injury datasets.	138
5.15	Migratory mode probabilities estimated via IMM-URTS algorithm from the severe injury data.	139
5.16	The chemoattractant concentration field estimated from the severe injury data. The colour bars are normalised to scale 0-80 a.u..	140
5.17	Computed mode transition frequencies for the severe injury datasets.	140
5.18	Migratory mode probabilities of neutrophils performing reverse migration	142
5.19	Comparison of cell velocity histograms for actively migrating and randomly moving neutrophils during the resolution stage of the inflammatory response	143
5.20	The chemoattractant concentration field estimated from the reverse migration data.	144
5.21	Computed mode transition frequencies for the reverse migration datasets.	144
6.1	Example time-lapse images of the PH-Akt-GFP neutrophils in the zebrafish body.	153
6.2	Example of image processing results for a migrating neutrophil.	159
6.3	Vertex bases and projections of local velocities	160
6.4	Correlation plots of the local normal velocity and the local concentration of PIP ₃ over the entire cell boundary over all tracking period.	161

6.5	Normal velocity histograms for Cell 6.	163
6.6	Smoothing of the extracted fluorescence intensity over the cell boundary at time t	164
6.7	Cell properties extracted from the imaging data.	165
6.8	Estimation results for a motile neutrophil.	165
6.9	Image processing results for a polarised HL-60 cell.	167
6.10	Quantitative data extracted from the images of a polarised HL-60 cell.	168
6.11	Estimation results for a polarised HL-60 cell.	168
6.12	Example of the cell membrane classification based on the local curvature and local velocity components.	170
6.13	The effect of 2-D microscopy on the single-cell scale.	172
C.1	An example of the basis located at a vertex of the polygon.	193
C.2	Approaches to calculating the area of non-convex polygons.	193

List of Tables

3.1	Tuning parameters of the dynamical model.	72
3.2	Statistics of the parameter MLEs obtained via the EM algorithm employing URTS state estimation over 100 simulations.	74
3.3	Normalised total variance of the estimation results obtained using the model with varying process noise parameters.	75
3.4	Normalised total variance of the estimation results obtained using the model with varying measurement noise parameters.	75
4.1	Estimation results for Experiment I over 100 simulations.	103
4.2	Selected estimation results for Experiment II over 100 simulations. . .	106
4.3	Selected estimation results for Experiment III over 100 simulations. . .	109
4.4	Selected estimation results for Experiment IV over 100 simulations. . .	109
4.5	Selected estimation results for Experiment V over 100 simulations. . .	111
4.6	Selected estimation results for Experiment VI over 100 simulations. . .	111
4.7	Volumes of confidence regions for N MC simulations.	113
4.8	Normalised total variance for N MC simulations.	115
4.9	Mode transition frequencies computed on the final iteration of the EM algorithm.	116
5.1	Alternative methods for the environment inference.	128
5.2	Tuning parameters of the hybrid model.	129
5.3	Volumes of confidence regions and normalised total MLE variance for the normal injury datasets.	134
5.4	Volumes of confidence regions and normalised total MLE variance for the mild injury datasets.	136

5.5	Volumes of confidence regions and normalised total MLE variance for the nick injury datasets.	138
5.6	Volumes of confidence regions and normalised total MLE variance for the severe injury datasets.	140
5.7	Volumes of confidence regions and normalised total MLE variance for the reverse migration datasets.	144
6.1	Mann-Whitney U-test results.	163
6.2	Estimation results for PH-Akt-GFP neutrophils.	164
6.3	Estimation results for a polarised cell.	166

List of Algorithms

2.1	Kalman filter	45
2.2	Two-filter smoother	46
2.3	Backward recursion of the RTS smoother	47
2.4	Cross-covariance computation	48
2.5	Forward time recursion of the IMM filter	54
3.1	Expectation-Maximisation algorithm	64
3.2	Extended forward-backward algorithm	68
3.3	Unscented forward-backward algorithm	69
4.1	Backward recursion of the IMM RTS smoother	97
6.1	Extraction of local fluorescence intensity along the cell boundary . .	156
C.1	Computation of polygon vertex bases	192
C.2	Computation of the polygon area with standard triangulation	193

Nomenclature

List of Symbols

t	continuous time.
T	time step.
$k = 1, \dots, K$	object index.
$t = 1, \dots, T$	discrete time index.
$j = 1, \dots, N_m$	mode index.
$h = 1, \dots, N_b$	basis function index.
i	iteration number of the EM algorithm.
$A^j \triangleq A(M^j)$	arbitrary matrix corresponding to the model M^j .
A^\top	matrix transpose.
$A^\dagger \triangleq (A^\top A)^{-1} A^\top$	Moore-Penrose pseudoinverse.
\mathbb{I}_n	identity matrix of size $n \times n$.
$\mathbb{O}_{n \times m}$	zero matrix of size $n \times m$.
$\delta(\cdot)$	Kronecker delta function.
$\nabla(\cdot)$	gradient operator.
$p(\cdot)$	probability density function.
$p(\cdot \cdot)$	conditional probability density function.
$P(\cdot)$	probability function.
$P(\cdot \cdot)$	conditional probability function.
$\mathbb{E}[\cdot]$	expectation function.
$\mathbb{E}_{\{\cdot\}}[\cdot] \triangleq \mathbb{E}[\cdot \cdot]$	conditional expectation function.
$z \sim p(z)$	variable is distributed according to $p(z)$.
$\mathcal{N}(\mu, \Sigma)$	Gaussian distribution with mean μ and covariance Σ .
$\beta^k(s)$	B-spline of order k for variable s .
$\beta(s_1, \dots, s_\nu)$	multivariate B-spline defined in ν dimensions.
d_h^ν	support knot of the h -th B-spline in the ν -th dimension.
$D_\nu^k \triangleq \{d_h^\nu, \dots, d_{h+k+1}^\nu\}$	local support of a k -th order B-spline in the ν -th dimension.
\mathbf{z}_t^k	arbitrary vector associated with the k -th object at time t .
$\mathbf{z}^k \triangleq \{\mathbf{z}_t^k\}_{t=1}^{T_k}$	collection over index t , associated with the k -th object.
$\mathbf{z} \triangleq \{\mathbf{z}^k\}_{k=1}^K$	collection over all time indices for all objects.

\mathbf{x}_t^k	continuous-valued state of the k-th object at time t.
\mathbf{m}_t^k	discrete-valued state of the k-th object at time t.
\mathbf{y}_t^k	observation available for the k-th object at time t.
\mathbf{u}_t^k	deterministic input of the k-th object at time t.
$\mathbf{w}_t \sim \mathcal{N}(0, Q_w)$	realisation of the process noise at time t.
$\mathbf{v}_t \sim \mathcal{N}(0, R_v)$	realisation of the measurement noise at time t.
\mathbb{R}^r	a set of real numbers of the dimension r .
$\mathbb{M} = \{M^1, \dots, M^{N_m}\}$	a set of candidate models of size N_m .
A_t	state transition matrix at time t.
B_t	control gain at time t.
G_t	process noise gain at time t.
C_t	observation matrix matrix at time t.
H_t	measurement noise gain at time t.
\mathcal{X}	hidden data set consisting of continuous-valued states.
\mathcal{M}	hidden data set consisting of discrete-valued states.
\mathcal{Y}	incomplete data set.
\mathcal{Z}	complete data set.
θ	unknown parameter.
Θ	vector of unknown parameters.
$\mathcal{L}(\Theta) \triangleq p(\mathcal{Z} \Theta)$	likelihood function of the parameter vector Θ .
$\log \mathcal{L}(\Theta)$	log likelihood function in the EM algorithm.
$\mathcal{Q}(\Theta, \hat{\Theta}^i)$	lower bound the log likelihood function.
$\hat{\Theta}$	parameter vector estimate.
$\hat{\mathbf{x}}_{t t-1}$	predicted state estimate.
$\hat{\mathbf{x}}_{t t}$	filtered state estimate.
$\hat{\mathbf{x}}_{t T^*}, T^* > t$	smoothed state estimate.
$\mathbf{P}_{t }$	covariance of the state estimate.
K_t	Kalman gain.
J_t	RTS smoother gain.
$\pi_j \triangleq p(M^j)$	prior probability of the model M^j .
$\mu_t^j \triangleq p(M^j)$	probability of the model M^j at time t.
$\phi_{lj} \triangleq p(M^j M^l)$	probability of transition from M^l to M^j .
$\Gamma(t)$	closed curve evolving in discrete time.
$ds \in \Gamma(t)$	segment of a continuous curve.
$s^k \in \Gamma(t), k = 1, \dots, K$	node on a discretised curve.
$\bar{\nu}(s^k)$	local outward pointing normal of the discretised curve.
$\bar{\tau}(s^k)$	local tangent of the discretised curve.
$\kappa(s^k)$	local curvature of the discretised curve.
$\mathcal{F}(s^k)$	local force acting normally to the curve.
\mathcal{A}_t	area enclosed by the curve at time t.

List of Acronyms

<i>pdf</i>	Probability density function.
<i>pmf</i>	Probability mass function.
a.u.	Arbitrary unit.
ABC	Approximate Bayesian computation.
Akt	Protein kinase B.
AR(<i>k</i>)	Autoregressive model of order <i>k</i> .
BEN	Biased excitable network.
ChA	Chemoattractant.
CV	Coefficient of variance.
DAG	Directed acyclic graph.
ECM	Expectation - conditional maximisation (algorithm).
EKF	Extended Kalman filter.
EM	Expectation-Maximisation (algorithm).
ERTS	Extended Rauch-Tung-Striebel (smoother).
GFP	Green fluorescent protein.
GPB	Generalised pseudo-Bayesian (state estimator).
GPC	G-protein coupled (receptor).
HMM	Hidden Markov model.
IMM	Interacting multiple-model (state estimator).
JMS	Jump Markov system.
KF	Kalman filter.
KS	Keller-Segel (model).
LEGI	Local excitation global inhibition.
LS	Least squares (estimation).
MA	Moving average.
MAP	Maximum <i>a posteriori</i> (estimation).
MC	Monte Carlo (method).
MCMC	Makrov chain Monte Carlo (method).
ML	Maximum likelihood (estimation).
MLE	Maximum likelihood estimate.
MM	Multiple model (state estimator).
NN	Nearest neighbour (method).
O-U	Ornstein-Uhlenbeck (process).
ODE	Ordinary differential equation.
PDE	Partial differential equation.
PI3k	Phosphatidylinositol-4,5-bisphosphate 3-kinase.
PIP ₃	Phosphatidylinositol (3,4,5)-trisphosphate.
RD	Reaction-diffusion (system).
RGB	Red-green-blue (colourmap).

RMSE	Root Mean Square Error.
RTS	Rauch-Tung-Striebel (smoother).
RW	Random walk.
SLAM	Simultaneous localisation and mapping.
SMC	Sequential Monte Carlo (method).
SSM	State space model.
UKF	Unscented Kalman filter.
URTS	Unscented Rauch-Tung-Striebel (smoother).
VB	Variational Bayes (method).
VBEM	Variational Bayes Expectation-Maximisation (algorithm).

1 | Introduction

1.1 Background

Inflammation is a fundamental defence reaction of the host immune system to infection or injury. It is initiated by leukocyte migration from the bloodstream to the damaged tissue where they kill bacteria and pathogens. Neutrophils are the most abundant type of leukocytes that are essential for combating various infectious threats [1]. Highly sensitive to the tissue “distress signals”, neutrophils are rapidly recruited to the wound area where they exert two critical functions: amplify the inflammatory signal to recruit other types of immune cells and eliminate bacteria via multiple defence mechanisms.

However, a malfunction in either neutrophil deployment or deactivation may have a devastating impact on the host. An overzealous neutrophilic response is accompanied by the release of multiple cytotoxic mediators that can damage healthy tissue. The neutrophil-induced tissue damage is characteristic for rheumatoid arthritis, pulmonary disease, and many autoimmune disorders. On the other hand, an impaired mechanism of neutrophil removal from the damaged area may lead to chronic inflammatory diseases. That is why the moderation of recruitment and timely removal of these cells from the damaged tissue is crucial for treatment of various inflammation-mediated diseases. Given the critical role the neutrophil migration plays in the inflammatory response, many therapeutic strategies target this process. To be able to manipulate this process precisely, one needs not only to understand but also to be able to predict the behaviour of these immune cells in various settings.

Since the early nineteenth century, when cell biology was postulated as a discipline, the cell function has been studied experimentally by means of observation. Therefore, the success of research in this field remains closely connected to the advances in experimental assays. For a long time, the understanding of neutrophil function and role in inflammation has been informed by the experimental studies that were conducted in artificial environments (*in vitro*). Dramatic improvements in genomics and microscopy lead to the development of several animal models of inflammation where the immune response can be observed running its natural course

inside the living organism (*in vivo*). The most promising *in vivo* model is the zebrafish. The combination of novel imaging techniques, transparency of zebrafish larval form, and fluorescent labelling of immune cells in transgenic animals helped to elucidate various mechanisms of neutrophil migration. Whilst *in vivo* experiments undoubtedly have moved the research in this topic forward, the need for quantitative evaluation of developed hypotheses is clear.

It is well known that the activity of neutrophils during the immune response is characterised by continuous interaction with their environment. At initial stages of inflammation, the cells are guided towards the target site by multiple external mediators, called chemoattractants. This directed migration, called chemotaxis, has been widely studied in various cell types. Despite the intensive research in this field, the principal mechanism that connects sensing the environment with cell locomotion remains undefined. In the light of this problem, two aspects require investigation: behavioural modes of cells moving in the chemoattractant environment, and intracellular signalling pathways responsible for cell locomotion in response to changes in the environment.

At the final stages of inflammation, after the infectious threat has been resolved, neutrophils are deactivated and removed from the wound. Recent *in vivo* experiments have revealed that some neutrophils migrate away from the damaged site [2], [3]; this migration may be induced by the interaction with the environment or, on the contrary, by the loss of cell's sensitivity to the environment caused by the receptor internalisation. The inability to assess the chemoattractant (or chemorepellent) concentration in *in vivo* studies makes it impossible to reject either of the hypotheses.

This thesis attempts to analyse different aspects of cell-environment interaction during the inflammatory process through the medium of mathematical modelling and system identification. The role of mathematical modelling in biological studies has become widely appreciated in recent years as it allows to represent a complex interconnected system with simple prototypes that account for the key interactions within the system. The predictive capacity of the developed model depends largely on how well it is linked to the data. Thus, the ambition of data-driven modelling is not to merely simulate cell motion but to identify mechanistic models which could represent its innate processes with satisfying accuracy.

Rapid development of data-driven modelling in biology is directed by three key processes: refinement of measurement technologies, continuous improvement and development of models representing complex biological processes, and fitting of the existing models to available experimental data. Both the task of model fitting and the model development can be viewed as a parameter estimation problem and solved using classical system identification tools. System identification is the discipline that

deals with the process of creating mathematical models of a dynamical system from the available data. It is constituted by the selection of the appropriate model structure and the estimation of unknown parameters in the constructed model. Because both the model structure and the measurements are subject to uncertainty, the inferences about the unknown parameters are of stochastic nature and can be obtained via various statistical methods. This thesis focuses on the dynamical modelling of neutrophil migration and employs a combination of statistical inference methods to link the developed models to the experimental data.

1.2 Systems perspective on neutrophil migration

The inflammatory response that evolves over long periods of time can be viewed as a dynamical process and investigated through the agency of systems theory. The dynamical system considered in this work is the population of activated pro-inflammatory neutrophils migrating in the tissue towards the wound, where each individual neutrophil is an element of the high-level system. The environment in which the system functions consists of the chemokine concentrations that drive the cell population towards and away from the wound area. Based on the experimental evidence that shows that neutrophil migration is actively regulated by these concentrations it is considered here that the evolution of the system is not only affected by the environment but explicitly linked to it.

In order for the described system to achieve the aim of efficiently removing the infectious threat and resolving the inflammation process, all of its elements must function correctly. If one is to consider a therapeutic treatment of inflammatory response as a controller for a dynamical system, the need for the formal description of the interaction between its constituent parts becomes evident. Thus, the development of a comprehensive mathematical model of neutrophil migration driven by various chemokine environments can be viewed as a part of the design process of a controller that would regulate that particular interaction.

A systematic approach to the model development involves a cycle of the three key processes: hypothesis development, experiment design and implementation, and testing the hypothesis against the experimental data. In biological research, the experimental design usually dictates the level of detail of the developed model. This particular study is dealing with a relatively novel type of experiment, where the direct observation of the system's environment is not possible. The model development in this thesis is concerned with the following questions:

- What can be inferred about the underlying environment from the observed migratory patterns of the neutrophil population?
- Do all neutrophils always respond to the external environment during the

recruitment stage of inflammation?

- Do all neutrophils always respond to the external environment during the resolution stage of inflammation?
- How is the influence of the environment translated into the locomotion of an individual neutrophil?

The objectives of this thesis are developed in accordance to the posed questions.

1.3 Aims and objectives

The principal aim of this thesis is to characterise the interaction between the migrating neutrophils and their environment during different stages of inflammation. This research is motivated by new challenges presented by the novel microscopy methods. In particular, it seeks to recover information about the external environment, that is typically not measured *in vivo*, based on the observed cell behaviour. The experimental data is obtained by collaborators from the transgenic zebrafish whose immune responses are similar to human ones.

The influence of the environment on neutrophils is assessed on various scales. The first part of the thesis deals with the macro-scale level of cell population but utilises the models of individual cell dynamics. The final chapter deals with the sub-cellular processes that occur as a result of spacial environment sensing. The analysis is conducted from a system-theoretic point of view and utilises a state-space framework to interpret the temporally resolved experimental data. The developed state-space models are fitted to the experimental data via various identification techniques.

The following objectives are associated with the established aim:

- Develop a phenomenological model of cell dynamics that would incorporate a parametrised model of the global environment.
- Develop a hybrid model of cell dynamics to account for the homogeneous behaviour of the migrating neutrophils.
- Create an estimation framework that is able to estimate simultaneously the hidden environment and cell dynamic properties with minimum prior assumptions made about the environment.
- Analyse the interaction between migrating neutrophils and their environment during different stages of inflammation by applying the developed estimation framework to the population *in vivo* data collected from zebrafish model of inflammation.
- Quantify and analyse the relationship between sub-cellular concentrations and the deformation of the cell boundary of a migrating cell by applying a classical

estimation method to the sub-cellular *in vivo* data collected from zebrafish model of inflammation.

1.4 Thesis overview

Below is the summary of the thesis structure and the key contributions of the chapters.

Chapter 2 consists of two parts. The first part provides a brief overview of neutrophil function during the immune response and discusses state of the art of the mathematical modelling in the study of chemotaxis. Models at different scales are reviewed alongside with recent developments in experimental procedures. The role and the challenges of evidencing the mathematical hypotheses with biological data are discussed. The second part surveys popular methods of statistical inference and provides the systems context for the data-driven analysis of neutrophil chemotaxis.

Chapter 3 proposes a dynamical model of cell migration that incorporates the influence of the unknown chemoattractant concentration for the quantitative analysis of the cell migration data obtained from the *in vivo* model of inflammation. The model builds on the potential field paradigm of object-environment interaction. The parametrised model of the hidden chemoattractant concentration is obtained via basis function decomposition with unknown scaling parameters. The estimation of these parameters is performed simultaneously with the cell state estimation in the form of an approximate Expectation-Maximisation (EM) algorithm. The developed algorithm is applied to several datasets of neutrophils observed in the zebrafish larvae during the recruitment stage of the inflammatory response. The Chapter makes the following contributions:

- A novel state space model of cell dynamics which includes the parametrised model of the chemoattractant environment.
- An inference framework that predicts the chemoattractant concentration based on tracking data. The estimation is performed under the assumptions that all cells at all times interact with the environment.

Chapter 4 proposes an approximate maximum likelihood framework for estimation of the hidden environment that drives the migration of multiple objects with hybrid dynamics. The hybrid model is a direct extension of the model introduced in Chapter 3. Along with explicitly linking the object dynamics with the external environment, the new model accounts for the heterogeneous behaviour of the migrating object which is common in many engineering and life systems applications. The joint state-parameter estimation of the proposed model is performed via the EM algorithm that utilises a multiple model smoother during the expectation step.

The performance of the proposed estimation framework is demonstrated on several simulation examples with varying migration patterns and model mismatches. The main contribution of this Chapter is summarised as:

- An approximate maximum-likelihood method for estimation of the hidden environment driving multiple objects, based on hybrid representation of object dynamics.
- An implementation of the interacting multiple model (IMM) unscented algorithm.

Chapter 5 demonstrates the application of the developed framework to the neutrophil *in vivo* data collected from zebrafish larva at different stages of inflammation. A hybrid system consisting of three state space models is proposed to represent the dynamics of an individual cell in order to relax the assumption that all neutrophils are responding to the environment at all times. The first model accounts for the cell behaviour when it is driven by the external stimuli, the second model represents the randomly moving cell, and the third model describes the behaviour of a stationary cell. Two types of the multiple model state estimation algorithms are used in the estimation procedure. The key contributions are the following:

- A novel hybrid model of cell dynamics that accounts for the heterogeneous interactions with its environment.
- Simultaneous inference of the hidden chemoattractant concentration field and cell behavioural modes from the cell migration patterns observed in zebrafish.
- Quantitative evidence that neutrophils randomly diffuse from the wound during the inflammation resolution.

Chapter 6 analyses the process of neutrophil chemotaxis on a finer scale. It investigates morphodynamics of the migrating neutrophil and proposes a novel framework of quantifying the internal concentrations. A dynamical model used in the chapter builds on the current understanding of the PIP_3 ¹ signalling mediator that governs changes of the cell boundary. A series of image-processing techniques are applied to the time-lapse images of the migrating neutrophils to quantify the key factors described by the selected model. The correlation analysis of the extracted data is performed to assess whether the observed sub-cellular concentration maybe considered as the sole or the most prominent contributor to the cell membrane evolution. The resultant state space model linking cell boundary evolution and concentrations of sub-cellular activators is tested against the *in vitro* imaging data obtained by observing the immortalised mammalian cells with pro-inflammatory proteins tagged with green fluorescent protein (GFP). The contributions of this Chapter are listed below:

¹ Phosphatidylinositol (3,4,5)-trisphosphate

- An image-processing module for quantification of intracellular concentrations and geometric properties of an individual cell.
- Quantitative evidence that the sub-cellular mediator PIP_3 is not the dominant local activator in pseudopod formation process.

Chapter 7 provides a summary of the results achieved in the thesis and proposes several possible directions for future work.

1.5 Disseminated results

The contributions of this thesis have been disseminated in the following works

- A. Kadochnikova, Y. Zhang, V. Kadiramanathan. “A dynamical systems modelling framework for eukaryotic cell migration analysis”. An abstract in proceedings of SEB Symposium “*Bridging scales in models of cell to tissue behaviour*”, Oxford, United Kingdom, 13-15 September 2016.
- A. Kadochnikova, H.M. Isles, S.A. Renshaw, V. Kadiramanathan. “Estimation of Hidden Chemoattractant Field from Observed Cell Migration Patterns”. A peer-reviewed paper in *Proceedings of 18th IFAC Symposium on System Identification SYSID 2018*, Stockholm, Sweden, 9-11 July 2018.
- H.M. Isles, C. Muir, A. Kadochnikova, C.A. Loynes, V. Kadiramanathan, P.M. Elks, S.A. Renshaw. “Non-apoptotic pioneer neutrophils initiate a swarming response in a zebrafish tissue injury model” submitted to EMBO Reports, 2019.

The following journal papers are currently in development:

- A. Kadochnikova, V. Kadiramanathan. “An Approximate Maximum Likelihood Framework for Estimating the Environment Driving multiple objects with Hybrid Dynamics”.
- A. Kadochnikova, H.M. Isles, S.A. Renshaw, V. Kadiramanathan. “Inference of the External Stimuli Environments from Heterogeneous Behaviour of Migrating Neutrophils in Zebrafish Model of Inflammation”.

2 | Mathematical modelling of cell migration

This chapter contains an overview of relevant experimental and theoretical research of directed immune cell migration. The concepts presented in this part of the thesis serve as a foundation for the analysis conducted in further chapters. Section 2.1 outlines an example scenario of the inflammatory response and explains the role of neutrophil chemotaxis in this process. Section 2.2 is a survey of state of the art in mathematical modelling of chemotaxis on different scales. Section 2.3 provides a system-theoretic perspective on time series analysis and defines the model structure that will be utilised throughout the thesis. Sections 2.4 and 2.5 overview commonly used methods of parameter and state estimation, respectively.

2.1 Background on inflammation

Inflammation is the innate immune system response to infectious threats or injuries. It is mediated by the migration of white blood cells, or *leukocytes*, to and from the damaged tissue. From the leukocyte-centred point of view, the inflammatory process can be divided into two phases: the *recruitment* during which these immune cells migrate towards the injury site, and the *resolution* during which leukocytes are removed from the damaged tissue after eradicating alien cells. Both phases are complex multi-stage processes governed by a variety of overlapping signals that incite cell motion. Although the insistent study elucidated the key events that take place during the recruitment, recent improvements in the experimental procedures reveal new aspects of this process that are yet to be understood [1]. The progress in investigating inflammation resolution for a long time was hindered by the lack of technological capacity to reproduce this process in artificial conditions. Whilst novel microscopy available nowadays allows to observe the natural course of this process, the level of involvement and the exact role of leukocytes in the inflammation resolution remain unknown [2].

2.1.1 Neutrophils in the inflammatory response

Neutrophils constitute a major class of mammalian white blood cells and play a vital role in inflammatory responses. They are deservedly called the first responders of the innate immune system as they migrate towards the damaged tissue within minutes post-injury. Early recruited neutrophils promote recruitment of other immune cells such as phagocytes [4], but also possess phagocytic mechanisms for removal of pathogens and microbial debris from the tissue [5]. Although neutrophils function in cooperation with other immune cells, they outnumber any other participants of the process and define the way the inflammation evolves.

In the absence of external threats to the host, the majority of neutrophils are found in the bone marrow and less than 2% in blood circulation in inactive state [6]. Once the immune system detects a microbial threat or tissue damage, neutrophils activate, leave the bone marrow, and enter the circulation that quickly delivers them closer to the damaged site. They then transmigrate through the endothelial layer of blood vessels to the tissue and enter the pro-inflammatory state to start their migration towards the target area [4].

In the tissue, neutrophil migration towards the wound site is tightly regulated by stimulating chemokines deployed by the damaged tissue, commonly called **chemoattractants (ChA)**. Activated neutrophils exert different types of receptors to sense the local ChA concentration. Depending on the perceived chemotactic environment, neutrophils can commit two distinct types of movement. In uniform attractant concentrations, cells undergo chemokinesis - the type of movement characterised by the increase in cell speed while the choice of direction remains random. In spatially changing ChA concentrations one can observe **chemotaxis** - the directed cell migration. Chemotaxing cells alternate between periods of persistent motion, called **chemotactic runs** [7], in a direction along the steepest concentration gradient, and periods of random motility, during which neutrophils are thought to assess the concentration gradient at their location [8]. Being the quickest eukaryotes, during chemotactic runs neutrophils can develop speeds up to 10-20 $\mu\text{m}/\text{min}$.

Having arrived to the target area, neutrophils eliminate the bacterial or infectious threat by means of phagocytosis (chase and engulfment of individual microbes [9]) or degranulation (release of toxic molecules [10]). Degranulated neutrophils also leave extracellular traps that deactivate and kill the pathogens by delivering large numbers of antimicrobial molecules to the target [11]. Neutrophils are known to amplify the pro-inflammatory signal to other leukocytes by releasing into the tissue various mediators that serve as attractants for other cell types [12].

During the inflammation resolution, surviving neutrophils can be removed from the wound site through one of the following processes:

- Apoptosis (cell death);
- Engulfment by macrophages¹;
- Reverse migration within the tissue;
- Reverse transmigration back into the circulation.

The first two processes can be characterised as passive. Neutrophil apoptosis prevails during the resolution phase and plays an important role in the dampening of the inflammatory response. Dying neutrophils release anti-inflammatory chemokines that slow down the recruitment of new leukocytes until it eventually ceases [13]. Interactions between macrophages and neutrophils have been examined in various animal models [14]. Recent advances in cell imaging revealed that neutrophils can actively migrate away from the wound site [3]. Ability to regulate this active process therapeutically could advance the treatment of many inflammation-mediated diseases. The migration back to circulation is normally driven by mechanical cues and has been observed in extraordinary conditions [15], therefore it is not considered to be of primary importance for the inflammation resolution stage.

The exact mechanism of neutrophil reverse migration remains unknown but there exist several hypotheses about the way it progresses. Firstly it was assumed that cells react to directional cues regulating their removal from the injury site via fugetaxis. In this process, opposite to chemotaxis, cells are driven away from high concentrations of repelling agents called *chemorepellents* [16]. Recent analysis of the experimental data suggests that this retrograde movement is rather a stochastic redistribution of cells than a directed migration [17]. There is an increasing trend of papers speculating that neutrophils lose sensitivity to local gradients owing to the receptor internalisation [14], however more evidence is needed to determine which mechanism is more prevalent.

Neutrophil malfunction at any stage of inflammation can lead to fatal damage to the host: if neutrophils fail to react during the recruitment phase, infection will spread further into the host body, while failure in neutrophil deactivation during the resolution phase can cause passer-by damage to healthy tissue cells and lead to chronic inflammation, which in turn contributes to the development of numerous disorders [18]. Therapeutic importance of neutrophil regulation is obvious, but despite decades of extensive study many aspects of their migratory mechanism remain unknown.

¹ Macrophages are a specialised type of white blood cells involved in the detection, phagocytosis and destruction of bacteria and other pathogens.

2.1.2 The mechanism of chemotaxis

In the literature, the chemotaxis is viewed as a result of complex coordination of motility, gradient sensing, and polarisation [19]. Motility is the main driver of the crawling motion of stimulated neutrophils [20]. A motile neutrophil undergoes severe morphological changes by extending and retracting actin-rich protrusions, also called *pseudopods*. It must be noted that the motility occurs independently from gradient sensing, as neutrophils exhibit spontaneous protrusion extension even in uniform concentrations [21].

Neutrophils are extremely receptive to the environment - they are able to detect the spatial change in chemoattractant concentrations as small as 1% across the cell length [22], [23]. The environment sensing is facilitated by the activation of G-protein-coupled (GPC) receptors located on the cell membrane [24]. They “measure” local concentrations of the attractant by capturing nearby ligands. The receptor-ligand binding initiates several signalling cascades that translate the external gradient into the intracellular gradient that mediates cell polarisation [25]. One of the critical cascades for protrusion development starts with activation of PI3-kinase (PI3k)² and its product, PIP₃³, that boosts actin polymerisation in the growing pseudopod [26]. Other cascades lead to the recruitment of GPC receptors from the rear to the front of the cell, and send myosin and other motor proteins involved in membrane contraction to the rear thus inhibiting the retraction of the new pseudopod [27]. As a result, the neutrophil assumes an elongated shape with two well defined domains: the *leading edge* and the *trailing edge* (see Figure 2.1). This process is called polarisation.

During the polarisation process, the dominant pseudopod develops into the leading edge and defines the direction of cell movement. Large number of receptors at the front create the self-amplifying loop of pro-migratory signalling, leading to persistent polarisation that ultimately results into a chemotactic run [12]. This process goes on until either the direction of the chemotactic gradient changes direction or receptor desensitisation occurs. The neutrophil with internalised receptors slows down, but never fully rounds up as it keeps extending pseudopods to find the direction for the next run. These protrusions are regulated by multiple activators and inhibited by receptor occupancy [28]. The pseudopod extension is illustrated in frames 2-7 of Figure 2.2, while the beginning of the chemotactic run is shown in frames 7-12.

² The phosphatidylinositide 3-kinases (PI3k) are a family of enzymes that are involved in cell motility, growth and directional sensing. In presence of the external stimuli, PI3k is activated after a GPC receptors start binding to external ligands.

³ Phosphatidylinositol (3,4,5)-trisphosphate (PIP₃) is the second messenger PI3k product that resides in the membrane and activates multiple downstream components, including the protein kinase (Akt).

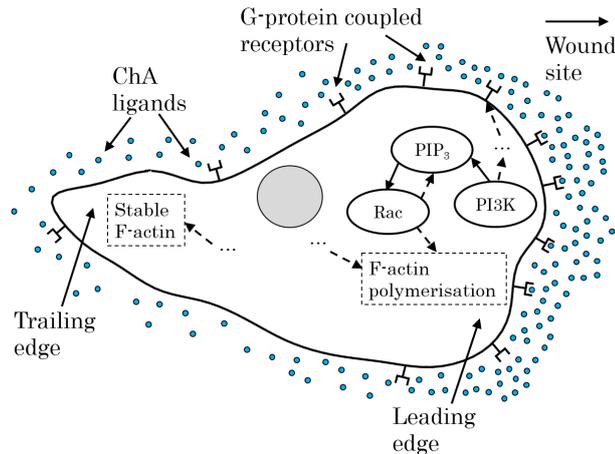


Figure 2.1: A polarised neutrophil senses the environment via GPC receptors clustered at the leading edge. Engagement of these receptors with external stimuli initiates several signalling cascades that lead to f-actin polymerisation and further protrusion of the cell membrane. The PIP_3 activity is inhibited in the trailing edge whose retraction is mediated by a different signalling cascade descending from PI3k.

Although the polarisation process is usually linked to the spatial sensing, it can be viewed as a result of the locomotion itself [21]. Once exposed to a gradient, a cell reorients by turning existing leading edge in line with the gradient, because new pseudopods are usually formed near the existing leading edge. If the gradient is steeper at the opposite side from the leading edge, neutrophils make a U-turn rather than extend new pseudopod at the rear [29].

2.1.3 Role of *in vivo* experiments in chemotaxis research

Our current understanding of the way neutrophils mediate the inflammatory response is evidenced by qualitative experimental research. For a long time the main experimental material for studying neutrophil function has been a cell purified from human blood [30] observed *in vitro*, *i.e.* in an artificial environment. The extraction of neutrophils from blood often leads to spontaneous apoptosis, which is beneficial for investigation of this process but creates extra difficulties in the observation of cell behaviour over long periods of time. The problem of short life span was initially solved by developing immortalized neutrophil-like cell lines, where tendencies observed in genetically modified cultures would later undergo the validation for primary neutrophils. This, however, has proven to be both time consuming and inaccurate. For example, modified cells are defective in apoptosis and inapplicable for testing hypotheses that account for the cell death [31].

Knowledge about the role of neutrophils in the inflammation process has progressed dramatically since the introduction of *in vivo* experiments, where cells are

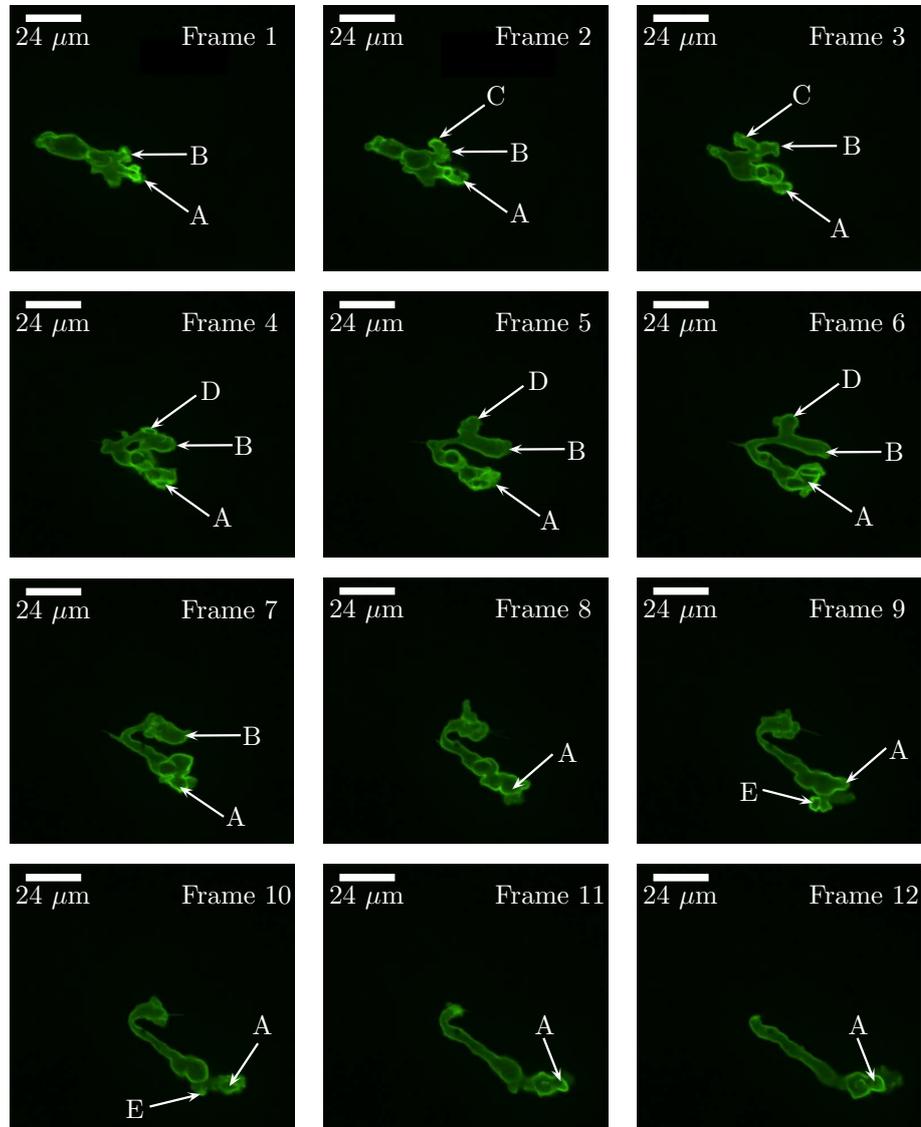


Figure 2.2: Morphological changes of the migrating PH-Akt-GFP neutrophil observed *in vivo*. Visible is the GFP-tagged pleckstrin homology (PH) domain of the Akt protein kinase that binds to PIP_3 and gets translocated to the membrane. The GFP intensity on each image indicates the concentration of PIP_3 and actin polymerisation. The highest level of actin polymerisation is maintained in the leading edge A. Pseudopods B, C, and E form near the leading edge in search for the highest ChA concentration, whilst the secondary pseudopod D branches from B in frame 4, severely stretching the cell material. When the signalling in pseudopods B, C, D is inhibited by the growing number of occupied receptors, the membrane contracts and pseudopods eventually collapse.

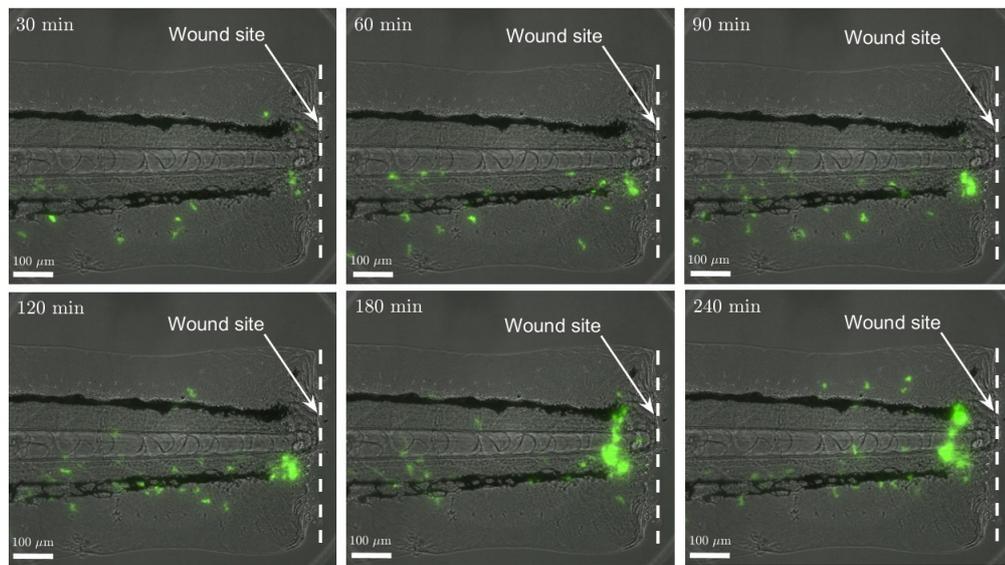


Figure 2.3: Time-lapse *in vivo* images: neutrophils labelled with a green fluorescent protein (GFP) are visible in the transparent zebrafish embryo as they arrive to the wound site during inflammation recruitment stage. The dashed line indicates the location of the tail fin transection. The process is observed starting from 30 minutes to 4 hours post the injury.

observed in their natural environment in the living organisms such as drosophila flies [32], zebrafish [33], and mice [34]. In these animal models innate immune system reaction is triggered by an injury after which recruited leukocytes are observed in real time using novel microscopy techniques. Various limitations of modelling human inflammation have been reported for both mice [35] and drosophila [36]. On the other hand, the zebrafish model has been steadily gaining popularity as it best resembles innate immune response of humans [37]. The biophysical and morphological properties of zebrafish immune cells are similar to those of human leukocytes, and the overall course of the inflammation has similar kinetics with mammalian immune responses [31].

The optical transparency of zebrafish larval form, combined with extensive genomic resources available for genetic manipulation, makes it an outstanding model of inflammation. In inflammation-specific line the neutrophils are labelled with the green fluorescent protein (GFP) or with photoconvertible Kaede protein that changes colour from green to red after exposure to violet light [38]. The inflammatory response is initiated by the tail fin transection of the fish larva. The fluorescent neutrophils are then tracked directly in the fish using time-lapse imaging technology (see Figure 2.3). Development of Kaede zebrafish line has fostered numerous studies of the inflammation resolution stage, revealing longer life span of neutrophils in the tissue and their potentially crucial role in anti-inflammatory processes

[39]. Furthermore, GFP markers can be applied to a certain protein within the cell such as GFP-labelled PH-domain of Akt in Figure 2.2, allowing the investigation of individual signalling pathways regulating neutrophil motility and polarisation [40].

The unique insight provided by the *in vivo* microscopy comes at a price of numerous shortcomings of a passive experiment. Although the inflammation process is actively triggered by the injury, it is impossible to manipulate the course of the inflammation or customise the experiment in any way. Furthermore, only limited measurements can be made along the process, while some important aspects of inflammation - such as the underlying environment driving neutrophil chemotaxis - cannot be measured directly. Thus, there arises a need for a rigorous methodology that can quantify voluminous *in vivo* data in a systematic and principled way; and the most straightforward way to perform quantitative analysis is through the medium of mathematical modelling.

2.2 Mathematical models of chemotaxis

Experimental studies of chemotaxis have been greatly complemented by the mathematical modelling as it provides a myriad of tools for interpreting quantitative data. The benefits of using mathematical modelling in biological applications are threefold [41]:

- Models that are closely linked to the biological data are instrumental in development of new experimental techniques and detection of potential outliers.
- Data-driven models help to rule out biological hypotheses at a fraction of cost and time required to run multiple laboratory experiments.
- Speculative models introduce additional detail to the processes that are currently observed. Although these models are often impossible to validate using currently available biological data, they greatly contribute to mechanistic understanding of complex processes that cannot yet be reached experimentally.

Complexity of a mathematical model depends on the level of detail of the experimental data that is being characterised. Traditionally, the modelling process progresses from high-level phenomenological models with limited number of parameters to more detailed ones. When a model on a certain descriptive level conflicts with the data, it usually means that the potential of modelling on this level has been exhausted. In this case one can either a) test a different hypothesis on the same descriptive level, or b) switch to a more detailed descriptive level by introducing extra components to the model, if sufficient knowledge about these components is available or can be inferred from the data.

Steady progress in microscopy and signal processing methods led to two import-

ant level-down transitions in the modelling of chemotaxis. Initially, the modelling literature was dominated by macro-scale models that describe global trends in the cell population behaviour [42]. The first shift was from the level of cell population to the whole-cell level of modelling, where mechanisms of the single cell locomotion are examined without association to molecular processes inside the cell [43]. In view of recent *in vivo* studies, it has become apparent that the whole-cell models fail to capture the intrinsic complexity of the locomotory mechanism of migrating cells. This led to the development of models describing the interdependent signalling and molecular processes inside the migrating cell, thus constituting the second transition to the subcellular level of modelling. In present days, when computational costs of high-order simulations are not as much of a concern, subcellular models enjoy ever-increasing interest from the cell biology community [41].

This section provides an overview of relevant mathematical models⁴ used to describe chemotaxis of the population level, the single cell level, and the subcellular level. Furthermore, models of cell morphology are discussed separately as they constitute an intermediate descriptive level that provides a connection between the subcellular signalling and the whole cell locomotion.

2.2.1 Population level

Stochastic nature of the migrating cell behaviour resembles Brownian motion of particles with external bias introduced by the acting environment. Thus, the directed motion of the cell can be mathematically described by the biased **random walk (RW)** [44], while the behaviour of cell population is viewed as a consequence of irregular movement of its members. The irregularity of individual motion impedes precise assessment of the attractant gradient by an individual, yet the average behaviour of cells in the population reflects it with remarkable accuracy.

The first approach to population modelling deals with the population as a whole and directly describes changes in the population density in space and time. The first macroscopic model of chemotaxis dates back to the pioneering work of Keller and Segel in application to slime molds [42]. The original **Keller-Segel (KS)** model consists of four strongly coupled partial differential equations (PDEs) describing the evolution of densities and the mean flux of two populations: attractant molecules and migrating cells. A more widely used KS model can be obtained by reducing the original system to two equations under the quasi-steady-state assumption [45]:

$$\frac{\partial u}{\partial t} = \nabla(D_u \nabla u) - \nabla(\chi(u, v)u \nabla v); \quad (2.2.1a)$$

⁴ Note that all mathematical models presented in this section preserve the notation of original papers and may contradict the notation adopted in the thesis. All parameters are defined in the text.

$$\frac{\partial v}{\partial t} = D_v \nabla^2 v - f(u, v)v, \quad (2.2.1b)$$

where $u(s, t)$ is the density of migrating cells and $v(s, t)$ is the density of attractant molecules at position s and time t , and where D_u denotes cell diffusion coefficient, χ is the coefficient of chemotactic sensitivity. In (2.2.1b), $f(u, v)$ is a function describing attractant degradation and $D_v \gg D_u$ denotes the coefficient of attractant diffusion. The **reaction-diffusion (RD)** system (2.2.1) describes slow diffusion of cells with a drift proportional to the gradient of the ChA concentration ∇v , and fast diffusion of the attractant molecules accompanied by the decay that can be influenced by the migrating cells or some external mediator.

Lauffenburger et al. in [46] account for the fact that the aggregation-diffusion process described above is also affected by fluctuations in cell population. In their model, equation (2.2.1b) remains the same, while two extra terms are added to (2.2.1a):

$$\frac{\partial u}{\partial t} = \nabla(D_u(u, v)\nabla u) - \nabla(\chi(u, v)u\nabla v) + g(u, v) - h(u, v), \quad (2.2.2)$$

where function $g(u, v)$ describes the rate of cell population growth, and function $h(u, v)$ describes the rate of cell death. Many mathematical descriptions of cell chemotaxis are derived from this modified model. Differences between transformed models mainly consist of the forms of functions $f(u, v)$, $g(u, v)$, and $h(u, v)$ [7].

An alternative to describing the evolution of densities is stochastically modelling the trajectory of each individual in the population and deriving the limit RD system to describe the population evolution. In the work that anticipated KS model by nearly two decades, Patlak derives the PDE that governs population densities from the random walk model with external bias [44]. Patlak is the first to coin the idea of directional persistence by introducing the non-independence in the particle direction between two successive steps, which suitably describes the behaviour of some chemotaxing cells. Alt extends this idea in [47] by considering the alternations between persistent random walk and reorientation. The resultant population model describes the evolution of density as a function of the position, time, orientation, and the time of starting the directed motion. Further improvements in explicit derivation of the population model from stochastic equations are recently reviewed in [48]. Microscopic approach to the population modelling requires good approximation of the dynamics of an individual cell. Several stochastic models of individual cell migration are discussed below.

2.2.2 Individual cell level

Majority of single-cell models rely on the similarity of the migrating cell behaviour to the Brownian motion, the process that for a long time has been treated statistically

based on fundamental works [49], [50]. One of the first attempts to model individual cell migration rigorously utilises Ornstein and Uhlenbeck (O-U) process [43]. It describes the velocity of a massive particle committing random walk in presence of friction:

$$\dot{v}(t) = -\beta v(t) + \sqrt{\sigma} \mathbf{W}(t); \quad (2.2.3a)$$

$$x(t) = \int_0^t v(\tau) d\tau, \quad (2.2.3b)$$

where $v(t)$ is the particle velocity and $x(t)$ is its position at time t . In (2.2.3a), the first term denotes the deterministic drag that opposes the motion of the particle with the rate β , and the second term represents random disturbances, where $\mathbf{W}(t)$ is the Wiener process scaled by volatility σ . This behaviour is similar to the random crawling neutrophils perform when not affected by the ChA environment. Various extensions to the model above have been introduced over the years to account for the environment influence.

Stokes et al. in their study of microvessel endothelial cell migration extend the O-U process to account for the directional bias in the cell velocity in reaction to the ChA gradient [51]. This behaviour is described by an additional term in the equation of cell velocity change:

$$\dot{v}(t) = -\beta v(t) + \sqrt{\sigma} \mathbf{W}(t) + \tilde{\Psi}(t); \quad (2.2.4)$$

where β and σ are the same motility parameters as in the O-U process, $\mathbf{W}(t)$ is the Wiener process, and $\tilde{\Psi}(t)$ is a drift function. According to (2.2.4), cell acceleration is the result of the chemotactic bias, random fluctuations, and resistance to motion. The drift function is calculated as the product of cell responsiveness to the attractant κ , magnitude of the chemotactic gradient $\nabla\psi$, and its orientation with respect to the gradient, $\phi(t)$:

$$\tilde{\Psi}(t) = \kappa \nabla\psi \sin \left| \frac{\phi(t)}{2} \right|. \quad (2.2.5)$$

This model is connected to the KS population model of chemotaxis, as the authors claim that chemotactic responsiveness “is related, although probably not in a trivial way, to the cell population chemotaxis coefficient χ ”. The ChA gradient $\nabla\psi$ is viewed as a force that accelerates the particle while the friction force $-\beta v(t)$ opposes the motion. Whilst this representation is a fundamental step in the progress of modelling the environment-driven cell migration, it is constructed under the assumption that the migrating cell is interacting with the environment at all times. This assumption contradicts recent experimental findings, particularly observed periods of random motility in neutrophils and slime molds [52].

A unifying model that combines two observed modes of neutrophil behaviour (chemotactic runs and random motility) treats the ChA gradient as a turning mo-

ment affecting cell orientation rather than acceleration [25], [53]. The model mathematically links the directional randomness of neutrophil migration to receptor-ligand binding fluctuations. Although the model yields important perspective on chemotaxis in the context of receptor sensitivity and occupation, it is created under the assumption of persistent cell polarity. This means that cells have defined front and rear at any moment, the assumption that may not hold the test against the experimental data (recall Figure 2.2).

A more complex hybrid model consisting of five types of random walk is proposed and linked to the *in vivo* neutrophil data in [54]. The following random walks are considered: Brownian motion, biased RW, persistent RW, and two distinct biased-persistent RWs. The type of random walk depending on values of three parameters: bias b , persistence p , and a random variable w that defines whether a cell commits biased or persistent motion. The cell state at time t is determined by a step length $s(t)$ (similar with cell displacement) and the turning angle $\phi(t)$, which is defined as the angle between cell reference axis and the direction of motion

$$s(t) = \sqrt{dt} \times \omega(t), \text{ where} \quad (2.2.6a)$$

$$\omega(t) \sim \mathcal{N}^+(0, 1), \quad (2.2.6b)$$

$$\phi(t) \sim \mathcal{N}(\mu, \sigma) \quad (2.2.6c)$$

where $\mathcal{N}^+(0, 1)$ is a normal distribution truncated at 0, while $\phi(t)$ follows the wrapped Gaussian distribution whose parameters μ and σ depend on the type of the random walk. If the cell is performing a biased RW, the mean $\mu = \beta$, where β is the direction to the wound, while for the persistent motion $\mu = \phi(t - 1)$. The model (2.2.6) describes the cell that at each time moves a distance $s(t)$ in the direction $\phi(t)$ that is determined by the type of RW the cell is currently performing. Although this model does not explicitly link cell motion with locally acting environment, it captures an emerging trend of employing hybrid models to describe non-trivial cell motion. Similar idea is explored in Chapter 3, where a hybrid model incorporates the parametrised model of the environment.

Random walk modelling treats the migrating cell as a particle and therefore only focuses on the direction and the displacement of the cell. A single cell is viewed as a material point whose mass is concentrated in the cell centroid or the leading edge [43], [54], or as a fixed domain [25], [53]. Although various random walk models have been successfully used to characterise the migration of different types of cells, their capacity is restricted by the following:

- The whole-cell approach does not consider subcellular signalling mechanisms that determine directional change of cell movement.
- The fast changing shape of neutrophils falls beyond the scope of cell centroid and leading edge models.

These limitations indicate that the complexity of neutrophil inner mechanisms requires modelling in greater detail.

2.2.3 Subcellular level

For a long time intracellular process regulating cell directed migration has been considered to be strictly sequential [8], [55]. Snyderman and Goetzl in [23] describe two distinct models accounting for the environment sensing mechanism in the migratory process. The first model involves temporal sensing which means that a cell first receives information about the environment, then reorients itself along the selected direction and moves in a straight line for a certain period of time. This type of movement can be observed in bacteria cultures. The second approach is more appropriate for modelling neutrophil intracellular processes because considers spatial sensing which involves morphological changes of a cell exposed to a stimulating gradient.

Since the subcellular processes are usually accompanied by the redistribution of various chemical concentrations, they are commonly represented by the RD systems [56]. A simple two-agent model is given by

$$\frac{\partial A}{\partial t} = D_a \nabla^2 A + f_a(A, B), \quad (2.2.7a)$$

$$\frac{\partial B}{\partial t} = D_b \nabla^2 B + f_b(A, B), \quad (2.2.7b)$$

where A is an autocatalytic activator and B is the inhibitor, $D_a < D_b$ are the diffusion coefficients for activator and inhibitor, respectively, and functions f_a, f_b represent the interactions that lead to the pattern formation [57]. The activator and the inhibitor describe the series of feedback loops between the receptors and the mediators of actin polymerisation (see Figure 2.4a). The accumulation of activator corresponds to the increasing activity of protrusion-promoting proteins (PI3k) and lipids (PIP₃), while the inhibitor may correspond to various mediators that can degrade PIP₃ and actin polymers. Some models consider global inhibition ($D_b \rightarrow \infty$), meaning that after the cell has identified the direction of the gradient an started moving, the signalling from surface receptors stops for some time.

The first model of the dynamic cell polarisation introduced by Meinhardt in [58] is based on the RD system with two antagonists to activation. For a discretised cell boundary composed of K elements, the local evolution of species at the k -th node is described as follows

$$\frac{\partial A_k}{\partial t} = D_a \frac{S_k(A_k^2/B + b_a)}{(S + C_k)(1 + S_a A_k^2)} - r_a A_k, \quad (2.2.8a)$$

$$\frac{\partial B}{\partial t} = b_b \sum_{k=1}^n A_k/n - r_b B, \quad (2.2.8b)$$

$$\frac{\partial C_k}{\partial t} = b_c A_k - r_c C_k, \quad (2.2.8c)$$

where A is the activator, B is the global, rapidly distributed inhibitor, and where C is the local inhibitor, r_a, r_b, r_c are the corresponding rates of decay, and b_a, b_b, b_c are the corresponding rates of production. S is the external signal that translates the information about the local environment (see Figure 2.4b). In this model, global inhibition generates the pattern, while local inhibition is only present at local maxima and deactivates it after some time, facilitating continuous pattern change [57]. The Meinhardt model accounts for many phenomena in the process of polarisation, including strong amplification of the external signal and fast reaction to the changes in the environment [56].

Another popular model of gradient sensing is the *local excitation, global inhibition (LEGI)* model [28], [59], where the external signal S has similar effect on both the local agent E and the slowly acting global inhibitor I , which jointly regulate the downstream signal to polarisation regulator R that can be active (R) or inactive ($R_T - R$) [60]:

$$\frac{\partial E}{\partial t} = -k_{-E}E + k_E S, \quad (2.2.9a)$$

$$\frac{\partial I}{\partial t} = -k_{-I}I + k_I S + D_I \nabla^2 I, \quad (2.2.9b)$$

$$\frac{\partial R}{\partial t} = -k_{-R}IR - k_R(R_T - R)E, \quad (2.2.9c)$$

where R_T denotes the total concentration of the regulator that is assumed to be constant. In this model, the local excitation represents fast self-amplifying generation of enzymes that initiate the growth of pseudopods, while the global inhibition reflects a slowly changing average level of receptor occupancy [61]. The original LEGI model represents only the sensing mechanism and does not account for the amplification of the external signal. In subsequent works [62], the same authors develop a model that couples the LEGI module with the two-agent RD called the biased excitable signalling network (BEN). In the new model, illustrated in Figure 2.4d, the LEGI mechanism models the environment sensing, while BEN accounts for the amplification of the external signal. The excitable network is represented by a RD system similar to (2.2.7). A more recent publication [60] introduces an additional RD system to model spontaneous polarisation in absence of the external gradient gradient.

Although RDS models account for multiple aspects of cell polarisation in presence of external gradients, most of them assume existence of a fast-diffusing global inhibitor. Their applicability for neutrophils is limited by the fact that no candidates for the global inhibitor have been proposed yet. Irimia *et al.* develop an adaptive-control model that combines spatial sensing and polarisation without relying on global inhibition [30]. A system of coupled equations, illustrated in Figure 2.4e,

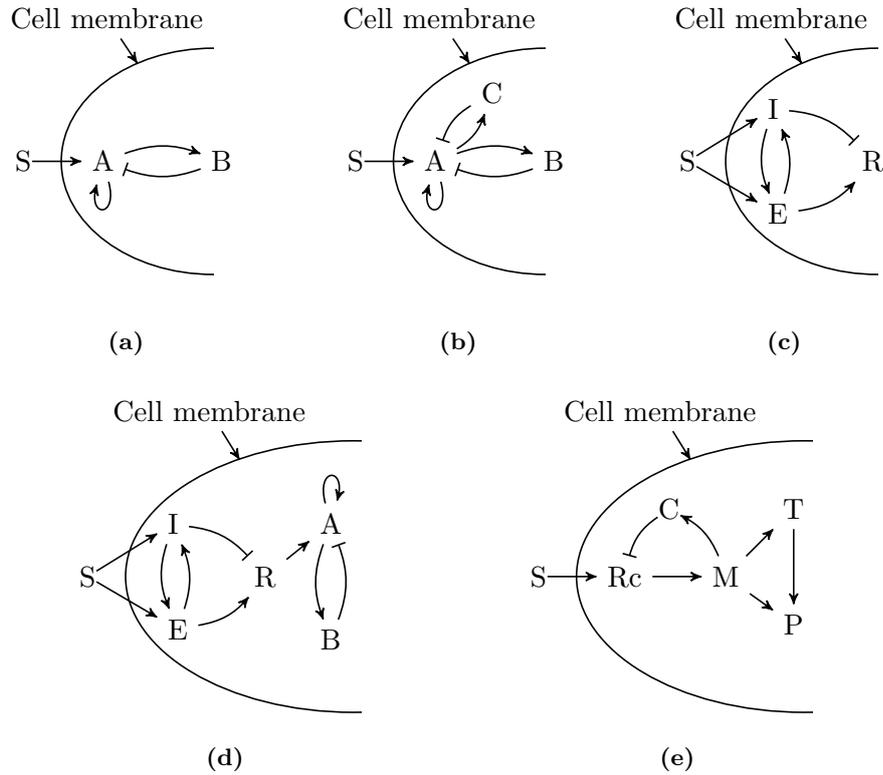


Figure 2.4: Schematic representation of various signalling models. **(a)** RDS with local autocatalytic activator (A) and global, fast-diffusing inhibitor (B) reacting to the external signal (S) [57]. **(b)** RDS with added local inhibitor (C) accounts for the dynamic pattern formation [58]. **(c)** LEGI mechanism of environment sensing: both local excitation (E) and global inhibition (I) are induced by the external signal and regulate downstream signal (R) [59]. **(d)** LEGI-BEN: a combination of LEGI module with the biased excitable network accounts for the signal amplification and polarisation [60], [62]. **(e)** Adaptive control model is an alternative to RD models. The external signal (S) induces the activation of the membrane receptors (Rs), which in turn initiate the activation of the response mediator (M) and the local inhibitor (C). The mediators amplify the external signal and initiate the redistribution of microtubules (T), which leads to the formation of pseudopods (P) [30].

thoroughly describes behaviour of multiple mediators within the cell and accounts for the temporal sensing under the assumption that receptors are situated at the tip of each protrusion. Results of the modelling work are tested on the data from the microfluidic device. The findings suggest that the key role in intracellular signalling processes may be allotted to a feedback between receptors and the actin network. Together with a simplistic model of local attractant concentration it provides necessary connection between interior processes of the cell and its environment.

The majority of computational studies of polarisation and environment sensing have been carried out under the assumption of fixed cell boundary, normally

represented as an unwrapped surface or a sphere. The role of cell geometry in the subcellular processes and vice versa is just starting to be appreciated in the literature [56]. The following subsection will focus on several models that take into account the changing cell domain.

2.2.4 Cell morphodynamics

One of the outstanding questions in modelling the cell motility is how to connect multiple subcellular events and processes with evolving shape of the chemotaxing cell [63]. Several recent theoretical works consider the deforming cell boundary [60], [64]–[66]. Usually, one of the RD systems described above is utilised to account for the subcellular signalling, while the evolving cell membrane is represented by a discretised closed hypersurface $\Gamma(t)$. The hypersurface is obtained by numerically approximating the cell membrane by a finite number of nodes or intervals (elements) whose motion is linked to the chemical patterns produced by the RD system.

The motion of each node is governed locally by the interaction of forces acting normal to the boundary:

$$\mathcal{F} = (\mathcal{F}_{\text{pro}} + \mathcal{F}_{\text{ten}} + \mathcal{F}_{\text{vol}} + \mathcal{F}_{\text{ext}} + \mathcal{F}_{\text{visc}} + \mathcal{F}_{\text{b}})\nu, \quad (2.2.10)$$

where ν is an outward pointing normal to the curve Γ at the node $s \in \Gamma$, and the acting forces are the following:

- **Protrusive force.** It is commonly assumed that \mathcal{F}_{pro} is proportional to the concentration of chemical species active along the cell boundary (for example, cross-linked filaments [67]). For N_a species with local concentrations $a_i(s)$ at the node s on the cell boundary, the corresponding force is expressed as follows

$$\mathcal{F}_{\text{pro}}(s) = \sum_i^{N_a} (k_a)_i a_i(s), \quad (2.2.11)$$

where the sign of the coefficient $(k_a)_i$ defines whether the species a_i promotes protrusion ($(k_a)_i > 0$) or retraction ($(k_a)_i < 0$) of the cell boundary.

- **Surface tension.** \mathcal{F}_{ten} corresponds to the surface energy that prevents the cell membrane from stretching. In the literature, it is incorporated in the model either as a global spatially invariant force that is a product of mean cell curvature \mathcal{K} and mean surface tension γ [65]:

$$\mathcal{F}_{\text{ten}} = \gamma\mathcal{K}, \quad (2.2.12)$$

or as a local force defined at discrete nodes x of curve $\Gamma(t)$ [60]

$$\mathcal{F}_{\text{ten}}(s) = \gamma(s)\kappa(s), \quad (2.2.13)$$

where $\gamma(s)$ is local tension and $\kappa(s)$ is local curvature.

- **Volume conservation.** There exists a pressure difference between exterior and interior of the cell that balances small volume changes caused by boundary evolution. In modelling literature, the corresponding tension force is either assumed to be a hard constant [65], or dependant on change in cell area in 2-D or cell volume in 3-D [60]

$$\mathcal{F}_{\text{vol}} = k_{\text{vol}}(\mathcal{A}_t - \mathcal{A}_0), \quad (2.2.14)$$

where k_{vol} is a constant coefficient, \mathcal{A}_t is the cell area at time t , and \mathcal{A}_0 is the initial cell area. This force is spatially constant.

- **Viscous force.** $\mathcal{F}_{\text{visc}}$ opposes the cell motion

$$\mathcal{F}_{\text{visc}}(s) = -k_{\text{visc}}v_\nu, \quad (2.2.15)$$

where v_ν denotes the local material velocity of the cell boundary that is directed normal to the boundary.

- **External force.** \mathcal{F}_{ext} accounts for possible interactions with the medium in which the cell moves. It can represent interaction with an obstacle or another cell. Majority of modelling studies to date do not account for this type of interaction.
- **Resistance to bending.** The main component of the cell membrane resists excessive bending and accounts for the contributing force

$$\mathcal{F}_b = k_b(\Delta_\Gamma \kappa + \kappa |\nabla_\Gamma|^2 - \frac{1}{2}\kappa), \quad (2.2.16)$$

where κ is the mean curvature, ∇_Γ is the surface gradient, and Δ_Γ is the Laplace-Beltrami operator. See [65] and references therein for the derivation of this model.

Several works couple the simplified force equation with the Meinhardt model. In [64], [68], a 2-D pseudopod-centred model is constructed, where each node on the cell boundary can either move outward subject to the protrusive force (2.2.11), or inward subject to the tension force (2.2.13). The protrusive force is proportional to the local activator level from the system (2.2.8) modified for the evolving domain. Such model implies no direct connection between the pseudopod formation and the centralised processing of external signal. Protrusions can emerge at any point of the cell surface based on local stimuli concentration and receptor occupancy [69]. The model is generalised to the 3-D case in [66].

The authors of [65] also utilise the Meinhardt model, but consider all the aforementioned forces in both 2-D and 3-D cases. The model is constructed under the assumption that the force balance holds, *e.g.* $\mathcal{F} = 0$, and the equation describing the evolution of the cell boundary is obtained by substituting the constituent terms of (2.2.10) by (2.2.11)-(2.2.15). Then the local velocity normal to an individual

segment of the discretised closed curve (or the surface) can be expressed as follows

$$v_\nu = \left(\sum_i \left\{ \frac{(k_a)_i}{k_{\text{visc}}} \mathbf{a}_i \right\} + \frac{\gamma}{k_{\text{visc}}} \kappa + \frac{k_{\text{vol}}}{k_{\text{visc}}} (\mathcal{A}_t - \mathcal{A}_0) + \frac{k_b}{k_{\text{visc}}} (\Delta_\Gamma \kappa + \kappa |\nabla_\Gamma|^2 - \frac{1}{2} \kappa) \right) \nu. \quad (2.2.17)$$

This computational model is constructed under the assumption that the tangent movement of cell membrane is accounted for in surface tension and volume conservation terms, *e.g.* $v_\tau = 0$.

On the other hand, Iglesias *et. al.* in [60] couple the polarised LEGI-BEN model with the evolution equation that includes the protrusive force (2.2.11), local tension (2.2.13), and volume conservation force (2.2.14). The protrusive force is a result of a downstream signal amplified through the BEN. The polarity module introduces an additional pair of feedbacks that translate the effect of this force on the future signalling. The local positive feedback ensures that the next protrusion is more likely to occur next to the existing one, while the global negative feedback occasionally inhibits the activity in the cell and prevents the overzealous polarisation.

One of the most challenging aspects of modelling cell morphodynamics is the representation of the deforming cell domain. Several techniques commonly used in 2-D modelling are reviewed in [41]. The interest of this thesis lies in coupling the velocity of individual nodes on the cell boundary and the local protrusive forces, so the direct modelling of the cell bulk and boundary are not considered.

2.3 Systems approach to the modelling of biological processes

A major challenge presented by the modelling process of any biological system is handling both measurement and model uncertainties. While the former are the result of measurement errors and are external to the system of interest, the latter are inherent to the system and can be subsumed under two broad categories. The first category includes random processes that may influence the evolution of the system. In the models of cell motility described above these processes are represented by stochastic terms in the dynamics. The model uncertainty of the second type is caused by our inability to predict which model terms are significant to describe the evolution of the underlying system. This uncertainty is particularly hard to reduce when modelling life systems, as typically not only the model parameters are unspecified, but some of the signals cannot be directly observed. The abundance of models reviewed in the previous section is the direct consequence of this uncertainty.

From a system-theoretic point of view, a cell interacting with its environment and undergoing changes as a result of this interaction can be viewed as a dynamical

system. Then the cell's perception of the local chemoattractant is the input signal of the system, while its velocity, orientation, and polarity are the characteristic quantities of the cell that evolve in time. The cell properties that can be measured experimentally, *e.g.* cell position, are the observables of the system. The construction of the dynamical model from observable signals is a fundamental problem in systems theory that relies heavily on testing various hypotheses about the system of interest against the available observables. Despite the importance of this problem, it is often overlooked in the cell modelling literature, where mathematical models are presented without validation or linkage to the experimental data.

System identification provides appropriate machinery for quantifying model uncertainty and linking models to the data. The purpose of system identification is to construct the mathematical model that would best describe the dynamical system by processing its observable input and output signals. The "accuracy" of the model is assessed in terms of some criterion and, coincidentally, the identification process often consists of combining methods of statistical inference and optimisation. Before we move on to outline the principal methods of system identification in sections below, a brief introduction of some important notions and assumptions is in order.

2.3.1 State space representation

Mathematical characterisation of complex biological processes often amounts to interpreting temporally resolved experimental data at hand. The *time series* structure of the data arises naturally from the format of many biological experiments, where the process of interest is observed periodically. Systems theory offers a unified way to analyse time series by introducing the notion of *state* - a minimum size vector of variables that fully summarise the past of the system and determine its future evolution [70]. A full set of all possible states the dynamical system can assume is called *state space*.

In the state space approach, the underlying process is described by two time series running in parallel. The first time series is a sequence of unobservable states $\mathbf{x}_0, \dots, \mathbf{x}_T$ that captures the evolution of the system, while the other one consists of observations $\mathbf{y}_0, \dots, \mathbf{y}_T$ associated with these states. It is assumed that the transition from \mathbf{x}_t to \mathbf{x}_{t+1} happens in a definable way, and the process of obtaining \mathbf{y}_t is known. A generic mathematical model that specifies both the state evolution and the observation process in discrete time is called the *discrete time state space*

model and is formulated as follows:

$$\underbrace{\begin{bmatrix} x_1(t+1) \\ x_2(t+1) \\ \vdots \\ x_n(t+1) \end{bmatrix}}_{\mathbf{x}_{t+1}} = \underbrace{\begin{bmatrix} f_1(x_1(t), \dots, x_n, u_1(t), \dots, u_m(t)) \\ f_2(x_1(t), \dots, x_n, u_1(t), \dots, u_m(t)) \\ \vdots \\ f_n(x_1(t), \dots, x_n, u_1(t), \dots, u_m(t)) \end{bmatrix}}_{f(\mathbf{x}_t, \mathbf{u}_t)}, \quad (2.3.1a)$$

$$\underbrace{\begin{bmatrix} y_1(t) \\ y_2(t) \\ \vdots \\ y_p(t) \end{bmatrix}}_{\mathbf{y}_t} = \underbrace{\begin{bmatrix} h_1(x_1(t), \dots, x_n, u_1(t), \dots, u_m(t)) \\ h_2(x_1(t), \dots, x_n, u_1(t), \dots, u_m(t)) \\ \vdots \\ h_n(x_1(t), \dots, x_n, u_1(t), \dots, u_m(t)) \end{bmatrix}}_{h(\mathbf{x}_t, \mathbf{u}_t)}, \quad (2.3.1b)$$

where $t = 0, 1, \dots, T$ is the discrete time index, $\mathbf{x}_t = [x_1 \ x_2 \ \dots \ x_n]^\top \in \mathbb{R}^n$ is the unobserved state of the system, also referred in literature as *hidden* or *latent*, $\mathbf{u}_t = [u_1 \ u_2 \ \dots \ u_m]^\top \in \mathbb{R}^m$ is the input vector consisting of either deterministic or random components, and $\mathbf{y}_t = [y_1 \ y_2 \ \dots \ y_p]^\top \in \mathbb{R}^p$ is the observation. State evolution is described by the function of dynamics, f , and the observation process is described by the measurement function h .

Throughout the thesis, the discrete time model (2.3.1) will be referred to simply as state state model or SSM. Continuous case is not discussed in further sections.

2.3.2 Preliminaries

All problems discussed in this thesis will be defined by the following assumptions.

Assumption 2.3.1: *All SSMs of the underlying processes are linear with respect to the state \mathbf{x}_t and deterministic input $\mathbf{u}_t \forall t$.*

$$f(\mathbf{x}_t, \mathbf{u}_t) = A_t \mathbf{x}_t + B_t \mathbf{u}_t, \quad (2.3.2a)$$

$$h(\mathbf{x}_t, \mathbf{u}_t) = C_t \mathbf{x}_t + D_t \mathbf{u}_t. \quad (2.3.2b)$$

The first equation of system (2.3.2) is known as the discrete dynamics equation with the following parameters:

$A_t \in \mathbb{R}^{n \times n}$ is the state transition matrix,

$B_t \in \mathbb{R}^{n \times m}$ is the control gain.

The second equation is known as the discrete measurement equation with the following parameters:

$C_t \in \mathbb{R}^{p \times n}$ is the measurement matrix,

$D_t \in \mathbb{R}^{p \times m}$ is input-to-measurement gain.

Assumption 2.3.2: *The state evolution happens in presence of the input disturbance, or the **process noise**, represented by a white zero mean Gaussian process ω_t , where whiteness property means that ω_{t_1} is independent of ω_{t_2} at all times.*

$$\mathbf{w}_t \sim \mathcal{N}(0, Q_\omega(t)), \quad (2.3.3a)$$

$$\mathbb{E} \left[\mathbf{w}_{t_1} \mathbf{w}_{t_2}^\top \right] = Q_{t_1} \delta_{t_1, t_2} \quad \forall t_1 \neq t_2, \quad (2.3.3b)$$

where δ_{t_1, t_2} is the Kronecker delta function and where the discrete time noise covariance Q_t is defined directly.

Assumption 2.3.3: *The measurement process happens in presence of the **measurement noise**, represented by a white zero mean Gaussian process \mathbf{v}_t independent of \mathbf{w}_t at all times.*

$$\mathbf{v}_t \sim \mathcal{N}(0, R_v(t)), \quad (2.3.4a)$$

$$\mathbb{E} \left[\mathbf{v}_{t_1} \mathbf{v}_{t_2}^\top \right] = R_{t_1} \delta_{t_1, t_2} \quad \forall t_1 \neq t_2, \quad (2.3.4b)$$

$$\mathbb{E} \left[\mathbf{w}_{t_1} \mathbf{v}_{t_2}^\top \right] = 0 \quad \forall t_1 \neq t_2, \quad (2.3.4c)$$

where the discrete time measurement noise covariance R_t is defined directly.

In presence of the disturbance, not the state itself but its **probability density function (pdf)** uniquely determines the evolution of the underlying process [71]. The dynamics of the stochastic system is governed by the probability laws specifying conditional dependence between hidden states and the measurements. The whiteness of noise in this case is a particularly important property, as it guarantees that the evolution of the deterministic system state is a **Markov process**, *i.e.* the pdf of the state at some time t , conditioned on the past states until the time $t-1$, depends only on the most recent state \mathbf{x}_{t-1} :

$$p(\mathbf{x}_t \mid \mathbf{x}_{0:t-1}, \mathbf{u}_{t-1}) = p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, \mathbf{u}_{t-1}). \quad (2.3.5)$$

Assumption 2.3.4: *The initial state x_0 conforms to a Gaussian distribution with a prior mean $\bar{\mathbf{x}}_0$:*

$$\mathbf{x}_0 \sim \pi(\mathbf{x}_0) = \mathcal{N}(\bar{\mathbf{x}}_0, \mathbf{P}_0), \quad (2.3.6)$$

which is independent from both noise sequences:

$$\mathbb{E} \left[\mathbf{x}_0 \mathbf{w}_t^\top \right] = 0 \quad \forall t, \quad (2.3.7a)$$

$$\mathbb{E} \left[\mathbf{x}_0 \mathbf{v}_t^\top \right] = 0 \quad \forall t. \quad (2.3.7b)$$

The resulting SSM defined by the assumptions above is called a **discrete time linear stochastic system** and is described by the following equations

$$\mathbf{x}_{t+1} = A_t \mathbf{x}_t + B_t \mathbf{u}_t + G_t \mathbf{w}_t \quad (2.3.8a)$$

$$\mathbf{y}_t = C_t \mathbf{x}_t + D_t \mathbf{u}_t + H_t \mathbf{v}_t, \quad (2.3.8b)$$

where $\mathbf{w}_t \sim \mathcal{N}(0, Q_t)$, $Q_t \in \mathbb{R}^{n_w \times n_w}$ is the Gaussian white noise that enters the system via the process noise gain $G_t \in \mathbb{R}^{n \times n_w}$, and where $\mathbf{v}_t \sim \mathcal{N}(0, R_t)$, $Q_t \in \mathbb{R}^{n_v \times n_v}$ is the Gaussian white noise that enters the system via the measurement noise gain $H_t \in \mathbb{R}^{p \times n_v}$. Since the relationship between hidden states is Markovian, (2.3.8) is also referred to as a **hidden Markov model (HMM)**. A diagram specifying causal relationship between variables of the system (2.3.8) is presented in Figure 2.5.

Typically, hidden states cannot be deduced directly from the observations, but they can be recovered given the model structure. However, as was discussed in the beginning of this section, it is very common in life systems applications to have limited idea about model parameters. Depending on whether states or parameters (or both) are of interest, there are three broad categories of estimation problems associated with stochastic SSMs:

- Reconstruction of the hidden states $\mathbf{x} = \{\mathbf{x}_t\}_{t=0}^T$ from the set of available measurements $\mathbf{y} = \{\mathbf{y}_t\}_{t=0}^T$ is called **state estimation** problem.
- The problem of estimating a single parameter or an ensemble of parameters, *e.g.* process matrices in (2.3.8), from a set of observations \mathbf{y} is called **parameter estimation** problem or **parameter learning** problem.
- In some cases, especially at early stages of the modelling process, both true states of the systems and model parameters are unknown and must be inferred given \mathbf{y} . This problem is referred to as **joint state-parameter estimation**.

Further sections explore common estimation methods dealing with these problems.

2.4 Parameter estimation

The function that infers a property of the underlying process from the available data is called an estimator. It is formally defined as follows [72]:

Definition 2.4.1 (The estimator): *The estimator of the unknown variable θ given a set of observations*

$$\mathbf{y}_t = f(t, \theta, \mathbf{v}_t), \quad t = 0, \dots, T, \quad (2.4.1)$$

made in presence of the measurement noise $\mathbf{v}_t \sim \mathcal{N}(0, R_v)$ is a function of $T+1$ discrete-time observations

$$\hat{\theta} \triangleq \hat{\theta}[T, \mathcal{Y}], \quad (2.4.2)$$

where \mathcal{Y} is the set of observations associated with θ

$$\mathcal{Y} = \{\mathbf{y}_t\}_{t=0}^T. \quad (2.4.3)$$

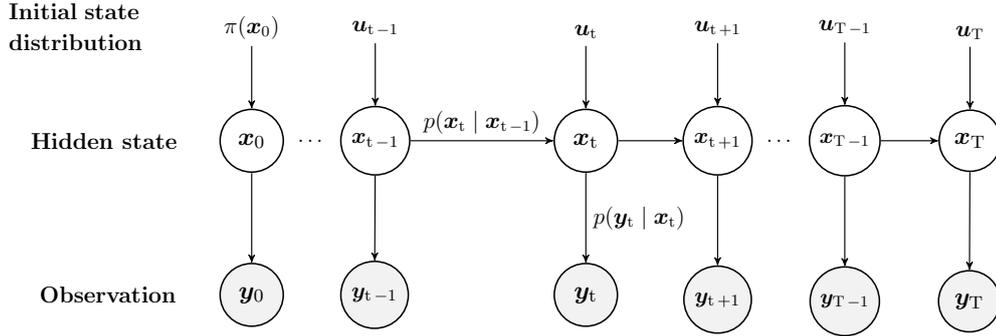


Figure 2.5: A directed graph of a hidden Markov model representing time series. Arrows represent causality between states \mathbf{x}_t and observations \mathbf{y}_t . In Markov processes, the state \mathbf{x}_t is conditionally independent on all past states given the previous one \mathbf{x}_{t-1} . The observation \mathbf{y}_t is dependent only on the current state \mathbf{x}_t .

The unknown variable can be either scalar, denoted by θ , or vector valued, denoted by Θ , and may include model parameters, hidden states, or future data that is yet to be observed [73]. The task of parameter estimation from the data collected in presence of measurement noise is commonly formulated as a problem of statistical inference [74].

There are two commonly used approaches to statistical inference that arise from different interpretations of probability. The first one is Bayesian approach, that treats probability as a reasonable expectation that is corresponding to some rational belief about the unknown parameters [75]. The uncertainty of the model is the result of one's limited knowledge about it, thus all the unknowns are uncertainties that can be quantified by probability distributions. This approach finds the parameter set $\hat{\Theta}(\mathcal{Y})$ that minimises the errors between the available data \mathcal{Y} and the predicted model behaviour.

An alternative to Bayesian is the frequentist approach [76], that asserts that the probability of an event is asymptotically close to the frequency of its occurrence as the number of independent trials approaches infinity. In this setting, the uncertainty of the model arises from the limited number of observations, while the unknown parameters are treated as constants. Frequentist methods aim to find the parameter set $\hat{\Theta}(\mathcal{Y})$ that has the highest probability of producing the available observation set. In other words, the frequentist approach fixes the underlying parameter and assesses variability of the data, while in Bayesian methods the data is fixed and the uncertainty of the parameter is assessed.

Despite the ongoing 250-year long debate about which approach is superior, it is widely recognised at the methodological level that both Bayesian and frequentist approaches are legitimate and may compliment each other in application to certain inference problems [77]. The following subsections will outline basic concept of both

approaches and demonstrate the connection between them.

2.4.1 Bayesian estimation

Bayesian estimation [78] deals with the model uncertainty by treating its unknown parameters as random variables. As such, the parameters have a *pdf* that should be taken into account in the estimation process. In this approach, the knowledge about the hidden parameters available prior to receiving any observations is summarised in the prior distribution normally referred to simply as *prior*. This distribution is then combined with the information coming from the available data to produce a posterior distribution of the unknown parameters.

In Bayesian approach, the propagation of distributions relies on the Bayes' rule [78]:

$$p(\Theta | \mathcal{Y}) = \frac{p(\mathcal{Y} | \Theta)\pi(\Theta)}{p(\mathcal{Y})}, \quad (2.4.4)$$

where Θ denotes a set of hidden properties that are of interest, $\pi(\Theta)$ is their prior distribution, and \mathcal{Y} is a set of observations associated with Θ , called the *evidence* of Θ . $p(\mathcal{Y})$ denotes the probability of observing \mathcal{Y} that can be rewritten as follows

$$p(\mathcal{Y}) = \int_{-\infty}^{\infty} p(\mathcal{Y} | \Theta)\pi(\Theta)d\Theta. \quad (2.4.5)$$

Although the evidence does not give the exact idea about Θ , there exists the probability density $p(\mathcal{Y} | \Theta)$ that is a measure of likelihood that the observation set is generated by the hidden parameters, called the *likelihood function*. Finally, $p(\Theta | \mathcal{Y})$ denotes the sought *posterior distribution* of the unknown parameters given the evidence \mathcal{Y} . Usually, calculating the observation probability $p(\mathcal{Y})$ is not necessary, and it can be treated as a normalising constant for the posterior distribution:

$$p(\Theta | \mathcal{Y}) = \frac{1}{c}p(\mathcal{Y} | \Theta)\pi(\Theta). \quad (2.4.6)$$

In practical applications, knowing the parameter distribution alone may not be particularly meaningful and it is necessary to obtain a finite-dimensional estimate of Θ , called the *point-wise estimate*. Bayesian methods find a point-wise estimate $\hat{\Theta}(\mathcal{Y})$ by minimising some *loss function* that quantifies the discrepancy between the true value of the unknown parameter and its estimate [79]. The most common types of loss functions are the quadratic loss function [80]:

$$l(\Theta, \hat{\Theta}) = (\Theta - \hat{\Theta})^\top L(\Theta - \hat{\Theta}), \quad (2.4.7)$$

and the 0-1 loss function [81]:

$$l(\Theta, \hat{\Theta}) = \begin{cases} 0, & \text{for } |\Theta - \hat{\Theta}| < \epsilon, \\ 1, & \text{otherwise.} \end{cases} \quad (2.4.8)$$

Since Θ is a random variable, it makes sense to consider statistics of the selected loss function. The average loss associated with $\hat{\Theta}(\mathcal{Y})$ is called the *risk* and is defined as the expectation of the loss function over the density of \mathcal{Y} :

$$R(\Theta, \hat{\Theta}) = \int_{\mathcal{Y}} l(\Theta, \hat{\Theta})p(\mathcal{Y}, \Theta)d\mathcal{Y}. \quad (2.4.9)$$

Bayesian estimation averages the risk over the available prior distribution of Θ .

Definition 2.4.2 (Generic Bayes estimator): *The Bayes estimator minimises the expected risk marginalised over the prior distribution of the unknown parameter:*

$$\hat{\Theta}_B = \arg \min_{\hat{\Theta}} \int_{\Theta} R(\Theta, \hat{\Theta})d\Theta = \arg \min_{\hat{\Theta}} \int_{\Theta} \int_{\mathcal{Y}} l(\Theta, \hat{\Theta})p(\mathcal{Y}, \Theta)d\mathcal{Y}d\Theta. \quad (2.4.10)$$

Most of Bayesian methods require knowledge of the generative model for \mathcal{Y} in order to compute the likelihood function. However, recently developed numerical methods offer ways to circumvent direct evaluation of the likelihood for scenarios when the mathematical properties of the likelihood function are not fully known. One class of these methods bypasses exact calculation of the likelihood by simulating from it using Markov Chain Monte Carlo (MCMC) method. An overview of MCMC-based algorithms is provided in [82]. Another popular class of algorithms is called Approximate Bayesian Computation (ABC) and is based on rejection sampling. Instead of evaluating the likelihood, an ABC rejection sampler draws a parameter sample from prior $\pi(\Theta)$ and simulates the measurement set with each value in the sample. The sampled parameter values are accepted or rejected after assessment of the distance between some summary statistics of simulated and real measurements. The fundamental idea behind the ABC and its extensions are reviewed in [83].

2.4.2 Maximum Likelihood estimation

In cases when the prior $\pi(\Theta)$ is unavailable, the estimation relies solely on the parameter likelihood function introduced in the subsection above. It is defined as a *pdf* of the observations conditioned on the unknown parameter

$$\mathcal{L}(\Theta) \triangleq p(\mathcal{Y} | \Theta). \quad (2.4.11)$$

The aim of the *Maximum Likelihood (ML)* estimation is to find the value of Θ that has the highest likelihood of generating the data \mathcal{Y} . As a frequentist approach, ML estimation does not associate any probabilities with the unknown parameter and outputs a point-wise estimate that maximises the likelihood function in the form (2.4.11).

Definition 2.4.3 (Maximum likelihood estimator):

$$\hat{\Theta}_{ML} = \arg \max_{\Theta} \mathcal{L}(\Theta). \quad (2.4.12)$$

Since the data is treated as an ensemble of independent random variables, its probability can be computed as a product of probabilities associated with each observation

$$\mathcal{L}(\Theta) = \prod_{t=0}^T p(\mathbf{y}_t | \Theta). \quad (2.4.13)$$

As all individual densities $p(\mathbf{y}_t | \Theta)$ are less than 1, the product (2.4.13) will quickly become very small, thus calculating the product for the large sample is infeasible, making logarithms are used to avoid some of the unnecessary computations. Monotonicity of the logarithm guarantees that maximising full likelihood function is equivalent to maximising the *log likelihood* in the form

$$\log \mathcal{L}(\Theta) = \sum_{t=0}^T \log p(\mathbf{y}_t | \Theta), \quad (2.4.14)$$

which is especially convenient with Gaussian assumptions, as (2.4.14) becomes the sum of products of estimation errors and their covariances.

The *Maximum Likelihood estimate (MLE)* can be obtained by maximising (2.4.14):

$$\frac{\partial \log \mathcal{L}(\Theta)}{\partial \Theta} = 0. \quad (2.4.15)$$

If the model is linear with respect to the unknown parameter set Θ and the log likelihood is tractable, the analytical solution can be easily obtained. In most cases, however, the log likelihood (2.4.14) must be maximised numerically. A number of available numerical techniques for solving (2.4.15), such as Newton-Raphson method and its modifications, are discussed in [84].

It is important to note that, while the true value of the parameter Θ^* is an unknown constant, the MLE $\hat{\Theta}_{ML}$ is a function of the set of random measurements \mathcal{Y} and therefore is a random variable. Furthermore, under the regularity conditions listed in the Theorem 2.6 of [85], the MLE is proven to be asymptotically normal as $n \rightarrow \infty$

$$\hat{\Theta}_n \sim \mathcal{N}\left(\Theta^*, \frac{1}{\mathcal{I}(\Theta^*)}\right) \quad (2.4.16)$$

where n denotes the size of the measurement set \mathcal{Y} and the covariance of the random variable is expressed as the reciprocal of the *expected Fisher information*, $\mathcal{I}(\Theta^*)$, which represents the amount of information about the parameter contained in the data \mathcal{Y} . It is often expressed as the expected negative Hessian of the log likelihood evaluated at the true value of the parameter

$$\mathcal{I}(\Theta^*) = \mathbb{E} \left[-\frac{\partial^2 \log \mathcal{L}(\Theta)}{\partial \Theta^2} \right] \Big|_{\Theta=\Theta^*} \quad (2.4.17)$$

In practice, the expected information is hard to compute, and the *observed Fisher information* is used instead. It is given by the negative Hessian of the

log likelihood evaluated at the current MLE value

$$\mathcal{J}(\hat{\Theta}_n) = -\frac{\partial^2 \log \mathcal{L}(\Theta)}{\partial \Theta^2} \Big|_{\Theta=\hat{\Theta}_n}. \quad (2.4.18)$$

It is often argued in the literature that the **observed information** is a more practical measure of the data variability (for the frequentist justification see [86]). Given a large enough data sample, $\mathcal{J}(\hat{\Theta}_n)$ can be used to approximate $\mathcal{I}(\Theta^*)$. This is because with the growing number of observations ($n \rightarrow \infty$), as the estimate approaches the true parameter value, the expectation of the observed information approaches the expected information.

$$\mathbb{E} \left[\mathcal{J}(\hat{\Theta}_n) \right] \rightarrow \mathcal{I}(\Theta^*). \quad (2.4.19)$$

Predictable behaviour of the MLE with the growing number of observations enables one to express the uncertainty of the estimator in terms of **confidence regions**. For the Gaussian distribution (2.4.16), the volume of the confidence region is given by

$$V_\alpha = \frac{2\pi^{p/2}}{p\Gamma(\frac{p}{2})} (\chi_{p,\alpha}^2)^{p/2} |\Sigma|^{1/2}, \quad (2.4.20)$$

where α denotes the level of confidence, p is the dimension of Θ , $\Gamma(\cdot)$ is the gamma function, and where the $\chi_{p,\alpha}^2$ is the Chi-Square probability associated with the given confidence level and the dimension p . The volume of the confidence region is proportional to the square root of the generalised variance Σ which is approximated by the inverse of $\mathcal{J}(\hat{\Theta}_n)$ [87]. In other words, the more "informative" is the data sample at hand, the smaller is the volume of the confidence region (width of the interval in univariate case). Note that confidence region is a random variable in a sense that given the realisation of the sample \mathcal{Y} , it will either contain Θ^* or not [73].

Remark 2.4.1: *Confidence regions are inferences about the reliability of the estimation procedure rather than about the parameter uncertainty.*

For example, the 95% confidence region will not contain 95% of MLE realisations, but it will contain the true value of the parameter with the certainty of 95%.

2.4.3 Maximum a posteriori estimation

Maximum a posteriori (MAP) estimation is often regarded as a hybrid between Bayesian and frequentist approaches, as it is a point-wise estimate that utilises prior knowledge about the underlying parameters. MAP estimator is usually described in the Bayesian setting, as it maximises the posterior distribution of the unknown parameter.

Definition 2.4.4 (Maximum a posteriori estimator):

$$\hat{\Theta}_{MAP} = \arg \max_{\Theta} p(\Theta | \mathcal{Y}). \quad (2.4.21)$$

It is commonly accepted that the MAP estimator is the limit of the Bayes estimator (see def. 2.4.2) with the hit-or-miss loss function (2.4.8) [88], which allows to subsume it under the category of Bayesian methods. The recent paper defines additional conditions for this assertion to hold in the general setting [89]. On the other hand, if one is to consider the MAP estimator with diffuse prior for the parameter

$$\pi(\Theta) = \epsilon \quad \text{for } |\Theta| < \frac{1}{2\epsilon}, \quad (2.4.22)$$

where $\epsilon > 0$ but small, from the definition of data probability (2.4.5) it follows that

$$p(\mathcal{Y}) = \int_{-\infty}^{\infty} p(\mathcal{Y} | \Theta) \pi(\Theta) d\Theta = \epsilon \int_{-\infty}^{\infty} p(\mathcal{Y} | \Theta) d\Theta = \epsilon g(\mathcal{Y}), \quad (2.4.23)$$

where $g(\mathcal{Y})$ is not a function of Θ . Inserting the denominator into the Bayes formula (2.4.4) renders

$$p(\Theta | \mathcal{Y}) = \frac{p(\mathcal{Y} | \Theta) \epsilon}{\epsilon g(\mathcal{Y})} = \frac{1}{\epsilon} p(\mathcal{Y} | \Theta), \quad (2.4.24)$$

meaning that a MAP estimate with a diffuse prior *pdf* is proportional to the likelihood function, and thus coincides with the MLE [72]. This property of the MAP estimator provides a unifying view on Bayesian and frequentist estimation.

Remark 2.4.2: *The ML estimate is a Bayesian MAP estimate with complete ignorance of prior.*

In view of this connection between the two approaches, the choice of the method depends on the expertise of an individual researcher and is often defined by such aspects as availability of the model structure, possibility to manipulate the experiment design, or one's confidence in the prior information about the unknowns. Admittedly, Bayesian methods are more effective for small data samples or when data is processed as it comes in. However, the trade-off of relying on the prior assumptions is the risk of collapsing into subjectivity. In absence of a unified procedure for choosing a prior, different Bayesian estimators may arrive to contradicting conclusions about the same data. The Maximum likelihood approach, on the other hand, provides an objective solution that is achieved in a consistent and rigorous way with minimum assumptions in cases when no prior information is available. This thesis utilises ML methods for parameter learning and Bayesian methods for state estimation.

2.4.4 Least Squares estimation

Another common approach to non-random parameter estimation is the **Least Squares (LS)** method. First introduced by Gauss in 1809 [90], it pre-dates the ML estim-

ator but can be considered its special case under the zero-mean Gaussian assumption. The simplest LS algorithm makes no assumption about the measurement noise $\mathbf{v} \sim \mathcal{N}(0, R_v)$. It considers a linear input-output model given by

$$\mathbf{y}_t = \Theta \mathbf{x}_t + \mathbf{v}_t, \quad t = 0, \dots, T, \quad (2.4.25)$$

where both the input \mathbf{x}_t and output \mathbf{y}_t are known, and where Θ and \mathbf{v}_t are unknown. For any selection of $\hat{\Theta}$ there exists a difference between the predicted output and available output, called the *residual*:

$$\epsilon_t = \mathbf{y}_t - \hat{\mathbf{y}}_t. \quad (2.4.26)$$

The LS estimate is obtained by minimising the sum of squared residuals.

Definition 2.4.5 (Ordinary Least Squares estimator):

$$\hat{\Theta}_{LS} = \arg \min_{\Theta} \sum_t^T \epsilon_t^\top \epsilon_t. \quad (2.4.27)$$

The criterion in (2.4.27) can be rewritten as a quadratic cost function

$$l(\hat{\Theta}) = \left(\mathbf{y}_t - \hat{\Theta} \mathbf{x}_t \right)^\top R_v^{-1} \left(\mathbf{y}_t - \hat{\Theta} \mathbf{x}_t \right), \quad (2.4.28)$$

minimising which is the equivalent of maximising the log likelihood function for a linear SSM with zero-mean Gaussian noise.

Despite the apparent simplicity of the LS estimator, it has found multiple applications in modified forms. For instance, a generalised Least Squares (GLS) algorithm derived in [91] is often employed for estimation in presence of correlated measurement errors. The regularised Least Squares (RLS) estimator can be utilised in situations when the model suffers from bad generalisation or the number of estimated parameters is unknown [92]. Furthermore, the recursive LS estimator [72] served as foundation for the development of the most ubiquitous recursive state estimator that will be discussed in the following section.

2.4.5 Parameter inference from the incomplete data

The *Expectation-Maximisation (EM)* algorithm [93] is a broadly used iterative method of obtaining MLEs in situations when only the absence of some data impedes straightforward ML estimation [84]. Under the assumption that there exists a complete data set $\mathcal{Z} = \{\mathcal{X}, \mathcal{Y}\}$ that consists of the observable *incomplete data* \mathcal{Y} , and the *missing data* \mathcal{X} , the algorithm converts the parameter estimation into two coupled problems: recovering the expected value of the complete-data log likelihood function, and optimising the obtained function with respect to the unknown

parameters Θ . The two problems are solved iteratively and on each iteration the expectation is conditioned on the previously obtained parameter estimate.

Let the $\mathcal{Z} = \{\mathcal{Y}, \mathcal{X}\}$ be a complete data set with a joint density $p(\mathcal{Z} | \Theta)$. Then log likelihood function of the parameter is defined as

$$\log \mathcal{L}(\Theta) = \log p(\mathcal{Z} | \Theta) = \log p(\mathcal{X} | \mathcal{Y}, \Theta). \quad (2.4.29)$$

As $\log \mathcal{L}(\Theta)$ is not tractable without the exact knowledge of \mathcal{X} , at each iteration i it is replaced by its lower bound which is obtained by introducing marginal distribution over hidden data, $p(\mathcal{X} | \mathcal{Y}, \Theta^i) > 0$, where Θ^i is the current parameter estimate. Then (??) takes the following form

$$\log \mathcal{L}(\Theta) = \log \int_{\mathcal{X}} p(\mathcal{X} | \mathcal{Y}, \Theta^i) \frac{p(\mathcal{Y}, \mathcal{X} | \Theta)}{p(\mathcal{X} | \mathcal{Y}, \Theta^i)} d\{\mathcal{X}\}. \quad (2.4.30)$$

Note that $\int_{\mathcal{X}} p(\mathcal{X} | \mathcal{Y}, \Theta) d\mathcal{X} = 1$. Using Jensen's inequality B.2, one can establish that

$$\log \mathcal{L}(\Theta) \geq \int_{\mathcal{X}} p(\mathcal{X} | \mathcal{Y}, \Theta^i) \log \frac{p(\mathcal{X}, \mathcal{Y} | \Theta)}{p(\mathcal{X} | \mathcal{Y}, \Theta^i)} d\{\mathcal{X}\}, \quad (2.4.31)$$

where the terms can be regrouped as follows

$$\begin{aligned} \log \mathcal{L}(\Theta) &\geq \int_{\mathcal{X}} p(\mathcal{X} | \mathcal{Y}, \Theta^i) \log p(\mathcal{X}, \mathcal{Y} | \Theta) d\{\mathcal{X}\} - \\ &\quad - \int_{\mathcal{X}} p(\mathcal{X} | \mathcal{Y}, \Theta^i) \log p(\mathcal{X} | \mathcal{Y}, \Theta^i) d\{\mathcal{X}\}. \end{aligned} \quad (2.4.32)$$

The second term in (2.4.32) is not a function of Θ and can be disregarded in maximising the log likelihood. Noting the definition of the expected value B.1, yields the lower bound that must be evaluated:

$$\mathcal{Q}(\Theta, \hat{\Theta}^{i-1}) = \mathbb{E} \left[\log \mathcal{L}(\Theta) | \mathcal{Y}, \hat{\Theta}^{i-1} \right]. \quad (2.4.33)$$

The computation of (2.4.33), commonly called the \mathcal{Q} -function, constitutes the **expectation step (E-step)** of the algorithm. When processing HMMs, this step usually amounts to computing the posterior marginals of the hidden states \mathcal{X} given observations \mathcal{Y} via the **forward-backward algorithm**. Then the **maximisation step (M-step)** amounts to obtaining the MLE that maximises the \mathcal{Q} -function:

$$\hat{\Theta}^i = \arg \max_{\Theta} \mathcal{Q}(\Theta, \hat{\Theta}^{i-1}).$$

The EM algorithm is usually initialised with an arbitrary estimate $\hat{\Theta}^0$. The two steps defined above are then repeated until some convergence criterion is satisfied. Convergence properties of the EM algorithm are first discussed in [93] and established for a more generic case in [94].

Theorem 2.4.1: *The generic EM algorithm generates a non-descending sequence of log-likelihood function values*

$$\log \mathcal{L}(\Theta^i) \geq \log \mathcal{L}(\Theta^{i-1}),$$

that converges to a stationary point $\log \mathcal{L}(\Theta^*)$.

Proof. From (4.3.3) one can observe that $\log \mathcal{L}(\Theta^i)$ is lower bounded by $F(\Theta | \Theta^i)$ expressed as

$$F(\Theta | \Theta^i) = \mathcal{Q}(\Theta | \Theta^i) - H(\Theta | \Theta^i),$$

where

$$H(\Theta | \Theta^i) = \mathbb{E}[\log p(\mathcal{X} | \mathcal{Y}, \Theta) | \mathcal{Y}, \Theta^i].$$

The difference between lower bound values in two successive iterations is given by

$$\begin{aligned} F(\Theta | \Theta^i) - F(\Theta | \Theta^{i-1}) &= [\mathcal{Q}(\Theta^i | \Theta^i) - \mathcal{Q}(\Theta^{i-1} | \Theta^{i-1})] - \\ &\quad - [H(\Theta^i | \Theta^{i-1}) - H(\Theta^{i-1} | \Theta^{i-1})], \end{aligned} \quad (2.4.34)$$

where $\mathcal{Q}(\Theta^i | \Theta^i) - \mathcal{Q}(\Theta^{i-1} | \Theta^{i-1}) \geq 0$ by definition of the maximisation step. The second term of (2.4.34) is the Kullback-Leibler distance between two distributions with densities $p(\mathcal{X} | \mathcal{Y}, \Theta^{i+1})$ and $p(\mathcal{X} | \mathcal{Y}, \Theta^i)$. As a consequence of Jensen's inequality B.2, $H(\Theta^i | \Theta^{i-1}) \leq H(\Theta^{i-1} | \Theta^{i-1})$ with equality reachable if and only if $p(\mathcal{X} | \mathcal{Y}, \Theta^{i+1}) = p(\mathcal{X} | \mathcal{Y}, \Theta^i)$ almost everywhere. Hence, the difference defined by (2.4.34) is non-negative and $\log \mathcal{L}(\Theta^i)$ is non-descending function of i .

The function $\mathcal{Q}(\Theta^i | \Theta^{i-1})$ is continuous in both arguments, satisfying condition of Theorem 2 in [94], which assures convergence of log-likelihood function to a local maximum $\log \mathcal{L}(\Theta^*)$, completing the proof. \blacksquare

Similarly with Newton-type methods, global convergence of the EM cannot be guaranteed and the estimation relies on the initial value selection. This is compensated by many advantages of the EM algorithm over numerical methods such as its numerical stability, low computational costs per iteration, and simplicity of monitoring convergence that arises from monotone increase in likelihood with each iteration [84].

Although here the EM algorithm is presented from the ML perspective, it can be utilised to produce a MAP estimate given an informative prior *pdf* $\pi(\Theta)$ [95]. The MAP estimate of Θ corresponds to the maximum of the log posterior density given by

$$\log p(\Theta | \mathcal{Y}) = \log \mathcal{L}(\Theta) + \log \pi(\Theta). \quad (2.4.35)$$

Effectively, the modified algorithm consists of the computation and maximisation of the augmented \mathcal{Q} -function in the following form:

$$\mathbb{E}[\log p(\Theta | \mathcal{Y}) | \mathcal{Y}, \hat{\Theta}^{i-1}] = \mathcal{Q}(\Theta, \hat{\Theta}^{i-1}) + \mathbb{E}[\log \pi(\Theta) | \mathcal{Y}, \hat{\Theta}^{i-1}]. \quad (2.4.36)$$

The presence of the second term in the right hand side of (2.4.36) as a Bayesian prior almost always makes the maximised function more concave than that in the ML case [84].

2.4.6 The expected log likelihood for a state space model

The problem of joint state-parameter estimation of the model (2.3.8) can be formulated as an incomplete-data problem [96], where the available measurement sequence $\mathbf{y} = \{\mathbf{y}\}_0^T$ are regarded as the incomplete data \mathcal{Y} , while hidden states $\mathbf{x} = \{\mathbf{x}\}_0^T$ constitute the missing data \mathcal{X} that must be recovered along with unknown parameters. Assume that initial state of the system and all its dynamical matrices (considered to be time-invariant for this derivation) are unknown:

$$\Theta = \{A, B, C, D, \bar{\mathbf{x}}_0, \mathbf{P}_0\}. \quad (2.4.37)$$

In the general EM framework, the complete data defined above consists of two subsets \mathcal{X} , \mathcal{Y} and a mappings from \mathcal{X} onto \mathcal{Y} . It follows that

$$p(\mathcal{Z} | \Theta) = p(\mathcal{Y} | \mathcal{X}, \Theta)p(\mathcal{X} | \Theta), \quad (2.4.38)$$

In state space context, two mappings emerge: one from the hidden data \mathcal{X} on the observed data \mathcal{Y} , and another one from \mathcal{X} onto itself. Under the Markovian assumptions adopted in Section 2.3, the right hand side in (2.4.38) can be substituted by the product of individual *pdfs*

$$p(\mathcal{Z} | \Theta) = \pi(\mathbf{x}_0 | \Theta) \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{x}_{t-1}, \Theta) \prod_{t=0}^T p(\mathbf{y}_t | \mathbf{x}_t, \Theta), \quad (2.4.39)$$

where $\pi(\mathbf{x}_0 | \Theta)$ is the initial state *pdf*, $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \Theta)$ is the state update *pdf*, and where $p(\mathbf{y}_t | \mathbf{x}_t, \Theta)$ is the observation emission *pdf*. Using the properties of the log function and the definition (2.4.29), the log likelihood function can be factorised as follows

$$\begin{aligned} \log \mathcal{L}(\Theta) &= \underbrace{\log \pi(\mathbf{x}_0 | \bar{\mathbf{x}}_0, \mathbf{P}_0)}_{\mathcal{L}_0(\Theta)} + \underbrace{\sum_{t=1}^T \log p(\mathbf{x}_t | \mathbf{x}_{t-1}, A, B, Q_w)}_{\mathcal{L}_X(\Theta)} + \\ &\quad + \underbrace{\sum_{t=0}^T \log p(\mathbf{y}_t | \mathbf{x}_t, C, R_v)}_{\mathcal{L}_Y(\Theta)}, \end{aligned} \quad (2.4.40)$$

Given the above expression for the log-likelihood, the Q -function is divided into three independent terms, each depending on the small portion of parameters:

$$\begin{aligned} \mathcal{Q}(\Theta, \hat{\Theta}^i) &= \mathbb{E} [\log \mathcal{L}(\Theta) | \mathcal{Y}, \hat{\Theta}^i] = \mathbb{E} [\log \mathcal{L}_0(\bar{\mathbf{x}}_0, \mathbf{P}_0) | \mathcal{Y}, \hat{\Theta}^i] + \\ &\quad + \mathbb{E} [\log \mathcal{L}_X(A, B, Q_w) | \mathcal{Y}, \hat{\Theta}^i] + [\log \mathcal{L}_Y(C, R_v) | \mathcal{Y}, \hat{\Theta}^i], \end{aligned} \quad (2.4.41)$$

which ensures a straightforward of the maximisation step. Further expanding each *pdf* (see, for example, [97]) shows that evaluation of (2.4.41) via the forward-backward algorithm depends on the following conditional expectations

$$\mathbb{E} \left[\mathbf{x}_t \mid \mathcal{Y}, \hat{\Theta}^i \right]; \quad (2.4.42a)$$

$$\mathbb{E} \left[\mathbf{x}_t \mathbf{x}_t^\top \mid \mathcal{Y}, \hat{\Theta}^i \right]; \quad (2.4.42b)$$

$$\mathbb{E} \left[\mathbf{x}_t \mathbf{x}_{t-1}^\top \mid \mathcal{Y}, \hat{\Theta}^i \right]. \quad (2.4.42c)$$

Hence, the computation of the expected log-likelihood function $Q(\Theta, \hat{\Theta}^i)$ requires evaluation of the state expectations given the set of observations \mathcal{Y} , which constitutes the problem of MAP state estimation. The common framework for solving this problem is presented in the following section.

2.5 State estimation

The purpose of state estimation is to recover the hidden state \mathbf{x}_t of a dynamical system which is observed through a set of noisy measurements $\mathbf{y} = \{y_t\}_{t=1}^{T^*}$. Depending on the time span of the measurement set, state estimation methods are divided in three classes:

- ***prediction***, where the state at time t is estimated given the measurements available until some time in the past $t^* < t$;
- ***filtering***, where the state at time t is estimated given all preceding and the current measurement.
- ***smoothing***, where the state at time t is estimated given a full set of measurements available until some subsequent time $T^* > t$.

Statistically optimal filtering and smoothing algorithms can be formulated from the Bayesian perspective.

2.5.1 Optimal state estimation

Consider a discrete time stochastic SSM in the form (2.3.8) with hidden state \mathbf{x}_t and observation vector \mathbf{y}_t that comply fully to the distributions

$$\mathbf{x}_t \sim p(\mathbf{x}_t \mid \mathbf{x}_{t-1}) \quad (2.5.1)$$

$$\mathbf{y}_t \sim p(\mathbf{y}_t \mid \mathbf{x}_t). \quad (2.5.2)$$

In the Bayesian framework, the purpose of filtering is to approximate the posterior distribution of $(\mathbf{x}_t$ given only preceding measurements. Filtering distributions are

computed recursively using the *prediction-correction* algorithm [98]. First, the predictive distribution of the hidden state \mathbf{x}_t is obtained in the following form

$$p(\mathbf{x}_t | \mathbf{y}_{0:t-1}) = \int p(\mathbf{x}_t | \mathbf{x}_{t-1})p(\mathbf{x}_{t-1} | \mathbf{y}_{0:t-1})d\mathbf{x}_{t-1} \quad (2.5.3)$$

generates the predicted distribution given the measurement sequence up to time $t-1$. When the new observation \mathbf{y}_t arrives at time t , it is incorporated at the correction step to produce the filtering (posterior) distribution using Bayes' rule (2.4.4) modified as follows

$$p(\mathbf{x}_t | \mathbf{y}_{0:t}) = \frac{p(\mathbf{y}_t | \mathbf{x}_t)p(\mathbf{x}_t | \mathbf{y}_{0:t-1})}{p(\mathbf{y}_t | \mathbf{y}_{0:t-1})}, \quad (2.5.4)$$

where the observation distribution is marginalised over the prior *pdf* of the hidden state

$$p(\mathbf{y}_t | \mathbf{y}_{0:t-1}) = \int p(\mathbf{y}_t | \mathbf{x}_t)p(\mathbf{x}_t | \mathbf{y}_{0:t-1})d\mathbf{x}_t. \quad (2.5.5)$$

In optimal smoothing, the probability density of the hidden state at a prior time t is "retrodicted" based on the full history of observations until some subsequent time $T^* > t$. The corresponding smoothing distributions can be computed in two forms:

1) Two-filter smoother [99]:

$$p(\mathbf{x}_t | \mathbf{y}_{1:T^*}) \propto p(\mathbf{x}_t | \mathbf{y}_{0:t})p(\mathbf{y}_{t+1:T^*} | \mathbf{x}_t), \quad (2.5.6)$$

where the first *pdf* on the right hand side obtained directly from the Bayesian filter, and the second term is computed via the backward-time filter. The applicability of this smoother rarely goes beyond the linear case. In more general cases calculation of additional probability densities is required to ensure the backward-time filter is normalisable [81].

2) Forward-backward smoother [100]:

$$p(\mathbf{x}_t | \mathbf{y}_{1:T^*}) \propto p(\mathbf{x}_t | \mathbf{y}_{0:t}) \int \frac{p(\mathbf{x}_{t+1} | \mathbf{x}_t)p(\mathbf{x}_{t+1} | \mathbf{y}_{1:T^*})}{p(\mathbf{x}_{t+1} | \mathbf{y}_{0:t})}d\mathbf{x}_{t+1}, \quad (2.5.7)$$

where the first term in the numerator is the filtering distribution of the state at time t , and the second term is the predicted distribution at time $t+1$ given by (2.5.3). While the smoothing algorithm proceeds backward in time starting at T^* , each recursion is implemented in forward time, making the implementation much more straightforward and easily applicable to nonlinear cases.

The formal equations presented in this subsection are generally computationally intractable and must be approximated numerically or evaluated for a specific distribution. The special case of optimal estimator for the linear Gaussian system is derived in the following subsection.

2.5.2 Filtering in linear Gaussian case

In the Bayesian framework, the model (2.3.8) is fully described by the following distributions

$$p(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; A_{t-1}\mathbf{x}_{t-1} + B_{t-1}\mathbf{u}_{t-1}, Q_{t-1}), \quad (2.5.8a)$$

$$p(\mathbf{y}_t | \mathbf{x}_t) = \mathcal{N}(\mathbf{y}_t; C_t\mathbf{x}_t, R_t), \quad (2.5.8b)$$

$$\pi(\mathbf{x}_0) = \mathcal{N}(\bar{\mathbf{x}}_0, \mathbf{P}_0). \quad (2.5.8c)$$

Given the model (2.5.8), filtering equations presented in the section above can be solved in the closed form and the resulting distributions will be Gaussian

$$p(\mathbf{x}_t | \mathbf{y}_{0:t-1}) = \mathcal{N}(\hat{\mathbf{x}}_{t|t-1}, \mathbf{P}_{t|t-1}), \quad (2.5.9a)$$

$$p(\mathbf{y}_t | \mathbf{y}_{0:t}) = \mathcal{N}(\tilde{\mathbf{y}}_{t|t}, \mathbf{S}_{t|t}), \quad (2.5.9b)$$

$$p(\mathbf{x}_t | \mathbf{y}_{0:t}) = \mathcal{N}(\hat{\mathbf{x}}_{t|t}, \mathbf{P}_{t|t}), \quad (2.5.9c)$$

where $\hat{\mathbf{x}}_{t|t-1}$ and $\hat{\mathbf{x}}_{t|t}$ denote the predicted and filtered estimates, respectively, and where $\tilde{\mathbf{y}}_{t|t}$ is the predicted measurement residual, also referred to as innovation. The resultant estimator is called the *Kalman filter (KF)* [101]. It is a widely used recursive state estimator that has found numerous applications in control [102], aerospace [103], and computational biology [104]. The historical importance of this algorithm is also notable, as the publication of Kalman's seminal paper coincided with the fundamental shift in the systems theory from input-output models to the state space representation [90].

Although originally the Kalman filter was presented as a LS approach, it can be derived as a point-wise Bayesian estimator $\hat{\mathbf{x}}_t$ with a quadratic loss function of the form

$$l(\mathbf{x}_t, \hat{\mathbf{x}}_t) = (\mathbf{x}_t - \hat{\mathbf{x}}_t)^\top L_t (\mathbf{x}_t - \hat{\mathbf{x}}_t), \quad (2.5.10)$$

where L_t is an arbitrary non-singular positive-definite symmetric matrix for all t [88]. Then the Bayesian risk is expressed as follows

$$R(\hat{\mathbf{x}}_t, p(\mathbf{x}_t)) = \int_{\mathbf{x}} R(\mathbf{x}_t, \hat{\mathbf{x}}_t) d\mathbf{x}_t = \int_{\mathbf{x}} (\mathbf{x}_t - \hat{\mathbf{x}}_t)^\top L_t (\mathbf{x}_t - \hat{\mathbf{x}}_t) p(\mathbf{x}_t | \mathbf{y}_{0:t-1}) d\mathbf{x}_t. \quad (2.5.11)$$

Under Gaussian assumptions, the estimate can be calculated in the closed form by taking the derivative of the risk with respect to \mathbf{x}_t and setting it equal to zero

$$\frac{\partial R(\hat{\mathbf{x}}_t, p(\mathbf{x}_t))}{\partial \mathbf{x}_t} = 2L_t \int_{\mathbf{x}} (\mathbf{x}_t - \hat{\mathbf{x}}_t) p(\mathbf{x}_t | \mathbf{y}_{0:t-1}) d\mathbf{x}_t = 0. \quad (2.5.12)$$

Solution of (2.5.12) is the expectation of \mathbf{x}_t given the preceding set of measurements $\mathbf{y}_{0:t-1}$ and is independent of the matrix L_t :

$$\hat{\mathbf{x}}_{t|t-1} = \int_{\mathbf{x}} \mathbf{x}_t p(\mathbf{x}_t | \mathbf{y}_{0:t-1}) d\mathbf{x} = \mathbb{E}[\mathbf{x}_t | \mathbf{y}_{0:t-1}]. \quad (2.5.13)$$

The corresponding covariance matrix is defined as

$$\begin{aligned} \mathbf{P}_{t|t-1} &= \int_{\mathbf{x}} (\mathbf{x}_t - \hat{\mathbf{x}}_t)(\mathbf{x}_t - \hat{\mathbf{x}}_t)^\top p(\mathbf{x}_t | \mathbf{y}_{0:t-1}) d\mathbf{x} = \\ &= \mathbb{E} \left[(\mathbf{x}_t - \hat{\mathbf{x}}_t)(\mathbf{x}_t - \hat{\mathbf{x}}_t)^\top | \mathbf{y}_{0:t-1} \right]. \end{aligned} \quad (2.5.14)$$

Expectations correspond to the predictive *pdf* in Bayesian filter and are called **prior state estimate** and **prior state covariance**. Under linear-Gaussian assumptions (2.3.1-2.3.2) they can be calculates follows

$$\hat{\mathbf{x}}_{t|t-1} = A_t \hat{\mathbf{x}}_{t-1|t-1} + B_t \mathbf{u}_t, \quad (2.5.15)$$

$$\mathbf{P}_{t|t-1} = A_t \mathbf{P}_{t-1|t-1} A_t^\top + Q_t. \quad (2.5.16)$$

The likelihood of \mathbf{x}_t in (2.5.4) is Gaussian by definition (2.5.8). What remains is to evaluate the denominator of (2.5.4) which will also be Gaussian under the given assumptions

$$p(\mathbf{y}_t | \mathbf{y}_{0:t}) = \mathcal{N}(\tilde{\mathbf{y}}_t, \mathbf{S}_{t|t}), \quad (2.5.17)$$

where $\tilde{\mathbf{y}}_t$ denotes the residual between the predicted observation and the available observation, and where $\mathbf{S}_{t|t}$ is the corresponding residual covariance:

$$\tilde{\mathbf{y}}_t = \mathbf{y}_t - C_t \mathbb{E}[\mathbf{x}_t | \mathbf{y}_{0:t-1}] = \mathbf{y}_t - C_t \hat{\mathbf{x}}_{t|t-1}, \quad (2.5.18)$$

$$\mathbf{S}_{t|t} = \mathbb{E} \left[(\tilde{\mathbf{y}}_t - C_t \hat{\mathbf{x}}_{t|t-1})(\tilde{\mathbf{y}}_t - C_t \hat{\mathbf{x}}_{t|t-1})^\top | \mathbf{y}_{0:t-1} \right] = C_t \mathbf{P}_{t|t-1} C_t^\top + R_t. \quad (2.5.19)$$

Substituting all terms in (2.5.4) with (2.5.18) and (2.5.19) renders

$$p(\mathbf{x}_t | \mathbf{y}_{0:t}) = \frac{|C_t \mathbf{P}_{t|t-1} C_t^\top + R_t|^{\frac{1}{2}}}{|2\pi \mathbf{P}_{t|t-1} R_t|^{\frac{1}{2}}} \exp^{-\frac{1}{2}\zeta}, \quad (2.5.20)$$

where

$$\begin{aligned} \zeta &= (\mathbf{x}_t - \hat{\mathbf{x}}_{t|t-1})^\top (\mathbf{P}_{t|t-1})^{-1} (\mathbf{x}_t - \hat{\mathbf{x}}_{t|t-1}) + \\ &+ (\mathbf{y}_t - C_t \mathbf{x}_t)^\top (R_t)^{-1} (\mathbf{y}_t - C_t \mathbf{x}_t) - \\ &- (\mathbf{y}_t - C_t \hat{\mathbf{x}}_{t|t-1})^\top \left(C_t \mathbf{P}_{t|t-1} C_t^\top + R_t \right)^{-1} (\mathbf{y}_t - C_t \hat{\mathbf{x}}_{t|t-1}). \end{aligned} \quad (2.5.21)$$

The exponent can be represented in the quadratic form

$$\zeta = (\mathbf{x}_t - \hat{\mathbf{x}}_{t|t})^\top (\mathbf{P}_{t|t})^{-1} (\mathbf{x}_t - \hat{\mathbf{x}}_{t|t}). \quad (2.5.22)$$

Simple matrix manipulations render the expressions for the **posterior state estimate** and the **posterior covariance** in the following form:

$$\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + K_t \tilde{\mathbf{y}}_t, \quad (2.5.23)$$

$$\mathbf{P}_{t|t} = \mathbf{P}_{t|t-1} - K_t \mathbf{S}_{t|t} K_t^\top. \quad (2.5.24)$$

where K_t is the optimal filter (Kalman) gain given by

$$K_t = \mathbf{P}_{t|t-1} C_t^\top (\mathbf{S}_{t|t})^{-1}. \quad (2.5.25)$$

An intuitive interpretation of the Kalman gain is offered in [72]: if viewed as a scalar, K_t is “proportional” to the variance of the prediction, and “inverse proportional” to the measurement variance. As such, a small gain yields a response to inaccurate observations, while a large gain corresponds to inaccurate predictions. The full filtering recursion is summarised in the Algorithm 2.1.

Algorithm 2.1 Kalman filter

Input: Measurement vector $\mathbf{y}_{1:T}$; model matrices; state noise variance, Q_t ; measurement noise variance, R_t ; initial state estimate, $\hat{\mathbf{x}}_0$; and initial covariance, \mathbf{P}_0 .

Output: Sequence of filtered state estimates, $\{\hat{\mathbf{x}}_{t|t}\}_{t=0}^T$; filtered covariances, $\{\mathbf{P}_{t|t}\}_{t=0}^T$.

- 1: **for** $t \leftarrow 1, T$ **do**
 - 2: Compute prior state estimate $\hat{\mathbf{x}}_{t|t-1} = A_t \hat{\mathbf{x}}_{t-1|t-1} + B_t \mathbf{u}_t$;
 - 3: Compute prior covariance $\mathbf{P}_{t|t-1} = A_t \mathbf{P}_{t-1|t-1} A_t^\top + Q_t$;
 - 4: Compute residual $\tilde{\mathbf{y}}_t = \mathbf{y}_t - C_t \hat{\mathbf{x}}_{t|t-1}$;
 - 5: Compute residual covariance $\mathbf{S}_{t|t} = C_t \mathbf{P}_{t|t-1} C_t^\top + R_t$;
 - 6: Compute Kalman gain $K_t = \mathbf{P}_{t|t-1} C_t^\top (\mathbf{S}_{t|t})^{-1}$;
 - 7: Compute posterior state estimate $\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + K_t \tilde{\mathbf{y}}_t$;
 - 8: Compute posterior covariance $\mathbf{P}_{t|t} = \mathbf{P}_{t|t-1} - K_t \mathbf{S}_{t|t} K_t^\top$;
 - 9: **end for**
-

Remark 2.5.1: *All of the state estimates are implicitly conditioned on the model structure of the dynamical system. In Kalman filter the residual $\tilde{\mathbf{y}}_t$ quantifies goodness of fit of the model to the observations. Thus the pdf of the residual (2.5.17) can be viewed as a **marginal likelihood** of the system model M :*

$$\mathcal{L}(M) \triangleq p(\mathbf{y}_t | M, \mathbf{y}_{0:t}) = \mathcal{N}(\tilde{\mathbf{y}}_t, \mathbf{S}_{t|t}). \quad (2.5.26)$$

This interpretation of the model likelihood is utilised in further sections to discriminate between various candidate models of the dynamical process.

2.5.3 Smoothing in linear Gaussian case

Similarly with the Kalman filter, there exist closed form smoothing solutions for the model (2.3.8) both in two-filter and forward-backward forms. The resultant smoothing distribution is normal

$$p(\mathbf{x}_t | \mathbf{y}_{1:T^*}) = \mathcal{N}(\hat{\mathbf{x}}_{t|T^*}, \mathbf{P}_{t|T^*}), \quad (2.5.27)$$

For the sake of brevity, this subsection only includes the resultant linear equations.

The linear case of the two-filter smoother consists of two Kalman filters [70], one running in direct time, while the other is applied to the reversed observation sequence $\mathbf{y}_{0:T^*}^-$ such that

$$\mathbf{y}_0^-, \dots, \mathbf{y}_{T^*}^- = \mathbf{y}_{T^*}, \dots, \mathbf{y}_0. \quad (2.5.28)$$

Given that A_t is a non-singular matrix, filtered estimates can be obtained for a system with reverse time dynamics defined as

$$\mathbf{x}_{t+1}^- = A_t^- \mathbf{x}_t^- + B_t^- \mathbf{u}_t^- + G_t^- \mathbf{w}_t^-, \quad (2.5.29a)$$

$$\mathbf{y}_t^- = C_t^- \mathbf{x}_t^- + H_t^- \mathbf{v}_t^-, \quad (2.5.29b)$$

where \mathbf{w}_t^- and \mathbf{v}_t^- are backward forced noises [105]. The justification for using this form for the reverse time system is provided in [70].

Then the smoothed state estimate is a weighted sum of two filtering estimates

$$\hat{\mathbf{x}}_{t|T^*} = \mathbf{P}_{t|T^*} \left((\mathbf{P}_{t|t})^{-1} \hat{\mathbf{x}}_{t|t} + (\mathbf{P}_{t|t}^-)^{-1} \hat{\mathbf{x}}_{t|t}^- \right), \quad (2.5.30)$$

where the smoothed covariance $\mathbf{P}_{t|T^*}$ is given by

$$\mathbf{P}_{t|T^*} = \mathbf{P}_{t|T^*} = \left((\mathbf{P}_{t|t})^{-1} + (\mathbf{P}_{t|t}^-)^{-1} \right)^{-1}. \quad (2.5.31)$$

The complete recursion is summaries in Algorithm 2.2. The limitation of this algorithm is evident as a it requires non-singularity of the state transition matrix and of all covariances. Information filters can be used to avoid taking the inverses [102].

Algorithm 2.2 Two-filter smoother

Input: Measurement vector $\mathbf{y}_{1:T^*}$; model matrices, state noise variance, Q_w , and measurement noise variance, R_v , fro forward and reverse time systems; initial state estimates, $\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_0^-$; and initial covariances, $\mathbf{P}_0, \mathbf{P}_0^-$.

Output: Sequence of smoothed state estimates, $\{\hat{\mathbf{x}}_{t|T^*}\}_{t=0}^{T^*}$; smoothed covariances, $\{\mathbf{P}_{t|T^*}\}_{t=0}^{T^*}$.

- 1: Execute Kalman filter for the forward-time system (2.3.8);
 - 2: Execute Kalman filter for the backward-time system (2.5.29);
 - 3: **for** $t \leftarrow 0, T$ **do**
 - 4: Merge covariances $\mathbf{P}_{t|T^*} = \left((\mathbf{P}_{t|t})^{-1} + (\mathbf{P}_{t|t}^-)^{-1} \right)^{-1}$;
 - 5: Merge estimates $\hat{\mathbf{x}}_{t|T^*} = \mathbf{P}_{t|T^*} \left((\mathbf{P}_{t|t})^{-1} \hat{\mathbf{x}}_{t|t} + (\mathbf{P}_{t|t}^-)^{-1} \hat{\mathbf{x}}_{t|t}^- \right)$;
 - 6: **end for**
-

The linear Gaussian case of the forward-backward algorithm is the Rauch-Tung-Striebel (RTS) smoother [106]. The forward pass of the smoother is equivalent to the regular KF, while the backward recursion is implemented in backward time as follows. The smoother is initialised at the time of final observation T^* :

$$\hat{\mathbf{x}}_{T^*} = \hat{\mathbf{x}}_{T^*|T^*}, \quad (2.5.32a)$$

$$\mathbf{P}_{T^*} = \mathbf{P}_{T^*|T^*}. \quad (2.5.32b)$$

Then, at each time $t < T^*$, the prediction is made about the future state $t+1$ in exact form of the Kalman filter prediction step:

$$\hat{\mathbf{x}}_{t+1|t} = A_t \hat{\mathbf{x}}_{t|t} + B_t \mathbf{u}_t, \quad (2.5.33)$$

$$\mathbf{P}_{t+1|t} = A_t \mathbf{P}_{t|t} A_t^\top + Q_t. \quad (2.5.34)$$

The smoothing step accounts for the discrepancy between the prediction of the future state and its smoothed estimate:

$$\hat{\mathbf{x}}_{t|T^*} = \hat{\mathbf{x}}_{t|t} + J_t \left(\hat{\mathbf{x}}_{t+1|T^*}^s - \hat{\mathbf{x}}_{t+1|t} \right), \quad (2.5.35)$$

$$\mathbf{P}_{t|T^*} = \mathbf{P}_{t|t} + J_t \left(\mathbf{P}_{t+1|T^*} - \mathbf{P}_{t+1|t} \right) J_t^\top, \quad (2.5.36)$$

where J_t is the smoothing gain defined as

$$J_t = \mathbf{P}_{t|t} A_t^\top \left(\mathbf{P}_{t+1|t} \right)^{-1}. \quad (2.5.37)$$

The resultant estimation recursion is summarised in the Algorithm 2.3.

Algorithm 2.3 Backward recursion of the RTS smoother

Input: Sequence of filtered state estimates, $\{\hat{\mathbf{x}}_{t|t}\}_{t=0}^{T^*}$; filtered covariances, $\{\mathbf{P}_{t|t}\}_{t=0}^{T^*}$; state transition matrices, A_t .

Output: Sequence of smoothed state estimates, $\{\hat{\mathbf{x}}_{t|T^*}\}_{t=0}^{T^*}$; smoothed covariances, $\{\mathbf{P}_{t|T^*}\}_{t=0}^{T^*}$.

- 1: *Initialise:* $\hat{\mathbf{x}}_{T^*} = \hat{\mathbf{x}}_{T^*|T^*}$, $\mathbf{P}_{T^*} = \mathbf{P}_{T^*|T^*}$;
 - 2: **for** $t \leftarrow T^* - 1, 0$ **do**
 - 3: Calculate the smoothing gain $J_t = \mathbf{P}_{t|t} A_t^\top \left(\mathbf{P}_{t+1|t} \right)^{-1}$;
 - 4: Smooth state estimate $\hat{\mathbf{x}}_{t|T^*} = \hat{\mathbf{x}}_{t|t} + J_t \left(\hat{\mathbf{x}}_{t+1|T^*} - \hat{\mathbf{x}}_{t+1|t} \right)$;
 - 5: Smooth covariance $\mathbf{P}_{t|T^*} = \mathbf{P}_{t|t} + J_t \left(\mathbf{P}_{t+1|T^*} - \mathbf{P}_{t+1|t} \right) J_t^\top$;
 - 6: **end for**
-

Depending on how many future observations are available (where is T^* located with respect to t), the smoothers of both types can be further classified as follows:

- **Fixed-point smoothing** estimates the previous state given the newest observation, *i.e.* $T^* = t + 1$.
- **Fixed-lag smoothing** estimated the state with a fixed time delay τ from the newest observation, *i.e.* $T^* = t + \tau$.
- **Fixed-interval smoothing** estimates the states in presence of the full observation set, *i.e.* $T^* = T$, where $t = 0, 1, \dots, T$.

Figure 2.6 illustrates the span of measurement sets utilised by each type of smoother.

Note that the fixed-interval smoother outputs the expectation and covariance of the system state given the full observation sequence:

$$\hat{\mathbf{x}}_{t|T} = \mathbb{E} [\mathbf{x}_t | \mathbf{y}_{0:T}] = \mathbb{E} [\mathbf{x}_t | \mathcal{Y}], \quad (2.5.38)$$

$$\mathbf{P}_{t|T} = \mathbb{E} \left[(\mathbf{x}_t - \hat{\mathbf{x}}_{t|T})(\mathbf{x}_t - \hat{\mathbf{x}}_{t|T})^\top | \mathcal{Y} \right], \quad (2.5.39)$$

which can be combined to produce the following

$$\mathbf{P}_{t|T} + \hat{\mathbf{x}}_{t|T}(\hat{\mathbf{x}}_{t|T})^\top = \mathbb{E} \left[\mathbf{x}_t \mathbf{x}_t^\top | \mathbf{y}_{0:T} \right] = \mathbb{E} \left[\mathbf{x}_t \mathbf{x}_t^\top | \mathcal{Y} \right], \quad (2.5.40)$$

thus providing the first and second expectation in (2.4.42) necessary for performing the E-step in the §2.4.6. The third expectation (2.4.42c) is given by

$$\mathbf{P}_{t,t-1|T} + \hat{\mathbf{x}}_{t|T}(\hat{\mathbf{x}}_{t-1|T})^\top = \mathbb{E} \left[\mathbf{x}_t \mathbf{x}_{t-1}^\top | \mathbf{y}_{0:T} \right] = \mathbb{E} \left[\mathbf{x}_t \mathbf{x}_{t-1}^\top | \mathcal{Y} \right], \quad (2.5.41)$$

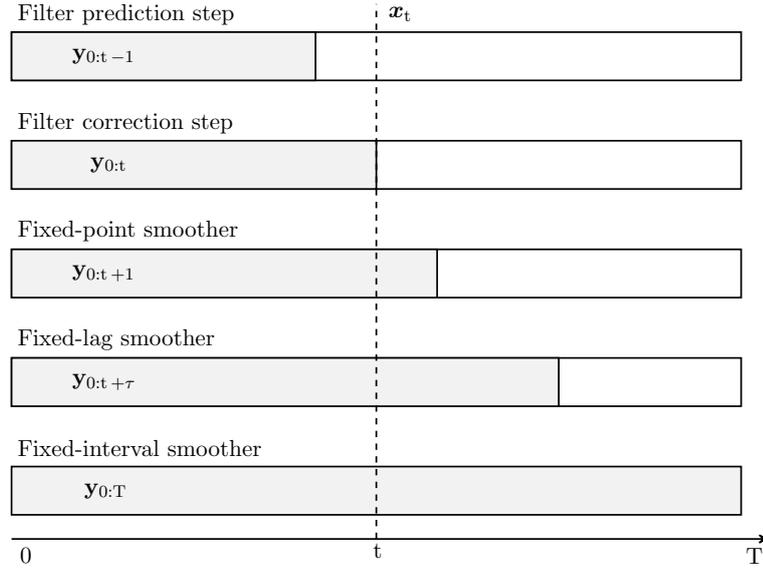


Figure 2.6: Observation sets utilised in filtering and smoothing algorithms.

where $\mathbf{P}_{t,t-1|T}$ denotes the cross-covariance of two subsequent states [107]. It is computed in backward-time recursion of the RTS smoother as shown in the Algorithm 2.4.

Algorithm 2.4 Cross-covariance computation

Input: Sequence of filtered covariances, $\{\mathbf{P}_{t|t}\}_{t=0}^T$; sequence of smoothing gains, $\{J_t\}_{t=0}^T$; system matrices at final time, A_T, C_T , Kalman gain at final time, K_T .

Output: Sequence of cross-covariances, $\{\mathbf{P}_{t,t-1|T}\}_{t=1}^T$.

- 1: *Initialise:* $\mathbf{P}_{T,T-1|T} = (\mathbb{I} - K_T C_T) A_T \mathbf{P}_{T-1|T-1}$;
 - 2: **for** $t \leftarrow T-1, 1$ **do**
 - 3: $\mathbf{P}_{t,t-1|T} = \mathbf{P}_{t|t} J_{t-1}^T + J_{t-1} (\mathbf{P}_{t+1,t|T}) J_{t-1}^T$;
 - 4: **end for**
-

In consequence, the forward-backward procedure that constitutes the expectation step of the EM algorithm for linear Gaussian SSMs is equivalent to the combination of the Kalman filter and the RTS smoother. It provides all necessary quantities for the evaluation of the expected log likelihood function in the EM algorithm in §2.4.6, ensuring a closed form solution of the maximisation step for linear Gaussian systems.

2.5.4 Suboptimal state estimation

In practical applications, the model of interest is often non-linear and the Kalman filter in its original form is not applicable. Derivation an optimal non-linear filter requires complete knowledge of the conditional *pdf* $p(\mathbf{x}_t | \mathbf{x}_{t-1})$ which is hardly achievable, thus multiple suboptimal filtering solutions have been sought. This subsection briefly discusses the most prominent extensions of Kalman algorithm in

applications to non-linear systems of the general form (2.3.1).

The *extended Kalman filter (EKF)* was developed in application to navigation and guidance systems. It is based on assuming the knowledge of the approximate solution and describing the deviation from the actual model by linear equations. This is done via the first-order Taylor series expansion of the given non-linear model functions f and h [108]. The resultant Jacobians $\frac{\partial f}{\partial \mathbf{x}}|_{\mathbf{x}=\hat{\mathbf{x}}}$, $\frac{\partial h}{\partial \mathbf{x}}|_{\mathbf{x}=\hat{\mathbf{x}}}$ evaluated at the current estimate are then utilised in the classical filter recursion. In case of non-additive noise, extra Jacobians $\frac{\partial f}{\partial \mathbf{w}}|_{\mathbf{w}=0}$, $\frac{\partial h}{\partial \mathbf{v}}|_{\mathbf{v}=0}$ are evaluated and used to calculate covariances [81]. Simplicity of the EKF makes it attractive in engineering applications, but local linearisation restricts applicability of this filter only to dynamical models with differentiable functions f and h .

The *unscented Kalman filter (UKF)* [109], [110] builds on a rather intuitive notion that it is easier to approximate a Gaussian distribution than an arbitrary non-linear SSM. In the filter, the distribution $p(\mathbf{x}_t | \mathbf{y}_{t-1})$ is approximated by deterministically choosing a set of sigma points that capture its mean and covariance. These points are then propagated through the non-linear functions f and h , to produce a new distribution with updated mean and covariance. The UKF has repeatedly outperformed the EKF as it can capture the moments of the distribution after a non-linear transformation [111], neither does it require differentiability of the model function, although it is computationally demanding. The complementing unscented RTS smoother is derived in [112].

Aside from EKF and UKF that are the most commonly used for non-linear state estimation, there exist a number of filtering algorithms, such as statistically linearised filter [113], cubature Kalman filter [114], Gauss-Hermite Kalman filter [115], each with their merits and varying applicability. However, all suboptimal solutions rely on some level of approximation, be it linearisation (EKF), fitting to a curve (cubature KF), or deterministic moment approximation (UKF), which may lead to divergence of an estimator in cases of significant non-linearity or in situations where the Gaussian approximation is not appropriate. Particle filters are the sequential Monte Carlo (SMC) methods that numerically approximate Bayesian filtering equations [116]. Unlike the UKF that relies on the deterministic selection of sigma-points, particle filters utilise random sampling from the posterior distribution. Particle filters do not require that distribution to be Gaussian and therefore are applicable to a wider class of models, though this flexibility comes at a price of high computational costs and bad scalability for high dimensional problems. Non-Gaussian situations are not considered in the following chapters.

2.5.5 State estimation of jump Markov systems

Another useful category of dynamical models are the *hybrid systems* that exhibit both continuous and discrete dynamic behaviour. A generic hybrid system is described by a finite set of discrete states, or *modes*, $\mathbb{M} = \{M^j\}_{j=1}^{N_m}$, each associated with a different SSM governing the evolution of the continuous-valued state \mathbf{x}_t .

Remark 2.5.2: *In the hybrid systems context, a continuous-valued state at some discrete time $t \in \{0, 1, 2, 3, \dots, T\}$ is defined as $\mathbf{x}_t \in \mathbb{X}_t$, where \mathbb{X}_t is an uncountable set.*

For example, a discrete-time hybrid system composed of linear SSMs is given by

$$\mathbf{x}_t = A(\mathbf{m}_t)\mathbf{x}_{t-1} + B(\mathbf{m}_t)\mathbf{u}_{t-1} + G(\mathbf{m}_t)\mathbf{w}_{t-1} \quad (2.5.42a)$$

$$\mathbf{y}_t = C(\mathbf{m}_t)\mathbf{x}_t + D(\mathbf{m}_t)\mathbf{u}_t + H(\mathbf{m}_t)\mathbf{v}_t, \quad (2.5.42b)$$

where $\mathbf{w}_t \sim \mathcal{N}(0, Q(\mathbf{m}_t))$, $\mathbf{v}_t \sim \mathcal{N}(0, R(\mathbf{m}_t))$, and where system matrices correspond to the current mode in effect $\mathbf{m}_t \in \mathbb{M}$. The mode switching process is usually assumed to be left-continuous, *i.e.* the mode \mathbf{m}_{t+1} comes into effect at time t^+ . Causal relationship between modes, states, and observations of a hybrid system (2.5.42) is demonstrated in Figure 2.7. Alternative structures of hybrid systems are examined in [117]. The systems where mode also depends on the continuous-valued state are not considered here.

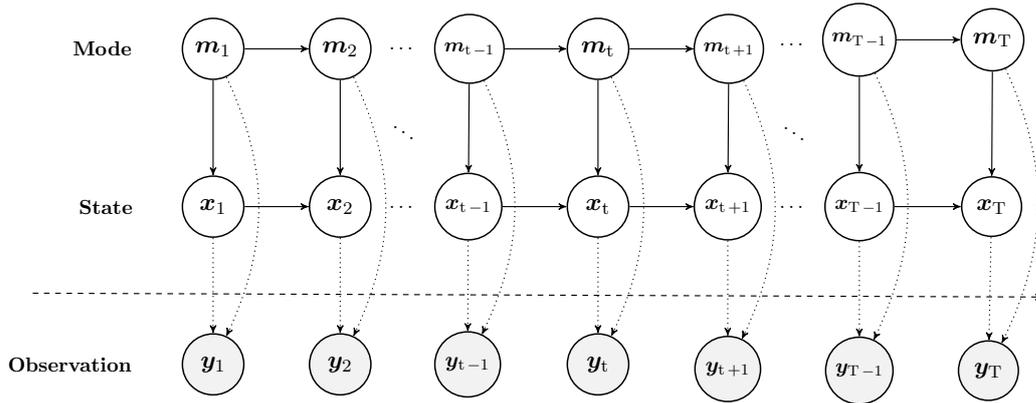


Figure 2.7: A directed acyclic graph of a generic left-continuous hybrid model. The discrete state (mode) \mathbf{m}_t affects the state and the observation at time t . Both discrete and continuous-valued states are hidden.

Jump Markov systems (JMS) constitute a class of hybrid systems in which the evolution of the discrete state \mathbf{m}_t is governed by a finite state Markov chain [72], [118] described by the set of initial probabilities $\Pi = \{\pi_j \geq 0\}_{j=1}^{N_M}$ and a transition

probability matrix

$$\Phi = \{\phi_{jl} \geq 0\}_{j,l=1}^{N_M} \quad (2.5.43)$$

where each element denotes an individual transition probability

$$\phi_{lj} \triangleq P(\mathbf{m}_t = M^l \mid \mathbf{m}_{t-1} = M^j). \quad (2.5.44)$$

The **multiple model (MM)** approach to hybrid system estimation involves running a bank of mode-matched filters [72]. At each time t the exact posterior *pdf* of the continuous-valued state comes as a mixture of mode-conditioned posteriors produced by model-matched filters

$$p(\mathbf{x}_t \mid \mathbf{y}_{0:t}) = \sum_{j=1}^{N_m} \mu_t^j p(\mathbf{x}_t \mid \mathbf{m}_t = M^j, \mathbf{y}_{0:t}), \quad (2.5.45)$$

where the mixing coefficients are given by the posterior mode probabilities

$$\mu_t^j \triangleq P(\mathbf{m}_t = M^j \mid \mathbf{y}_{0:t}); \quad (2.5.46)$$

and where, under assumptions 2.3.2-2.3.3, the filtered posterior is Gaussian

$$p(\mathbf{x}_t \mid \mathbf{m}_t = M^j, \mathbf{y}_{0:t}) = \mathcal{N}(\hat{\mathbf{x}}_{t|t}^j, \mathbf{P}_{t|t}^j). \quad (2.5.47)$$

The mode probability at each time μ_t^j is directly linked to the model likelihood $\mathcal{L}(M^j)$ that is produced by the filter as described in Remark 2.5.1.

The exact MM estimation of jumping systems proves impractical, as keeping exhaustive history of each model leads to exponentially increasing computational complexity (see Figure 2.8). Multiple methods that have been developed to make JMS estimation feasible rely on either pruning or merging techniques [119]. Pruning is an *ad hoc* solution that cuts off unlikely mode histories at each time. While it is an efficient strategy in change point detection, pruning algorithms are hardly applicable in situations where the true system differs from all of the hypothesised modes. Contrarily, merging achieves the tractability by replacing the mixture (2.5.45) with a single Gaussian that would summarise the "combined" history of all possible modes:

$$p(\mathbf{x}_t \mid \mathbf{y}_{0:t}) \approx \mathcal{N}(\hat{\mathbf{x}}_{t|t}, \mathbf{P}_{t|t}), \quad (2.5.48)$$

where the combined state estimate and its covariance are obtained via moment matching:

$$\hat{\mathbf{x}}_{t|t} = \sum_{j=1}^{N_m} \mu_{t|t}^j \hat{\mathbf{x}}_{t|t}^j \quad (2.5.49a)$$

$$\mathbf{P}_{t|t} = \sum_{j=1}^{N_m} \mu_{t|t}^j \left\{ \mathbf{P}_{t|t}^j + (\hat{\mathbf{x}}_{t|t}^j - \hat{\mathbf{x}}_{t|t})(\hat{\mathbf{x}}_{t|t}^j - \hat{\mathbf{x}}_{t|t})^\top \right\}. \quad (2.5.49b)$$

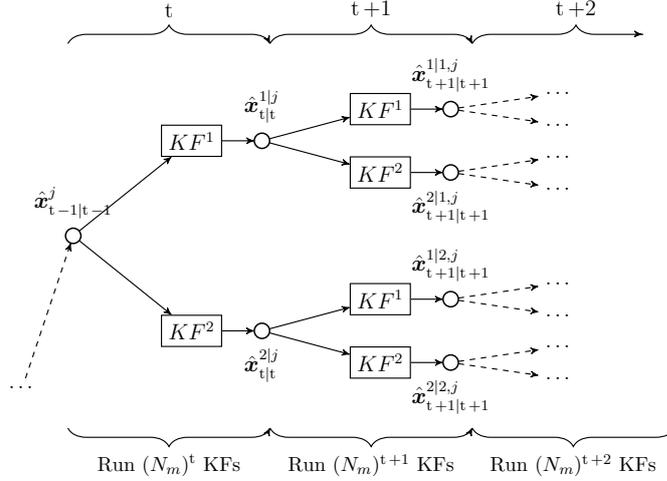
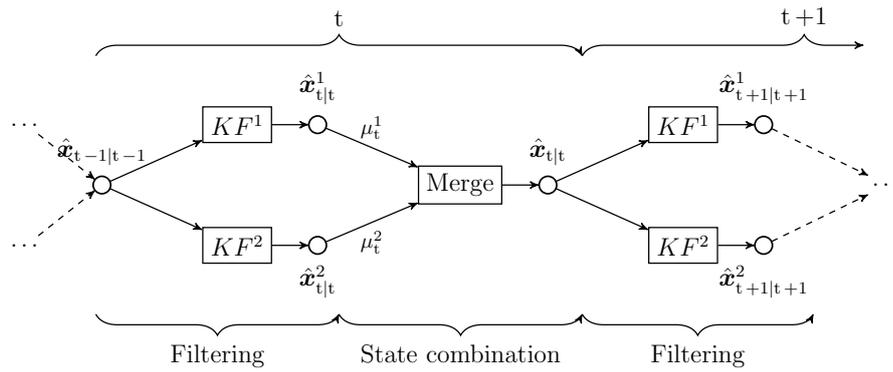


Figure 2.8: The number of model-matched Kalman filters (KF^j) required for the exact estimation of the jump Markov system grows exponentially.

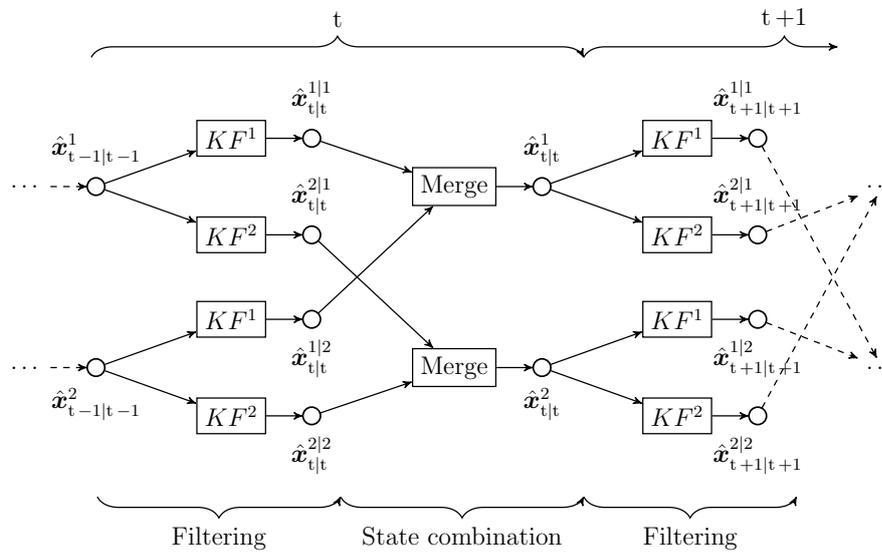
Generalised pseudo-Bayesian (GPB) algorithms are a class of algorithms that utilise the merging procedure to combine estimates that have different history at previous n sampling periods, where n is defined by the order of the algorithm. The first order GPB (GPB1) stores mode-conditioned estimates from the last sampling period only. The second order algorithm (GPB2) stores all previous estimates and all their possible histories, *i.e.* all modes they could have arrived from. While the GPB2 clearly outperforms GPB1 in terms of accuracy, it requires running $(N_m)^2$ filters at each iteration [119]. The recursions of GPB1 and GPB2 estimation of a JMS with two modes are illustrated in Figures 2.9a and 2.9b, respectively.

The most preferred method for JMS state estimation is the **interacting multiple model (IMM)** algorithm [120], [121]. It is principally similar to the GPB2, but has one crucial advantage. Instead of running $(N_m)^2$ model-matched filters for all possible posteriors, it first produces the mixed posteriors using the updated mode probabilities as mixing weights. This procedure is called the interaction step (see steps 3-4 of Algorithm 2.5). As demonstrated in Figure 2.9c, the IMM framework employs the merging procedure twice: once during the interaction step to produce mixed initial conditions for the model-matched filters, and once after the filtering step to produce the combined estimate. The interaction step ensures that the algorithm utilises information from two subsequent sampling periods while running only N_m filters, thus providing a compromise between the under-performing GPB1 and the computationally expensive GPB2.

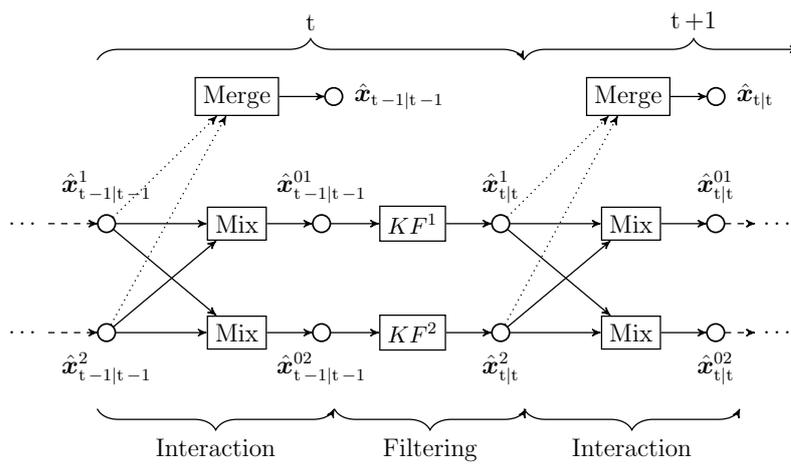
Several smoothing counterparts to forward IMM filters have been developed to improve the estimation accuracy in tracking applications. The two-filter IMM approach [122], [123] combines all possible estimates of forward and backward-time models, and is therefore restricted to the systems where time-reversal is possible.



(a) GPB1



(b) GPB2



(c) IMM

Figure 2.9: Multiple model algorithms for a jumping system with two modes.

Algorithm 2.5 Forward time recursion of the IMM filter

Input: Measurements sequence, \mathbf{y} ; hypothesised models, \mathbb{M} ; initial model probabilities, Π ; mode transition probabilities, Φ .

Output: Sequence of mode-conditioned state estimates, $\hat{\mathbf{x}}_{0:T}^j$, $j = 1, \dots, N_m$; Sequence of merged state estimates, $\hat{\mathbf{x}}_{0:T}$, and their covariances, $P_{0:T}$; estimated model probabilities, $\hat{\mu}_{0:T}$.

1: Initialise the algorithm with $\mu_{0|0}^j, \hat{\mathbf{x}}_{0|0}^j, \mathbf{P}_{0|0}^j$ $j = 1, \dots, N_m$;

2: **for** $t \leftarrow 1, T$ **do**

3: **for** $j \leftarrow 1, N_m$ **do**

4: Calculate mixing probabilities $\mu_{t-1|t-1}^{0j} = \frac{1}{c_j} \phi_{ij} \mu_{t-1|t-1}^j$, where $c_j = \sum_{i=1}^{N_m} \phi_{ij} \mu_{t-1|t-1}^i$;

5: Compute mixed initial conditions:

$$\hat{\mathbf{x}}_{t-1|t-1}^{0j} = \sum_{j=1}^{N_m} \mu_{t-1|t-1}^{0j} \hat{\mathbf{x}}_{t-1|t-1}^j$$

$$\mathbf{P}_{t-1|t-1}^{0j} = \sum_{j=1}^{N_m} \mu_{t-1|t-1}^{0j} \left\{ \mathbf{P}_{t-1|t-1}^{0j} + (\hat{\mathbf{x}}_{t-1|t-1}^j - \hat{\mathbf{x}}_{t-1|t-1}^{0j})(\hat{\mathbf{x}}_{t-1|t-1}^j - \hat{\mathbf{x}}_{t-1|t-1}^{0j})^\top \right\}$$

6: Run KF matched with the mode M^j to produce $\hat{\mathbf{x}}_{t|t}^j, \mathbf{P}_{t|t}^j$;

7: Compute mode likelihood $\mathcal{L}_{t|t}^j = p(\mathbf{y}_t | M^j, \mathbf{y}_{0:t})$;

8: **end for**

9: Update mode probabilities $\mu_{t|t}^j = \frac{1}{c} \mathcal{L}_{t|t}^j c_j$, where $c = \sum_{j=1}^{N_m} \mathcal{L}_{t|t}^j c_j$;

10: Merge mode-conditioned states and covariances:

$$\hat{\mathbf{x}}_{t|t} = \sum_{j=1}^{N_m} \mu_{t|t}^j \hat{\mathbf{x}}_{t|t}^j$$

$$\mathbf{P}_{t|t} = \sum_{j=1}^{N_m} \mu_{t|t}^j \left\{ \mathbf{P}_{t|t}^j + (\hat{\mathbf{x}}_{t|t}^j - \hat{\mathbf{x}}_{t|t})(\hat{\mathbf{x}}_{t|t}^j - \hat{\mathbf{x}}_{t|t})^\top \right\};$$

11: **end for**

The fixed-lag IMM smoother proposed in [124] utilises state augmentation to provide filtered and smoothed estimates simultaneously. However, the method builds on the assumption that the mode does not change within the lag period. The first RTS-type MM algorithm presented in [125] uses Bayesian retrodiction during the backward recursion and utilises $(N_m)^2$ smoothers on each iteration, which renders it inferior to the RTS-type IMM algorithm recently proposed in [126]. The RTS IMM estimator first proceeds in forward time using IMM filter, then utilises mode-conditioned filtered estimates during the backward-time recursion. The merged smoothed estimate produced by the algorithm on each iteration can be computed separately and has no direct effect on the smoothing process.

Although the algorithms discussed in this subsection build on linear-Gaussian assumptions for candidate models, the multiple model framework is not limited to strictly linear cases. Modular structure of IMM is particularly accommodating to non-linear model-matched filters, although such modification adds an additional level of approximation. A range of non-linear multiple model methods will be discussed in the following chapters in application to a specific model structure.

3 Inference of the hidden chemoattractant field from observed cell migration patterns

This chapter introduces the problem of estimating a spatially-varying chemoattractant (ChA) concentration field from observed migration patterns of cell populations. The analysis is motivated by the shortcomings of novel *in vivo* microscopy techniques that do not have the capacity for quantifying the ChA environment. The problem, once introduced, is formulated within a state space framework that incorporates a parametrised model of the hidden environment in the cell dynamics. An approximate Expectation-Maximisation (EM) solution is proposed for the joint state-parameter estimation and applied to several *in vivo* datasets obtained from transgenic zebrafish. Selected results of this chapter have been presented in [127].

3.1 Background

Neutrophilic chemotaxis is an essential part of immune system response to external threats. The recruitment phase of the inflammation is defined by rapid cell migration towards the damaged tissue driven by the extracellular chemoattractant concentration, here referred to as the environment. Therefore the recruitment process can be manipulated in two ways: by either changing the ChA environment, or changing the way cells respond to it. However, in order to create robust cost-effective therapeutic strategies the researchers must have a clear understanding on the way the immune cells perceive the chemoattractant.

Rapid progress of the experimental procedures, and development of *in vivo* assays in particular, has informed numerous theoretical studies of neutrophil dynamics. At the same time, new microscopy methods pose a novel problem for the theorists. Whereas in artificial circumstances the concentration of the attractant is tightly regulated and is available for measurement, it is impossible to directly observe it in the living organism (see Figure 3.1). Processing of *in vivo* data presents a new challenge in quantifying the unobserved environment. Because the environ-

ment influences migrating cells, it is possible to recover the information about the hidden chemoattractant concentration from the observed migration patterns of the cell population. The primary goal of this chapter is to develop a dynamical model that links the hidden global environment with the trajectory of an individual cell. This link is then exploited in the estimation of the hidden ChA concentration governing neutrophil recruitment in zebrafish larvae. An extension of the EM algorithm discussed in §2.4.6 is utilised for the estimation.

It has been observed that the migrating cells alternate between directed motion towards the higher concentration of external stimuli and random walk. An individual neutrophil migration can thus be characterised as a stochastic dynamical process driven by an external environment. Various dynamical models describing the movement of an individual cell have been proposed, each incorporating the influence of the chemoattractant differently. In [51] the field gradient acts on a cell as a force, contributing to its acceleration. In [25], [53] the environment affects the orientation of the cell. In [128], neutrophil velocity is directly proportional to the gradient. This chapter utilises the model of [51] because it can be easily transformed into a state-space representation and incorporate the influence of the global environment as an input vector. The resultant model is presented in Section 3.2.

The problem of ChA field inference addressed here is similar to the ones solved by [129] and [128]. Authors of the former tackle the environment estimation problem by introducing three different phenomenological models of the hidden concentration field and discriminating between them within the Bayesian rejection algorithm. The proposed solution requires making prior assumptions about the shape of the field thus limiting its inference to several generic models. The latter work introduces a more flexible methodology in terms of field representation that relies on basis function decomposition of the chemoattractant environment. However, the authors solve the problem deterministically by decoupling field parameter learning and cell state estimation.

The main contribution of this chapter is an approximate Maximum Likelihood framework that can estimate the hidden chemoattractant concentration field and recover cell states simultaneously from spatiotemporally resolved *in vivo* data. The algorithm is derived in Section 3.3. The solution employs the artificial potential field method of modelling object interaction with the global environment, a popular reactive mechanism that is often used to model the environment in wide range of applications, including algorithms of swarm formation for ground-based robots [130], [131] or UAVs [132]. In this approach, the underlying environment is modelled as a static potential field \mathcal{U} driving object movement. This translates to the cell migration as follows: a cell that moves through the chemoattractant concentration field is subject to the attractive force that is proportional to the field gradient at the current

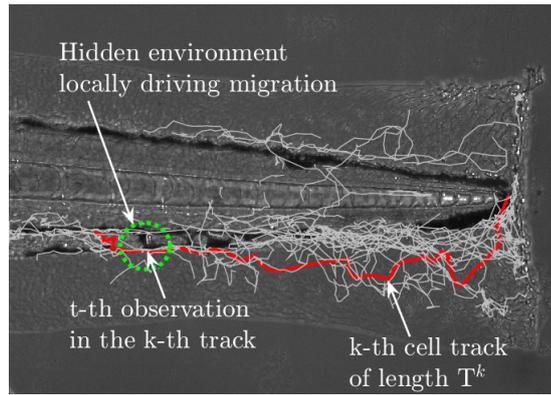


Figure 3.1: An example of tracking data obtained from *in vivo* observation of the injured zebrafish larva. Two key aspects of chemotaxis are unobserved: cell velocities and ChA concentration environment.

cell position. The underlying field is parametrised via basis function decomposition which provides the flexibility necessary for inference of complex shapes without changing the model structure.

The framework is applied to several datasets to recover the ChA concentration generated by different types of injury. The combination of ML method and basis function decomposition means that no prior assumptions about the shape of concentration environment are required within the selected parametrisation. The estimation results are presented and analysed in Section 3.4. All datasets are processed with the same tuning parameter values, meaning that the inferences made about the environment are made based solely on the available tracking data.

3.2 Methods

3.2.1 Data acquisition

All animal experiments were performed by the collaborators at the Department of Infection, Immunity & Cardiovascular Disease according to legislation and guidelines detailed in the Animals (Scientific Procedures) Act 1986. Ethical approval was given by the University of Sheffield Local Ethical Review Panel and experiments were fully approved by the Home Office (Project license PPL 70/8178). The neutrophil-specific fluorescent Tg(mpx:GFP)i114 zebrafish line, also referred to as mpx:GFP, was used for all experiments. Adult zebrafish were raised in the Bateson Centre at The University of Sheffield in UK Home-Office approved aquaria. All zebrafish were maintained according to standard protocols in [133].

All experiments were conducted on 3 days post fertilisation (dpf) mpx:GFP embryos anaesthetised by immersion in E3 containing 4.2% tricaine (Sigma-Aldrich).

The inflammatory response was activated by tail fin transection using a sterile scalpel blade as described previously in [33], or by a tail fin nick by a sterile needle as described in [1]. 3dpf mpx:GFP embryos were mounted in 0.7% low melting point agarose (Sigma-Aldrich) containing 4.2% tricaine for imaging immediately after tail transection. Time lapse imaging was performed during 0.5-5 hours post injury with 30 second sampling period using an Eclipse TE2000-U fluorescence microscope with a Andor Zyla 5.5 camera (Nikon).

Fast sampled imaging allows one to track individual cells. The Z-stacked images obtained from the video microscopy were compressed into one maximum intensity projection layer in within the NIS Elements software (Nikon). The 2-D neutrophil trajectory data was then extracted from the compressed time lapse images by applying a threshold to detect GFP neutrophils and linking their centroid positions via the tracking tool within the NIS Elements. The tracking results superimposed on brightfield images of fish bodies are shown in Figure 3.2.

Different injury methods demonstrated in the figure can induce inflammatory responses of varying severity. Normally, the tail fin transection is performed near the end of the circulation loop without damaging it or the notochord (the circulation defined by the black line and the spinal chord defined by the bright line are intact in Figure 3.2a). In the absence of inflammation inhibitors, normal injury induces continuous recruitment of immune cells for 5-6 hours post injury, during which neutrophils get uniformly distributed along the wound site. Occasionally, one can observe a more concentrated swarm-like formation near the wound, as in two larvae in Figure 3.2b. A mild injury can be triggered by either a tail fin transection further away from the circulation (see Figure 3.2c), or by a small cut with a needle below the notochord (as in Figure 3.2d). It can be seen that fewer neutrophils are activated by the mild injury, and in case of the nick injury the neutrophils are often delivered to the wound site by the circulation rather than travelling in the tissue. This results in rather short tracks concentrated around the wound. A severe injury is performed by cutting through the notochord and the circulation loop (severed circulation loops can be seen in both larvae in Figure 3.2e). A rapid inflammatory response induced by the severe injury is characterised by a more “chaotic” neutrophil behaviour, which may be caused by the fact that the ChA ligands quickly diffuse within the tissue and their global concentration gradient becomes more shallow.

For each fish larva, the observation set \mathcal{Y} consists of K tracks:

$$\mathcal{Y} = \{\mathbf{y}^k\}_{k=1}^K,$$

where each track \mathbf{y}^k is a collection of observed positions of the k -th cell ordered in time:

$$\mathbf{y}^k = \{\mathbf{y}_t^k\}_{t=0}^{T^k},$$

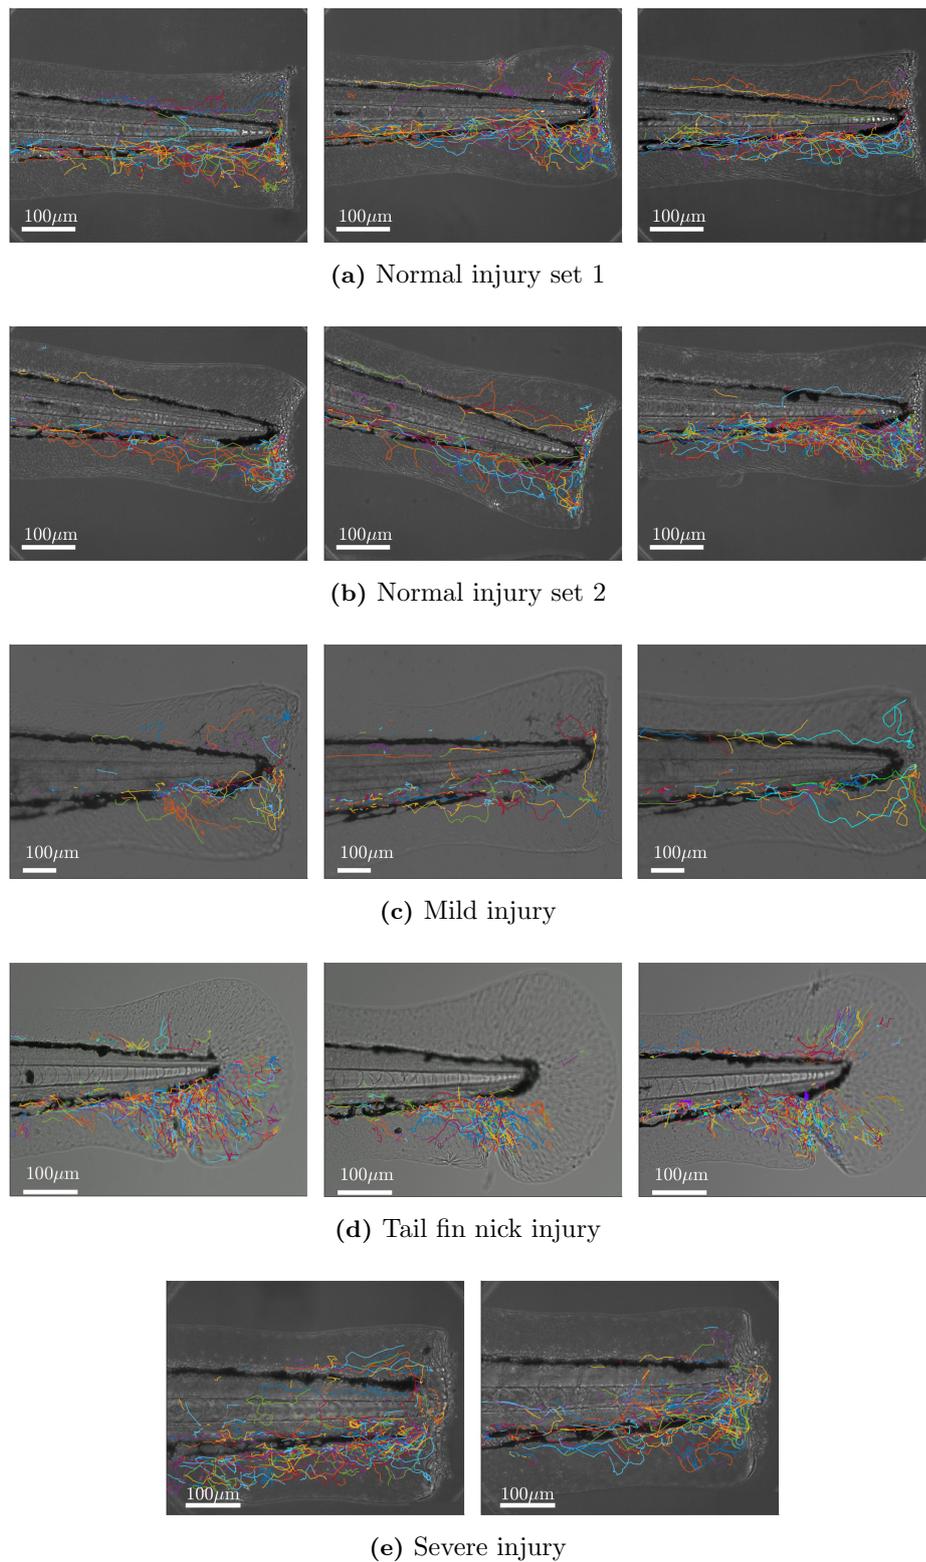


Figure 3.2: The tail fin transection initiates recruitment of neutrophils that are tracked for 5 hours post injury. (a)-(b) Normal injury is characterised by lengthy tracks with notable preference in direction. (c-d) A minor injury is inflicted on the fish by either a smaller fin transection or a tail fin nick with a needle and is characterised by numerous short tracks. (e) Severe injury activates more neutrophils but the recruited cells do not maintain the directionality and diffuse within the tissue.

where T^k is the length of the individual track, and an individual measurement at time t contains the observed cell position $\mathbf{y}_t^k = [\bar{s}_x, \bar{s}_y]_t^\top$.

3.2.2 Model of neutrophil dynamics

In this chapter, the problem of environment inference is solved under the following key assumptions.

Assumption 3.2.1: *The ChA environment \mathcal{U} is time-invariant.*

Assumption 3.2.2: *The global environment acts on an individual cell as a potential field.*

Assumption 3.2.3: *All neutrophils migrate in response to the acting environment.*

The dynamics of the k -th cell are described by a standard discrete time SSM:

$$\mathbf{x}_t^k = A\mathbf{x}_{t-1}^k + B\mathbf{u}_{t-1}^k + G\mathbf{w}_{t-1}^k, \quad (3.2.1a)$$

$$\mathbf{y}_t^k = C\mathbf{x}_t^k + \mathbf{v}_t^k, \quad (3.2.1b)$$

where the state vector at sample time t contains position and velocity of the cell centroid on the two-dimensional image

$$\mathbf{x}_t^k = [s_x, s_y, v_x, v_y]_t^\top, \quad (3.2.2)$$

and where the process noise $\mathbf{w}_t \sim \mathcal{N}(0, Q_w)$, $Q_w \in \mathbb{R}^{2 \times 2}$ and the measurement noise $\mathbf{v}_t \sim \mathcal{N}(0, R_v)$, $R_v \in \mathbb{R}^{2 \times 2}$ are i.i.d. Gaussian sequences. System matrices A, B, G , and C are considered to be time-invariant and are defined as follows

$$A = \begin{bmatrix} \mathbb{I}_2 & T\mathbb{I}_2 \\ \mathbb{O}_2 & \mathbb{I}_2 - T\rho\mathbb{I}_2 \end{bmatrix}, \quad (3.2.3a)$$

$$B = \begin{bmatrix} \mathbb{O}_2 & T\mathbb{I}_2 \end{bmatrix}^\top, \quad (3.2.3b)$$

$$G = \begin{bmatrix} \mathbb{O}_2 & T\mathbb{I}_2 \end{bmatrix}^\top, \quad (3.2.3c)$$

$$C = \begin{bmatrix} \mathbb{I}_2 & \mathbb{O}_2 \end{bmatrix}^\top, \quad (3.2.3d)$$

where T is the time increment, \mathbb{I}_2 is an identity matrix and \mathbb{O}_2 is a zero matrix of size 2×2 . Noise variances Q_w and R_v are defined as

$$Q_w = \sigma_w^2 \mathbb{I}_2; \quad R_v = \sigma_v^2 \mathbb{I}_2. \quad (3.2.4)$$

The presented model is the Euler–Maruyama discretisation of the biased random walk with resistance to the environment described in [51]. In the transition

matrix A , the term $\mathbb{I}_2 - T\rho\mathbb{I}_2$ corresponds to the reversion to mean in the O-U process describing velocity of a large Brownian particle in continuous time. Random fluctuations of neutrophil acceleration that are modelled as a Wiener process in [51] are approximated here by the Gaussian noise because the process is sampled with a constant time increment. The rate of reversion to mean ρ , and variances σ_w and σ_v are the tuning parameters that are specified prior to the estimation procedure.

Under the potential field paradigm, the deterministic input term in (3.2.1a) corresponds to the influence of the environment that can be described as follows

$$\mathbf{u}_t^k = \mathbf{u}_t^k(s_x, s_y) = \nabla \mathcal{U}(s_x, s_y). \quad (3.2.5)$$

where ∇ denotes the gradient and where $\mathcal{U}(s_x, s_y)$ is the ChA concentration at the current position of the migrating cell. In the interest of the posed inference problem, the potential function \mathcal{U} must be defined globally.

3.2.3 Chemoattractant field decomposition

Direct estimation of the global ChA concentration from the localised tracking data poses an infinite-dimensional problem, thus necessitating a simplified parametrised model of the environment. A natural way to reduce the order of the model is to represent a complex function by a linear combination of isotropic basis functions. As any other smooth surface, the chemoattractant concentration field can be approximated to an arbitrary degree of accuracy via the following decomposition

$$\mathcal{U}(s_x, s_y) = \sum_{h=1}^{N_b} \beta_h(s_x, s_y) \theta_h; \quad (3.2.6)$$

where $\beta_h(s_x, s_y)$ is a bivariate basis function defined in 2-D space and θ_h is the corresponding scaling coefficient that defines the magnitude of this function (Figure 3.3a). Since only one scaling parameter is assigned to each function (shown in Figure 3.3b), overall inference of such a model is easy to implement and has computational costs proportional to the size of the basis function grid. Note that the decomposition in the form (3.2.6) poses no restrictions on the kernel of basis functions. For example, in [128] the environment $\mathcal{U}(s_x, s_y)$ is approximated by a composition of overlapping Gaussian basis functions defined as follows

$$\beta_h(s) = \exp\left[-\frac{1}{2}(s - c_h)^T \Sigma_{\beta_h}^{-1} (s - c_h)\right], \quad (3.2.7)$$

where $c_h \in \mathbb{R}^2$ denotes the centre of h -th basis function, and $s = [s_x, s_y]^T$ denotes the position of a cell. Matrix Σ_{β_h} defines width of each basis function in two spatial dimensions. In a simple case, symmetric basis functions are considered, and hence diagonal elements of Σ_{β_h} are equal.

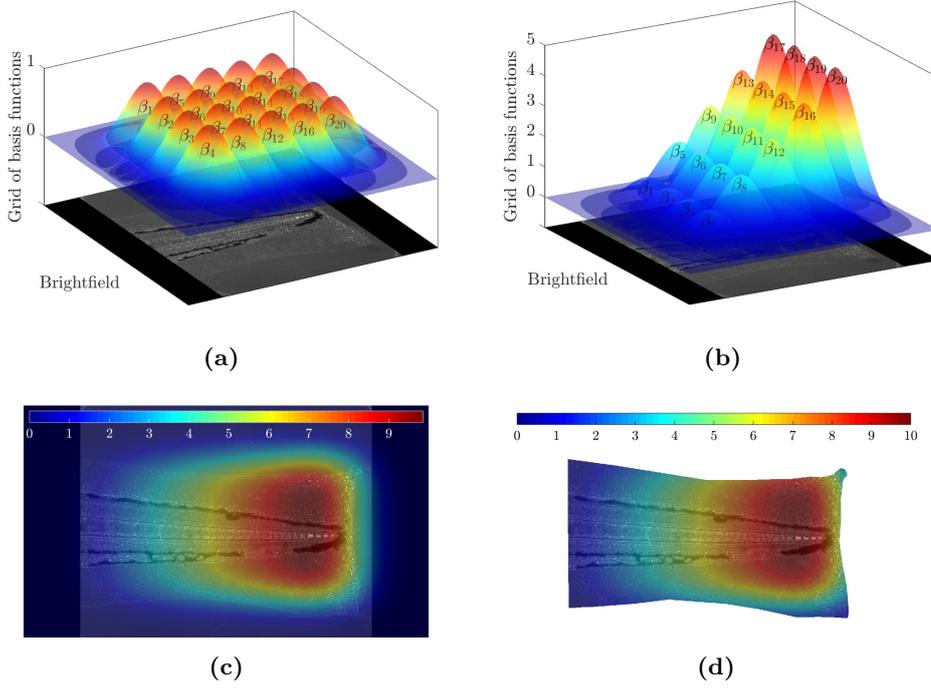


Figure 3.3: Decomposition of the environment with a grid of basis functions. **(a)** A 5×4 grid of uniformly spaced tensor product cubic B-splines is placed over the image of the fish body. Each b-spline β_h is a function of two coordinates, s_x and s_y . **(b)** Each b-spline is associated with a scaling parameter θ_h that regulates its magnitude. **(c)** Even a small number of b-splines guaranties smooth approximation of complex shapes. **(d)** The masked image illustrates the ChA environment within the fish body.

This thesis utilises another common type of multivariate basis functions called the tensor product B-splines

$$\beta_h(s_x, s_y) = \beta_l^4(s_x)\beta_m^4(s_y), \quad (3.2.8)$$

where $l \times m = N_b$ and where $\beta_l^4(s_x)$ and $\beta_m^4(s_x)$ are cubic B-splines constructed for a uniform knot sequence in each spatial dimension. The base of each B-spline is defined by a rectangle with sides $D_l^x = \{d_l^x, \dots, d_{l+5}^x\}$ for $\beta_l^4(s_x)$ and $D_m^y = \{d_m^y, \dots, d_{m+5}^y\}$ for $\beta_m^4(s_y)$. The tensor product formalism is discussed in chapter 7 of [134]. Its main advantage is that all properties of univariate B-splines extend to the multivariate case:

- $\beta_h(s_x, s_y)$ is a non-negative polynomial on the compact support containing knots $D_{l,x}, D_{m,y}$.
- $\beta_h(s_x, s_y)$ is equal to zero everywhere outside the area
- $\beta_h(s_x, s_y)$ is differentiable $k - 1$ times where k is the B-spline order and the derivative can be calculated analytically.

These properties, as well as inherent smoothness and flexibility make B-splines particularly attractive in applied problems that may have high computational complexity. Although the estimation framework introduced below is not restricted to the particular type of basis functions as long as they are differentiable, the aim of this work is to not only approximate a global function of unknown order but also provide the consistent gradients: in areas with no tracking data the field is expected to be a flat surface for approximation of which local polynomial representations are more suitable. This is why the low order B-splines are selected to parametrise the environment. Moreover, the B-splines have no additional tuning parameters (such as width of the Gaussian), meaning that for any given grid of B-splines and the image of any size the locations of the control points can be computed automatically.

This chapter utilises a coarse grid of tensor product B-splines to infer the global environment based on localised tracking data. The resultant model of individual cell dynamics can be rewritten as

$$\mathbf{x}_t^k = A\mathbf{x}_{t-1}^k + B\varphi_{t-1}^k(s_x, s_y)\Theta + G\mathbf{w}_{t-1}^k, \quad (3.2.9)$$

where the term $\varphi_t^k(s_x, s_y) = \mu\nabla\mathcal{B}(s_x, s_y)$ describes the gradient of superposition of basis functions at the current cell location. The chemotactic sensitivity coefficient μ is arbitrary set to one, while the gradient of the superposition of splines is presented as follows

$$\nabla\mathcal{B}(s_x, s_y) = \begin{bmatrix} \frac{\partial\beta_1(s_x, s_y)}{\partial s_x} & \frac{\partial\beta_2(s_x, s_y)}{\partial s_x} & \cdots & \frac{\partial\beta_h(s_x, s_y)}{\partial s_x} & \cdots & \frac{\partial\beta_{N_b}(s_x, s_y)}{\partial s_x} \\ \frac{\partial\beta_1(s_x, s_y)}{\partial s_y} & \frac{\partial\beta_2(s_x, s_y)}{\partial s_y} & \cdots & \frac{\partial\beta_h(s_x, s_y)}{\partial s_y} & \cdots & \frac{\partial\beta_{N_b}(s_x, s_y)}{\partial s_y} \end{bmatrix}. \quad (3.2.10)$$

In (3.2.9), Θ denotes a vector of the corresponding scaling parameters that are unknown

$$\Theta = [\theta_1, \theta_2, \dots, \theta_h, \dots, \theta_{N_b}]^\top. \quad (3.2.11)$$

It can be seen that the developed model is linear with respect to the unknown parameter vector Θ . However, since the basis functions depend on cell positions, the dynamics are non-linear with respect to the cell state. The measurement equation (3.2.1b) remains unchanged.

The following estimation objectives are associated with the model (3.2.9):

- Estimate all hidden states $\mathcal{X} = \{\{\mathbf{x}_t^k\}_{t=0}^{T_k}\}_{k=1}^K$ conditional on the given cell tracks \mathcal{Y} and the prediction of the underlying environment (here amounts to the estimated parameter vector $\hat{\Theta}$).
- Estimate the value of scaling coefficients $\hat{\Theta}$ given the complete data set $\mathcal{Z} = \{\mathcal{X}, \mathcal{Y}\}$ and estimated cell states \mathcal{X} .

The following section presents the formulation of the chemoattractant field inference problem within the ML framework.

3.3 Chemoattractant concentration inference

The problem of joint state-parameter estimation in the presence of the whole observation set \mathcal{Y} can be straightforwardly solved via the EM algorithm discussed in the previous chapter and summarised in Algorithm 3.1. It is a popular approach to estimating the parameters of generic linear SSMs [135] that treats all unobserved states as the hidden data \mathcal{X} . The algorithm efficiently separates two estimation problems and solves them iteratively until convergence. The fact that the ML approach does not require any *a priori* knowledge about unobserved data makes it particularly applicable to biological systems, where it is desirable to minimise the number of assumptions about the parametrised model [136]. The constituent steps of Algorithm 3.1 are discussed in the following subsections.

Algorithm 3.1 Expectation-Maximisation algorithm

Input: The set of observations \mathcal{Y} ; the dynamics function, $f(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_t)$; the measurement function $h(\mathbf{x}_t, \mathbf{u}_t, \mathbf{v}_t)$; state noise variance, Q_w ; measurement noise variance, R_v ; initial state estimate for each cell, $\bar{\mathbf{x}}_0^k$; and initial covariance for each cell, \mathbf{P}_0^k .

Output: The MLE of the unknown parameter vector, $\hat{\Theta}^*$.

- 1: Initialise with $\hat{\Theta}^0$;
 - 2: **while** not converged **do**
 - 3: E-step: compute the expected log-likelihood function $\mathcal{Q}(\Theta, \hat{\Theta}^i) = \mathbb{E} [\log p(\mathcal{X}, \mathcal{Y} | \Theta) | \mathcal{Y}, \hat{\Theta}^i]$;
 - 4: M-step: compute ML estimates $\hat{\Theta}^{i+1} = \arg \max_{\Theta} \mathcal{Q}(\Theta, \hat{\Theta}^i)$;
 - 5: Check convergence condition;
 - 6: **end while**
-

3.3.1 The log likelihood function

For K cells observed across all sampling times there exists complete data set $\mathcal{Z} = \{\mathcal{X}, \mathcal{Y}\}$ with joint probability defined as follows

$$\mathcal{L}(\Theta) = p(\mathcal{Z} | \Theta) = \prod_{k=1}^K \left[p(\mathbf{x}_0^k) \prod_{t=1}^{T_k} p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \Theta) \prod_{t=0}^{T_k} p(\mathbf{y}_t^k | \mathbf{x}_t^k) \right], \quad (3.3.1)$$

where individual probability densities are

$$p(\mathbf{x}_0^k) = \mathcal{N}(\mu_0, \Pi), \quad (3.3.2a)$$

$$p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \Theta) = \mathcal{N}((G)^\dagger (\mathbf{x}_t^k - A\mathbf{x}_{t-1}^k - B\phi(C\mathbf{x}_{t-1}^k)\Theta), \Sigma_w), \quad (3.3.2b)$$

$$p(\mathbf{y}_t^k | \mathbf{x}_t^k) = \mathcal{N}(C\mathbf{x}_t^k, R_v). \quad (3.3.2c)$$

where

$$\Sigma_w \triangleq \{(G)^\dagger\}^\top (Q_w)^{-1} (G)^\dagger \quad (3.3.3)$$

is the process noise variance updated to account for the process noise gain. The Moore-Penrose pseudoinverse

$$(G)^\dagger \triangleq (G^\top G)^{-1} G^\top$$

is used in case the noise gain G is a singular matrix. Taking the logarithm of (3.3.1) leads to the following expression for the log likelihood

$$\log \mathcal{L}(\Theta) = \sum_{k=1}^K \left[p(\mathbf{x}_0^k) + \sum_{t=1}^{T_k} \log p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \Theta) + \sum_{t=1}^{T_k} \log p(\mathbf{y}_t^k | \mathbf{x}_t^k) \right], \quad (3.3.4)$$

where only the term corresponding to (3.3.2b) is a function of Θ , which significantly simplifies the computation of the M-step:

$$\mathcal{Q}(\Theta, \hat{\Theta}^i) = \mathbb{E} [\log \mathcal{L}(\Theta)] = \mathbb{E} \left[\sum_{k=1}^K \sum_{t=1}^{T_k} \log p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \Theta) | \mathcal{Y}, \hat{\Theta}^i \right] + c, \quad (3.3.5)$$

where c includes all terms that are independent of Θ .

3.3.2 The forecasting step

In the linear Gaussian case, both the expectation and maximisation step can be solved analytically, as the \mathcal{Q} -function becomes the quadratic form and depends only on the first two central moments of the state \mathbf{x}_t^k and the unknown parameters Θ . In case of non-linear dynamics considered here, a tractable approximation is required for both steps of the algorithm. There exists a multitude of non-linear state estimators that can be utilised in the expectations step. Further complications arise during the maximisation step, since the maximised \mathcal{Q} -function may depend on high order moments of the state (in case of B-spline basis) or involve non-polynomial dependencies (in case of Gaussian basis). Explicit computation of (3.3.2b) becomes infeasible.

Numerical optimisation of the log likelihood may be computationally expensive and difficult to implement. Possible workarounds for this problem range from basis function decomposition of the dynamics $f(\mathbf{x}_t^k)$ and the observer $h(\mathbf{x}_t^k)$ [137] to using partial M-steps that increase the lower bound. Such algorithms as the Expectation Conditional Maximisation (ECM) algorithm [138] utilise this principle, but are liable to slow convergence as each CM step only guarantees to increase the log likelihood but does not maximise it [139]. Another solution is to approximate the non-linear function by the first-order Taylor series expansion, similar to the one employed in the EKF.

This work circumvents the problem of numerical optimisation by introducing another layer of hidden states that are algebraically linked to the expected state of the neutrophil at each time

$$\mathbf{s}_t^k = C \hat{\mathbf{x}}_{t|T^k}^k, \quad t = 0, \dots, T^k, k = 1, \dots, K, \quad (3.3.6)$$

so the augmented state \mathbf{s}_t^k can be interpreted as a forecast of the cell position that is utilised in computing the gradient of basis functions $\phi(\mathbf{s}_t^k) = \mu \nabla \mathcal{B}(s_x, s_y)$ as shown

in (3.2.10). Similar state augmentation is utilised in [140] to generate the forecast distribution in on-line particle filtering. In this case a simpler, mean field type approximation is used to replace the inherent stochasticity of the dynamical process with determinism [141], [142]: each forecast position \mathbf{s}_t^k behaves as an independent real-valued constant random variable that is deterministically linked to the state. The added state conforms to the degenerate distribution with the following probability mass function (*pmf*):

$$p(\mathbf{s}_t^k | \mathbf{x}_t^k) = \begin{cases} 1, & \mathbf{s}_t^k = C\hat{\mathbf{x}}_t^k, \\ 0 & \text{otherwise,} \end{cases} \quad (3.3.7)$$

which ensures that the complete-data likelihood function only takes non-zero value when \mathbf{s}_t^k is equal to the expected position of the cell. In the mean field approach, independence of \mathbf{s}_t^k means that the joint conditional *pdf* is factorisable

$$p(\mathbf{x}_t^k, \mathbf{s}_t^k | \mathbf{x}_{t-1}^k, \mathbf{s}_{t-1}^k) = p(\mathbf{s}_t^k | \mathbf{x}_t^k)p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \mathbf{s}_{t-1}^k), \quad (3.3.8)$$

so the addition of the forecasting step does not affect the cost function of the estimation algorithm.

Since the entire observation set \mathcal{Y} is available prior to the estimation procedure, cell states can be estimated via one of the non-linear smoothers discussed in §2.5.4. An additional, “forecasting” step (3.3.6) may be implemented after all expected state values have been obtained. The basis grid gradient is then calculated in the forecast positions, $\phi_t^k = \phi(\mathbf{s}_t^k)$. For the convenience of parameter estimation discussed in the following subsection, after the expectation and forecasting steps are complete the state transition model may be written in the following form

$$\mathbf{x}_t^k = A\mathbf{x}_{t-1}^k + B\phi(\mathbf{s}_{t-1}^k) + G\mathbf{w}_{t-1}^k, \quad (3.3.9)$$

meaning that the *pdf* of the state transition is approximated as

$$\begin{aligned} p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \Theta) &\approx p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \mathbf{s}_{t-1}^k, \Theta) = \\ &= \mathcal{N}((G)^\dagger(\mathbf{x}_t^k - A\mathbf{x}_{t-1}^k - B\phi(\mathbf{s}_{t-1}^k))\Theta, \Sigma_{\mathbf{w}}). \end{aligned} \quad (3.3.10)$$

Hence, as long as the sequence of forecast cell positions $\mathcal{S} = \{\{\mathbf{s}_t^k\}_{t=1}^{T_k}\}_{k=1}^K$ is obtained prior to the maximisation step, the \mathcal{Q} -function can be approximated as

$$\mathcal{Q}(\Theta, \hat{\Theta}^i) \approx \mathbb{E} \left[\sum_{k=1}^K \sum_{t=1}^{T_k} \log p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \mathbf{s}_{t-1}^k, \Theta) \right] + c. \quad (3.3.11)$$

The estimation procedure is then simplified to cycle between three steps:

$$\begin{aligned} \text{E-step:} & \quad \mathcal{X} | \hat{\Theta}^{i-1}, (\mathcal{S}) \\ \text{Forecast:} & \quad \mathcal{S} | \mathcal{X} \end{aligned}$$

$$\text{M-step: } \hat{\Theta}^i \mid \mathcal{X}, \mathcal{S}.$$

The intermediate step effectively separates the variables and reduces the maximisation step to optimising (3.3.11) in the closed form.

3.3.3 The maximisation step

Given the augmented state introduced above, the gradient of basis functions can be computed for all cell positions prior to the maximisation step

$$\varphi_t^k = \varphi(\mathbf{s}_t^k), \quad t = 1, \dots, T^k, k = 1, \dots, K, \quad (3.3.13)$$

making the Q -function a quadratic form of the state that can be maximised analytically.

Lemma 3.3.1: *Suppose that all constituent terms of (3.3.11) are estimated during the expectation step. Then the new MLE of the unknown parameter vector is given by*

$$\begin{aligned} \hat{\Theta}^{i+1} &= \left(\sum_{k=1}^K \sum_{t=1}^{T_k} (\varphi_{t-1}^k)^\top \Sigma_{\mathbf{w}}^{-1} \varphi_{t-1}^k \right)^{-1} \times \\ &\times \left[\sum_{k=1}^K \sum_{t=1}^{T_k} (\varphi_{t-1}^k)^\top \Sigma_{\mathbf{w}}^{-1} \left(\mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_t^k] - A \mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_{t-1}^k] \right) \right], \end{aligned} \quad (3.3.14)$$

where

$$\mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_{t-1}^k] \triangleq \mathbb{E}[\mathbf{x}_{t-1}^k \mid \mathbf{y}^k, \hat{\Theta}^i], \quad (3.3.15a)$$

$$\mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_t^k] \triangleq \mathbb{E}[\mathbf{x}_t^k \mid \mathbf{y}^k, \hat{\Theta}^i] \quad (3.3.15b)$$

are expectations of cell full states at two subsequent times conditioned on the observed cell tracks and the current parameter estimate $\hat{\Theta}^i$.

Proof. See Appendix A. ■

Note that because of the model structure and the adopted approximation, the maximisation step relies only on the first moments of the conditional hidden state distribution at two subsequent time instances. These expectations, along with augmented states \mathbf{s}_t^k , are computed during the expectation step.

3.3.4 The expectation step

As has been shown in §2.4.6, the state space modelling framework allows one to substitute evaluation of the full Q -function with computation of state expectations

Algorithm 3.2 Extended forward-backward algorithm

Input: Measurement vector $\mathbf{y}_{1:T}$; the dynamics function, $f(\mathbf{x}, \mathbf{u}, \mathbf{w})$; the measurement function $h(\mathbf{x}, \mathbf{u}, \mathbf{v})$; state noise variance, $Q_{\mathbf{w}}$; measurement noise variance, $R_{\mathbf{v}}$; initial state estimate, $\bar{\mathbf{x}}_0$; and initial covariance, \mathbf{P}_0 .

Output: Sequence of smoothed state estimates, $\{\hat{\mathbf{x}}_{t|T}\}_{t=1}^T$; and the corresponding covariances, $\{\mathbf{P}_{t|T}\}_{t=1}^T$.

- 1: *Forward recursion:* Initialise with $\bar{\mathbf{x}}_0$ and \mathbf{P}_0 ;
- 2: **for** $t \leftarrow 1, T$ **do**
- 3: Compute the Jacobian of the dynamics function $F_t = \frac{\partial f}{\partial \mathbf{x}}|_{\mathbf{x}=\hat{\mathbf{x}}_{t|t}}$;
- 4: Compute the Jacobian of the measurement function $H_t = \frac{\partial h}{\partial \mathbf{x}}|_{\mathbf{x}=\hat{\mathbf{x}}_{t|t}}$;
- 5: Run the standard Kalman filter (see Algorithm 2.1) recursion with F_t as a transition matrix and H_t as an observation matrix;
- 6: *Backward recursion:* Initialise with $\bar{\mathbf{x}}_{T|T}$ and $\mathbf{P}_{T|T}$ obtained in the forward pass;
- 7: **end for**
- 8: **for** $t \leftarrow T-1, 1$ **do**
- 9: Compute the Jacobian of the dynamics function $F_t = \frac{\partial f}{\partial \mathbf{x}}|_{\mathbf{x}=\hat{\mathbf{x}}_{t|t}}$
- 10: Run the RTS smoother recursion with F_t as a transition matrix as described in Algorithm 2.3.
- 11: **end for**

conditioned on the current parameter estimates and available measurements. The maximisation step presented above requires estimation of cell states at two successive sampling periods:

$$\mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_{t-1}^k] = \mathbb{E}[\mathbf{x}_{t-1}^k | \mathbf{y}^k, \hat{\Theta}^i] = \hat{\mathbf{x}}_{t-1|T^k}^k, \quad (3.3.16)$$

$$\mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_t^k] = \mathbb{E}[\mathbf{x}_t^k | \mathbf{y}^k, \hat{\Theta}^i] = \hat{\mathbf{x}}_{t|T^k}^k, \quad (3.3.17)$$

where $\hat{\mathbf{x}}_{t-1|T^k}^k$ and $\hat{\mathbf{x}}_{t|T^k}^k$ are produced by the smoothing algorithm. Since the set of full tracks \mathcal{Y} is available *a priori*, the state space model described by (3.3.9) and (3.2.1b) is used in a fixed-interval smoother to recover full cell states. This chapter utilises the forward-backward scheme in the form of the RTS smoother.

Recall that the basis functions depend on the current position s_t^k of the k -th cell at time t , which makes the matrix ϕ_t^k a non-linear function of the cell hidden state, rendering the SSM non-linear as well. As has been discussed in the previous chapter, there exist a number of approximate methods for non-linear state estimation. Two most popular algorithms for non-linear estimation combine either the extended Kalman filter (EKF) and extended RTS smoother (ERTS), as summarised for one object ($K = 1$) in Algorithm 3.2, or the unscented Kalman filter (UKF) and unscented RTS smoother (URTS), as shown in the Algorithm 3.3. Alternatively, one can utilise the forecast cell positions computed on the previous iteration of the EM algorithm, and estimate full cell states using the linear RTS smoother.

After the state estimation procedure with one of the smoothing algorithms, the forecast cell positions can be easily obtained by passing state estimates through the noiseless measurement equation

$$\hat{\mathbf{s}}_t^k = C\hat{\mathbf{x}}_t^k, \quad t = 1, \dots, T^k, k = 1, \dots, K.$$

Algorithm 3.3 Unscented forward-backward algorithm

Input: Measurement vector $\mathbf{y}_{1:T}$; the dynamics function, $f(\mathbf{x}, \mathbf{u}, \mathbf{w})$; the measurement function $h(\mathbf{x}, \mathbf{u}, \mathbf{v})$; state noise variance, Q_w ; measurement noise variance, R_v ; initial state estimate, $\hat{\mathbf{x}}_0$; and initial covariance, \mathbf{P}_0 ; state vector length, n ; filter parameters, λ, α, β .

Output: Sequence of smoothed state estimates, $\{\hat{\mathbf{x}}_{t|T}\}_{t=1}^T$; and their covariances, $\{\mathbf{P}_{t|T}\}_{t=1}^T$.

1: *Forward recursion:* initialise with $\hat{\mathbf{x}}_0$ and \mathbf{P}_0 ;

2: **for** $t \leftarrow 1, T$ **do**

3: Generate $2n$ sigma points $\tilde{\mathbf{x}}_{t-1}$ representative of the distribution $N(\hat{\mathbf{x}}_{t-1|t-1}, \mathbf{P}_{t-1|t-1})$:

$$\begin{aligned}\tilde{\mathbf{X}}_{0,t-1} &= \hat{\mathbf{x}}_{t-1|t-1}, \\ \tilde{\mathbf{X}}_{k,t-1} &= \hat{\mathbf{x}}_{t-1|t-1} + (\sqrt{(n+\lambda)}\sqrt{\mathbf{P}_{t-1|t-1}})_k, \quad k = 1, \dots, n; \\ \tilde{\mathbf{X}}_{k,t-1} &= \hat{\mathbf{x}}_{t-1|t-1} + (\sqrt{(n+\lambda)}\sqrt{\mathbf{P}_{t-1|t-1}})_{k-n}, \quad k = n+1, \dots, 2n;\end{aligned}$$

where n is the size of the state vector $\mathbf{x}_{t|t}$.

4: Compute the weights:

$$\begin{aligned}W_0^m &= \frac{\lambda}{n+\lambda}; \\ W_0^c &= \frac{\lambda}{n+\lambda} + (1 - \alpha^2 + \beta); \\ W_k^m &= W_k^c = \frac{1}{2(n+\lambda)}, \quad k = 1, \dots, 2n;\end{aligned}$$

5: Propagate the sigma-points through the dynamics to get $\tilde{\mathbf{X}}_{k,t} = f(\tilde{\mathbf{X}}_{k,t-1}, \mathbf{u}_{t-1})$, $k = 0, \dots, 2n$;

6: Compute prior mean $\hat{\mathbf{x}}_{t|t-1} = \sum_{k=0}^{2n} W_k^m \tilde{\mathbf{X}}_{k,t}$;

7: Compute prior covariance $\mathbf{P}_{t|t-1} = \sum_{k=0}^{2n} W_k^c (\tilde{\mathbf{X}}_{k,t} - \hat{\mathbf{x}}_{t|t-1})(\tilde{\mathbf{X}}_{k,t} - \hat{\mathbf{x}}_{t|t-1})^\top + Q_w$;

8: Predict the measurement $\tilde{\mathbf{Y}}_{k,t} = \sum_{k=0}^{2n} W_k^m h(\tilde{\mathbf{X}}_{k,t}, \mathbf{u}_t)$;

9: Compute residual covariance $\mathbf{S}_t = \sum_{k=0}^{2n} W_k^c (\tilde{\mathbf{Y}}_{k,t} - \mathbf{y}_t)(\tilde{\mathbf{Y}}_{k,t} - \mathbf{y}_t)^\top + R_v$;

10: Compute cross-covariance $\mathbf{P}_{t,t|t-1} = \sum_{k=0}^{2n} W_k^c (\tilde{\mathbf{Y}}_{k,t} - \mathbf{y}_t)(\tilde{\mathbf{X}}_{k,t} - \hat{\mathbf{x}}_{t|t-1})^\top$;

11: Compute Kalman gain $\mathbf{K}_t = \mathbf{P}_{t,t|t-1}(\mathbf{S}_t)^{-1}$;

12: Complete the regular measurement update to acquire $\hat{\mathbf{x}}_{t|t}$ and $\mathbf{P}_{t|t}$ as described in Algorithm 2.1;

13: **end for**

14: *Backward recursion:* initialise with $\hat{\mathbf{x}}_{T|T}$ and $\mathbf{P}_{T|T}$ obtained during the forward pass;

15: **for** $t \leftarrow T-1, 1$ **do**

16: Augment the state and covariance:

$$\tilde{\mathbf{x}}_{t|t} = \begin{bmatrix} \hat{\mathbf{x}}_{t|t} \\ 0 \end{bmatrix}; \quad \tilde{\mathbf{P}}_{t|t} = \begin{bmatrix} \mathbf{P}_{t|t} & 0 \\ 0 & Q_w \end{bmatrix}.$$

17: Repeat steps 3-7 for the augmented state $\tilde{\mathbf{x}}_{t|t}$ to obtain $\tilde{\mathbf{X}}_{t+1|t}$, $\hat{\tilde{\mathbf{x}}}_{t+1|t}$, and $\tilde{\mathbf{P}}_{t+1|t}$;

18: Compute cross-covariance $\mathbf{P}_{t+1,t|t} = \sum_{k=0}^{2n} W_k^c (\tilde{\mathbf{X}}_{k,t} - \hat{\tilde{\mathbf{x}}}_{t+1|t})(\tilde{\mathbf{X}}_{k,t+1} - \hat{\tilde{\mathbf{x}}}_{t+1|t})^\top$;

19: Compute the smoother gain $\mathbf{J}_t = \mathbf{P}_{t+1,t|t}(\tilde{\mathbf{P}}_{t+1|t})^{-1}$;

20: Obtain the smoothed state $\hat{\mathbf{x}}_{t|T} = \hat{\tilde{\mathbf{x}}}_{t+1|t} + \mathbf{J}_t(\hat{\tilde{\mathbf{x}}}_{t+1|T} - \hat{\tilde{\mathbf{x}}}_{t+1|t})$;

21: Obtain the smoothed covariance $\mathbf{P}_{t|T} = \tilde{\mathbf{P}}_{t+1|T} + \mathbf{J}_t(\tilde{\mathbf{P}}_{t+1|T} - \tilde{\mathbf{P}}_{t+1|t})\mathbf{J}_t^\top$;

22: **end for**

Each of the computed forecast positions is then utilised to approximate the matrix φ_t^k necessary for the maximisation step using equation (3.3.13).

3.3.5 The estimation algorithm

The EM algorithm is initialised by calculating values of ϕ_t^k for all measured cell positions ($\hat{\varphi}_t^k)^0 = \varphi(\mathbf{y}_t^k)$. These matrices are then utilised in the estimation of full cell states with the assumption that there is no field acting on moving cells, *i.e.* all magnitude coefficients are set equal to zero $\Theta = 0$. The maximisation step is performed afterwards to obtain the initial parameter estimate $\hat{\Theta}^0$. Then the three steps described above are repeated until convergence as follows:

- i) E-step: compute the expected values of the cell state \mathbf{s}_t^k , $t = 1, \dots, T^k, k = 1, \dots, K$ via the URTS algorithm.
- ii) Forecasting: compute forecasts of object positions, \mathbf{s}_t^k , and approximate gradients of the basis function grid in those positions, $\hat{\varphi}_t^k$, for $t = 0, \dots, T^k, k = 1, \dots, K$.
- iii) M-step: maximise the \mathcal{Q} -function in closed form.

Assessing convergence of the likelihood function comes at a high computational cost as it requires computation of all terms included in the constant in (3.3.11) at each iteration of the algorithm. Instead, the following stopping criteria is used

$$\Delta\Theta^i = \frac{(\Theta^i - \Theta^{i-1})^\top (\Theta^i - \Theta^{i-1})}{(\Theta^i)^\top \Theta^i} < \epsilon. \quad (3.3.20)$$

The rate of convergence is highly dependent on the selection of the initial parameter vector estimate. As an EM solution, the algorithm described in this section is liable to converge to local maxima of the likelihood function. It must also be noted that the shape of the estimated chemoattractant field is only unique up to an additive constant since it is not the magnitude of the field that drives cell migration but its gradient.

In addition to the risk of converging to a neighbourhood of a local maxima, single-point estimation obscures the fact that there may exist other local maxima solutions worth investigating. As discussed in §2.4.2, the most straightforward way to assess the uncertainty of the ML estimator is via confidence regions. In incomplete-data problems it is impossible to compute the expected Fisher information as part of the information is inevitably lost to the missing data, but there exists a number of techniques to recover the observed information.

The observed Fisher information is not a direct by-product of the EM algorithm, but it conforms to the missing data principle [143]:

$$\mathcal{J}(\Theta | \mathcal{Y}) = \mathcal{I}_c(\Theta, \mathcal{Y}) - \mathcal{I}_m(\Theta, \mathcal{Y}), \quad (3.3.21)$$

where the first term on the right-hand side denotes the expected complete-data information and can be calculated as the expectation of the negative Hessian of the

complete log-likelihood function

$$\mathcal{I}_c(\Theta | \mathcal{Y}) = -\mathbb{E} \left[\frac{\partial^2 \log \mathcal{L}(\Theta)}{\partial \Theta^2} | \mathcal{Y}, \Theta \right], \quad (3.3.22)$$

and the second term denotes the quantity that can be interpreted as missing information that is lost as a consequence of observing only \mathcal{Y}

$$\mathcal{I}_m(\Theta | \mathcal{Y}) = -\mathbb{E} \left[\frac{\partial^2 \log g(\mathcal{X} | \mathcal{Y}, \Theta)}{\partial \Theta^2} | \mathcal{Y}, \Theta \right], \quad (3.3.23)$$

where $g(\mathcal{X} | \mathcal{Y}, \Theta) = p(\mathcal{Z} | \Theta)/p(\mathcal{Y} | \Theta)$ is the conditional density of \mathcal{Z} given \mathcal{Y} .

It is shown in [144] that both terms in (3.3.21) can be computed directly on the last iteration of the EM algorithm as the covariance of scores

$$\mathcal{I}_m(\Theta | \mathcal{Y}) = \text{cov}_{\Theta} \left\{ \mathbb{E} \left[\frac{\partial \log \mathcal{L}(\Theta)}{\partial \Theta} | \mathcal{Y} \right] \right\}. \quad (3.3.24)$$

Then the observed information is calculated as follows

$$\begin{aligned} \mathcal{J}(\Theta | \mathcal{Y}) = & -\mathbb{E} \left[\frac{\partial^2 \log \mathcal{L}(\Theta)}{\partial \Theta^2} | \mathcal{Y}, \Theta \right] - \mathbb{E} \left[\frac{\partial \log \mathcal{L}(\Theta)}{\partial \Theta} \frac{\partial \log \mathcal{L}(\Theta)^\top}{\partial \Theta} | \mathcal{Y}, \Theta \right] + \\ & + \left\{ \mathbb{E} \left[\frac{\partial \log \mathcal{L}(\Theta)}{\partial \Theta} | \mathcal{Y}, \Theta \right] \mathbb{E} \left[\frac{\partial \log \mathcal{L}(\Theta)}{\partial \Theta} | \mathcal{Y}, \Theta \right] \right\}, \end{aligned} \quad (3.3.25)$$

where the last term will be equal to zero if the EM reached the stationary point on the final iteration. The reciprocal of the observed information is fully defines the orientation and size of N_b -dimensional confidence ellipsoid [145]. The trace of this reciprocal is used in the thesis as a measure of total variance of the obtained MLE. In order to provide a unified metric for all experiments, the trace is divided by the dimension of Θ :

$$\text{Var}(\hat{\Theta}) = \text{tr}((\mathcal{J}(\Theta | \mathcal{Y}))^{-1}) / N_b.$$

3.4 Results

Given several data sets with different types of injury, we can see how the neutrophil recruitment process adjusts depending on the varying scale of infectious threats. The proposed estimation algorithm is used to process all datasets described above with similar parameters summarised in Table 3.1. The time step for each experiment is equal to the time increment between two frames in the video microscopy, and the measurement noise covariance matches the variance of the tracking module in NIS Elements software. The rest of the parameters used for this analysis are selected arbitrarily. All datasets from the same experiment are processed with the same tuning parameters.

Parameter	Normal injury	Mild injury	Nick injury	Severe injury
T	2 min	1.5 min	0.5 min	2 min
ρ	0.3	0.3	0.2	0.3
Q_w	$2\mathbb{I}_2$	$1\mathbb{I}_2$	$0.5\mathbb{I}_2$	$2\mathbb{I}_2$
R_v	$2\mathbb{I}_2$	$2\mathbb{I}_2$	$2\mathbb{I}_2$	$2\mathbb{I}_2$

Table 3.1: Tuning parameters of the dynamical model.

In the dataset with the normal injury, a grid of 5×4 basis functions has been created to parametrise the environment for the normal injury experiment, with the knot sequence generated automatically given the dimensions of the fish image. In datasets with other injury types, cell populations are usually confined to a smaller area in the fish body, therefore a sparser, 4×4 grid was used to approximate the environment for those datasets.

3.4.1 Simulation examples

The proposed estimation framework has been tested on the Monte Carlo simulation of tracking data generated by multiple objects migrating in response to the artificial potential field. The field is modelled using a 4×4 grid of overlapping tensor product cubic B-splines $\mathcal{B}(s_x, s_y)$ placed on the map of size 1000×1000 arbitrary units (a.u.). The magnitude of the h -th B-spline is defined by a corresponding coefficient θ_h . The vector of magnitude coefficients is assumed to be unknown for purposes of testing the algorithm. The gradient of the modelled field is illustrated in Figure 3.4. In

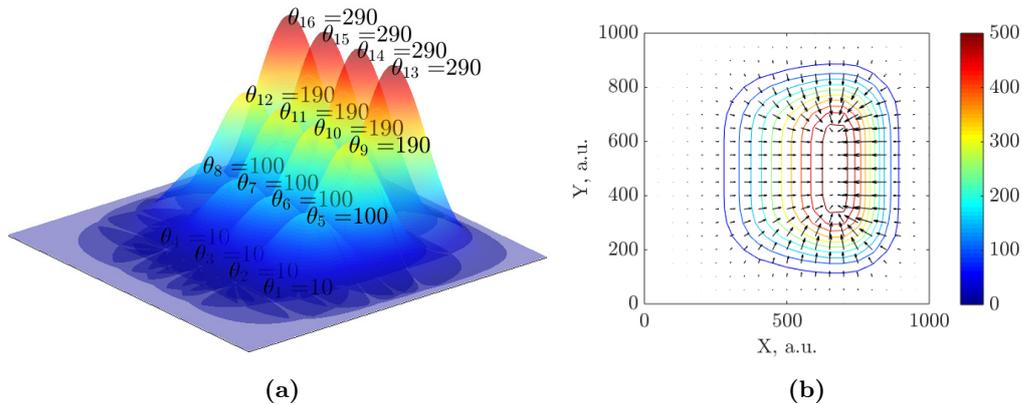


Figure 3.4: The potential field modelled with a 4×4 grid of tensor product B-splines. (a) Basis functions at each column are assigned with equal scaling parameters. (b) The gradient of the resultant modelled field. The colour bar corresponds to the magnitude of the constructed field.

each simulation, 100 tracks were generated by integrating the discrete-time SMM for 100 minutes with time increment T set to 1 min. Other tuning parameters have been selected to match the normal injury set (see Table 3.1).

Performance of the estimation framework is assessed by processing the simulation data with the assumption that all tuning parameters are fully known. An unscented forward-backward algorithm is used during the expectation step. The gradient of the mean estimated potential field that is shown in Figure 3.5a matches that of the modelled field, but the magnitude of the field is higher. This is the result of the initial choice of the parameter estimate $\hat{\Theta}^0$, as the bias of the mean field (see Figure 3.5b) is nearly invariant spatially and can be accounted for in an additive constant discussed above.

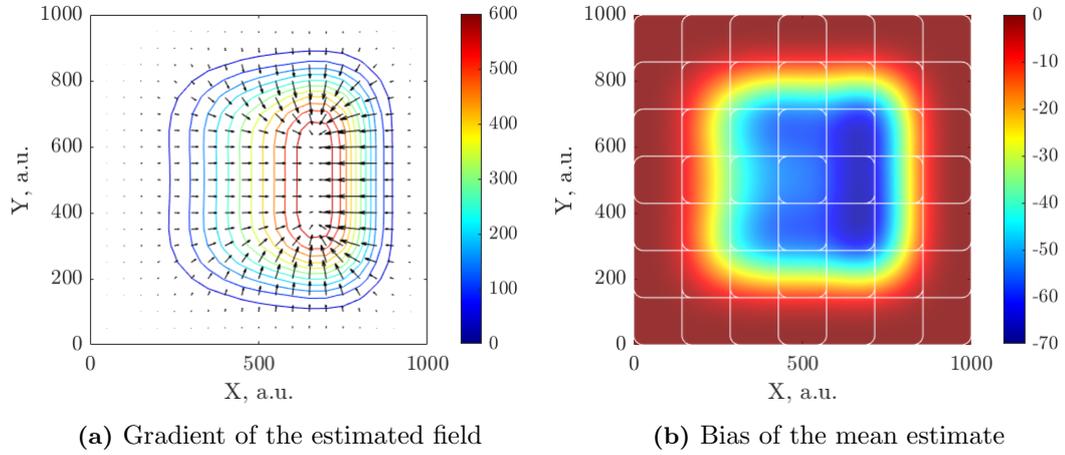


Figure 3.5: Field estimation results obtained from 100 MC simulations via the EM algorithm employing URTS smoother during the expectation step. **(a)** The gradient of the potential field constructed with the mean parameter estimates, $\nabla\mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the constructed field. **(b)** Total bias between the true field and the mean estimate, $\mathcal{U}(\Theta_{\text{true}}) - \mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the total bias.

The statistics of the parameter estimation results obtained using the correct model parameters are presented in Table 3.2. Since the scaling parameter values have varying magnitude it is more reasonable to compute the percentage (%) bias and the coefficient of variance (CV) for accessing the unbiasedness and the precision of individual MLEs. The %-bias computed as the percentage ratio of the total bias to the true parameter value:

$$\% \text{-bias} = \frac{\theta - \hat{\theta}}{\theta} \times 100\%.$$

The CV is the percentage ratio of standard deviation to the mean estimate:

$$\text{CV} = \frac{\sigma_{\hat{\theta}}}{\hat{\theta}} \times 100\%.$$

Parameter	True	Mean est.	St. dev.	CV (%)	%-bias
θ_1	10	32.680	20.878	63.887	226.801
θ_2	10	27.049	18.949	70.053	170.491
θ_3	10	27.485	18.533	67.430	174.847
θ_4	10	31.057	19.341	62.276	210.573
θ_5	100	129.034	12.608	9.771	29.034
θ_6	100	125.105	10.199	8.153	25.105
θ_7	100	123.581	10.487	8.486	23.581
θ_8	100	130.831	13.357	10.209	30.831
θ_9	190	216.838	9.232	4.258	14.125
θ_{10}	190	213.117	6.607	3.100	12.167
θ_{11}	190	213.588	6.963	3.260	12.415
θ_{12}	190	215.734	8.391	3.890	13.544
θ_{13}	290	325.707	9.721	2.985	12.313
θ_{14}	290	326.660	8.960	2.743	12.641
θ_{15}	290	325.258	8.219	2.527	12.158
θ_{16}	290	326.411	10.219	3.131	12.556

Table 3.2: Statistics of the parameter MLEs obtained via the EM algorithm employing URTS state estimation over 100 simulations.

The standard deviation is consistent for all elements of the parameter vector, meaning that all MLEs have similar dispersion from the mean value. Small deviation also indicates that the estimator is consistent over all MC simulations. Since the estimation was performed with correct tuning parameters, precision of the MLEs is high compared to the absolute parameter values (as indicated by small CV). The highest variance coefficient and %-bias are observed for the scaling parameters corresponding to B-splines at the left hand side of the map ($\theta_1 - \theta_4$). In this case such results can be explained by the choose of initial parameter values rather than poor estimation performance: recall absolute bias observed in Figure 3.5b is spatially constant, meaning that the %-bias will be large for small scaling coefficients.

While in this simulation study it is possible to compute statistics of the parameter MLEs from the Monte Carlo sample, for real-life applications of the ML estimator the measure of estimation uncertainty is provided by the reciprocal of Fisher information. In order to provide a point of reference for the estimation results presented in the following subsections, measures of the normalised total variance for the MLE vector obtained from a single Monte Carlo realisation of the tracking data are presented in Tables 3.3 and 3.4. Table 3.4 demonstrates the effect of the process

Q_w	$0.5 \mathbb{I}_2$	$2 \mathbb{I}_2$	$4 \mathbb{I}_2$	$8 \mathbb{I}_2$
$\text{Var}(\hat{\Theta})_{\text{ERTS}}$	41.315	165.147	330.235	660.349
$\text{Var}(\hat{\Theta})_{\text{URTS}}$	39.098	156.083	270.976	416.914

Table 3.3: Normalised total variance of the estimation results obtained using the model with varying process noise parameters. The measurement noise covariance is fixed at $R_v = 2 \mathbb{I}_2$.

R_v	$0.5 \mathbb{I}_2$	$2 \mathbb{I}_2$	$4 \mathbb{I}_2$	$8 \mathbb{I}_2$
$\text{Var}(\hat{\Theta})_{\text{ERTS}}$	165.058	165.147	165.198	165.256
$\text{Var}(\hat{\Theta})_{\text{URTS}}$	155.975	156.083	156.157	156.358

Table 3.4: Normalised total variance of the estimation results obtained using the model with varying measurement noise parameters. The process noise covariance is fixed at $Q_w = 2 \mathbb{I}_2$.

noise covariance on the dispersion of the parameters. Even though the smaller noise covariance $Q_w = 0.5 \mathbb{I}_2$ does not match the “true” setting, it results into a smaller scatter of the parameter MLEs. This may be explained by the fact that only the first moment of the cell state is explicitly involved in the computation of parameter estimates, and a smaller noise variance would result into a smoother mean estimate of the state while preserving the general trend of the true state. Higher process noises lead to the distortion of mean state estimates which in turn contribute to the higher dispersion of MLEs. Table 3.4, on the other hand, shows that the mismatch in the model of the measurement noise does not contribute to increase the estimate total variance. Note that the choice of the non-linear smoother employed in the expectation step also has no significant effect on the parameter estimation. However, it is worth investigating the accuracy of cell state estimation in more detail.

Performance of the state estimation stage is demonstrated on an example simulated track processed with ERTS and URTS smoothers on the last iteration of the EM algorithm in Figures 3.6 and 3.7, respectively. Mean values and $\pm 3\sigma$ confidence regions are presented for both filtering and smoothing procedures. While the confidence region for the extended framework is narrower than that of the unscented one, the UKF performs better than the EKF in the forward recursion, in particular with respect to the hidden states. Initial estimation errors of the EKF are corrected on the smoothing stage, but the URTS does produces similar results as UKF. The latter can be explained by accurate state estimation in the forward recursion of the unscented framework that does not leave much room for improvement.

Additional tests performed on the simulated tracking data have revealed that the ERTS framework is also sensitive to the tuning parameters Q_w and R_v , while

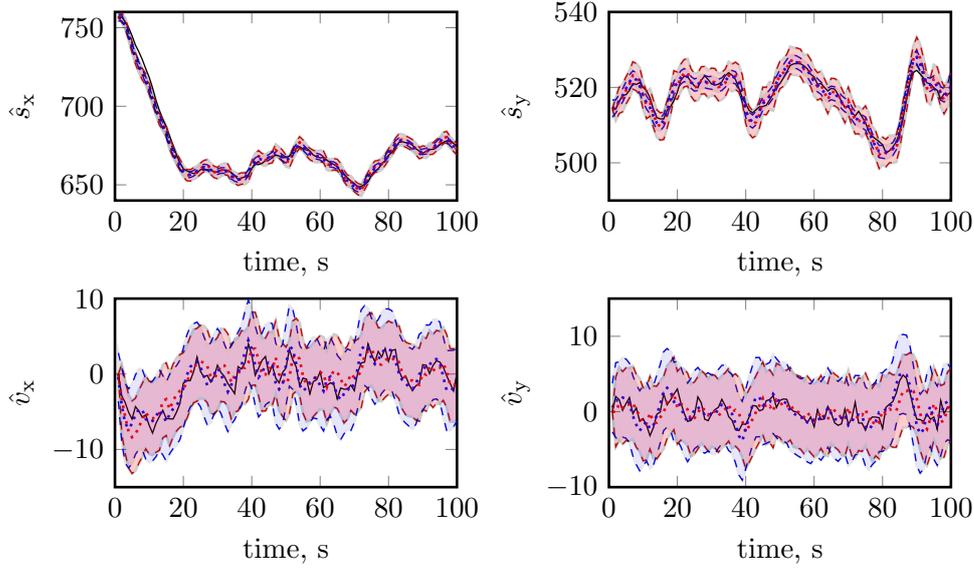


Figure 3.6: State estimation results obtained from the ERTS smoother for a simulated track. The true state (—), filtered state (.....) and its $\pm 3\sigma$ tube (■), smoothed state (.....) and its $\pm 3\sigma$ tube (■). Noise covariances used in the smoother match the parameters used to simulate the track: $Q_w = 2\mathbb{I}_2$, $R_v = 2\mathbb{I}_2$.

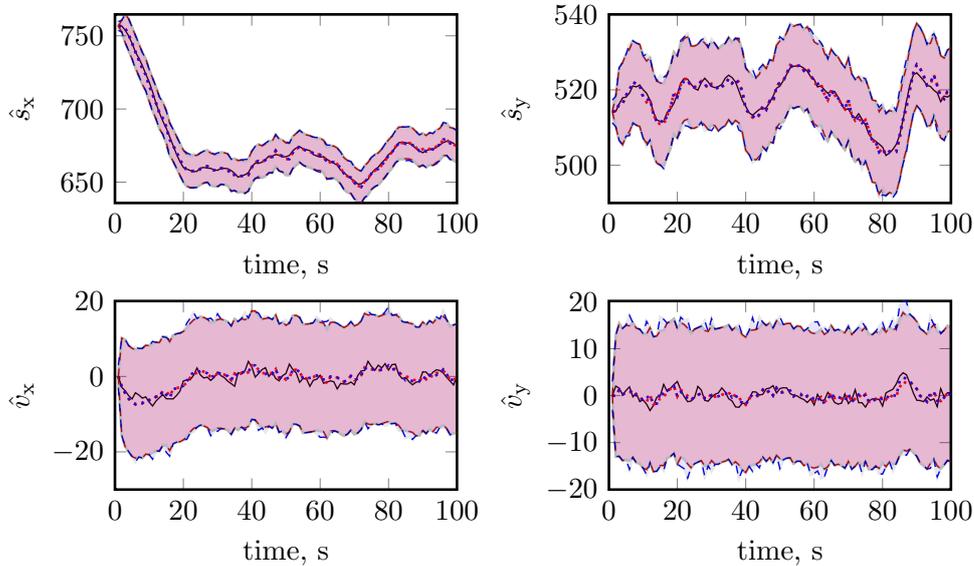


Figure 3.7: State estimation results obtained from the URTS smoother for a simulated track. The true state (—), filtered state (.....) and its $\pm 3\sigma$ tube (■), smoothed state (.....) and its $\pm 3\sigma$ tube (■). Noise covariances used in the smoother match the parameters used to simulate the track: $Q_w = 2\mathbb{I}_2$, $R_v = 2\mathbb{I}_2$.

the URTS scheme provides similar estimates regardless of the noise level, which indicates that it is more robust to the mismatch in the model of uncertainty. This is demonstrated in the plots of the Root Mean Square Error (RMSE) of state es-

timates for the example track processed with varying noise covariances (see Figures 3.8, 3.9). For both extended and unscented algorithms the mismatch in the process noise model has stronger effect on the estimation accuracy, particularly for the hidden states v_x , v_y . Notably, the extended filter has larger RMSE even when the estimation is performed with the correct levels of noise, but the accuracy is significantly improved in the smoother. It can be concluded from Figure 3.9 that the unscented filter copes better with the measurement noise mismatch: RMSE quickly stabilises on the same level as that of the filter with correct parameters. Since the true process and measurement noises of the dynamical system describing the migrating neutrophil are unknown, the URTS smoother is employed in the estimation framework while processing the experimental data.

3.4.2 Estimation results: cell velocities

First, the performance of various state estimation algorithms discussed in this chapter is compared in application to the neutrophil tracking data. Neutrophil tracks are processed with the EM algorithm employs different non-linear methods during the expectation step. The estimated states for one cell track in Fish 1 from the normal injury set are presented in Figure 3.10. It can be seen that the results obtained via the URTS smoother and the ERTS smoother are similar, potentially because the non-linearity is present only in the control input term, and the dynamical model reflects the fact that cell positions are measured with small observation errors. However, the extended smoother tends to smooth out small abrupt changes in direction, attributing them to the random fluctuation. The estimation framework based on the unscented transform attributes these “turns” to the dynamics and is generally more sensitive to small changes in velocity. The field estimation results presented in the following subsections are obtained by the framework that utilises the URTS algorithm.

The zebrafish in majority of the experiments are mounted on the microscope in such way so that the wound is aligned with the y -axis. This means that the directed migration of neutrophils towards the wound should be characterised by both median and mean shift in the horizontal velocity component. Example histograms of velocity components for various wound types are illustrated in Figure 3.11. Small but evident shift in the mean of the v_x component has been observed for all data sets, although there is little difference between the velocity shift caused by the normal wound and that caused by the severe wound. However, in zebrafish with severe injury the net mean shift in velocity along the X-axis is more notable compared to the normal injury data, but 25% of of all cells have v_x component close to zero, which brings the mean value to that of the zebrafish with normal injury. On the other hand, there is no positive median shift for any of the wound types. The observed results can be

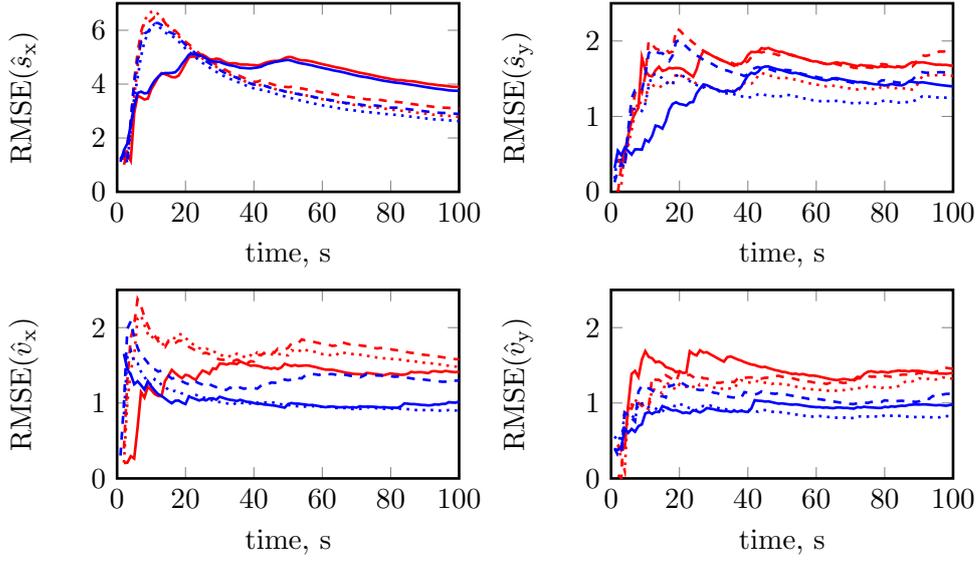


Figure 3.8: RMSEs of the state estimates obtained from EKF and ERTS with varying noise covariances. RMSE of filtered (—) and smoothed (—) estimates obtained with $Q_w = 2\mathbb{I}_2$, $R_v = 2\mathbb{I}_2$. RMSE of filtered (---) and smoothed (---) estimates obtained with $Q_w = 4\mathbb{I}_2$, $R_v = 2\mathbb{I}_2$. RMSE of filtered (⋯) and smoothed (⋯) estimates obtained with $Q_w = 2\mathbb{I}_2$, $R_v = 4\mathbb{I}_2$.

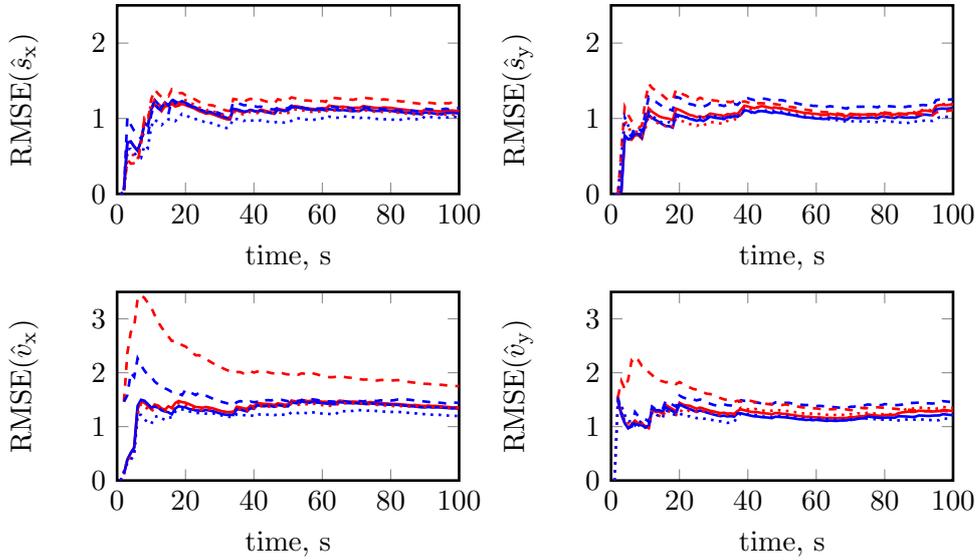


Figure 3.9: RMSEs of the state estimates obtained from UKF and URTS with varying noise covariances. RMSE of filtered (—) and smoothed (—) estimates obtained with $Q_w = 2\mathbb{I}_2$, $R_v = 2\mathbb{I}_2$. RMSE of filtered (---) and smoothed (---) estimates obtained with $Q_w = 4\mathbb{I}_2$, $R_v = 2\mathbb{I}_2$. RMSE of filtered (⋯) and smoothed (⋯) estimates obtained with $Q_w = 2\mathbb{I}_2$, $R_v = 4\mathbb{I}_2$.

interpreted as follows: while neutrophils increase their velocity while moving towards the wound, they are still equally likely to be migrating towards and away from the

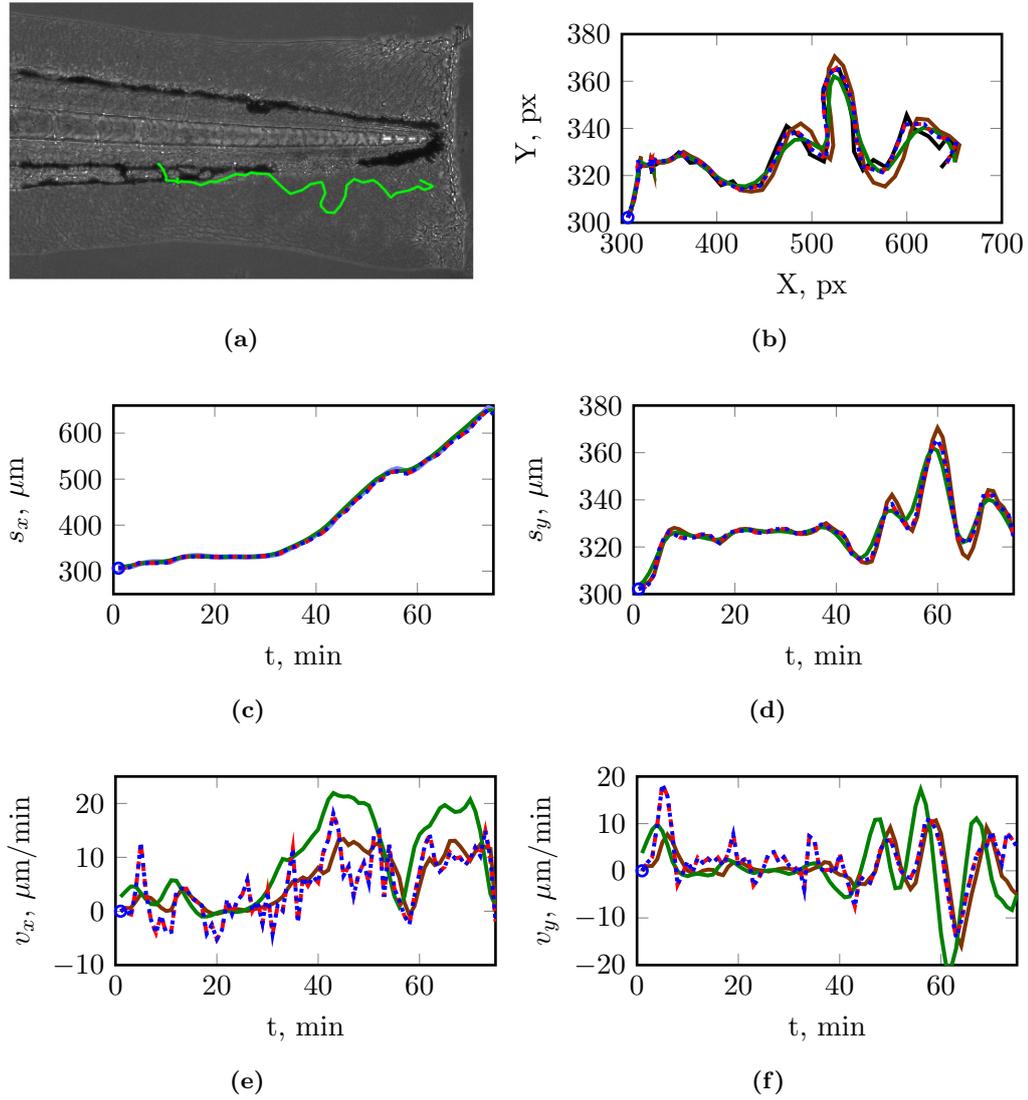


Figure 3.10: Comparison of state estimation results obtained from non-linear state estimation algorithms. (a) Observed cell track in the zebrafish larva. (b) The observed cell track (—), cell positions estimated by EKF (—), estimated by ERTS (—), estimated by UKF (---), and estimated by URTS (⋯). (c)-(d) Estimates of cell coordinates. (e)-(f) Estimates of velocity components.

wound site. This means that the initial assumption that all cells are driven by the ChA environment is falsified for the *in vivo* neutrophil data. The v_y component has similar distributions for all wound types, though the severe injury data has smaller variance.

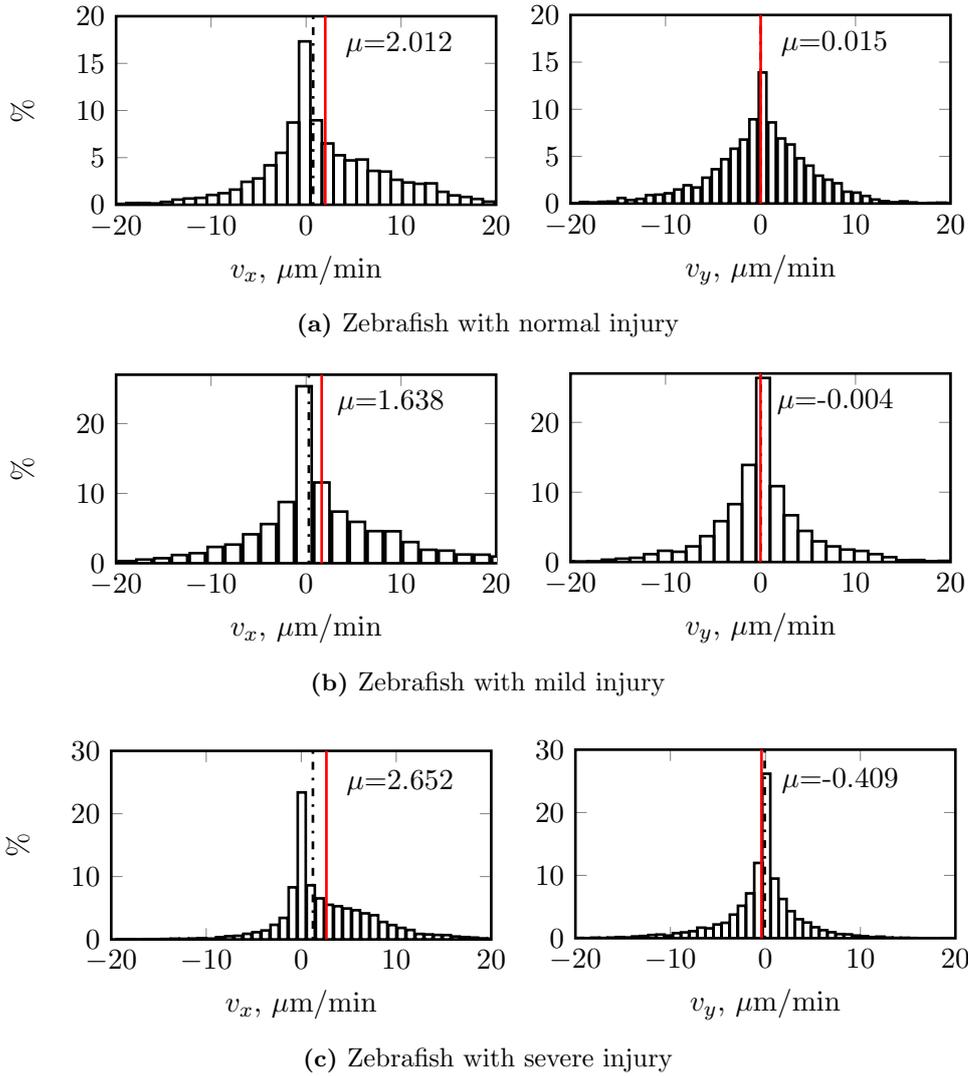


Figure 3.11: The example distributions of cell velocities estimated with UKF-URTS algorithm. (a) Fish 6 from the normal injury set. (b) Fish 1 from the mild injury set. (c) Fish 1 from the severe injury set.

3.4.3 Estimation results: normal injury

The estimated chemoattractant fields for the six zebrafish larvae from the first set are presented in Figure 3.12. Each result demonstrates the correlation between the total number of cells passing through the area and the estimated magnitude of the chemoattractant field in that area. The presented estimation results are adjusted post-estimation so that the smallest concentration is equal to zero. It can be seen that the estimated fields for Fish 1, 2, 3 and 6 conform to the well-established hypothesis that the chemoattractant is uniformly distributed along the injury with the highest concentration located at the injury site. Fish 4 and 5, however, demonstrate a different type of neutrophil behaviour. Contrary to the

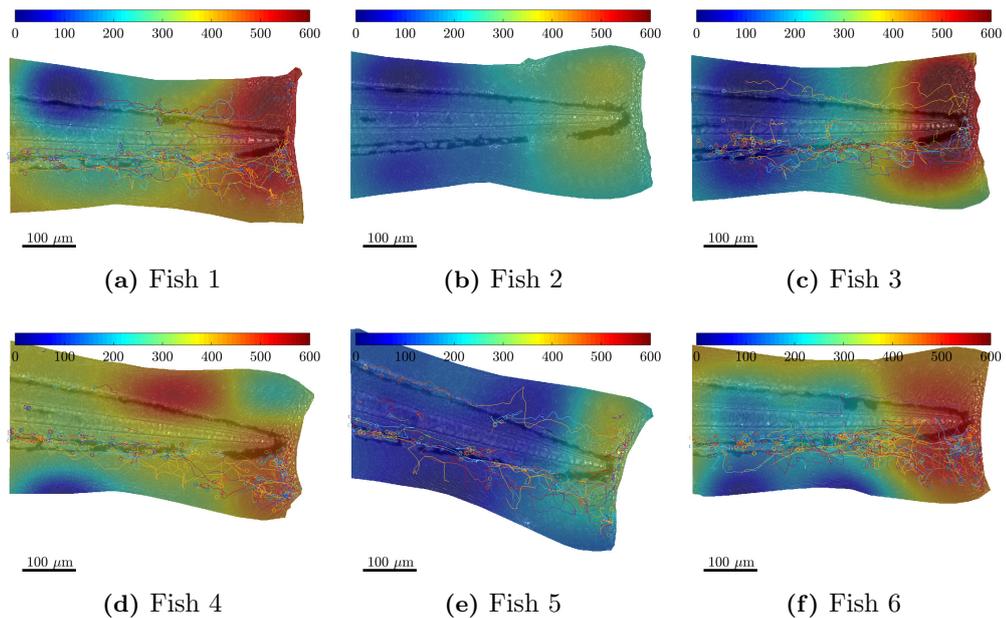


Figure 3.12: Inferred chemoattractant fields for the normal injury data. All colour bars are normalised to scale 0-600 a.u..

expected pattern, cells “swarm” towards the lower region of the injury site in Fish 4 and to the upper region in Fish 5. The swarming may be a response to the self-generated gradient recently observed in [146], that may be modulated predominantly by the LTB_4^1 mediator released by neutrophils themselves. In Fish 4, for which the lowest number of cell tracks is available, the magnitude of the inferred field peaks in the area with the most persistent cell tracks. The estimation results for Fish 4 and Fish 5 indicate that the future modelling and estimation work on population dynamics should take into account cell-to-cell interaction.

3.4.4 Estimation results: mild and severe injury

Results of the environment inference for other types of injury are summarised in Figure 3.13. In case of mild injury, the magnitude of chemoattractant concentrations ranges between 0 and 300 a.u., which is almost half of the range for several larva with normal injury. This is consistent with the fact that smaller wound releases fewer pro-inflammatory mediators which, in turn, attract fewer neutrophils. Sparsity of the neutrophil tracking data leads to sharp spatial changes in the estimated field that contradict the diffusive nature chemoattractants (see Figure 3.13b), and to unusually low concentration outliers near the wound site (see Figure 3.13f). This is the result of the environment decomposition with basis functions that have local support, the

¹ Leukotriene B4 is a lipid mediator secreted by various types of leukocytes and identified as a chemotactic agent.

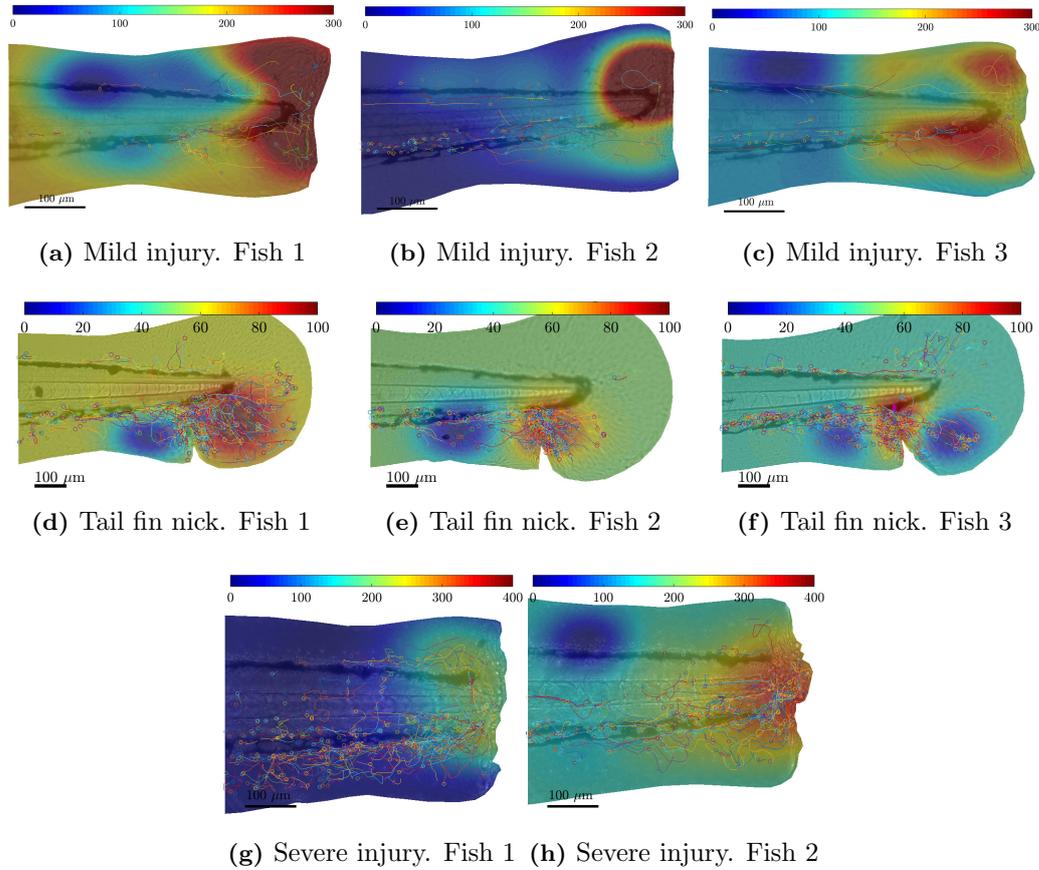


Figure 3.13: Inferred chemoattractant field for various wound types. (a-c) The mild injury. Colour bars are normalised to scale 0-300 a.u.. (d-f) The tail fin nick injury. Colour bars are normalised to scale 0-100 a.u.. (g-h) The severe injury. Colour bars are normalised to scale 0-400 a.u..

effect of which will be further investigated in Chapter 4. Despite individual outliers in the MLEs of scaling parameters, the framework successfully identifies the highest concentration in the tail fin nick injury without any additional information about the field shape. This proves that the selected environment model enhances the flexibility of the estimation framework.

The estimation results for the severe injury dataset also demonstrate smaller relative difference between the ChA concentration near the wound and at the fish body. This can be explained by the fact that the larger wound leads to the intensified generation of the pro-inflammatory mediators that rapidly diffuse away from the wound. Since the model parameter estimates obtained from the developed framework are only unique up to an additive constant, the magnitude of the global concentration field may be much higher but will have a reduced slope in comparison with the normal injury. The observed cell behaviour appears to support this claim: instead of rushing towards the wound, a large proportion of the activated neutrophils

remain in the fish body throughout the entire observation period. At the same time, the neutrophil presence is far more abundant than in the normal injury examples, which allows us to speculate that the recruited immune cells encounter the bacterial threat much deeper in the fish body. A clear outlier in the scaling parameter MLEs can be seen in Figure 3.13h at the area above the circulation. Similarly to the mild injury datasets, the magnitude of that area has not been estimated correctly because there are no cell tracks located there.

3.5 Discussion

This chapter presents a quantitative framework for the inference of global chemoattractant driving neutrophil recruitment that is normally not measured in *in vivo* experiments. Although based on several simplifying assumptions about the nature of cell-environment interaction, the proposed framework is the first attempt to solve the problems of the ChA field inference and the cell velocity estimation simultaneously without requiring any prior knowledge about the field shape.

The estimation procedure is derived from the classical EM algorithm with addition of an intermediate step that computes “forecasts” of cell positions that are utilised to solve the maximisation step in closed form. Such approximation is rather *ad hoc*, but its use is justified in this application because the noise-to-signal ratio in video microscopy of cell centroids is very small, meaning that the estimates of the cell positions obtained from the RTS smoother do not change significantly with iterations of the EM algorithm. Since the basis used for the decomposition of the potential field is a polynomial function of the cell position, it is technically possible to compute all expectations involved in the B-spline decomposition using moment generating functions for multivariate random variables [147]. However, explicit computation of high-order moments for every time point in the dataset will increase the computational complexity of the algorithm. Instead, several alternative solutions may be explored in the future, such as linearisation of the SSM via the first order Taylor series expansion or numerical optimisation of the expected log-likelihood.

The estimation results confirm the consensus that the highest concentration of the chemoattractant is located near the wound and show that the inference framework is able to identify complex concentration patterns, such as ones generated by the tail fin nick injury. The statistical analysis of the data from different wound types has revealed net bias in cell velocity along the X-axis, which in most experiments corresponds to the direction towards the injury. This indicates that even though some cells do not appear to maintain the persistent direction of their movement, on average they are drifting towards the wound site. The biased random walk can therefore be considered as an appropriate mathematical model of neutrophil

chemotaxis. However, lack of median shift in velocity distributions indicates that describing the dynamics of each neutrophil in the observed population by a single SSM implies that the all recruited cells uniformly respond to the environment. In reality, neutrophils are known to alternate between the directed motion to random motion that may be directed away from the wound. It is speculated in the literature that these “switches” are the result of neutrophil desensitisation that allows them discriminate between various sources of chemoattractant [148]. This complex behaviour requires a more detailed mathematical description. Furthermore, the model considered here has a pre-defined structure, where all dynamic matrices are known. Including the estimation of model parameters such as process noise variance σ_w and the rate of the reversion to mean ρ will improve the flexibility of the developed framework.

3.6 Summary

This chapter introduces an approximate Maximum Likelihood framework for the chemoattractant environment inference from neutrophil migration patterns observed *in vivo*. A novel dynamical model of neutrophil migration is proposed that builds on the potential field model of interaction between migrating cells and the hidden underlying environment. The global hidden environment is viewed as an inverse potential field in which the cells move up the gradient. It is parametrised via the decomposition with cardinal tensor product B-splines, such that the scaling parameters of individual B-splines are the unknown parameters of the model. Based on the assumption that cell movement is driven exclusively by the field gradient, the ML framework utilises the proposed model to process cell positions tracked *in vivo* to infer the chemoattractant environment and cell velocities simultaneously.

The performance of the framework and state estimation algorithms is demonstrated on MC simulation examples. Then the framework is applied to several datasets of neutrophil recruitment to the tail injury site observed in a transgenic zebrafish larvae. The estimation results conform with the established theory that neutrophils are driven up the ChA gradient to the wound. The environment inference results for several types of tail fin injury demonstrate that the steepness of the chemoattractant concentration field differs for wounds of varying severity, however it is impossible to determine whether this change in perception is the result of inhibited neutrophil sensitivity or the change in the concentration itself. Results for all experiments are obtained with the same field parametrisation and support the claim that the framework does not require any *a priori* knowledge about the shape of the environment.

A number of extensions for the proposed solution can be considered. For in-

stance, changes in the behaviour of an individual migrating cell fall beyond the scope of the linear state space representation. The model of neutrophil dynamics considered here implies that all cells are responding to the environment all the time. A more complex hybrid model can be employed to represent experimentally observed heterogeneous behaviour of migrating neutrophils, although lack of exact methods for state estimation of hybrid systems would significantly increase the complexity of the resulting algorithm. The following chapter aims to develop the environment estimation framework that would employ such a model.

4 | An approximate Maximum Likelihood framework for estimating the environment driving objects with hybrid dynamics

The inference solution proposed in the previous chapter utilises the state-space formulation of individual object dynamics to recover information about the hidden global environment. A similar problem is often solved as a part of simultaneous localisation and mapping (SLAM), monitoring of population migration, or in detecting the anomalies in crowd monitoring. In most of these applications, the observed object behaviour is prone to abrupt changes that can be caused by sensing faults (in robotics and SLAM), behavioural anomalies (in crowd or population monitoring), or physiological limitations (in migrating cells). For example, a robot solving the SLAM problem may have different sensors and therefore different regimes of interaction with its surroundings. In population monitoring, an individual that moves in a direction different from the majority may be characterised as someone who is not interacting with an artificial underlying environment. Finally, a migrating cell has a limited number of receptors that react to the chemoattractant gradient and may stop interacting with the environment once the receptor capacity has been saturated.

Considering a simple SSM in these cases may lead to errors in the environment estimation. This chapter considers a more complex model of individual object dynamics that can account for heterogeneous behaviour of the observed objects. This is achieved by introducing a discrete Markov chain with finite number of states that correspond to the possible regimes of object-environment interaction. The resultant model is a *jump Markov system (JMS)* that utilises the parametrisation of the hidden environment introduced in §3.2.3, so the unknown parameters incorporated in the input term of the JMS. The field inference problem is then solved via the EM algorithm extended to the case of a hybrid system. Several simulation examples are

presented to assess the performance of the estimation framework.

4.1 Background

The problem of tracking objects interacting with their environment is extensively studied in a variety of fields. Knowledge of object surroundings can be used to improve accuracy of target tracking [149], or to impose constraints in robotics path planning [150], [151]. A common assumption considered in literature is that a model of the environment is available and its parameters are known. However, a number of real-life applications pose an inverse problem of estimating the environment based on observed movement patterns of multiple manoeuvring objects. An example of such problem is presented in the previous chapter, where characterisation of cells migrating under the influence of global stimuli requires simultaneous state estimation of multiple tracked cells and learning model of the environment that governs cell movement.

This chapter considers the problem of environment learning conditioned on the tracking results, that is conceptually similar to the ones addressed in [128] and [152]. The former adopts a deterministic approach, while the latter proposes computationally expensive Bayesian solution based on Rao-Blackwellized particle filter. Both solutions rely heavily on the assumption that dynamics of moving objects is homogeneous, while this chapter deals with a new class of problems where object dynamics varies in time. This chapter introduces an estimation framework that employs a hybrid model that can better reflect the heterogeneous dynamics of moving objects.

Recall a Jump Markov system described by a finite set of hypothesised state space models (SSMs) $\mathbb{M} = \{M^1, M^2, \dots, M^{N_m}\}$, where N_m denotes the total number of models. The model in effect is chosen based on the discrete-valued state \mathbf{m}_t that takes one of the values in the set \mathbb{M} :

$$\mathbf{x}_t = A(\mathbf{m}_t)\mathbf{x}_{t-1} + B(\mathbf{m}_t)\mathbf{u}_{t-1} + G(\mathbf{m}_t)\mathbf{w}_{t-1} \quad (4.1.1a)$$

$$\mathbf{y}_t = C(\mathbf{m}_t)\mathbf{x}_t + \mathbf{v}_t, \quad (4.1.1b)$$

where $\mathbf{x}_t \in \mathbb{R}^{p \times 1}$ is the full state of the system, $\mathbf{u}_t \in \mathbb{R}^{q \times 1}$ is an input vector, and $\mathbf{y}_t \in \mathbb{R}^{r \times 1}$, where $r < p$, is state observation at time t . The state noise $\mathbf{w}_t \sim \mathcal{N}(0, Q_w(\mathbf{m}_t))$, $Q_w(\mathbf{m}_t) \in \mathbb{R}^{p \times p}$ and the measurement noise $\mathbf{v}_t \sim \mathcal{N}(0, R_v(\mathbf{m}_t))$, $R_v \in \mathbb{R}^{r \times r}$ are independent Gaussian sequences. $A(\mathbf{m}_t)$, $B(\mathbf{m}_t)$, $G(\mathbf{m}_t)$, and $C(\mathbf{m}_t)$, $Q_w(\mathbf{m}_t)$, $R_v(\mathbf{m}_t)$ are the system matrices of the model in effect $\mathbf{m}_t \in \mathbb{M}$. In the interest of introduced notation the models hereafter are referred to simply as modes. Jump systems are often employed to represent abrupt changes in the system dynamics that can be caused, amongst other reasons, by compon-

ent failures, measurement faults [153], or sudden changes in the environment [118]. Markovian switching is also relevant in modelling biological systems such as cells that transition through distinct states [154].

The framework developed here deals with multiple discrete-time JMSs interacting with the time-invariant attractive environment. The aim is to estimate parameters of the environment model using knowledge about how it influences objects moving within it. A parametrised model of the hidden environment is obtained via decomposition with overlapping basis functions as shown in Chapter 3. This approach allows the building of complex shapes with a simple grid of identical functions, and thus can be employed in a wide range of applications with varying settings. In the adopted potential field paradigm, influence of the environment on a moving object is incorporated as the input term \mathbf{u}_t^k of each hypothesised model in set \mathcal{M} . This results into a JMS that remains linear with respect to unknown parameters regardless of the mode, which allows one to pose the task of potential function inference as a problem of JMS parameter estimation.

The problem of parameter estimation of JMS is thought to be considered first in [155], where the truncated ML framework is proposed for estimation of the unknown transition probabilities. Another EM solution that relies on variational approximation is suggested in [117], but it only considers switching between the measurement models, while the state transition model remains mode-independent. Following the development of simulation-based methods, various versions of the EM algorithm have been introduced, each utilizing different approximations of the hidden data distribution. For example, parameter learning of JMSs using particle filters is considered in [156]. Other off-line and on-line estimation algorithms based on SMC methods are presented in [157] and [158], respectively.

This chapter is concerned with the case where object states are only partially observed and state sequence of the governing Markov chain is unknown. If the full set of observations is available, both the state sequence and parameters can be estimated off-line within the ML framework via the EM algorithm. However, the joint state-parameter estimation of JMS introduces a challenging problem: intractability of posterior probabilities of hidden data prevents the computation exact state estimates [72]. The most popular approach to state estimation of JMSs is to approximate a mixture model with a single Gaussian at each iteration via the IMM algorithm [121], [159] principle of which in forward time has been previously described in §2.5.5. The backward time recursion of the IMM RTS-type smoother utilised in the developed framework is described in §4.3.3.

The contribution of this chapter is a novel approximate Maximum Likelihood framework that infers a spatially varying hidden global environment from the observed movement of objects interacting with it in different modes. An off-line EM

algorithm for a set of JMSs is derived in order to separate joint state-parameter estimation problem into two coupled tasks and solve them iteratively. An approximate estimation of JMS states is performed during the expectation step in the form of the IMM smoother, which also provides mode probabilities and mode-conditioned state estimates that are instrumental in parameter learning. These quantities are exploited during the maximisation step to compute MLEs of environment model parameters in closed form. The inference framework takes into account all possible mode sequences of each moving object by constructing the log-likelihood function as a weighted sum of mode-conditioned probabilities. No prior knowledge about the environment is required, however, if available, it can be incorporated at the initialisation stage.

4.2 The problem statement

The formalised state-space model considered here consists of four components: a set of SSMs for object dynamics as defined in equation (4.1.1a), a model describing the observation process of object positions in the form (4.1.1b), a Markov chain for modelling mode transitions, and a parametrised model for the environment in which the objects move.

Consider K objects whose dynamics is described by a JMS composed of N_m SSMs

$$\mathbf{x}_t^k = A(\mathbf{m}_t^k)\mathbf{x}_{t-1}^k + B(\mathbf{m}_{t-1}^k)\mathbf{u}_t^k + G(\mathbf{m}_t^k)\mathbf{w}_{t-1}^k \quad (4.2.1a)$$

$$\mathbf{y}_t^k = C\mathbf{x}_t^k + \mathbf{v}_t^k, \quad (4.2.1b)$$

where the state \mathbf{x}_t^k consists of the spatial position and velocity projections on two axes

$$\mathbf{x}_t^k = [s_x, s_y, v_x, v_y]^\top,$$

and where the control input vector is defined as the gradient of the hidden environment at the current position of the cell

$$\mathbf{u}_t^k = \mathbf{u}_t^k(s_x, s_y) = \mu \nabla \mathcal{U}(s_x, s_y). \quad (4.2.2)$$

where μ is the chemoattractive sensitivity coefficient arbitrarily set to one. Similarly with the previous chapter, moving objects are influenced by the global spatially varying environment acting as an attractive potential field. Decomposition with bivariate tensor product cubic B-splines provides a parametrised model of the field

$$\mathcal{U}(s_x, s_y) = \mathcal{B}(s_x, s_y)\Theta = \sum_{h=1}^{N_b} \beta_h^4(s_x, s_y)\theta_h, \quad (4.2.3)$$

where $\beta_h^4(s_x, s_y)$ denotes a bivariate tensor product cubic B-spline at position (s_x, s_y) , and θ_h is a corresponding scaling parameter that defines its magnitude. This representation results in each candidate SSM being a linear function of the parameter vector

$$\mathbf{x}_t^k = A(\mathbf{m}_t^k)\mathbf{x}_{t-1}^k + B(\mathbf{m}_t^k)\phi(s_x, s_y)\Theta + G(\mathbf{m}_t^k)\mathbf{w}_{t-1}^k \quad (4.2.4a)$$

$$\mathbf{y}_t^k = C\mathbf{x}_t^k + \mathbf{v}_t^k, \quad (4.2.4b)$$

where $\phi(s_x, s_y) = \mu\nabla\mathcal{B}(s_x, s_y)$ describes the gradient of superposition of basis functions at the current cell location. Consistently with §3.2.3, the matrices in the control input term denote the following

$$\nabla\mathcal{B}(s_x, s_y) = \begin{bmatrix} \frac{\partial\beta_1(s_x, s_y)}{\partial s_x} & \frac{\partial\beta_2(s_x, s_y)}{\partial s_x} & \cdots & \frac{\partial\beta_h(s_x, s_y)}{\partial s_x} & \cdots & \frac{\partial\beta_{N_b}(s_x, s_y)}{\partial s_x} \\ \frac{\partial\beta_1(s_x, s_y)}{\partial s_y} & \frac{\partial\beta_2(s_x, s_y)}{\partial s_y} & \cdots & \frac{\partial\beta_h(s_x, s_y)}{\partial s_y} & \cdots & \frac{\partial\beta_{N_b}(s_x, s_y)}{\partial s_y} \end{bmatrix}, \quad (4.2.5)$$

$$\Theta = [\theta_1, \theta_2, \dots, \theta_h, \dots, \theta_{N_b}]^\top. \quad (4.2.6)$$

The model of an individual object considered in this chapter is defined by the following assumptions.

Assumption 4.2.1: *The mode transition process is left-continuous, i.e. the effect of model m_{t+1}^k starts at time t^+ .*

Assumption 4.2.2: *The initial state of the k -th object $\mathbf{x}_0^k \sim \mathcal{N}(\hat{\mathbf{x}}_0^k, \mathbf{P}_0^k)$ is independent from \mathbf{w}_t^k and \mathbf{v}_t^k for all t .*

Assumption 4.2.3: *The model parameters λ are known*

$$\lambda \triangleq \{A(M^j), B(M^j), G(M^j), Q(M^j), C, R, \Pi, \Phi, \hat{\mathbf{x}}_0^k, \mathbf{P}_0^k \quad \forall k \in \mathbb{K}\}, \quad (4.2.7)$$

where Π and Φ are initial and transition probability matrices, respectively.

Assumption 4.2.4: *The observation process is mode-independent, i.e measurement matrix C is not a function of the discrete state.*

The DAG illustrating casual relationships between states, modes, and observations under the adopted assumptions is presented in Figure 4.1. Transition between the modes is governed by a discrete-time finite homogeneous Markov chain with initial probabilities $\Pi = \{\pi_j \geq 0\}_{j=1}^{N_m}$. The evolution of the given JMS is fully characterised by the following quantities:

- Mode transition probability matrix $\Phi = \{\phi_{lj} \geq 0\}_{l,j=1}^{N_m}$ regulates the switching between states of Markov chain

$$\phi_{lj}(\mathbf{m}_t^k) \triangleq P(\mathbf{m}_t^k = M^j \mid \mathbf{m}_{t-1}^k = M^l), \quad (4.2.8)$$

where $M^j, M^l \in \mathbb{M}$.

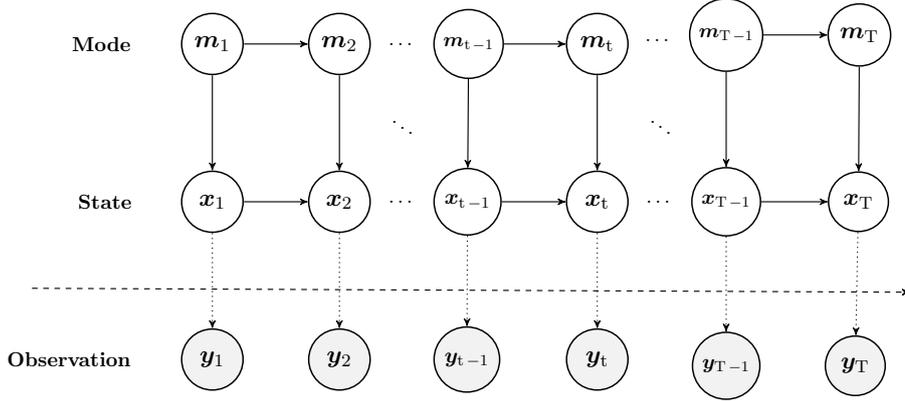


Figure 4.1: A DAG of the JMS considered in this chapter illustrates conditional dependence between variables. The observation \mathbf{y}_t depends only on the state \mathbf{x}_t , while state update is conditioned by the current system mode \mathbf{m}_t and the previous state \mathbf{x}_{t-1} . Note that in the system (4.2.1b) this relation is reflected in mode-independent measurement parameters $C(\mathbf{m}_t) = C = \text{const}$, $R_v(\mathbf{m}_t) = R_v = \text{const}$.

- Probability distribution $\Psi = \{\psi_j(\mathbf{x}_t^k)\}_{j=1}^{N_m}$ of arriving at state \mathbf{x}_t^k is conditioned on the previous state \mathbf{x}_{t-1}^k , and mode \mathbf{m}_t^k

$$\psi_j(\mathbf{x}_t^k) \triangleq p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \mathbf{m}_t^k = M^j, \Theta). \quad (4.2.9)$$

- Emission probability distribution $\Gamma = \{\gamma_j(\mathbf{y}_t^k)\}_{j=1}^{N_m}$ of observing \mathbf{y}_t^k is conditioned on the current state

$$\gamma_j(\mathbf{y}_t^k) \triangleq p(\mathbf{y}_t^k | \mathbf{x}_t^k). \quad (4.2.10)$$

For each object whose dynamics is described by equation (4.2.1a) the state \mathbf{x}_t^k is partially observed and the mode \mathbf{m}_t^k is unobserved. Therefore, the hidden data set includes full sequences of all states and modes $\{\mathcal{X}, \mathcal{M}\}$. A set of observations $\mathcal{Y} = \{\mathbf{y}^k\}_{k=1}^K$ is available for K objects.

Given the model defined above, the ML estimation of environment model parameters relies on distributions (4.2.8)-(4.2.10). For derivation of the EM-based estimation algorithm we consider the following estimation problems:

- *Discrete-valued state estimation*, where the probability of each mode in the set \mathcal{M} is computed given parameters Θ and the full set of observations \mathcal{Y} .
- *Continuous-valued state estimation*, where each hidden state in the set \mathcal{X} is estimated given parameters Θ and the full set of observations \mathcal{Y} .
- *Parameter inference*, where unknown parameters Θ are estimated given the complete data set $\mathcal{Z} = \{\mathcal{Y}, \mathcal{X}, \mathcal{M}\}$.

The first two objectives are coupled and should be solved simultaneously prior to the parameter inference.

4.3 The estimation framework

In the EM algorithm, MLEs of unknown parameters can be obtained by maximising the log-likelihood function of the complete data set \mathcal{Z} . For the system of interest the log-likelihood function is given by

$$\log \mathcal{L}(\Theta) = \log p(\mathcal{M}, \mathcal{X}, \mathcal{Y} | \Theta), \quad (4.3.1)$$

where \mathcal{Y} is the incomplete data and where sets \mathcal{M} and \mathcal{X} constitute the hidden data. One can obtain the lower bound on $\log \mathcal{L}(\Theta)$ by introducing the marginal distribution of the hidden data

$$q(\mathcal{M}, \mathcal{X}) \triangleq p(\mathcal{M}, \mathcal{X} | \mathcal{Y}, \hat{\Theta}^i) > 0,$$

where $\hat{\Theta}^i$ is the current parameter estimate [117]. Then (4.3.1) takes the following form

$$\log \mathcal{L}(\Theta) = \log \int_{\mathcal{M}, \mathcal{X}} q(\mathcal{M}, \mathcal{X}) \frac{p(\mathcal{M}, \mathcal{X}, \mathcal{Y} | \Theta)}{q(\mathcal{M}, \mathcal{X})} d\mathcal{M}, \mathcal{X}. \quad (4.3.2)$$

Using Jensen's inequality (B.2) one can establish

$$\log \mathcal{L}(\Theta) \geq \int_{\mathcal{M}, \mathcal{X}} q(\mathcal{M}, \mathcal{X}) \log \frac{p(\mathcal{M}, \mathcal{X}, \mathcal{Y} | \Theta)}{q(\mathcal{M}, \mathcal{X})} d(\mathcal{M}, \mathcal{X}). \quad (4.3.3)$$

Regroup terms:

$$\begin{aligned} \log \mathcal{L}(\Theta) &\geq \int_{\mathcal{M}, \mathcal{X}} q(\mathcal{M}, \mathcal{X}) \log p(\mathcal{M}, \mathcal{X}, \mathcal{Y} | \Theta) d(\mathcal{M}, \mathcal{X}) \\ &\quad - \int_{\mathcal{M}, \mathcal{X}} q(\mathcal{M}, \mathcal{X}) \log q(\mathcal{M}, \mathcal{X}) d(\mathcal{M}, \mathcal{X}), \end{aligned} \quad (4.3.4)$$

where the second term is not a function of Θ and can be disregarded on the maximisation step. The first term defines the \mathcal{Q} -function, the lower bound of log-likelihood that must be evaluated

$$\mathcal{Q}(\Theta, \hat{\Theta}^i) = \mathbb{E}_{\mathcal{M}, \mathcal{X}} [\log p(\mathcal{Y}, \mathcal{X}, \mathcal{M} | \Theta) | \mathcal{Y}, \hat{\Theta}^i], \quad (4.3.5)$$

where $\mathbb{E}[\cdot]$ denotes the expected value of the function. In case of a hybrid system, the lower bound is intractable because the number of possible mode sequences grows exponentially with time. A way to circumvent this problem is to approximate (4.3.5) using the law of iterated expectations defined in Appendix B.1:

$$\mathcal{Q}(\Theta, \hat{\Theta}^i) = \mathbb{E}_{\mathcal{M} | \mathcal{Y}, \hat{\Theta}^i} [\mathbb{E}_{\mathcal{X} | \mathcal{M}, \mathcal{Y}, \hat{\Theta}^i} \log p(\mathcal{Y}, \mathcal{X}, \mathcal{M} | \Theta) | \mathcal{Y}, \hat{\Theta}^i], \quad (4.3.6)$$

where the inner expectation is over the continuous-valued state conditioned on the mode sequence, and the outer expectation is over the mode sequence conditioned on the observation set [160], [161]. Equation (4.3.6) constitutes the expectation step of the algorithm. MLEs of unknown parameters are then obtained through the maximisation of computed expectation:

$$\hat{\Theta}^i = \arg \max_{\Theta} \mathcal{Q}(\Theta, \hat{\Theta}^{i-1}). \quad (4.3.7)$$

The two steps are iterated until convergence.

4.3.1 The likelihood function

For the model described above there exists a complete data set $\mathcal{Z} = \{\mathcal{M}, \mathcal{X}, \mathcal{Y}\}$ with joint probability density across all sampling points for K objects given by

$$p(\mathcal{Z} | \Theta) = \prod_{k=1}^K \left[\pi_k p(\mathbf{x}_0^k | \mathbf{m}_0^k) \prod_{t=1}^{T_k} p(\mathbf{m}_t^k | \mathbf{m}_{t-1}^k) \times \right. \\ \left. \times \prod_{t=1}^{T_k} p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \mathbf{m}_t^k, \Theta) \prod_{t=1}^{T_k} \prod_{j=1}^{T_k} p(\mathbf{y}_t^k | \mathbf{x}_t^k) \right], \quad (4.3.8)$$

where all individual *pdfs* are defined in the section above. Because of the specific model structure, only the state update *pdf* depends on the unknown parameter vector, which significantly simplifies evaluation of the expected log likelihood.

Lemma 4.3.1: *Suppose assumptions 4.2.1-4.2.4 hold, then the \mathcal{Q} -function for K JMSs satisfies*

$$\mathcal{Q}(\Theta, \hat{\Theta}^i) = \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \mathbb{E}[\log p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \mathbf{m}_t^k, \Theta) | \mathbf{m}_t^k, \mathbf{y}_t^k, \hat{\Theta}^i], \quad (4.3.9)$$

where $\mu_t^{k,j}$ denotes the mode association probability of the k -th object at time t conditioned on the observation sequence for the given object \mathbf{y}^k

$$\mu_t^{k,j} \triangleq P(\mathbf{m}_t^k = M^j | \mathbf{y}^k, \Theta) \geq 0. \quad (4.3.10)$$

Proof. Expanding the lower bound (4.3.6) for each moving object gives a superposition of simpler individual *pdfs*

$$\mathcal{Q}^k(\Theta, \hat{\Theta}^i) = \sum_{\mathbf{m}^k} \int_{\mathbf{x}^k} q(\mathbf{m}^k, \mathbf{x}^k) \log p(\mathbf{m}^k, \mathbf{x}^k, \mathbf{y}^k) d\mathbf{x}^k, \quad (4.3.11)$$

where the joint marginal distribution of the hidden data for an individual object is described by

$$q(\mathbf{m}^k, \mathbf{x}^k) = p(\mathbf{m}^k, \mathbf{x}^k | \mathbf{y}^k, \hat{\Theta}^i) \quad (4.3.12)$$

and where the sum $\sum_{\mathbf{m}^k} [\cdot]$ is over the mode sequence and the integral $\int_{\mathbf{x}^k} [\cdot]$ is over the state sequence of the k -th object. Then the lower bound of the joint log-likelihood function for K objects is a superposition of individual \mathcal{Q} -functions (4.3.11):

$$\mathcal{Q}(\Theta, \hat{\Theta}^i) = \sum_{k=1}^K \mathcal{Q}^k(\Theta, \hat{\Theta}^i) = \sum_{k=1}^K \left[\sum_{\mathbf{m}^k} \int_{\mathbf{x}^k} q(\mathbf{m}^k, \mathbf{x}^k) \log p(\mathbf{m}^k, \mathbf{x}^k, \mathbf{y}^k) d\mathbf{x}^k \right]. \quad (4.3.13)$$

From substituting the complete data *pdf* in (4.3.13) by (4.3.8) arises the following expression

$$\begin{aligned} \mathcal{Q}(\Theta, \hat{\Theta}^i) &= \sum_{k=1}^K \left[\sum_{\mathbf{m}^k} \int_{\mathbf{x}^k} \left\{ q(\mathbf{m}^k, \mathbf{x}^k) \log \pi_j \right\} d\mathbf{x}^k + \right. \\ &+ \sum_{\mathbf{m}^k} \int_{\mathbf{x}^k} \left\{ q(\mathbf{m}^k, \mathbf{x}^k) \log \psi_j(\mathbf{x}_0^k) \right\} d\mathbf{x}^k + \sum_{\mathbf{m}^k} \int_{\mathbf{x}^k} \left\{ q(\mathbf{m}^k, \mathbf{x}^k) \sum_{t=1}^{T_k} \log \varphi_{lj}(\mathbf{m}_t^k) \right\} d\mathbf{x}^k + \\ &+ \left. \sum_{\mathbf{m}^k} \int_{\mathbf{x}^k} \left\{ q(\mathbf{m}^k, \mathbf{x}^k) \sum_{t=1}^{T_k} \log \psi_j(\mathbf{x}_t^k) \right\} d\mathbf{x}^k + \sum_{\mathbf{m}^k} \int_{\mathbf{x}^k} \left\{ q(\mathbf{m}^k, \mathbf{x}^k) \sum_{t=1}^{T_k} \log \gamma_j(\mathbf{y}_t^k) \right\} d\mathbf{x}^k \right], \end{aligned} \quad (4.3.14)$$

Marginalising all unused variables in the hidden data *pdf* renders

$$\begin{aligned} \mathcal{Q}(\Theta, \hat{\Theta}^i) &= \sum_{k=1}^K \left[\sum_{\mathbf{m}_0^k} \left\{ q(\mathbf{m}_0^k) \log \pi_j \right\} + \sum_{\mathbf{m}_0^k} \int_{\mathbf{x}_0^k} \left\{ q(\mathbf{x}_0^k, \mathbf{m}_0^k) \log \psi_j(\mathbf{x}_0^k) \right\} d\mathbf{x}_0^k + \right. \\ &+ \sum_{t=1}^{T_k} \sum_{\mathbf{m}_t^k} \sum_{\mathbf{m}_{t-1}^k} \left\{ q(\mathbf{m}_t^k, \mathbf{m}_{t-1}^k) \log \varphi_{lj}(\mathbf{m}_t^k) \right\} + \sum_{t=1}^{T_k} \int_{\mathbf{x}_t^k} \left\{ q(\mathbf{x}_t^k) \log \gamma_j(\mathbf{y}_t^k) \right\} d\mathbf{x}_t^k + \\ &+ \left. \sum_{t=1}^{T_k} \sum_{\mathbf{m}_t^k} \int_{\mathbf{x}_t^k} \int_{\mathbf{x}_{t-1}^k} \left\{ q(\mathbf{x}_t^k, \mathbf{x}_{t-1}^k, \mathbf{m}_t^k) \log \psi_j(\mathbf{x}_t^k) \right\} d\mathbf{x}_t^k d\mathbf{x}_{t-1}^k \right], \end{aligned} \quad (4.3.15)$$

where $q(\cdot) = p(\cdot \mid \mathbf{y}^k, \hat{\Theta}^i)$ denotes the joint marginalised *pdf* of hidden variables. Given the definition of probability distributions (4.2.8)-(4.2.10) for the given system, only one term is a function of unknown parameters. Substituting the state update probability with (4.2.9) and expanding summation over the mode index leads to the \mathcal{Q} -function in the following form

$$\begin{aligned} \mathcal{Q}(\Theta, \hat{\Theta}^i) &= \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \int_{\mathbf{x}_t^k} \int_{\mathbf{x}_{t-1}^k} \left\{ q(\mathbf{x}_t^k, \mathbf{x}_{t-1}^k, \mathbf{m}_t^k) \times \right. \\ &\quad \left. \times \log p(\mathbf{x}_t^k \mid \mathbf{x}_{t-1}^k, \mathbf{m}_t^k, \Theta) \mid \mathbf{m}_t^k \Theta \right\} d\mathbf{x}_t^k d\mathbf{x}_{t-1}^k + c, \end{aligned} \quad (4.3.16)$$

where c denotes the sum of all terms that are independent of Θ . The joint probability density of hidden data can be partitioned as

$$\begin{aligned} q(\mathbf{x}_t^k, \mathbf{x}_{t-1}^k, \mathbf{m}_t^k) &= p(\mathbf{x}_t^k, \mathbf{x}_{t-1}^k, \mathbf{m}_t^k \mid \mathbf{y}^k, \hat{\Theta}^i) = \\ &= \mu_t^{k,j} p(\mathbf{x}_t^k, \mathbf{x}_{t-1}^k \mid \mathbf{m}_t^k, \mathbf{y}^k, \hat{\Theta}^i), \end{aligned} \quad (4.3.17)$$

Noting that the mode association probability (4.3.10) is not a function of \mathbf{x} , we can rearrange terms in (4.3.16) as follows

$$\begin{aligned} \mathcal{Q}(\Theta, \hat{\Theta}^i) &= \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \int_{\mathbf{x}_t^k} \int_{\mathbf{x}_{t-1}^k} \left\{ \log p(\mathbf{x}_t^k \mid \mathbf{x}_{t-1}^k, \mathbf{m}_t^k, \Theta) \mid \mathbf{m}_t^k, \Theta \right\} \times \\ &\quad \times p(\mathbf{x}_t^k, \mathbf{x}_{t-1}^k \mid \mathbf{m}_t^k, \mathbf{y}^k, \hat{\Theta}^i) \Big\} d\mathbf{x}_t^k d\mathbf{x}_{t-1}^k + c. \end{aligned} \quad (4.3.18)$$

The expression (4.3.18) conforms with the expectation of a distribution of two continuous random variables with a given joint probability density is defined in B.2 and renders a simplified form of the \mathcal{Q} -function (4.3.9), completing the proof. ■

4.3.2 The maximisation step

Suppose that all constituent terms of the joint probability density function are computed during the E-step.

Assumption 4.3.1: *The control input term φ_{t-1}^k is approximated prior to the maximisation step $\forall t, k$ by computing position forecasts as explained in §3.3.2.*

The above assumption makes use of the batch structure of the proposed algorithm and will be addressed in §4.3.4. Then (4.3.9) becomes a deterministic function of the parameter vector Θ , and an algebraic solution can be obtained.

Lemma 4.3.2: *The MLE of the unknown parameter vector $\hat{\Theta}^{i+1}$ that maximises the \mathcal{Q} -function is given by*

$$\begin{aligned} \hat{\Theta}^{i+1} = & \left(\sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \left\{ (\varphi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j B^j \varphi_{t-1}^k \right\} \right)^{-1} \times \\ & \times \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \left\{ (\varphi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j \mathbb{E}[\mathbf{x}_t^{k,j}] - \right. \\ & \left. - (\varphi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j A^j \mathbb{E}[\mathbf{x}_{t-1}^k] \right\} \end{aligned} \quad (4.3.19)$$

where

$$\Sigma_\omega^j \triangleq \{(G^j)^\dagger\}^\top (Q^j)^{-1} (G^j)^\dagger, \quad (4.3.20)$$

$$\mathbb{E}_{\mathbf{m}_t^k = M^j} [\mathbf{x}_{t-1}^k] \triangleq \mathbb{E}[\mathbf{x}_{t-1}^k \mid \mathbf{m}_t^k = M^j, \mathbf{y}^k, \hat{\Theta}^i], \quad (4.3.21)$$

$$\mathbb{E}_{\mathbf{m}_t^k = M^j} [\mathbf{x}_t^{k,j}] \triangleq \mathbb{E}[\mathbf{x}_t^k \mid \mathbf{m}_t^k = M^j, \mathbf{y}^k, \hat{\Theta}^i]. \quad (4.3.22)$$

Proof. See Appendix A. ■

4.3.3 The expectation step

Maximisation of the \mathcal{Q} -function in the form (4.3.9) requires the calculation of mode associated probabilities, $\mu_t^{k,j}$, $j = 1, \dots, N_m$, and conditional expectations of system state at two successive steps, $\mathbb{E}_{\mathbf{m}_t^k = M^j} [\mathbf{x}_{t-1}^k]$ and $\mathbb{E}_{\mathbf{m}_t^k = M^j} [\mathbf{x}_t^k]$. Adopted representation of the object dynamics is suitable for the state estimation via multiple model algorithms discussed in §2.5.5. Here the desired quantities are obtained by running the IMM filter for the forward-time recursion and the IMM smoother for the backward-time recursion. The filtering algorithm has been described in detail in

§2.5.5, and the smoothing procedure for the k -th object is summarised in Algorithm 4.1.

Assumption 4.2.1 implies independence of \mathbf{x}_{t-1}^k and \mathbf{m}_t^k . Then the expectations defined by (4.3.21) and (4.3.22) correspond to the following

$$\mathbb{E}[\mathbf{x}_{t-1}^k | \mathbf{m}_t^k = M^j, \mathbf{y}^k, \hat{\Theta}^i] = \mathbb{E}[\mathbf{x}_{t-1}^k | \mathbf{y}^k, \hat{\Theta}^i] = \hat{\mathbf{x}}_{t-1|T_k}^{k,s}; \quad (4.3.23)$$

$$\mathbb{E}[\mathbf{x}_t^k | \mathbf{m}_t^k = M^j, \mathbf{y}^k, \hat{\Theta}^i] = \hat{\mathbf{x}}_{t|T_k}^{k,j}, \quad (4.3.24)$$

where $\hat{\mathbf{x}}_{t-1|T_k}^{k,s}$ denotes the smoothed state estimate, calculated at the final step of the Algorithm 4.1, and $\hat{\mathbf{x}}_{t|T_k}^{k,j}$ is a mode-conditioned smoothed estimate, produced by the j -th mode-matched smoother.

Algorithm 4.1 Backward recursion of the IMM RTS smoother

Input: Sequence of mode-conditioned state estimates, $\hat{\mathbf{x}}_{0:T}^j$, and their covariances, $\mathbf{P}_{0:T}^j$; sequence of mode probabilities produced by IMM filter, $\mu_{0:T}^j$; hypothesised models, \mathcal{M} ; initial model probabilities, Π ; mode transition probabilities, Φ .

Output: Sequence of smoothed state estimates, $\hat{\mathbf{x}}_{0:T|T}$, and their covariances, $\mathbf{P}_{0:T|T}$; sequence of mode-conditioned smoothed state estimates, $\hat{\mathbf{x}}_{0:T}^j$, and their covariances, $\mathbf{P}_{0:T}^j$; estimated model probabilities, $\hat{\mu}_{0:T}$.

1: Initialise the algorithm with $\mu_{T|T}^j, \hat{\mathbf{x}}_{T|T}^j, \mathbf{P}_{T|T}^j \quad j = 1, \dots, N_m$;

2: **for** $t \leftarrow T-1, 0$ **do**

3: **for** $j \leftarrow 1, N_m$ **do**

4: Calculate mixing probabilities $\mu_{t+1|T}^{0j} = \frac{1}{d_j} \phi_{ij} \mu_{t+1|T}^j$, where $d_j = \sum_{i=1}^{N_m} \phi_{ij} \mu_{t+1|T}^i$;

5: Compute mixed initial conditions:

$$\begin{aligned} \hat{\mathbf{x}}_{t+1|T}^{0j} &= \sum_{j=1}^{N_m} \mu_{t+1|T}^{0j} \hat{\mathbf{x}}_{t+1|T}^j \\ \mathbf{P}_{t+1|T}^{0j} &= \sum_{j=1}^{N_m} \mu_{t+1|T}^{0j} \left\{ \mathbf{P}_{t+1|T}^{0j} + (\hat{\mathbf{x}}_{t+1|T}^j - \hat{\mathbf{x}}_{t+1|T}^{0j})(\hat{\mathbf{x}}_{t+1|T}^j - \hat{\mathbf{x}}_{t+1|T}^{0j})^\top \right\}; \end{aligned}$$

6: Run RTS smoother for a model M^j to produce $\hat{\mathbf{x}}_{t|T}^{s,j}, \mathbf{P}_{t|T}^j$;

7: **end for**

8: **for** $j \leftarrow 1, N_m$ **do**

9: **for** $i \leftarrow 1, N_m$ **do**

10: Compute mode-conditioned likelihoods $\mathcal{L}_{t+1|T}^{i|j} = \mathcal{N}(\hat{\mathbf{x}}_{t+1|T}^{s,i} - \hat{\mathbf{x}}_{t+1|T}^j, S^{j,i})$;

11: **end for**

12: Compute smoothed mode likelihood $\mathcal{L}_{t+1|T}^j = \sum_{i=1}^{N_m} \phi_{ij} \mathcal{L}_{t+1|T}^{i|j}$;

13: Update mode probability $\mu_{t|T}^{s,j} = \frac{1}{d} \mathcal{L}_{t+1|T}^j \mu_{t|T}^j$, where $d = \sum_{j=1}^{N_m} \mathcal{L}_{t+1|T}^j \mu_{t|T}^j$;

14: **end for**

15: Merge mode-conditioned states and covariances:

$$\begin{aligned} \hat{\mathbf{x}}_{t|T} &= \sum_{j=1}^{N_m} \mu_{t|T}^j \hat{\mathbf{x}}_{t|T}^j \\ \mathbf{P}_{t|T} &= \sum_{j=1}^{N_m} \mu_{t|T}^j \left\{ \mathbf{P}_{t|T}^j + (\hat{\mathbf{x}}_{t|T}^j - \hat{\mathbf{x}}_{t|T})(\hat{\mathbf{x}}_{t|T}^j - \hat{\mathbf{x}}_{t|T})^\top \right\}; \end{aligned}$$

16: **end for**

The implemented IMM algorithm is an RTS-type smoother [126] that does not require that the dynamical model to be reversible in time (unlike the smoother in [123]) and is applicable both linear and non-linear models using either unscented transform or Taylor expansion.

4.3.4 The approximate algorithm

The structure of the environment model adopted in this chapter leads to the hybrid system that is non-linear with respect to the object state. Thus, alongside with merging of mode histories, two additional approximations are utilised to ensure a computationally feasible solution to the problem of joint state-environment estimation. The first approximation is to employ banks of unscented filters and smoothers within the IMM framework during the expectation step. The second approximation is similar to the one introduced in §3.3.2. It is employed to separate the variables in the Q -function prior to the maximisation step by computing “forecasts” the object positions at each time $\mathbf{s}_t^k = [s_x, s_y]^\top$, $t = 1, \dots, T^k, k = 1, \dots, K$. The forecast is treated as an independent variable that is algebraically linked to its expected state

$$\mathbf{s}_t^k = C\hat{\mathbf{x}}_{t|T^k}^k. \quad (4.3.27)$$

Because of the batch nature of the algorithm, all forecasts can be computed after the expectation step. The gradient of the basis grid is then calculated in the forecast positions

$$\hat{\varphi}_t^k = \varphi(\mathbf{s}_t^k), \quad (4.3.28)$$

leading to the analytical solution of the maximisation problem that relies only on the first central moments of the object state at two subsequent time instances.

The resulting algorithm is initialised with computing the gradient of the basis function grid for each measured object position $\hat{\varphi}_t^k = \varphi(\mathbf{y}_t^k)$ and setting $\hat{\Theta} = 0$. Then, the initial velocity estimates are recovered from the observation set \mathcal{Y} by applying the unscented RTS smoother to the model structure that accounts for the object-environment interaction. The initial parameter estimate $\hat{\Theta}^0$ is obtained using the LS-type algorithm under the assumption that all objects were driven by the environment. The framework then iterates between the expectation, forecasting, and maximisation steps until the parameter convergence is achieved:

- i) E-step: compute the expected values of the continuous-valued state \mathbf{s}_t^k and mode probabilities $\mu_t^{k,j}$, $j = 1, \dots, N_m, t = 1, \dots, T^k, k = 1, \dots, K$ via IMM-RTS algorithm.
- ii) Forecasting: compute forecasts of object positions, \mathbf{s}_t^k , and approximate gradients of the basis function grid in those positions, $\hat{\varphi}_t^k$, for $t = 0, \dots, T^k, k = 1, \dots, K$.

iii) M-step: maximise the \mathcal{Q} -function in closed form.

As an ML estimator, the developed algorithm arrives at a point estimate of unknown parameters in a finite number of iterations but does not guarantee that the estimate corresponds to the global maximum of the log-likelihood.

The computational complexity of each iteration is linear with respect to the number of objects and the size of the observation set. By using the IMM framework we ensure that the expectation step will be of order $O(K T_k(N_m + N_m))$ at most, which is beneficial when dealing with large set of objects or long time series for individual objects. The maximisation step is of the same complexity as the LS algorithm, $O(N_b^3)$. The runtime of overall algorithm increases linearly with each iteration.

4.4 Simulation examples

The performance of the proposed algorithm is tested through Monte Carlo simulation of several scenarios inspired by various real-life phenomena of objects interacting with their environment. Similarly with the previous chapter, the acting environment is modelled using a 4×4 grid of overlapping tensor product cubic B-splines $\mathcal{B}(s_x, s_y)$ placed on the map of size 1000×1000 arbitrary units (a.u.). The scaling coefficients and the modelled field have been previously shown in Figure 3.4. Each scenario is simulated a hundred times with different realisations of noise and with the same field parameters. The field drives the motion of a hundred objects whose trajectories are modelled with time increment of 1 min for 100 minutes or until the object leaves the field of influence. Dynamics of each object are modelled by a JMS with a known transition probability matrix. It is assumed that at each time an object can be in one of two modes:

- M^1 : the driven mode in which the object responds to the acting environment;
- M^2 : the desensitized mode in which an object does not perceive the environment, *i.e.* the input gain associated with the acting environment is set to zero.

The dynamics of the k -th object moving in the field are described by the system (4.1.1) with a state vector

$$x_t^k = [s_x, s_y, v_x, v_y]^T,$$

where s_x, s_y denote the spatial position and v_x, v_y denote velocity projections on corresponding axes.

The following SSM matrices are used in the model:

$$A^1 = A^2 = \begin{bmatrix} \mathbb{I}_2 & T\mathbb{I}_2 \\ \mathbb{O}_2 & \mathbb{I}_2 - T\rho\mathbb{I}_2 \end{bmatrix}; \quad (4.4.1a)$$

$$B^1 = \begin{bmatrix} \mathbb{O}_2 & T\mathbb{I}_2 \end{bmatrix}^\top; \quad (4.4.1b)$$

$$B^2 = \begin{bmatrix} \mathbb{O}_2 & \mathbb{O}_2 \end{bmatrix}^\top, \quad (4.4.1c)$$

$$G^1 = G^2 = \begin{bmatrix} \mathbb{O}_2 & T\mathbb{I}_2 \end{bmatrix}^\top, \quad (4.4.1d)$$

$$C = \begin{bmatrix} \mathbb{I}_2 & \mathbb{O}_2 \end{bmatrix}, \quad (4.4.1e)$$

where $T = 1$ min is the time increment and \mathbb{I}_2 is the identity matrix and \mathbb{O}_2 is a zero matrix of size 2×2 . The rate of the reversion to mean has been set to $\rho = 0.3$.

The process noise covariances are set to $Q_w^1 = 2.5\mathbb{I}_2$ a.u.²/min⁴ for the field-driven mode, and $Q_w^2 = 2\mathbb{I}_2$ a.u.²/min⁴ for the desensitized mode, where \mathbb{I}_2 is a 2×2 identity matrix. The measurement noise variance is set to $R_v = \mathbb{I}_2$ a.u.². Initial velocity estimates are arbitrarily assumed to be zero with covariance $\mathbf{P}_0 = \mathbb{I}_2$ a.u.²/min⁴. Initial mode probabilities for the IMM filter are set to $\pi = [0.5, 0.5]$ to ensure complete uncertainty of the initially correct model. The mode transition probability matrix is set to

$$\begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{bmatrix} = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}. \quad (4.4.2)$$

It must be noted that the model parameters have been selected arbitrarily with the purpose to demonstrate the performance of the proposed estimation framework. Several cases of object movement patterns are generated using the hybrid model and the “true” underlying potential function in the interest of testing the estimation framework (see Figure 4.2). The case with uniformly distributed starting points is used to demonstrate the capabilities of the framework and analyse the estimator uncertainty. The rest of the examples are used to assess the quality of the global environment inference from the localised tracking data with presence of various model mismatches.

4.4.1 Experiment I: uniformly distributed starting points

This case of data generation is inspired by a phenomenon of directed cell migration in global chemotactic field. Recall the data described in the previous chapter, where cells tracks can start anywhere on the image as neutrophils enter the tissue from the circulation at multiple locations. The estimation is performed with the assumption that the model structure is known for both cell modes. Two aspects of the estimation procedure are assessed here: the probabilistic statement about the mode of each moving object, and the quality of the field parameter inference.

The results of the IMM estimator from the final iteration of the EM algorithm are presented in Figure 4.3. Shown are the twenty tracks from one Monte-Carlo

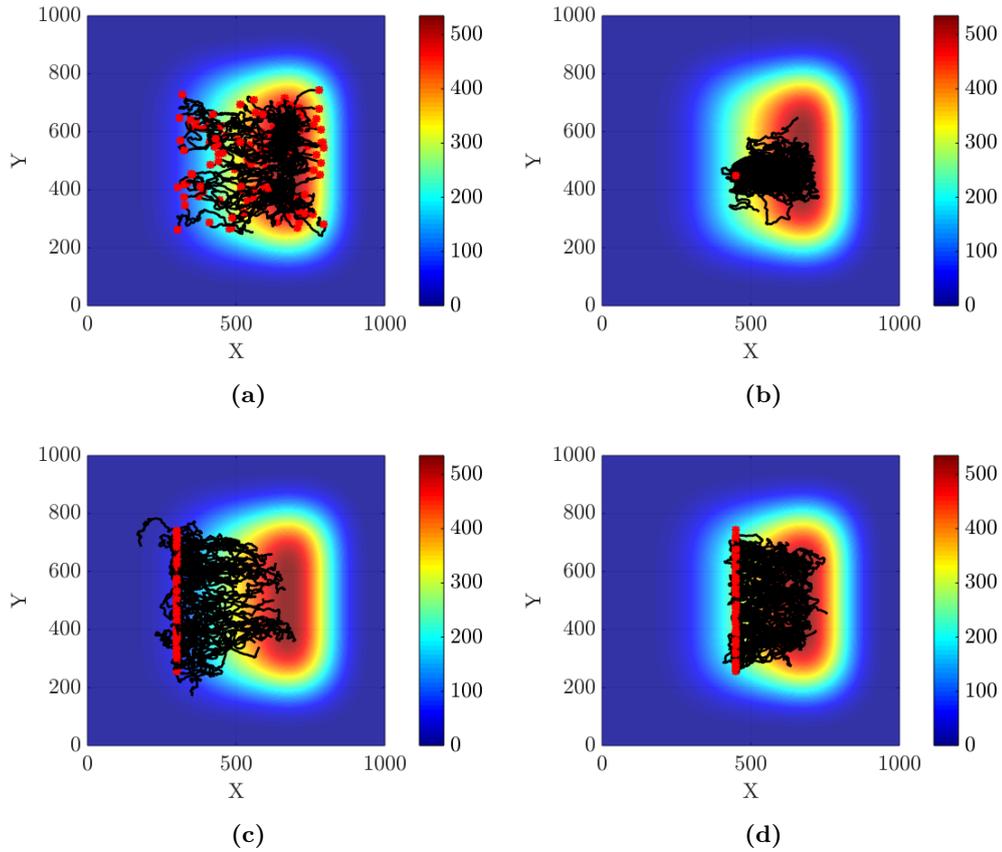


Figure 4.2: Examples of migration patterns generated by the interaction of N hybrid systems with the modelled potential field, where $N = 100$. Colour bars correspond to the magnitude of the constructed field. The starting points for each track is marked by the red asterisk (*). **(a)** Uniformly distributed starting points. **(b)** All objects start at the position $s_x = 450$ a.u., $s_y = 450$ a.u.. **(c)** All objects start at positions with $s_x = 300$ a.u.. **(d)** All objects start at positions with $s_x = 450$ a.u.

realisation plotted on top of the gradient of the estimated field. It can be seen that both IMM filter and smoother correctly identify that the segments of tracks that are aligned with the gradient direction and maintain persistence are more likely to have been generated by the object in the responsive mode M^1 .

The estimation framework is used to process 100 Monte Carlo realisations of the migrating pattern to obtain a sample of 100 parameter vector estimates $\hat{\Theta}$. This sample is then used to compute the statistics of the MLE distribution. The gradient of the potential field constructed with mean estimates of the scaling parameters is presented in Figure 4.3c. While the gradient field appears to be similar to the modelled one, the magnitude of the total bias between the true field and the mean estimate grows towards the right side of the map (see Figure 4.3d), meaning that most of the mean estimates overstate the value of the parameter. The bias in scaling parameter estimates can be caused by two factors. Firstly, the IMM algorithm

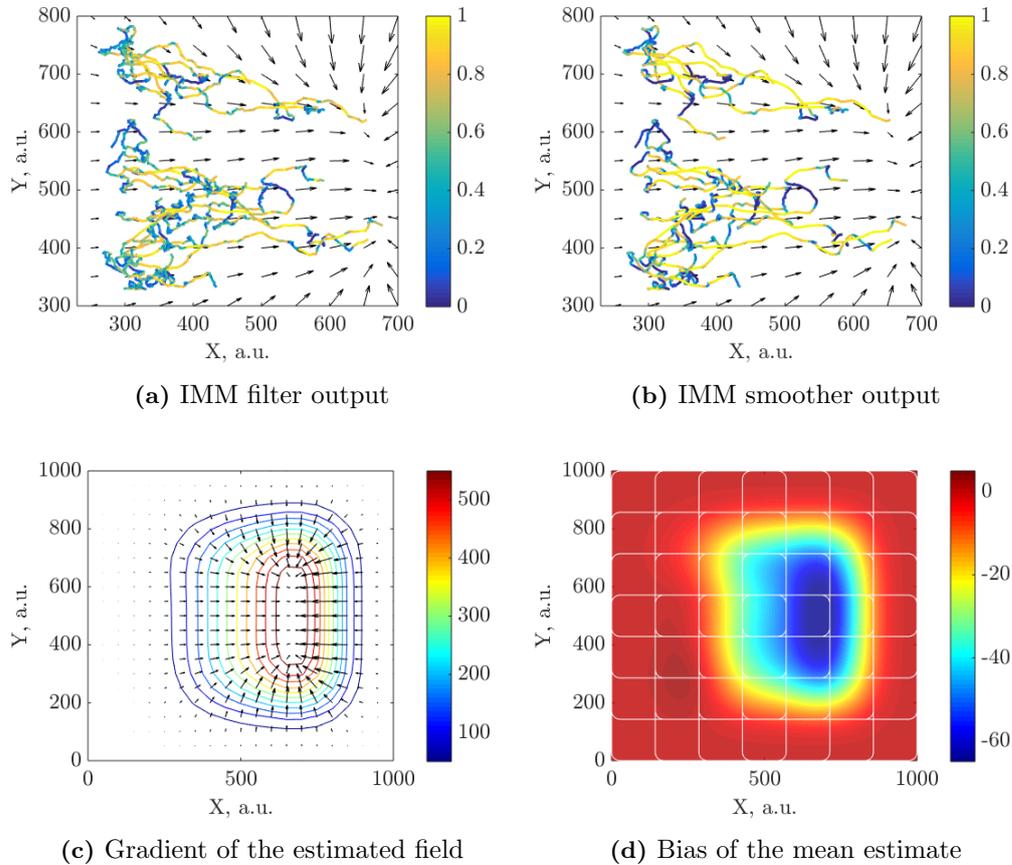


Figure 4.3: Estimation results for Experiment I. **(a)-(b)** Example tracks processed by the IMM algorithm plotted against the gradient of the estimated field from one Monte Carlo realisation. Each segment of an individual track between the neighbouring points s_{t-1} and s_t is colour-coded according to the probability $\mu_t^{k,1}$ of an object being in the responsive mode (M^1) at time t . The colour bars correspond to the probability of an object being in the responsive mode. **(c)** The gradient of the potential field constructed with the mean parameter estimates, $\nabla\mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the constructed field. **(d)** Total bias between the true field and the mean estimate, $\mathcal{U}(\Theta_{\text{true}}) - \mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the total bias.

utilised for the state estimation of the hybrid system may introduce the bias into the velocity estimate during the merging stage; secondly, the selected model of object-parameter interaction incorporates the information about the gradient only and not the magnitude of the underlying field, meaning that the gradient may be inferred to have been produced by the field with the same slope as the true one but with either higher or lower magnitude. The bias of the second type is spatially invariant and can be eliminated by introducing constraints or prior information about the field. In Figure 4.3d, however, the total bias is not constant meaning that the slope of the estimated field is not in full agreement with that of the true field.

The statistics of the parameter MLEs are presented in Table 4.1. It can be seen that standard deviations of individual estimates $\theta_5 - \theta_{16}$ are similar irrespective of the parameter value. The variability of these MLEs is rather small compared to the parameter value, meaning that the estimation is consistent from one MC realisation to another. Both CV and %-bias are the highest for the parameters in the left corners of the map where the true values of scaling parameters are the smallest (θ_1 and θ_4). This can be explained by the fact that fewer tracks pass through the outer area of the potential field, and even those objects whose tracks originate at the left hand side of the map quickly migrate away from that area along the gradient.

Parameter	True	Mean est.	St. dev.	CV (%)	%-bias
θ_1	10	6.026	58.344	968.264	39.744
θ_2	10	8.981	35.181	391.732	10.189
θ_3	10	9.000	27.805	308.935	9.997
θ_4	10	14.843	29.052	195.729	48.428
θ_5	100	114.095	29.795	26.114	14.095
θ_6	100	118.349	18.131	15.320	18.349
θ_7	100	115.129	17.821	15.479	15.129
θ_8	100	120.766	15.990	13.241	20.766
θ_9	190	209.359	20.268	9.681	10.189
θ_{10}	190	211.904	13.764	6.495	11.528
θ_{11}	190	210.064	11.689	5.564	12.322
θ_{12}	190	213.411	14.823	6.945	10.873
θ_{13}	290	321.531	21.109	6.565	10.873
θ_{14}	290	327.303	15.010	4.586	12.863
θ_{15}	290	329.106	13.209	4.014	13.485
θ_{16}	290	318.868	15.298	4.798	9.954

Table 4.1: Estimation results for Experiment I over 100 simulations.

The characteristics presented in the table above are computed using the sample of MLE estimates obtained using the Monte Carlo approach. However, majority of real-life applications do not allow for repetitive experiments, thus other methods to assess the estimation uncertainty are often sought. The advantage of the proposed framework is that it can provide interval estimates of the scaling parameters in the form of confidence regions (see Figure 4.4). As has been discussed in §2.4.2, the orientation and the volume of a N_b -dimensional confidence ellipsoid are defined by the reciprocal of the observed Fisher information that can be computed directly in the EM algorithm as shown in §4.3.4. The eigenvalues of this matrix indicate the variability of the parameter estimates along the principal axes. The provided

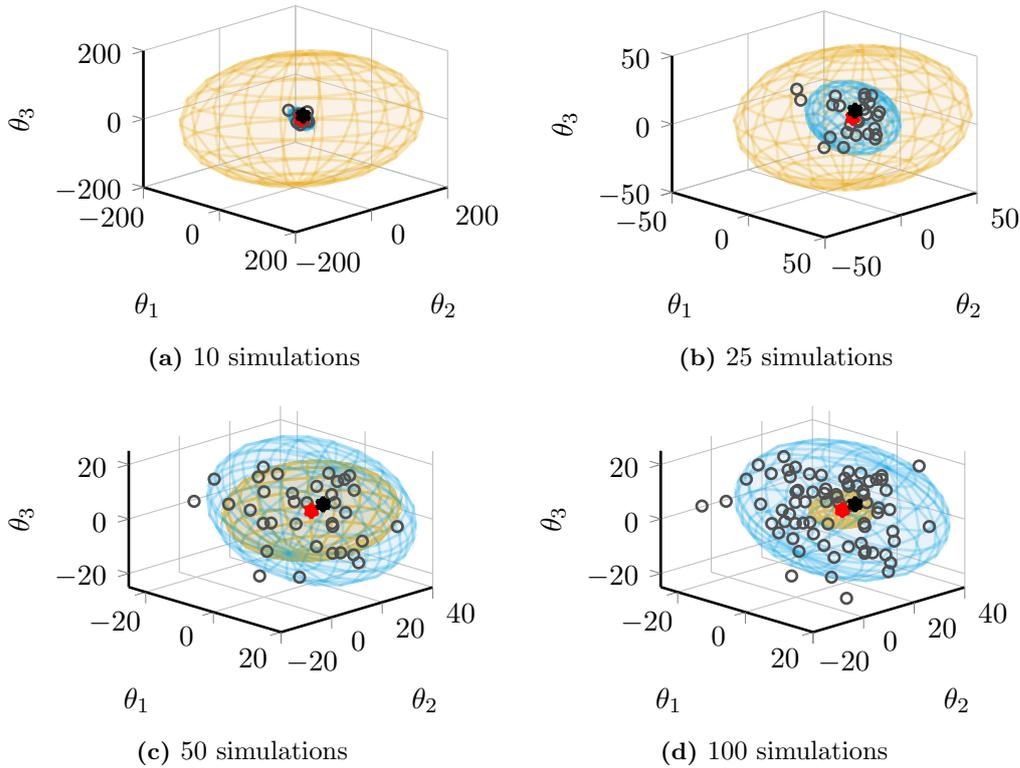


Figure 4.4: Example estimation results and confidence regions. The 95% confidence ellipsoid (orange wireframe) centred at the mean estimate (\bullet) changes volume with decreasing number of experiments but maintains the orientation. The mean estimate is located close to the true parameter value (\blacklozenge). The size of the 95% covariance ellipsoid (blue wireframe) is defined by the dispersion of the estimates (\circ).

examples clearly demonstrate the correspondence between the volume of the confidence region and the size of the observation set that amounts to the number of Monte Carlo simulations. While the variance ellipsoid maintains nearly the same size, the volume of the confidence region decreases dramatically when there is more data available. In further experiments, the volumes of the 95-% confidence regions will be compared to assess the variability of the parameter vector MLEs conditioned on the available tracking data and the selected parametrisation. In addition, the total variance normalised by the length of the parameter vector is provided as a unified measure of MLE scatter.

The rate of convergence of the EM algorithm for a hybrid system is impossible to determine analytically, therefore here it is examined by letting the estimation algorithm run for 10 iterations. The rate of parameter estimate convergence is presented on the log scale in Figure 4.5. It can be seen that the convergence threshold selected for this framework is typically reached within 3-4 iterations.

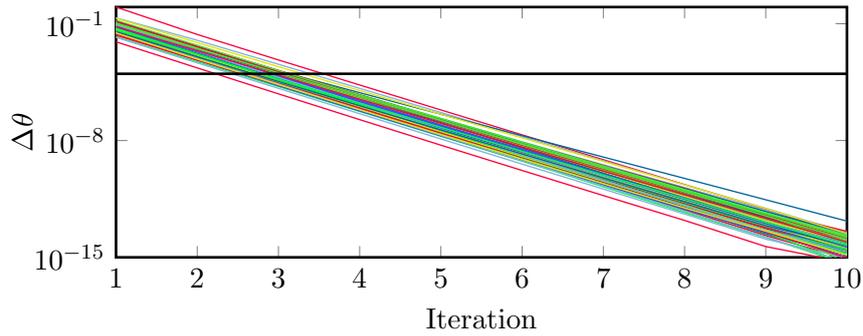


Figure 4.5: Parameter estimate convergence for 100 simulations in log scale. The black line shows the convergence threshold selected for testing the algorithm.

4.4.2 Experiment II: single starting point

This example is relevant to robotics applications where a swarm of robots is deployed at some initial position, as well as to processing some in vitro neutrophil data obtained from experimental assays where cells are released in to the chemical environment via a particular entrance. All objects in this scenario enter the potential field at the same point on the map as shown in Figure 4.2b. The main concern that may arise in this situation is that all object tracks are confined to a certain area and do not inform the estimation framework about the potential function in remote locations on the map. The estimation is performed under the assumption that all the dynamical matrices are known in order to assess how the localisation of the tracking data can affect field estimation results, hence the state estimation results are omitted for this case as they are similar to the previous scenario.

The mean field parameters estimated over 100 realisations of the tracking data are used to reconstruct the gradient field in Figure 4.6a. The magenta point on the map indicates the starting position of all objects, $s_x = 450$, $s_y = 450$. The contour lines show that, apart from the area in the top left corner of the map, the inferred underlying field is of higher magnitude than the true field. The bias of the estimated field from the true field is shown in Figure 4.6b. It can be seen that the bias surface is flat almost everywhere, meaning that the bias in MLEs of θ_5 – θ_{16} can be attributed to the selected initial estimate. On the other hand, the estimate of the parameter corresponding to B-spline β_4 , whose local support is highlighted in Figure 4.6b, is much further from the true value than the rest of the MLEs. This can be explained by the fact that there are no object tracks passing through this area of the map, because the single starting point is located near the centre of the map. Therefore, no information about the contribution of that area of the field to the directed migration of the population can be extracted by the proposed inference framework.

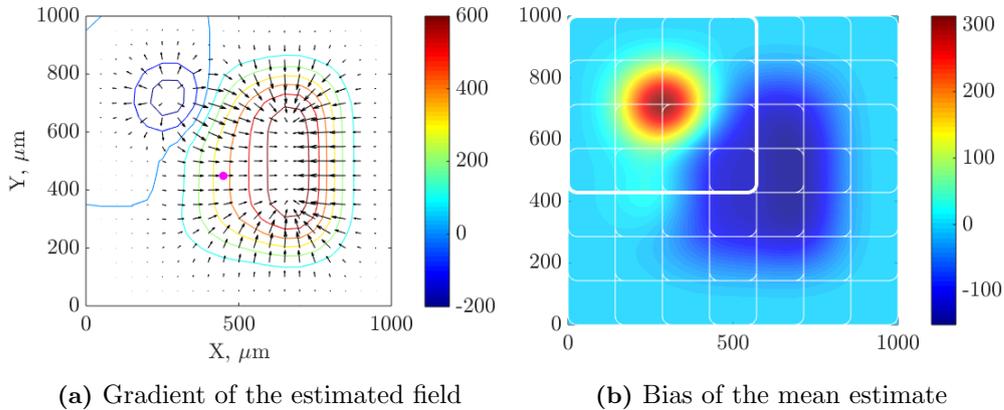


Figure 4.6: Estimation results for Experiment II. **(a)** The gradient of the potential field constructed with the mean parameter estimates, $\nabla\mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the constructed field. The magenta dot (\bullet) indicates the starting point of all generated tracks. **(b)** Total bias between the true field and the mean estimate, $\mathcal{U}(\Theta_{\text{true}}) - \mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the total bias.

The statistics of the selected estimates associated with B-splines in different areas of the map are summarised in Table 4.2. It can be seen that the lack of tracking information leads to loss of precision and increased bias in MLEs of the scaling parameters that correspond to the basis functions located in the left hand side of the map. While the mean estimates of θ_1 and θ_8 are much closer to the true value than θ_4 , their corresponding CV indicate poor estimation precision. This example reveals the drawback of using the field decomposition with local support when processing highly localised data. However, despite the bias in individual MLEs, the inference framework is able to locally reproduce the slope of the potential field that has been used to generate the object tracks.

Parameter	True	Mean est.	St. dev.	CV (%)	%-bias
θ_1	10	21.140	504.673	2387.372	111.402
θ_4	10	-303.403	2937.143	968.056	—
θ_8	100	85.6469	227.501	382.385	14.353
θ_{12}	190	253.982	111.009	43.707	33.67
θ_{16}	290	348.453	106.682	30.616	20.156

Table 4.2: Selected estimation results for Experiment II over 100 simulations.

4.4.3 Experiments III and IV: starting points on the line

Another application that may require simultaneous learning of object behavioural modes and the environment concerns the development of computer vision systems that determine the speeding in the traffic flow, or disruption in the crowd flow in public spaces. The abnormal behaviour in the crowd motion is usually detected as an individual deviation from the crowd treated as a whole population [162]. However, this approach requires first creating a sufficient model for the patterns of the crowd behaviour, and thus only works in situations when the population behaviour is uniform. While these scenarios are common in motorway traffic control, monitoring of the pedestrian flow deals with heterogeneous patterns.

In cases of non-uniform population behaviour the abnormality identification relies on the individual object tracking [163], [164]. For example, detecting an individual crossing the road outside of the designated pedestrian pathway using population models is hindered by the fact that this "abnormality" is a common occurrence and can be identified as a normal pattern during the training process training the computer vision system. Instead, this problem can be interpreted as a problem of identifying objects that are not responding to the artificial global environment. The framework derived here can be used to infer such artificial environment.

The tracking data for these experiments has been generated to replicate the situations when all objects enter the map from one side. Two cases are considered: i) all objects enter the map at the area with the smallest magnitude as demonstrated in Figure 4.2c); ii) all object enter the map at the central area (see Figure 4.2d). Note that the colour bars are not normalised across two cases in order to highlight the magnitude of the bias surface for individual experiments.

The estimation results for both scenarios are presented in Figure 4.7. The gradient of the estimated field in the Experiment III is in good agreement with the true gradient almost everywhere, though its magnitude is inferred to be lower. Spatially constant bias in Figure 4.7b is likely the result of the choice of the initial estimate $\hat{\Theta}^0$. The coefficients corresponding to the right column of B-splines are slightly overstated by the framework but the difference is not significant judging by the %-bias in Table 4.3. The standard deviation for MLEs of $\theta_{13} - \theta_{16}$ is greater than it is for other coefficients because only a limited number of tracks reaches highlighted area of local support in the right hand side of the map before terminating after 100 minutes.

When the object tracks originate in the middle of the map, the first column of B-splines covers the area with a rather limited number of data points (the local support β_1 and β_4 is highlighted in Figure 4.7d). Similarly with Experiment II, MLEs of corresponding scaling coefficients understated the parameter values. However, the

fact that all object tracks originate in that area is enough to sufficiently decrease the variance of the MLEs in comparison with the results from Experiment II. This is evidenced by the reduced CV for $\theta_1 - \theta_4$ in Table 4.4. The rest of the estimates are unbiased and have small variance, which allows us to conclude that the environment is estimated with a high degree of accuracy.

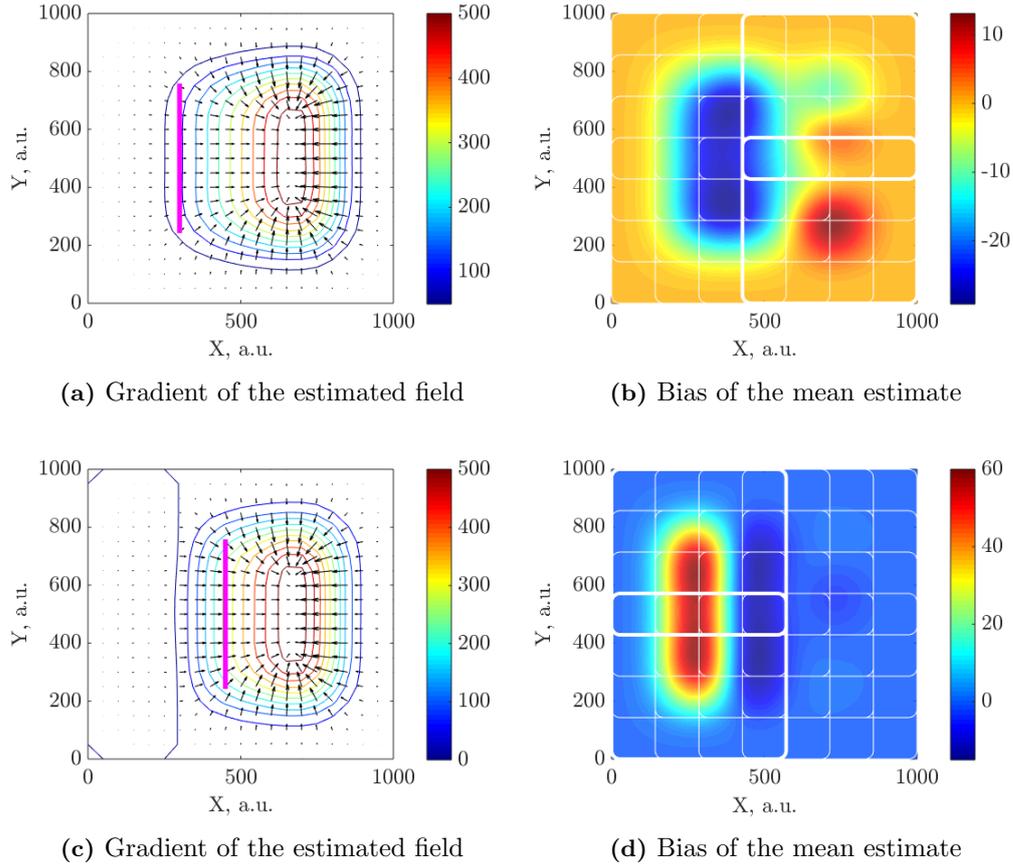


Figure 4.7: Estimation results for Experiments III and IV. (a) The gradient of the potential field constructed with the mean parameter estimates, $\nabla\mathcal{U}(\hat{\Theta})$. The magenta line (—) indicates the starting position of the generated tracks at $s_x = 300$ a.u.. (c) The gradient of the potential field constructed with the mean parameter estimates, $\nabla\mathcal{U}(\hat{\Theta})$. The magenta (—) line indicates the starting position of the generated tracks at $s_x = 450$ a.u.. (b,d) Total bias between the true field and the mean estimate, $\mathcal{U}(\Theta_{\text{true}}) - \mathcal{U}(\hat{\Theta})$. The colour bars correspond to the magnitude of the total bias.

4.4.4 Experiment V: dynamical model mismatch

The simulation examples presented above all rely on the fact that each SSM in the hybrid system modelling object dynamics is fully known. While this may be the case with some engineering applications, in life systems and the crowd monitoring,

Parameter	True	Mean est.	St. dev.	CV (%)	%-bias
θ_1	10	23.903	17.958	75.131	139.024
θ_4	10	21.638	19.385	89.586	116.380
θ_8	100	118.038	15.335	12.992	18.038
θ_{13}	290	274.599	74.027	26.958	5.311
θ_{16}	290	301.069	63.362	21.046	3.817

Table 4.3: Selected estimation results for Experiment III over 100 simulations.

Parameter	True	Mean est.	St. dev.	CV (%)	%-bias
θ_1	10	-35.767	88.752	248.140	457.670
θ_2	10	-30.893	75.748	245.197	408.928
θ_3	10	-30.143	79.534	263.856	401.430
θ_4	10	-37.000	87.530	236.566	470.003
θ_{12}	190	191.018	19.543	10.231	0.535
θ_{16}	290	288.495	27.647	9.583	0.519

Table 4.4: Selected estimation results for Experiment IV over 100 simulations.

the hybrid model is always an approximation of a true, more complex, dynamical system. This experiment examines the sensitivity of the estimation framework to the model mismatch in the object dynamics. The tracks from the experiment III were processed using a hybrid model that differs from the one used to generate them. In the estimation algorithm the responsive mode M^1 is described by the dynamics without reversion to mean ($\rho = 0$). The state transition matrix takes the following form:

$$A^1 = \begin{bmatrix} \mathbb{I}_2 & T \mathbb{I}_2 \\ \mathbb{O}_2 & \mathbb{I}_2 \end{bmatrix}, \quad (4.4.3)$$

while the rest of the model parameters remain the same. The estimation results are presented in Figure 4.8. The IMM estimation results reveal the apparent reason of poor estimation field estimation results. The IMM framework identifies the objects moving against the gradient to be in the responsive mode, which results in poor estimation accuracy of the underlying field. The total bias increases sharply in the right hand side area of the map where the true field magnitude is the largest.

Statistics of the selected parameter estimates presented in Table 4.5 demonstrate the severe effect of the dynamical model mismatch on the estimation results. It can be seen that the standard deviation of all parameters is small in comparison to previous experiments. In this case it small variance means that the estimation bias is consistent for all Monte Carlo realisations. The percentage bias is also the highest compared to previous experiments. This example demonstrates the consequences of

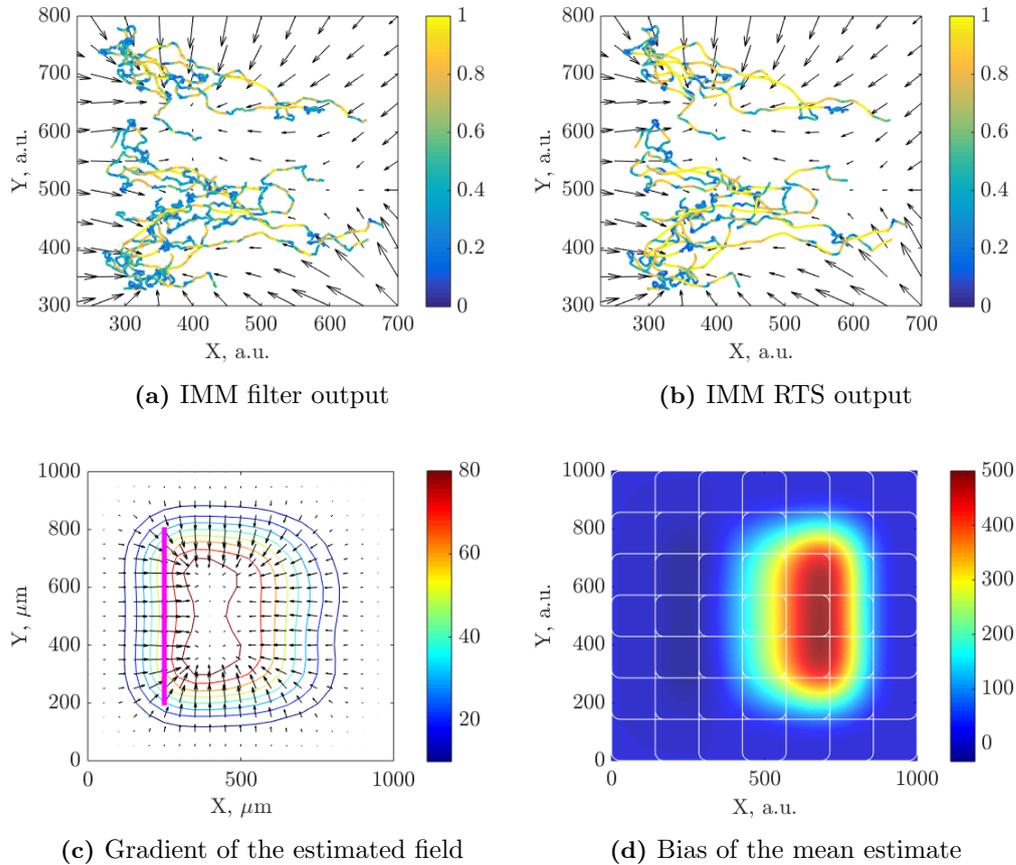


Figure 4.8: Estimation results w for Experiment V. **(a)-(b)** Example tracks processed by the IMM algorithm plotted against the gradient of the estimated field from one Monte Carlo realisation. Each segment of an individual track between the neighbouring points s_{t-1} and s_t is colour-coded according to the probability $\mu_t^{k,1}$ of an object being in the responsive mode (M^1) at time t . The colour bars correspond to the probability of an object being in the responsive mode. **(c)** The gradient of the potential field constructed with the mean parameter estimates, $\nabla\mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the constructed field. The magenta line (—) indicates the starting point of all generated tracks. **(d)** Total bias between the true field and the mean estimate, $\mathcal{U}(\Theta_{\text{true}}) - \mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the total bias.

field estimation with the wrong model structure.

4.4.5 Experiment VI: field model mismatch

One of the advantages of the proposed framework that has been pointed out earlier is that it does not require any knowledge about the structure of the underlying potential field. A way to test this claim is to assess the sensitivity of the estimation framework to the mismatch in the field model structure. This simulation example

Parameter	True	Mean est.	St. dev.	CV (%)	%-bias
θ_1	10	45.309	8.708	19.218	—
θ_4	10	46.170	8.688	18.817	57.935
θ_8	100	40.954	6.753	16.488	59.046
θ_{12}	190	35.792	11.009	30.758	81.162
θ_{16}	290	9.523	35.435	372.121	96.716

Table 4.5: Selected estimation results for Experiment V over 100 simulations.

utilises the data from the Experiment I to estimate the hidden field with the grid of 5×5 B-splines.

The gradient of the field constructed with mean parameter estimates and the bias of the inferred field from the true field are shown in Figure 4.9. In this case it is impossible to assess bias of individual parameter estimates, so the bias of the estimated field is calculated as the difference between the true potential field used to generate the tracks and the potential function reconstructed with the mean parameter estimates. It can be seen that the bias function is nearly flat at the right hand side of the map, meaning that the slope of the estimated field is in good alignment with the true slope. Again, the framework overestimates the magnitude of the underlying field almost everywhere. The scaling coefficient corresponding to the B-spline β_1 is the only parameter whose value is understated. This is because in the finer grid of B-splines the local support of an individual B-spline covers a smaller area than the one in, for example, Experiment I. Limited number of tracks passing through that area of the map leads to poor estimation accuracy. As can be seen in Table 4.6, the MLE of θ_1 also has the highest variance, while the other estimates are more precise. As shown in Figure 4.9a, the framework successfully infers the slope of the underlying field almost everywhere despite the increased dimension of the parameter vector.

Parameter	True	Mean est.	St. dev.	CV (%)	%-bias
θ_1	—	-36.811	314.097	853.270	—
θ_5	—	46.584	58.741	126.098	—
θ_{10}	—	114.326	28.955	25.327	—
θ_{15}	—	182.433	27.426	15.033	—
θ_{20}	—	264.317	22.924	8.673	—
θ_{25}	—	269.622	29.283	10.861	—

Table 4.6: Selected estimation results for Experiment VI over 100 simulations.

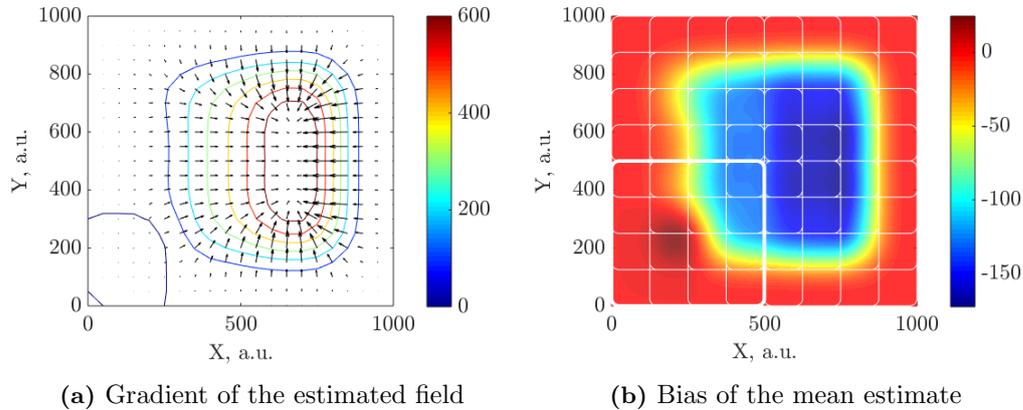


Figure 4.9: Estimation results for Experiments VI. **(a)** The gradient of the potential field constructed with the mean parameter estimates, $\nabla\mathcal{U}(\hat{\Theta})$. **(b)** Total bias between the true field and the mean estimate, $\mathcal{U}(\Theta_{\text{true}}) - \mathcal{U}(\hat{\Theta})$. The colour bar corresponds to the magnitude of the total bias.

4.5 Discussion

The potential field method of modelling the cell-environment interaction employed in this chapter leads to a hybrid model that is non-linear with respect to the object state. The E-step of the estimation algorithm therefore consists of running the IMM-URTS smoother to recover full object states at the reduced computational cost. The attractive feature of the IMM algorithm in the EM context is that it explicitly computes the probabilities of the object being in each mode at each time, which are used as weighing coefficients of the corresponding mode *pdf*. This means that the framework does not single out the most probable migratory mode and accounts for the estimation uncertainty that is usually caused by the fact that the true model of the moving object may be unknown. It must be noted that the price paid for computational efficiency is the suboptimality of the environment model estimation. Firstly, the interaction stage of the state estimator may introduce undesired bias in merged state estimates, especially if interacting state vectors have different lengths. Secondly, the intermediate step is introduced which involves approximation of object positions and allows to solve the optimisation problem in closed form. Although this approximation is considered suitable for the application of interest because of small noise-to-signal ratio in cell centroid microscopy, alternative approximation techniques must be considered in the future.

As an alternative to the IMM algorithm, a combination of optimal methods [165], [166] can be employed to guarantee unbiased estimation at a cost of the increased complexity, which will make the algorithm unsuitable for use with large data sets. It is also possible to utilise a particle filter instead of running a bank

of smoothers on each step. An EM scheme that utilises Rao-Blackwellised particle smoother during the expectation step is proposed in [167] for joint state-parameter estimation of non-linear JMSs. This scheme is particularly suitable for the inference problem considered here because Rao-blackwellisation marginalises discrete-valued states and explicitly computes their probabilities via an HMM filter. However, in applications where all candidate models approximate only a certain aspect of the “true” model which is unknown, the number of particles drastically reduces with each resampling step. This may result in posterior densities being represented by one particle, a known degeneracy problem of the particle-based methods [116].

The proposed framework utilises the ML method which is inherently a point-wise estimator. However, the measure of the estimation uncertainty is provided by the confidence regions of the mean estimate $\hat{\Theta}$. Volumes of the N_b -dimensional confidence ellipsoids obtained from the data samples of varying size are compared in Table 4.7. In the table, N denotes the number of Monte Carlo realisations that were used to compute the mean estimate of the parameter vector. As expected,

Experiment	$N = 1$	$N = 10$	$N = 25$	$N = 50$	$N = 100$
I	1.88×10^{26}	8.27×10^{17}	6.22×10^{14}	8.71×10^{12}	8.71×10^9
II	3.64×10^{32}	4.74×10^{23}	3.77×10^{20}	8.39×10^{17}	3.95×10^{15}
III	2.12×10^{26}	2.49×10^{18}	1.67×10^{15}	6.65×10^{12}	5.26×10^{10}
IV	4.09×10^{26}	1.31×10^{19}	8.38×10^{17}	3.36×10^{13}	1.35×10^{11}
V	4.87×10^{29}	5.45×10^{25}	3.55×10^{22}	1.37×10^{19}	5.01×10^{17}
VI	3.11×10^{42}	2.19×10^{30}	2.77×10^{25}	4.33×10^{21}	6.66×10^{17}

Table 4.7: Volumes of confidence regions for N MC simulations.

the estimation from the uniformly distributed tracking data in with full knowledge of the hybrid model has the smallest level of uncertainty. Restricting the tracking data to a particular starting line on the map III and IV has no significant impact on the estimation uncertainty. On the other hand, the effect of “less informative” localised tracking data in Experiment II is the increased uncertainty of parameter MLEs. It must be noted that the confidence region for the Experiment VI is of higher dimension ($N_b = 25$ compared to $N_b = 16$ for the other experiments), yet with the growing number of realisations its size reduces significantly. The uncertainty introduced by the model mismatch in Experiment V becomes evident with the increasing number of Monte Carlo realisations, as the volumes confidence regions for a hundred realisations are of similar order despite the difference in dimensions.

A better insight into the estimation uncertainty can be provided by treating the unknown parameters as random variables. This can be done by introducing a weak prior for Θ and extending the inference to fully Bayesian framework. The tractable

joint posterior of JMS states, modes, and unknown parameters in this case can be obtained via variational approximation [168]: the true posterior $p(\mathcal{X}, \mathcal{M}, \Theta | \mathcal{Y})$ is approximated by a parametrised distribution

$$\tilde{p}(\mathcal{X}, \mathcal{M}, \Theta) := \arg \min_{\eta} KL [\tilde{p}(\mathcal{X}, \mathcal{M}, \Theta) || p(\mathcal{X}, \mathcal{M}, \Theta)],$$

where $KL[\cdot]$ denotes Kullback–Leibler divergence between two distributions, an η denotes the set of hyper parameters that define “closeness” of the approximating posterior to the true one. The interdependencies of the true joint posterior can be efficiently decoupled via the mean-field approximation [169], [170]:

$$\tilde{p}(\mathcal{X}, \mathcal{M}, \Theta) = \tilde{p}(\mathcal{X})\tilde{p}(\mathcal{M})\tilde{p}(\Theta),$$

where $\tilde{p}(\cdot)$ denotes a variational posterior distribution of each random variable. Recursive variational Bayes (VB) scheme is similar to the EM algorithm and is often referred to as variational Bayes Expectation-Maximisation (VBEM) [171], [172], in which the expectation step is replaced by computing the variational posteriors $\tilde{p}(\mathcal{X})$ and $\tilde{p}(\mathcal{M})$ to maximise the lower bound of the log evidence rather than log likelihood function. Maximisation step then updates $\tilde{p}(\Theta)$ and η . Several recent studies utilise VB method for inference of linear JMS, either in parameter learning [117], [173], or in approximate state estimation [174]. The VB emerges as a versatile alternative to the IMM which allows to bypass the merging of mode histories by considering independent distributions of \mathcal{M} and \mathcal{X} . The comparative analysis in [117] shows that IMM and VB methods perform equally well in Gaussian case. The authors speculate that a combination of two methods will provide a robust inference for a wide class of hybrid models and propose using IMM framework to initialise VBEM recursion.

The anticipated advantage of extending the present framework to mean-field VB inference is that it will obviate the need for the forecasting step: while the classical EM relies only on the first central moments of state *pdfs*, VBEM considers full distributions and therefore incorporates the knowledge of higher order moments. Another attractive feature of variational methods is that, while fully Bayesian, they are deterministic and computationally cheap in comparison with the sampling-based methods discussed above. However, additional investigation is required to analyse the effect of decoupling state and mode of the JSM in variational distributions on the quality of the environment inference.

Based on the conducted experiments we can conclude that the quality of the environment inference is affected by the following three factors: the localisation of the available tracking data, the fit of the selected dynamical model to the data, and the choice of initial conditions for the estimation framework. Experiments II, III, and IV assess the estimation accuracy in scenarios where the tracks are

confined to an area that is smaller than the map approximated by the B-spline grid. Because the cardinal B-splines used for the environment decomposition have local support, the lack of tracking data within the support of an individual B-spline leads to poor estimation accuracy of the corresponding scaling parameter. This is best demonstrated in Figure 4.6b. The results indicate that the uniform support knot sequence may not be the best choice in applications with highly localised migration patterns. A potential solution is to utilise the hierarchical B-spline decomposition or a combination of coarse and fine basis function grids.

Values of the total variance for the MLE vector obtained from a various number Monte Carlo realisations are presented in Tables 4.8. The values presented in the Table are normalised over the size of the unknown parameter vector (N_b) to provide a unified measure of MLE variance per individual parameter. It can be noted that the dynamical model mismatch has the most significant effect on the scatter of the individual MLE, while for other experiments variance values are comparably similar for the same number of MC simulations. This means that the precision of the estimation procedure suffers the most from the lack of knowledge of the true dynamics. The increase in variance for Experiment VI from Experiment I can be explained by the fact that a larger number of unknown parameters is being estimated based on the observation set of the same size, thus rising the level of uncertainty.

Experiment	$N = 1$	$N = 10$	$N = 25$	$N = 50$	$N = 100$
I	169.91	18.81	7.68	3.81	1.82
II	287.91	35.35	14.17	6.95	3.39
III	256.62	29.65	11.97	5.96	2.91
IV	187.20	20.42	8.37	4.17	2.05
V	2.42×10^5	5.47×10^3	1.51×10^3	475.78	322.72
VI	459.33	51.87	20.89	9.99	4.93

Table 4.8: Normalised total variance for N MC simulations.

Because of the selected model structure, scaling parameter MLEs produced by the proposed framework are unique up to an additive constant, so the absence of the prior information and the choice of the initial estimate $\hat{\Theta}^0$ will contribute to the bias of the field estimator. This is demonstrated in Experiments II and VI, where in the right hand side of the map the total bias is spatially constant. This type of bias can only be eliminated by introducing some restrictions on the field magnitude prior to the estimation. The estimation with dynamical model mismatch is tested in Experiment V. In this case, the bias in the field estimate is caused by poor performance of the state estimation algorithm. When the SSM of the responsive migratory mode is wrong, the MM framework favours the random walk over the

wrong model of the responsive mode, thus reducing the number of data points that contribute to the knowledge about the field.

The choice of the function grid does not strongly influence the accuracy of the estimation, however with the growing number of unknown parameters the variance of the estimator increases dramatically. This is evidenced by the larger volume of the confidence regions for the Experiment VI in Table 4.7. Despite the bias in magnitude caused by the choice of the initial conditions, the slope of the potential field estimated in the Experiment VI matches that of the true field. This result supports the claim that the proposed framework does not require prior knowledge about the environment model, as long as there are no significant mismatches present in the model of object dynamics.

The estimation is carried out under the assumption that the mode transition probability matrix is known *a priori*. In the life systems applications this is an artificial parameter that cannot be evaluated from observations. It is possible to approximate the switching probability by the switching frequency which is computed here based on the estimation results of the IMM smoother at the final iteration of the EM algorithm. It is computed as the number of switches between the estimated most probable modes at two subsequent times related to the total number of switches. The results shown in Table 4.9 appear to be in good agreement with the transition matrix used to generate the tracking data, however it is not clear whether this is caused by the self-fulfilling nature of the IMM algorithm. A potential route of the future work is extending the proposed framework to include the identification of the transition frequency matrix. The estimation algorithm will have to be extended to the ECM form with two conditional maximisation steps: one for computing the environment model parameters and one for estimating the transition probability matrix as described, for instance, in [155].

Experiment	M^1 to M^1	M^2 to M^1	M^1 to M^2	M^2 to M^2
I	0.9127	0.0873	0.1412	0.8588
II	0.9135	0.0865	0.1302	0.8698
IV	0.8366	0.1634	0.1302	0.8698
II	0.9031	0.0959	0.1215	0.8785
V	0.8366	0.1634	0.1195	0.8805
VI	0.9105	0.0895	0.1441	0.8559

Table 4.9: Mode transition frequencies computed on the final iteration of the EM algorithm.

Finally, the proposed algorithm deals with a specific problem of quantifying the hidden environment based on object tracking data, while the model parameters of object dynamics are assumed to be known. The estimation framework can be

extended to infer model coefficients [156], [175]. Such modifications would improve the applicability of the proposed algorithm for cases where the system dynamics is largely undetermined.

4.6 Summary

This chapter introduces an approximate Maximum Likelihood framework for a novel problem, in which the potential field environment is inferred from only the observed positions of objects with hybrid dynamics. An approximate EM solution is derived, that involves state estimation of jump Markov non-linear system and identification of environment model parameters. The expectation step is approximated with the interacting multiple model algorithm that runs a bank of unscented Kalman filters for forward recursion and unscented RTS smoothers for backward recursion. The IMM procedure avoids incurring a prohibitive computational cost, while the unscented transform utilised by each smoother provides accurate estimates of cell positions and velocities. Additional set of variables is introduced to approximate object positions for computing the bases values, thus allowing to solve the maximisation step in the closed form.

Several simulation examples are presented to demonstrate the performance of the developed algorithm in various settings. The hybrid model utilised in the framework describes a dynamical system with two migratory modes, in one of which the system driven by the gradient of the environment while in another it is insensitive to its influence. The estimation results demonstrate that an accurate estimate of the global underlying environment can be obtained when the model of the responsive mode is close to the true dynamics and when the tracking data is available in every local support. The lack of prior information may lead to bias in the estimated field magnitude, however the slope of the potential field can be reconstructed accurately. The simulation example that utilises a finer grid of basis function also provides a precise (but biased) estimate, indicating that the estimation framework is scalable to higher dimensions of the unknown parameter vector or the larger volume of tracking data. The following chapter demonstrates a practical application of the developed framework to the problem of estimating global concentration of pro- and anti-inflammatory chemokines in the zebrafish body that drive immune cells towards and from the inflammation site.

5 | Characterisation of the environment influence on neutrophils during stages of inflammation

The analysis in Chapter 3 is performed under a rather limiting assumption that all neutrophils constantly respond to the environment throughout the whole observation period. It is well understood that the complex behaviour of chemotactic cells cannot be sufficiently described by a single state space model (SSM). Clear distinction between observed modes of neutrophil behaviour invites using a hybrid-like model of individual cell migration.

This study makes use of the framework proposed in Chapter 4 for the inference of the hidden chemoattractant environment with the assumption that a migrating cell can at some point stop responding to its surroundings and switch to pure random walk. Each neutrophil tracked *in vivo* is described with a hybrid model that accounts for heterogeneous behaviour of migrating cells evidenced by the experimental data. The hybrid model proposed in this chapter combines several dynamical models available in the literature, with three modes selected based on the observed behaviour of the immune cells. The EM algorithm adapted for the proposed hybrid model is applied to all datasets introduced in §3.2.1, as well as several datasets obtained during the inflammation resolution process.

5.1 Background

Neutrophil responsiveness to the environment during the recruitment stage of inflammation is of fundamental importance in designing effective therapeutic treatments of inflammation-mediated diseases. Majority of works that analyse the inflammation process on macro-scale consider a cell population as a uniform entity where every individual member conforms to the same mathematical model [7], [128], [129]. At the same time, multiple experimental studies suggest that migrating neutrophils

alternate between different types of motion. The two most reported types of chemotaxing cell behaviour are a) chemotactic runs during which the moving cell maintains its directionality, and b) periods of random motility that neutrophils appear to be searching for the direction of the steepest gradient in the ChA concentration [7]. In light of the shortcomings of *in vivo* microscopy discussed before, there arises a novel challenge of identifying which of the observed cells are responding to the hidden external environment. The results achieved in the preceding chapters allow us to account for this switching behaviour when solving the problem of the environment inference by using discrete-time finite-order Markov chains to model the evolution of neutrophil behaviour.

The first attempt to characterise neutrophil chemotaxis using Markov chain theory is described in [176]. The authors experimentally verify the Markovian nature of neutrophil locomotion by matching a five-state continuous-time Markov chain to the *in vitro* data. Four states correspond to the movement in the direction within a certain segment of angles with respect to the X-axis. The fifth state, called 0-state, describes a stationary cell that moves around a fixed position and extends multiple pseudopods in all directions. It is discovered that neutrophils frequently enter the 0-state between periods of directed locomotion. The authors note that neutrophils can randomly switch direction bypassing the zero-state and explain these switches by local changes in the otherwise monotonous ChA environment.

A more recent study uses a hybrid model to quantify the *in vivo* data from zebrafish neutrophils [54], and *drosophila* hemocytes [177]. In [54], the considered hybrid model includes five types of random walk (RW): pure RW, biased RW, persistent RW, and two types of biased-persistent RW. These models are linked to the tracking data via the ABC framework. In the designed estimation framework, an individual cell at each time instance is characterised by the length of its displacement and the direction angle that complies to a wrapped normal distribution. The mean and variance of this distribution depend on the type of RW to which the cell is committed at a given time. The parameters of distributions for each RW model are estimated based on the *in vivo* data obtained from four zebrafish with acute injury. Based on the estimation results, the authors suggest spatio-temporal dependence between bias and persistence in the direction of neutrophil migration towards the wound. The same authors employ the proposed model in further analysis of the *drosophila in vivo* data [177] that links the estimated directional bias to the ChA gradient produced by a circular wound. The attractant gradient is described by a standard 2-D diffusion model with three unknown parameters: total strength of the attractant source, the time of the attractant emission, and the rate of the diffusion. Although the study investigates more complex scenarios of conflicting gradients, it assumes perfect geometry of the wound and uniform process of ChA diffusion.

Neutrophil behaviour during the inflammation resolution stage poses a similar problem. Development of a novel zebrafish experimental assay revealed that the immune cells that survive during the resolution stage could migrate away from the wound within the tissue. It has been speculated that these neutrophils are actively driven away in the process called *fugetaxis* [178]. This process, also known as retrograde chemotaxis, is characterised by cell migration along the negative concentration gradient of the driving chemokines called chemorepellent [16]. However, the *fugetaxis* paradigm is widely disputed in the literature, as no candidate chemokines for the role of the repellent have been identified yet [178]. An alternative hypothesis has been proposed, that after spending a considerable amount of time at the inflammatory site with high ChA concentration, the neutrophils lose sensitivity to the environment and diffuse away from the wound.

The two hypotheses are tested on zebrafish inflammation resolution data in a series of works [17], [39], [179] followed by a thesis-long investigation of the reverse migration mechanisms [180]. The observed neutrophil positions are processed via the ABC framework that selects between two models: the pure diffusion model, which describes completely random migration of individual neutrophils, and the drift-diffusion model, that describes biased migration induced by the unobserved chemorepellent concentration. The estimation results strongly suggest that cells are randomly redistributed rather than actively led away from the wound site. Because of the specifics of the experimental data, the authors characterise neutrophil behaviour under the assumption that all cells in the observed population behave in the same way, and their estimation framework only makes a statement about the observed population as a whole.

This chapter builds on the idea of Markov switching by introducing multiple SSMs of cell migration that incorporate the parametrised model of the hidden environment. Switching between these models is governed by a three-state Markov chain, while the influence of the hidden environment is assumed to affect cell acceleration rather than its direction. The concepts proposed in previous chapters for the estimation procedure, in [54] for the single-cell model of chemotaxis, and in [179] for models of reverse migration are utilised to address the following questions:

- Is the observed cell movement at a given time determined by the interaction with its environment or randomly?
- What is the global environment that is influencing cell migration?

These questions are of interest when characterising the migration of cells both towards and away from the wound. The framework developed in the previous chapter is used to process the data sets obtained during two stages of inflammation. The novelty of this analysis is that it utilises a hybrid model that accounts for varying cell behaviour and makes a probabilistic statement about the migratory mode of an

individual neutrophil at each time as well as providing the information about the hidden global ChA environment that may be driving neutrophil migration without relying on any prior information about the shape of this environment.

5.2 Methods

5.2.1 Data acquisition

In addition to the inflammation recruitment data introduced before, this chapter deals with several datasets obtained during the inflammation resolution stage. The reverse migration process can be observed in Tg(mpx:GAL4;UAS:Kaede)i222 transgenic zebrafish line, in which neutrophils are tagged with a photoswitchable fluorescent protein kaede that converts from green to red fluorescence upon exposure to UV light [181].

Four larvae 3dpf were anaesthetised and subjected to the tail fin transection with a sterile scalpel. The fluorescence of kaede-tagged neutrophils was photoconverted using an UltraVIEWPhotoKinesis™ device (Perkin Elmer and Analytical Sciences) on an UltraVIEWVoX spinning disc confocal laser imaging system (Perkin Elmer). Photoconversion of kaede protein was performed using a 405nm laser at 40% energy using 120 cycles. The calibrating procedure of the device and the details of photoconversion are described in [182].

Time lapse imaging was performed during 6-16 hours post injury with 5 minute sampling period using an Eclipse TE2000-U fluorescence microscope with a Andor Zyla 5.5 camera (Nikon). Both recruited and departing neutrophils were tracked simultaneously using GFP and mCherry filters with 488 and 561 nm excitation respectively. Similarly with the previously described datasets, the 2-D neutrophil trajectory data was obtained by processing the Z-stacked images via the tracking tool within the NIS Elements (Nikon), where the threshold was set to detect the mCherry cells. Tracking results are shown against the brightfield images of zebrafish larvae in Figure 5.1.

5.2.2 Describing heterogeneous cell behaviour

The model of cell dynamics utilised in this chapter accounts for both random aspect of Brownian-like cell motion and the influence of the underlying field which is included as a deterministic term. The state \mathbf{x}_t^k consists of the spatial position of the cell and velocity projections on two axes

$$\mathbf{x}_t^k = [s_x, s_y, v_x, v_y]^\top,$$

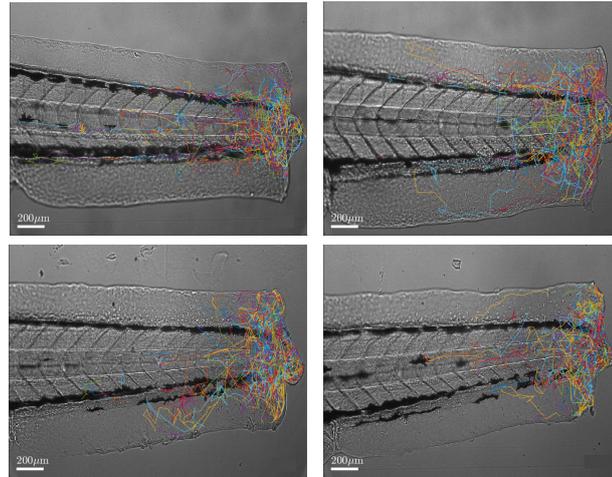


Figure 5.1: Tracking results from the inflammation resolution stage. Neutrophils can be seen slowly migrating away from the wound site. The resolution process was observed using video microscopy during 6-16 hours post injury with 5 minute intervals between frames.

In order to distinguish between behavioural modes of an individual cell and identify which of the observed migratory patterns may be the product of the cell-attractant interaction, we describe the dynamics of each moving cell by a hybrid model consisting of several SMMs, each describing a different type of cell behaviour that can be observed during the course of *in vivo* experiments (see Figure 5.2). Available centroid observations allow us to conclude that there are at least three types of neutrophil behaviour:

- Interacting with the environment, or responsive. This mode corresponds to the cells with more directionally persistent tracks. The state of the chemotaxing cell is described by the discrete-time O-U-type process with bias induced by the acting environment:

$$M^1 : \mathbf{x}_{t+1} = A(M^1)\mathbf{x}_{t+1} + B(M^1)\nabla U + G(M^1)\mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(0, Q(M^1)), \quad (5.2.1)$$

where the matrix $A(M^1)$ accounts for reversion to mean.

- Diffusing. This mode is used to describe cells that keep moving as if there is no gradient in the attractant concentration and commit Brownian-like motion. Dynamics of an individual cell in this mode is represented by the discrete-time O-U-type process:

$$M^2 : \mathbf{x}_{t+1} = A(M^2)\mathbf{x}_t + G(M^2)\mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(0, Q(M^2)). \quad (5.2.2)$$

- Stationary. Dead neutrophils remain fluorescent and are registered by tracking software. The tracks of stationary cells are normally lengthy and may corrupt the environment inference if these cells are considered as actively migrating.

On the other hand, some of the observed neutrophils that are waiting to be released into the bloodstream can be observed in the fish body. Behaviour of both of these cells can be described by a small magnitude RW. Dynamics of stationary cells is therefore described by the discrete-time O-U-type process with variance significantly smaller than that of the modes described above:

$$M^3 : \mathbf{x}_{t+1} = A(M^3)\mathbf{x}_t + G(M^3)\mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(0, Q(M^3)). \quad (5.2.3)$$

Modelling two types of cell behaviour by the same SSM means that the cells can switch from the stationary mode to other modes, which means that the mode evolution can be appropriately described by the ergodic Markov chain.

The ergodic Markov chain governing the mode switching is illustrated in Figure 5.2b. Combined with the basis function representation of underlying field utilised in previous chapters, the JMS describing the dynamics of the k -th cell is formulated as follows:

$$\mathbf{x}_{t+1} = A(M^j)\mathbf{x}_t + B(M^j)\varphi(\mathbf{s}_{t-1}^k) + G(M^j)\mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(0, Q(M^j)), \quad (5.2.4)$$

where the modes are described as follows:

$$\begin{aligned} M^1 : A &= \begin{bmatrix} \mathbb{I}_2 & T\mathbb{I}_2 \\ \mathbb{O}_2 & \mathbb{I}_2 - T\rho_1\mathbb{I}_2 \end{bmatrix}; B = \begin{bmatrix} \mathbb{O}_2 \\ T\mathbb{I}_2 \end{bmatrix}; G = \begin{bmatrix} \mathbb{O}_2 \\ T\mathbb{I}_2 \end{bmatrix}. \\ M^2 : A &= \begin{bmatrix} \mathbb{I}_2 & T\mathbb{I}_2 \\ \mathbb{O}_2 & \mathbb{I}_2 - T\rho_2\mathbb{I}_2 \end{bmatrix}; B = \begin{bmatrix} \mathbb{O}_2 \\ \mathbb{O}_2 \end{bmatrix}; G = \begin{bmatrix} \mathbb{O}_2 \\ T\mathbb{I}_2 \end{bmatrix}. \\ M^3 : A &= \begin{bmatrix} \mathbb{I}_2 & T\mathbb{I} \\ \mathbb{O}_2 & \mathbb{I}_2 - T\rho_3\mathbb{I}_2 \end{bmatrix}; B = \begin{bmatrix} \mathbb{O}_2 \\ \mathbb{O}_2 \end{bmatrix}; G = \begin{bmatrix} \mathbb{O}_2 \\ T\mathbb{I}_2 \end{bmatrix}. \end{aligned}$$

where T is the time increment, \mathbb{I} is an identity matrix and \mathbb{O} is a zero matrix of size 2×2 . The rate of reversion to mean ρ_j may be different for each mode. The standard deviation of the process noise also depends on the mode of the cell:

$$\sigma(M^3) \ll \sigma(M^1) \leq \sigma(M^2) \quad (5.2.5)$$

In (5.2.4) \mathbf{s}_{t-1}^k is the mode-independent noiseless measurement of cell position artificially introduced to simplify the estimation process

$$\mathbf{s}_t^k = C\mathbf{x}_t^k, \quad (5.2.6)$$

where

$$C = \begin{bmatrix} \mathbb{I}_2 & \mathbb{O}_2 \end{bmatrix}. \quad (5.2.7)$$

The observation process is also independent of the mode:

$$\mathbf{y}_t^k = C\mathbf{x}_t^k + \mathbf{v}_t^k, \quad (5.2.8)$$

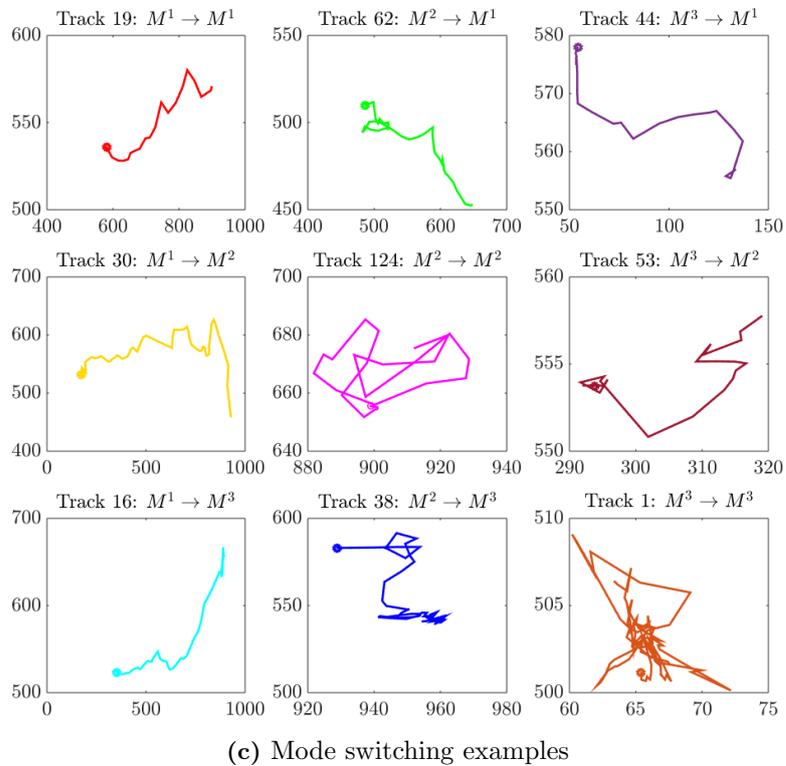
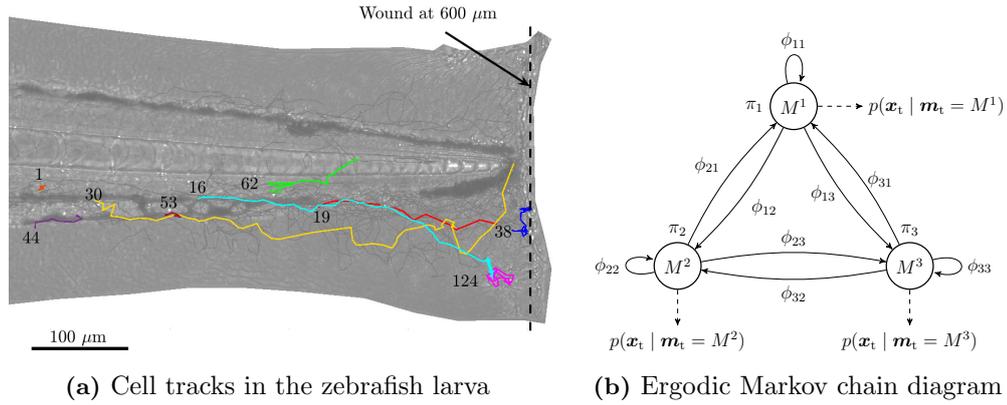


Figure 5.2: The examples of neutrophils switching between three proposed behavioural modes. **(a)** The tracks superimposed on the fish image are colour-coded to match the tracks in the table. **(b)** A three state ergodic Markov chain governs the mode switching. Each mode can be reached from every other mode in a single transaction. **(c)** Each of the selected cells demonstrates the switching between the described modes. The tracks in the first column from the left have clearly defined directionality, as can be seen on the fish image. The tracks in the central column demonstrate possible magnitude of random fluctuations in the directionality of desensitised cells. The tracks in the right column belong to the cells that were tracked before becoming activated and starting their migration. All of these cells are located in the fish body next to the spinal cord.

where the observation vector at sample time t contains positions of cell centroids on the two-dimensional image. Some of the tracked neutrophils that switch between the types of behaviour discussed above are shown in Figure 5.2c. The figure demonstrates that each mode can be reached from any other mode at any point throughout the recruitment process, which means that the switching process can be described by an ergodic HMM, making the JMS an appropriate hybrid model for describing neutrophil dynamics. Track 1, highlighted in the figure, belongs not to an apoptotic neutrophil but to the cell within the spinal cord that has not been released into the tissue. Considering such cells to be interacting with the environment will inevitably lead to the bias in the parameter estimation. The mode M^3 is introduced to account for these tracks and can be interpreted as a way of dealing with outliers.

The JMS defined above may be used in the inference of global chemorepellent concentration, with the sign of the B matrix in the responsive mode reflecting the fact that the neutrophils are moving in the direction opposite to the steepest gradient:

$$M^1 : A = \begin{bmatrix} \mathbb{I}_2 & T \mathbb{I}_2 \\ \mathbb{O}_2 & \mathbb{I}_2 - T \rho_1 \mathbb{I}_2 \end{bmatrix}; B = - \begin{bmatrix} \mathbb{O}_2 \\ T \mathbb{I}_2 \end{bmatrix}; G = \begin{bmatrix} \mathbb{O}_2 \\ T \mathbb{I}_2 \end{bmatrix}.$$

5.2.3 Environment inference with a hybrid model of neutrophil dynamics

The estimation algorithm utilised for the chemoattractant field inference is summarised in Figure 5.3. Similarly with the preceding chapters, the inference framework is implemented as an EM-like algorithm with an intermediate step where the forecast of cell positions is computed prior to the maximisation step. Although no prior information about the shape of the environment is required in this method, the assumptions are made about cell behavioural modes, particularly about the structure and tuning parameters of individual SSMs. Since the true dynamics of migrating cells are unknown, all hypothesised models must be considered at each time point when estimating the cell state. This leads to the problem of multiple model state estimation, that cannot be solved directly because of the prohibitive computational cost of keeping exhaustive mode histories. There exist several ways to decrease these costs, but two most common approaches achieve this by reducing the size of the mode history. One method relies on merging all modes from previous time instances, while the other "cuts off" the least probable modes at each time.

Merging. A number of methods that merge possible mode histories have been discussed in §2.5.5. The framework proposed in the previous chapter employs one such method - the IMM smoother. In the interest of the estimation objective, the IMM method is used to make a probabilistic statement about the cell behaviour at

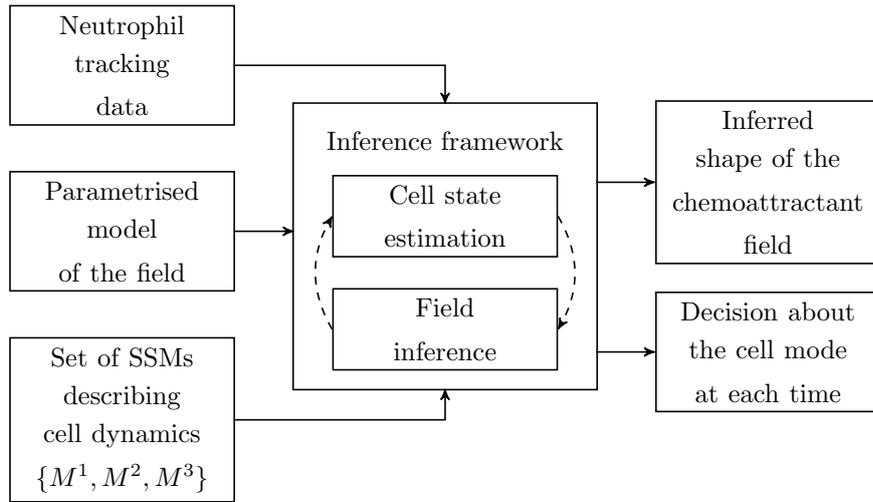


Figure 5.3: A schematic overview of the estimation framework utilised for the joint estimation of the global chemoattractant concentration field and the individual cell behaviour.

each time point. The estimated probabilities of an individual neutrophil being in a given mode are used as mixture weights in the log-likelihood function of the complete data set. This technique is preferable in tracking applications where the true dynamics of the object is unknown and does not match with any of the hypothesised models. In reality, a migrating cell can only be in one mode at any given time. In cases when the modes accurately describe possible behaviour of the migrating object, the merged state that is the output of the IMM algorithm is a weighted sum of all possible states and can thus be biased to some average value.

Pruning. A popular alternative to merging histories is reducing the number of sequences that are brought forward in time. This is achieved by pruning the posterior densities that correspond to the smaller model likelihoods after each recursion of the estimator [183], [184]. The simplest version of a pruning algorithm is the threshold-based pruning, that at each iteration of multiple-model filtering runs N_m filters to produce the likelihood of each mode, then only keep an arbitrary number $M < N_m$ of modes with highest probabilities, discarding the rest. The difficulty with pruning arise on the smoothing stage. While there exists a two-filter smoother that utilises pruning in both forward and backward-time recursions [185], it requires time-reversal of the dynamical model. On the other hand, an RTS-type smoother suffers from degeneracy as the number of hypotheses has been reduced because of pruning on the filtering stage. This chapter considers the combination of IMM forward filter and RTS smoother with M -best pruning, where $M = 1$. In other words, the smoother will follow the most probable mode sequence of an individual cell. Implementing the pruning procedure allows the computation of a simplified M-step derived in Lemma 3.3.1 that utilises only the sequence of states conditioned on the most probable mode

at each time. The estimation algorithms that employ merging and pruning methods are compared in Table 5.1.

Input: Observation set, \mathcal{Y} ; hypothesised models, \mathbb{M} ; measurement noise covariance, R ; initial model probabilities, Π ; mode transition probabilities, Φ ; convergence threshold, ϵ .
Output: Smoothed estimate of full hidden state sequence, \mathcal{X} ; estimated mode probabilities, $\hat{\mu}$; estimated model parameters, $\hat{\Theta}$.

Estimation with history merging

```

1: Initialise parameter vector  $\hat{\Theta}^0$ ;
2: while ( $i \leq i_{max}$ ) and  $\neg(\text{convergence})$  do
3:   for  $k \leftarrow 1, K$  do
4:     Run IMM-URTS algorithm to recover
       full states  $\hat{\mathbf{x}}^k$  and mode-associated probab-
       ilities  $\hat{\mu}$ ;
5:   for  $t \leftarrow 1, T$  do
6:     Obtain cell position forecast using
       merged states  $\hat{\mathbf{s}}_t^k = C\hat{\mathbf{x}}_t^k$ ;
7:   end for
8: end for
9: Estimate  $\hat{\Theta}^i$  using Lemma 4.3.2;
10: Check convergence of the parameter vector  $\hat{\Theta}$ ;
11: end while

```

Estimation with history pruning:

```

1: Initialise parameter vector  $\hat{\Theta}^0$ ;
2: while ( $i \leq i_{max}$ ) and  $\neg(\text{convergence})$  do
3:   for  $t \leftarrow 1, T$  do
4:     Run IMM filter to produce  $N_m$  mode-
       conditioned state pdfs;
5:   end for
6:   for  $t \leftarrow T, 1$  do
7:     Run RTS smoother with pruning at  $M = 1$ 
       to produce the most probable sequence of
       cell modes;
8:     Compute cell position forecast using the
       state conditioned on the most probable
       mode  $\hat{\mathbf{s}}_t^k = C\hat{\mathbf{x}}_t^{k,j}$ ;
9:   end for
10: Estimate  $\hat{\Theta}^i$  using Lemma 3.3.1;
11: Check convergence of the parameter vector  $\hat{\Theta}$ ;
12: end while

```

Table 5.1: Alternative methods for the environment inference.

5.3 Results

The estimation framework developed in the previous chapter is applied to the neutrophil tracking data from two stages of the inflammatory response. The following aspects of the observed cell migration are examined here: the probability of an individual neutrophil being in the responsive migratory mode, the mean bias in the velocity projections of neutrophils in different modes, mode switching frequencies, and the shape of the chemokine concentration environment perceived by the responsive cells. In addition, the estimation algorithm that employs history pruning during the expectation step is applied to the normal injury data in order to compare two alternative ways of approximating the exact mode history when the true model describing neutrophil dynamics is unknown.

The tuning parameters used for the processing of the data from various experiments are summarised in Table 5.2. In the table, the time increment of the dynamics model is equal to the interval between frames selected for an individual experiment. The measurement noise covariance is the same for all datasets because

all data was obtained using the same equipment. The process noise variances are selected arbitrarily to satisfy (5.2.5). Note that $Q_w(M^1)$ and $Q_w(M^2)$ are the same because the cell that is driven by the external gradient is assumed to be subjected to random disturbances of the same magnitude as those affecting a randomly moving cell. While the sensitivity of the estimation framework to the tuning parameters is not assessed in this work; it is important to note that as long as all datasets from one experiment are processed with the same settings the estimation results can be analysed in the same context. The mode transition probability matrix is arbitrarily

Parameter	Normal injury	Mild injury	Nick injury	Severe injury	Reverse migration
T	2 min	1.5 min	0.5 min	2 min	5 min
$\rho(M^1)$	$T/20$	$T/20$	$T/20$	$T/20$	$T/100$
$\rho(M^2)$	$T/10$	$T/10$	$T/10$	$T/10$	$T/50$
$\rho(M^3)$	$T/10$	$T/10$	$T/10$	$T/10$	$T/40$
$Q_w(M^1)$	$2 \mathbb{I}_2$	$2 \mathbb{I}_2$	$1 \mathbb{I}_2$	$3 \mathbb{I}_2$	$0.250 \mathbb{I}_2$
$Q_w(M^2)$	$2 \mathbb{I}_2$	$2 \mathbb{I}_2$	$1 \mathbb{I}_2$	$3 \mathbb{I}_2$	$0.250 \mathbb{I}_2$
$Q_w(M^3)$	$0.5 \mathbb{I}_2$	$0.05 \mathbb{I}_2$	$0.01 \mathbb{I}_2$	$0.5 \mathbb{I}_2$	$0.005 \mathbb{I}_2$
R_v	$2 \mathbb{I}_2$	$2 \mathbb{I}_2$	$2 \mathbb{I}_2$	$2 \mathbb{I}_2$	$2 \mathbb{I}_2$

Table 5.2: Tuning parameters of the hybrid model.

set to the following

$$\Phi = \begin{bmatrix} \phi_{11} & \phi_{12} & \phi_{13} \\ \phi_{21} & \phi_{22} & \phi_{23} \\ \phi_{31} & \phi_{32} & \phi_{33} \end{bmatrix} = \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix}. \quad (5.3.1)$$

The EM algorithm for all types of injury converges within 5 iterations. The convergence plots are omitted here. Volumes of the N_b -dimensional confidence regions are presented to provide the metric of the estimation uncertainty for various wound types. In this application, the confidence regions can also be interpreted as the measure of how informative the tracking data is in experiments with various wound types.

5.3.1 Chemoattractant concentration inference during neutrophil recruitment stage

The normal injury.

At each iteration of the estimation procedure neutrophil tracks are processed via the IMM-URTS smoother which identifies the probability of the cell being at

either of the modes. Examples of mode and state estimation results from the final iteration of the EM algorithm are presented in Figure 5.4. Neutrophil tracks in each instance are colour coded according to the estimated probability of being in a particular mode at each time. It can be observed that more persistent tracks in the the central area of the fish body belong to the cells that are identified as responding to the environment, while those closer to the wound site are inferred to be randomly moving. As expected, the cells that are more likely to be in the stationary mode are located either at the fish body or at the very edge of the injury. It can be speculated that the cells in the body were not activated throughout the inflammation recruitment, while the cells at the injury have undergone apoptosis.

The framework appears to make a clear distinction between different migratory modes: for example, the probabilities of being in the responsive mode are normally close to 1. Thus, these neutrophil tracks provide the most information about the chemotactic environment. The example histograms of the horizontal velocity component for cells that are identified to be in different modes are compared in Figure 5.5. The responsive cells are characterised by the clear bias of both mean and median horizontal velocity in all larvae. This means that the proposed model better captures the ChA-induced directional bias in comparison with the simple SSM considered in Chapter 3. The unusual bias of mean in the diffusing mode is observed in the histogram corresponding to Fish 4 in which the wound is not aligned with the Y-axis, thus some of the cells diffusing in the wound area have increased v_x . The stationary cells are characterised by the small variance of both velocity components. The histograms of vertical velocity components are not examined here as they are similar for all migratory modes.

Field inference results are presented in Figure 5.6. The hidden chemoattractant field in this case has been approximated by a 5×4 grid of B-splines. It can be seen that the results for Fish 1–3, 5 and 6 are consistent with the assumption that the ChA is uniformly distributed along the wound. It must be noted that Fish 1, 2 and 3 do not have outliers, while the results for Fish 4, 5 and 6 have sinks outside of the fish body. These results, once again, demonstrate the effect of using the B-spline grid with uniform local support. Nonetheless, for five out of six larvae the inferred environment is characterised by uniform concentration along the wound and a gradual decrease along the length of the fish body.

Selected results of applying an alternative algorithm that employed mode history pruning is presented in Figure 5.7. The plotted tracks illustrate a definitive decision about the cell behaviour at each time made by the pruning algorithm: each tracks in the figure is colour coded corresponding to the identified modes, rather than with the probability of each mode. The figure shows that the decisions made about the migratory mode of each neutrophil coincide with the results of the IMM

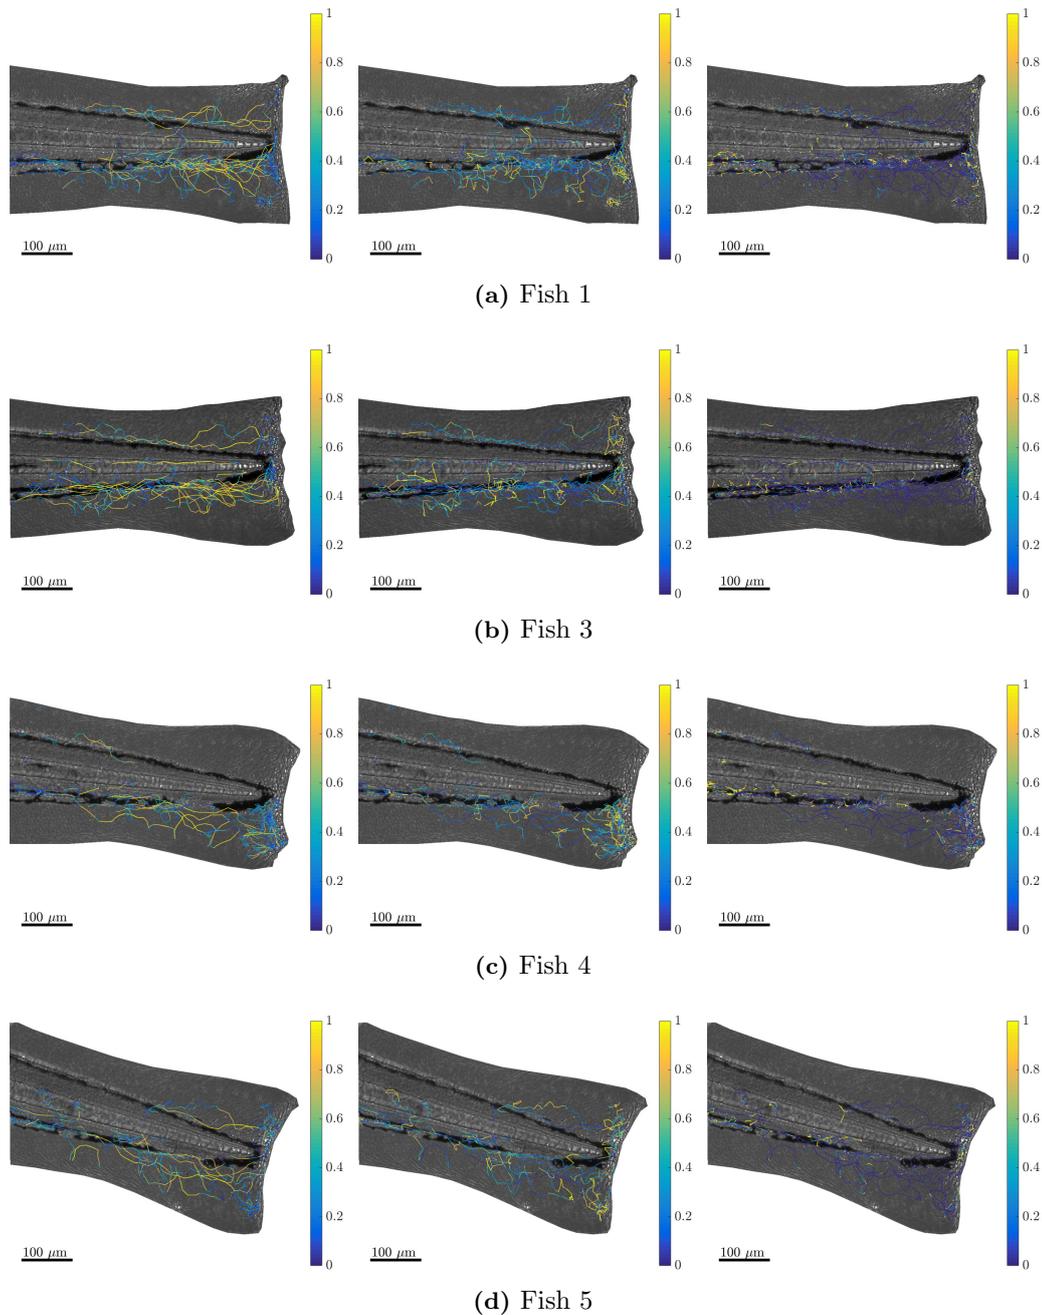


Figure 5.4: Examples of migratory mode probabilities estimated via IMM-URTS algorithm during the expectation step. Plots in each row are ordered from left to right as follows: probabilities of M^1 , probabilities of M^2 , probabilities of M^3 . The cell tracks are colour coded in accordance to the probability of the cell being at a given mode. Colour bars correspond to the probability of the particular mode.

algorithm. This is because the hybrid model describes three distinctive types of behaviour, so both frameworks discriminate between the migratory modes with high degree of certainty. Another reason for this similarity is that both algorithms utilise the IMM filter for forward estimation, and only differ at the smoothing procedure.

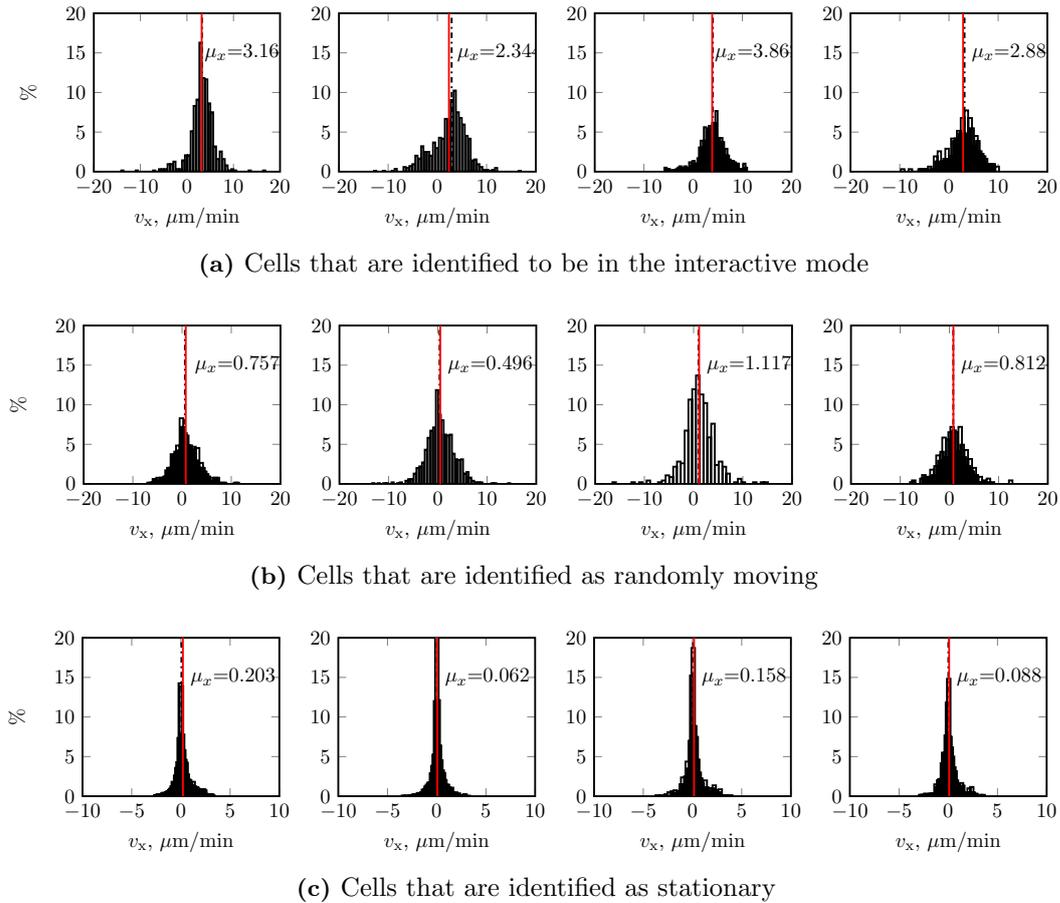


Figure 5.5: Comparison of cell velocity histograms for actively migrating and randomly moving neutrophils during the recruitment stage of the inflammatory response. Velocity components estimated with IMM-URTS algorithm. **(a)** Histograms of the horizontal velocity component corresponding to neutrophils that are more likely to be in the active mode. **(b)** Histograms of the horizontal velocity component of neutrophils that are more likely to perform random walk. **(c)** Histograms of the horizontal velocity component of neutrophils that are more probable to be stationary.

Field inference results obtained from the pruning algorithm are in good alignment with those obtained by using the IMM smoother but are characterised by the steeper slope of ChA concentration. Similar shapes of the estimated fields are the result of the fact the neutrophils that are identified by the framework as responding to the environment are assigned high probabilities (> 0.8), thus their contribution to the inference is close to the contribution of cell that are "definitely" responding to the field, according to the gradient. At the same time, even small probability of a cell being in a responsive mode carries the information about the environment which is utilised in the framework with the IMM smoother. This leads to the milder slope of the field estimated using the IMM-based approach. On the other hand, all cells with small probability of being in a responsive mode are discarded by the pruning

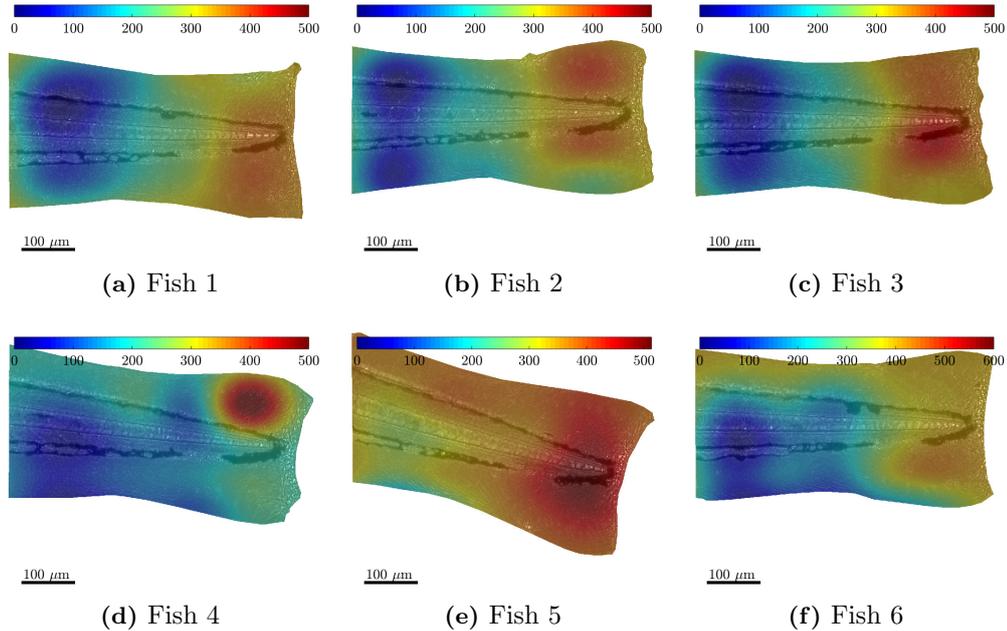


Figure 5.6: The chemoattractant concentration field estimated from the normal injury data via the framework with IMM algorithm. The colour bars are normalised to scale 0-500 a.u..

algorithm as diffusing, so their states are not utilised in the environment estimation step. This results in a steeper concentration gradient along the length of the fish body. Since the inference outcomes are rather similar for both algorithms, the further subsections will only utilise the framework that employs the IMM smoother for the cell state estimation.

Volumes of 20-dimensional confidence regions for model parameter estimates obtained from the normal injury data are presented in Table 5.3. While sizes of confidence ellipsoids for Fish 1-3 are similar, the highest levels of uncertainty correspond to Fish 4 and 5 where the body of a larva is not horizontally aligned on the image. This results in two of the B-splines being located completely out of the fish body with no tracking data available. High order of total variance values may be the result of the inaccurate model of cell dynamics or noise levels, as model mismatches in the previous chapter were resulting in similar levels of MLE dispersion. The computed mode transition frequencies illustrated in Figure 5.8 indicate that the stationary cells rarely switch to other migratory modes, and that responding neutrophils do not immediately become stationary. Instead, they tend to switch to the diffusion mode which seems to be intermediate between the other two, as switches from M^2 to M^3 and M^1 happen with similar frequency.

The mild injury.

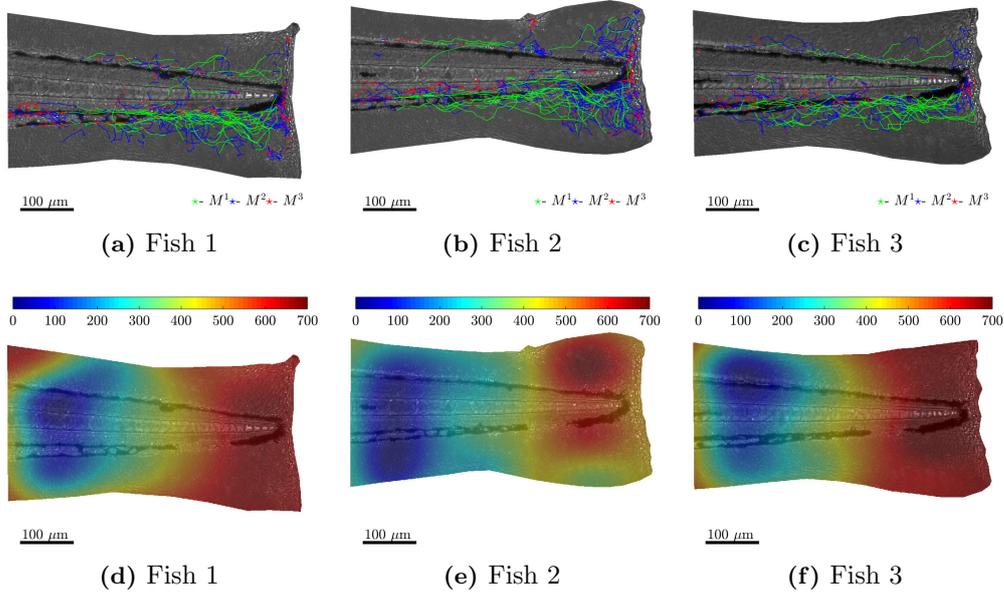


Figure 5.7: Estimation results obtained from the EM algorithm with mode history pruning. (a)-(c) Cell tracks colour-coded with the migratory mode decided by the pruning algorithm. (d)-(f) The chemoattractant concentration field estimated from the normal injury data with pruning of cell modes. The colour bars are normalised to scale 0-700 a.u..

Dataset	$V_{N_b=20}$	$\text{Var}(\hat{\Theta})$
Fish 1	1.36×10^{41}	3.90×10^3
Fish 2	8.87×10^{40}	1.50×10^3
Fish 3	2.32×10^{41}	1.28×10^3
Fish 4	8.53×10^{47}	5.90×10^3
Fish 5	2.12×10^{45}	7.75×10^3
Fish 6	2.68×10^{43}	2.41×10^3

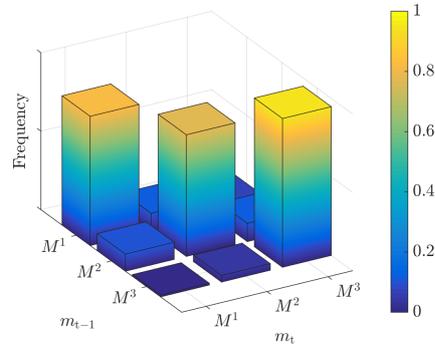


Table 5.3: Volumes of confidence regions and normalised total MLE variance for the normal injury datasets.

Figure 5.8: Computed mode transition frequencies for the normal injury datasets.

The datasets from this experiment are characterised by sparse but long neutrophil tracks with high degree of directional persistence. Mode probabilities are estimated by the IMM smoother at the final iteration of the EM algorithm are shown in Figure 5.9. Similarly with normal injury data, the framework distinguishes between modes with high degree of certainty: the cells that are identified to be in the responsive mode have probabilities close to 1. The decisions about the diffusing mode

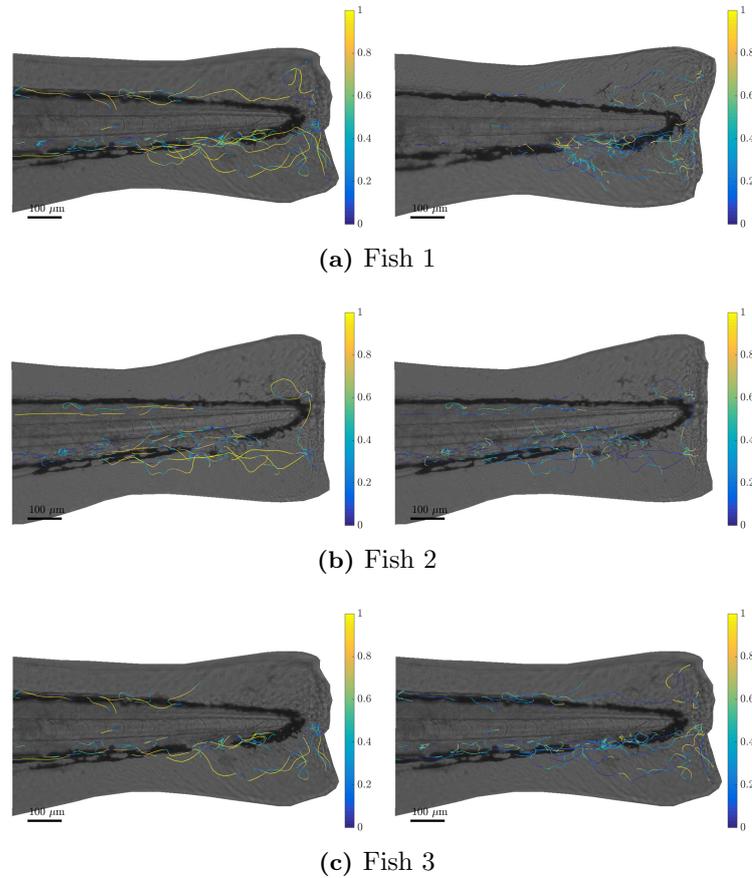


Figure 5.9: Migratory mode probabilities estimated via IMM-URTS algorithm from the mild injury data. Plots in each row are ordered from left to right as follows: probabilities of M^1 , probabilities of M^2 . The cell tracks are colour coded in accordance to the probability of the cell being at a given mode.

are not as conclusive, as the probability of being in M^2 for most of cell tracks does not go over 0.6. Estimated probabilities of the stationary mode are close to 1 as well. The results for M^3 are not illustrated in this figure.

Because of the sparsity of the tracking data, the underlying environment is parametrised by a 4×4 grid of basis functions. The inference results are shown in Figure 5.10. Some outlier peaks are observed for Fish 2 and 3 in the areas with no tracks available. It can be seen that the high concentration peaks correspond to the areas where a single responsive neutrophil is found. The slope of the estimated fields not as steep as that inferred from the normal injury data, but the diffusive nature of the global environment is not captured by the framework in this case.

Volumes of confidence regions for the parameter MLEs presented in Table 5.4 provide a measure of the estimation uncertainty. While they are much smaller in size than those obtained from the normal injury dataset, they correspond to the 16-dimensional hyperellipsoid, as opposed to the 20-dimensional region considered

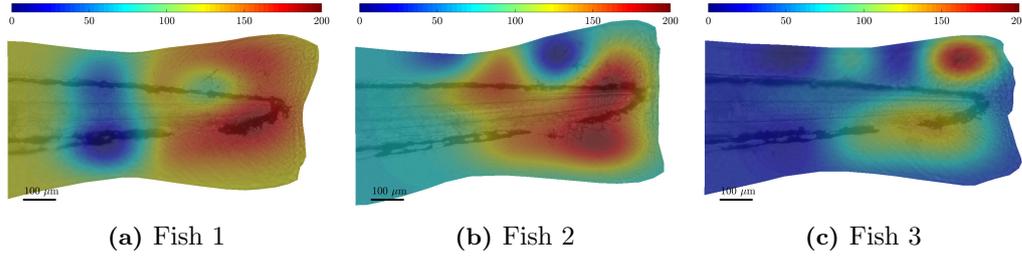


Figure 5.10: The chemoattractant concentration field estimated from the mild injury data. The colour bars are normalised to scale 0-200 a.u..

Dataset	$V_{N_b=16}$	$\text{Var}(\hat{\Theta})$
Fish 1	4.62×10^{30}	2.21×10^3
Fish 2	2.65×10^{30}	3.88×10^3
Fish 3	7.40×10^{29}	3.83×10^3

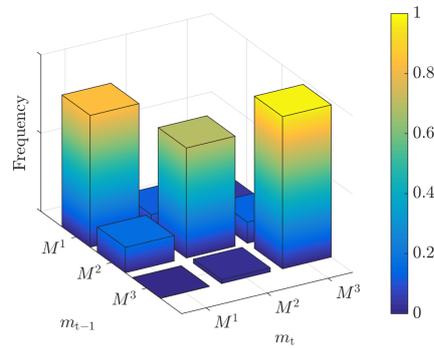


Table 5.4: Volumes of confidence regions and normalised total MLE variance for the mild injury datasets.

Figure 5.11: Computed mode transition frequencies for the mild injury datasets.

above. Variance values for this dataset are consistent with the results of the normal injury set. Estimated mode transition frequencies illustrated in 5.11 suggest similar behaviour to the normal injury case.

The tail fin nick injury.

The zebrafish with tail fin injury is characterised by numerous short tracks confined to the lower area of the fish body. The estimated mode probabilities of the observed neutrophils are illustrated in Figure 5.12. It can be seen from the figure that even the neutrophils that do not appear to migrate towards the wound are identified as responding to the environment, while a large proportion of neutrophils at the wound site are perceived by the framework as more probable to be diffusing. Very few cells are estimated to be stationary therefore the probabilities of this mode are not shown in the figure.

The global environment is approximated with a 4×4 grid to avoid outliers in the scaling parameter estimates. The inference results are presented in Figure 5.13. While in Fish 2 the estimated peak of chemoattractant concentration is located near the wound, Fish 1 and 3 have outlier peaks that match the area with several persistent cell tracks directed away from the wound. These cells are also

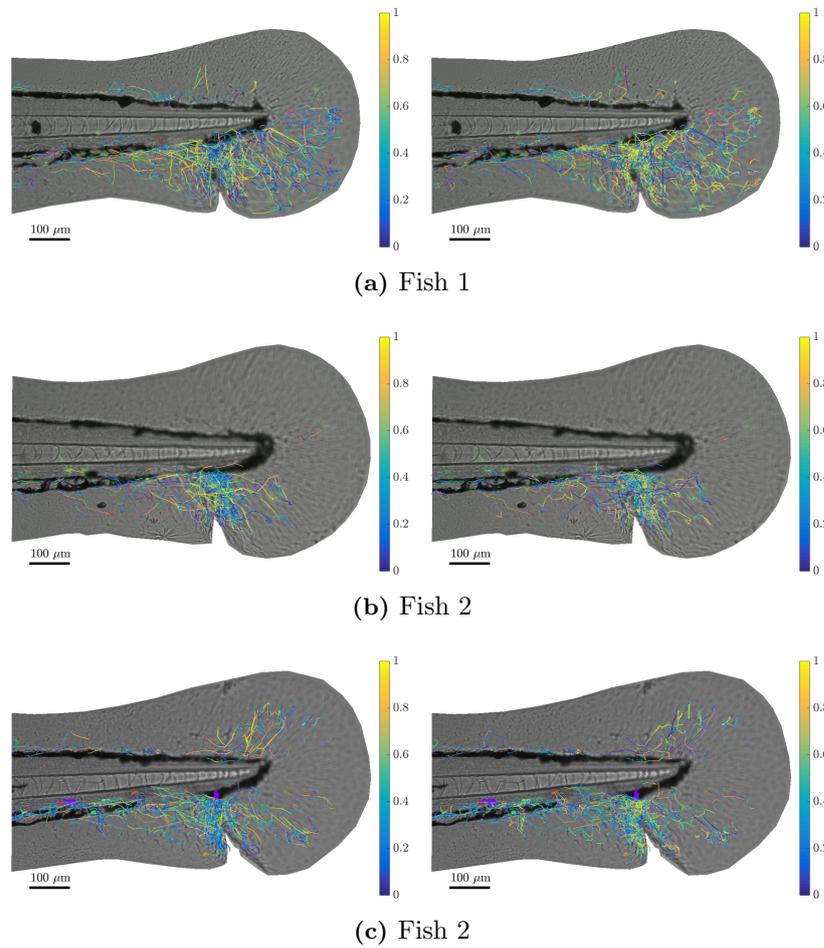


Figure 5.12: Neutrophil mode probabilities estimated via IMM-URTS algorithm from the tail fin nick injury data. Plots in each row are ordered from left to right as follows: probabilities of M^1 , probabilities of M^2 . The cell tracks are colour coded in accordance to the probability of the cell being at a given mode.

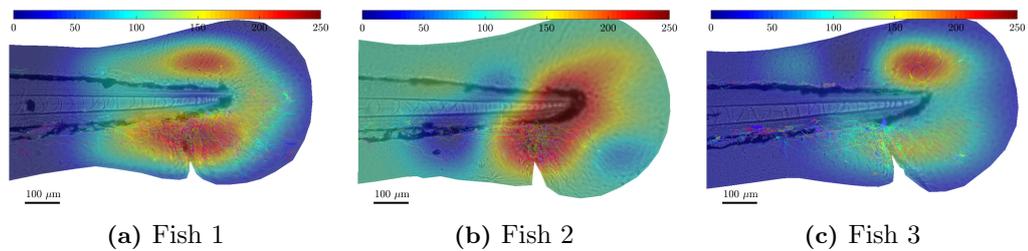


Figure 5.13: The chemoattractant concentration field estimated from the normal injury data. The colour bars are normalised to scale 0-250 a.u..

Dataset	$V_{N_b=16}$	$\text{Var}(\hat{\Theta})$
Fish 1	9.48×10^{29}	3.58×10^3
Fish 2	7.39×10^{29}	9.42×10^3
Fish 3	7.45×10^{29}	5.35×10^3

Table 5.5: Volumes of confidence regions and normalised total MLE variance for the nick injury datasets.

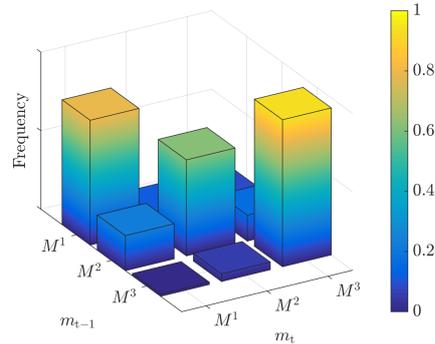


Figure 5.14: Computed mode transition frequencies for the tail fin nick injury datasets.

identified as responding to the ChA gradient, however they may be driven by another neutrophil-generated gradient. The framework utilised for this analysis does not differentiate between various sources of attractants, however it successfully picks up on the "unusual" neutrophil activity in the area remote from the wound, unlike the estimation framework based on the simple SSM (recall the estimation results for the same dataset in Figure 3.13).

The severe injury.

Mode probability estimates for the severe injury tracking data are illustrated in Figure 5.15. Interestingly, it is more difficult to distinguish between the responsive mode and the diffusing mode of the neutrophils in this case. The majority of the tracks do not maintain directionality, yet the ones that are closer to the wound site are identified as responsive with high probability. In addition, stationary neutrophils are scattered around the fish body, rather than being confined to the spinal cord or the wound site. However, in order to assess whether this unusual behaviour is a result of the higher magnitude of the ChA concentration in the fish body, it is necessary to conduct more experiments with severe injury.

The overwhelming number of randomly diffusing and stationary cells results into a rather small slope of the estimated field (see Figure. 5.16). While in Fish 1 the inferred ChA concentration appears to be uniformly diffused along the wound site, the results for Fish 2 suggest that there might be a local peak of chemoattractant at the wound site generated by neutrophils themselves. The slope of the estimated concentration field is rather moderate compared to the estimate obtained in Chapter 3, because only a portion of cell tracks inform the inference. However, because of the non-unique solution provided by the inference framework, the estimates shown in the figure are likely to be understating the actual magnitude of the attractant field that could be much higher. Confidence region volumes presented in Table 5.6

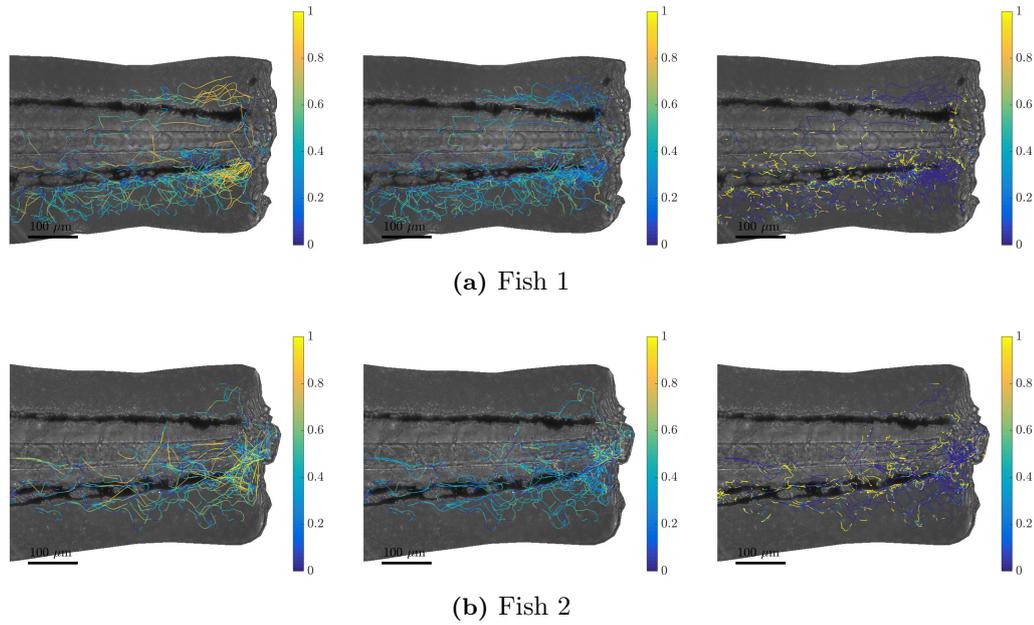


Figure 5.15: Migratory mode probabilities estimated via IMM-URTS algorithm from the severe injury data. Plots in each row are ordered from left to right as follows: probabilities of M^1 , probabilities of M^2 , probabilities of M^3 . The cell tracks are colour coded in accordance to the probability of the cell being at a given mode.

are consistent with the results obtained for the mild and tail nick injury data, as do the computed mode transition frequencies shown in Figure 5.17.

5.3.2 Chemorepellent concentration inference during inflammation resolution stage

Neutrophil tracks from the reverse migration dataset are concentrated in the tail fin area. As has been demonstrated in the simulation examples from the previous chapter, using a fine grid of the basis functions to parametrise the global environment in case of highly localised tracking data may lead to high variance estimates of individual scaling parameters. In order to avoid this, a sparse 4×4 grid of B-splines has been selected to approximate the hidden environment.

The output of the IMM smoother from the final iteration of the EM algorithm is presented in Figure 5.18. The colour bars of each plot correspond to the probability of being at a given mode. It can be observed that the more persistent tracks are assigned high probabilities (≥ 0.8) of being in the responsive mode. However, the number of such tracks is very limited, and most of them are confined to the central area of the fish body which means that they could have been removed from the tissue into the spinal cord. Overwhelming majority of neutrophils are identified as randomly migrating.

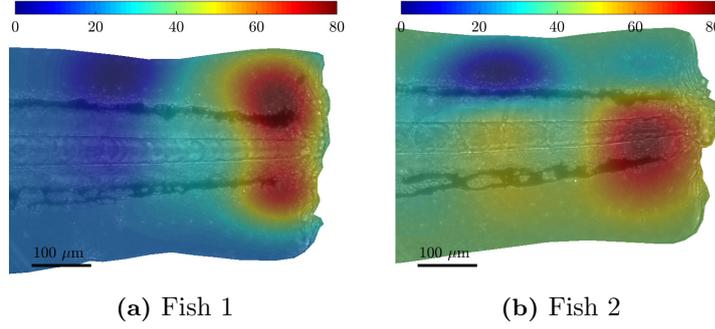


Figure 5.16: The chemoattractant concentration field estimated from the severe injury data. The colour bars are normalised to scale 0-80 a.u..

Dataset	$V_{N_b=16}$	$\text{Var}(\hat{\Theta})$
Fish 1	8.58×10^{28}	3.65×10^3
Fish 2	8.33×10^{27}	2.55×10^3

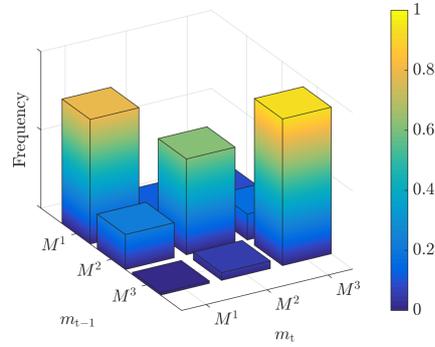


Table 5.6: Volumes of confidence regions and normalised total MLE variance for the severe injury datasets.

Figure 5.17: Computed mode transition frequencies for the severe injury datasets.

The histograms of the estimated horizontal velocity components for the neutrophils in the reverse migration dataset are shown in Figure 5.19. The plots in each row correspond to one of four larvae observed during the experiment. Note that the mean values of v_x are negative because the cells are migrating away from the wound. Unlike the case of active neutrophil recruitment, the distribution parameters of the horizontal velocity corresponding to the responsive mode are inconsistent from one zebrafish larva to another. In Fish 1 and 3, where more persistent tracks are observed, the mean estimate of v_x is comparable in magnitude with the mean of the recruited cells. In Fish 2 and 4, there is little difference between the responsive mode and the diffusing mode, however velocity distributions in the responsive mode appear to have larger variance. Note that for Fish 2 the bias of the mean v_x for the diffusing cells is larger than that of the actively migrating cells. The estimated velocities for the stationary cells are similar for all larvae.

Since very few neutrophils are identified as responsive, the tracking data that carries the information about the underlying chemorepellent concentration is rather

limited which leads to poor estimation accuracy. The environment inference results for the reverse migration data are presented in Figure 5.20. The inferred field for the Fish 1 and 3 demonstrates slight increase in concentration towards the wound area (ignoring the outlier peaks), however the gradient of the estimated field is rather shallow. Estimation results for Fish 2 and Fish 4 contradict the hypothesis of the chemorepellent concentration near the wound site. Because the EM algorithm estimates cell modes and the environment iteratively, the field inference results are directly connected with the number of responding neutrophils.

The neutrophils that are identified to be in the responsive mode are the main contributors to the information about the underlying environment, thus the variance of the field estimator is expected to be rather large. Volumes of the 16-dimensional confidence hyperellipsoids and total variances are presented in Table 5.7. It can be seen that the confidence region with the smallest volume corresponds to Fish 1 that has the largest number of persistent tracks in the dataset. The confidence regions indicate that Fish 3 and 4 are the least informative because the tracking data is confined to the tail fin. The total variance of parameter MLEs is inconsistent for different datasets. While for Fish 1 and 2 it is significantly smaller than for the previous datasets, a drastic increase for Fish 3 and 4 indicates large dispersion of MLEs.

5.4 Discussion

Contemporary therapeutic strategies for inflammation-mediated diseases have several directions for the development, including the following:

- Antagonising the pro-inflammatory signals in neutrophils by inhibiting participants of the PI3k signalling pathway [186].
- Inhibiting sensory activity on the neutrophil to the external chemokines [187], [188].

The first strategy motivates early inflammation resolution, while the latter reduces the sensitivity of neutrophils and limits the cell recruitment. Although both treatment strategies have enjoyed considerable success in some animal models, they rarely succeed in clinical trials because the mechanism of neutrophilic response to the environment is not fully understood [189], [190]. Identifying which intracellular mediators should be inhibited is a lengthy and expensive process especially when it involves *in vivo* experiments.

If a neutrophil has reduced sensory capacity, it may undergo several periods of desensitisation during the recruitment stage. It may be necessary to identify the length of the periods during which the cell remains responsive, as well as estimate

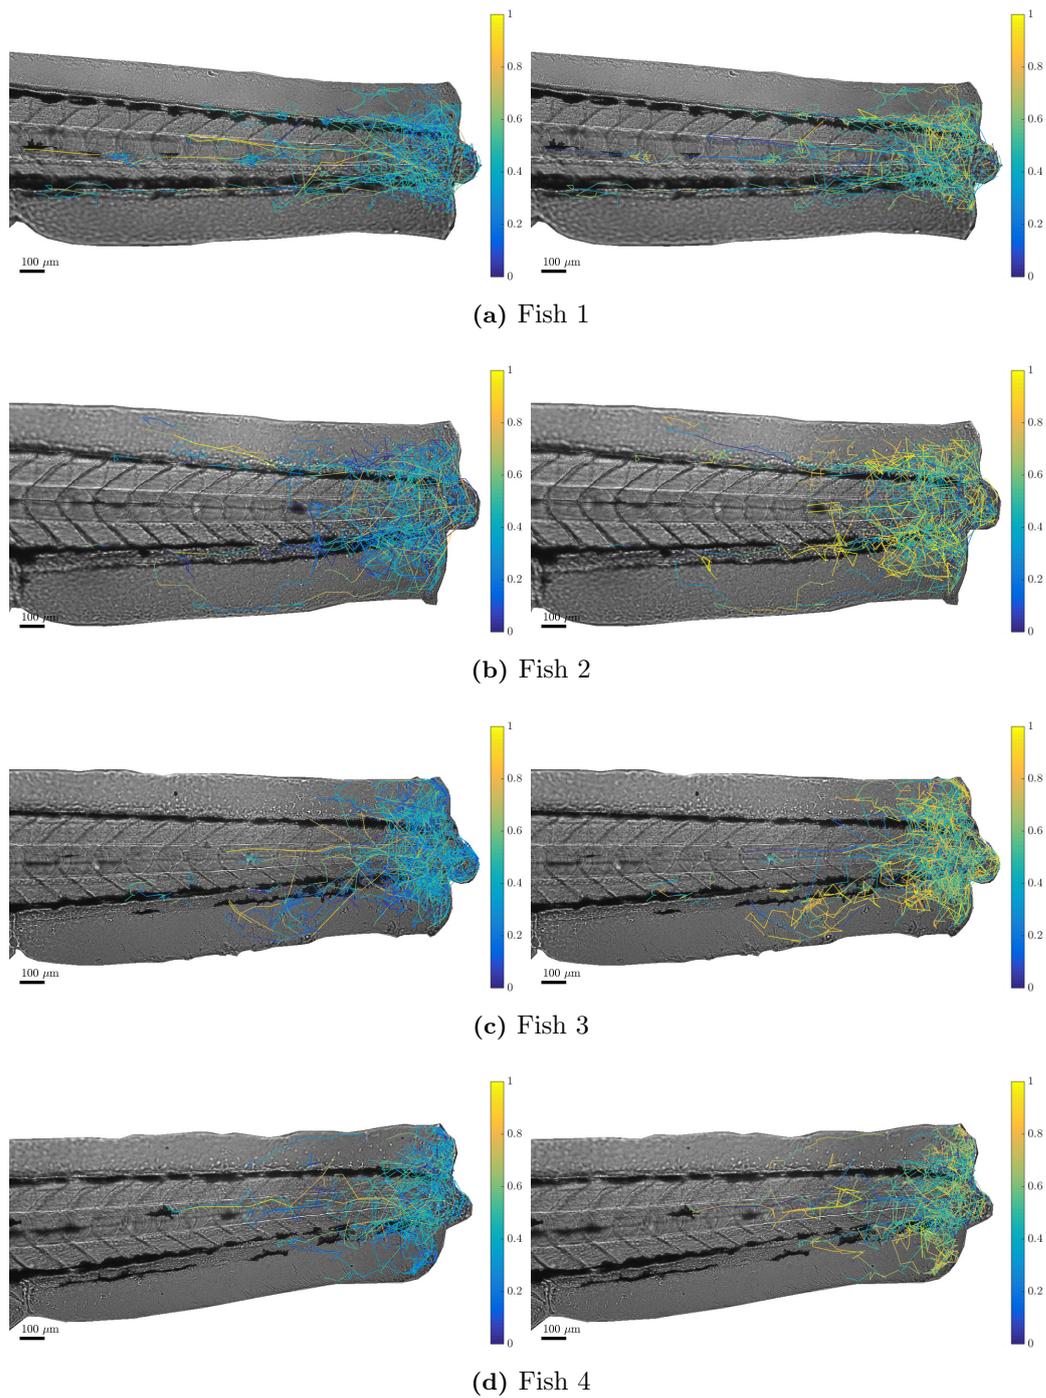


Figure 5.18: Migratory mode probabilities of neutrophils performing reverse migration. Plots in each row are ordered from left to right as follows: probabilities of M^1 , probabilities of M^2 . The cell tracks are colour coded in accordance to the probability of the cell being at a given mode.

the underlying concentration field that is driving it during such periods. The novelty of this analysis is that it does not restrict the neutrophils to one migratory mode while estimating the hidden environment. The hybrid model employed to describe

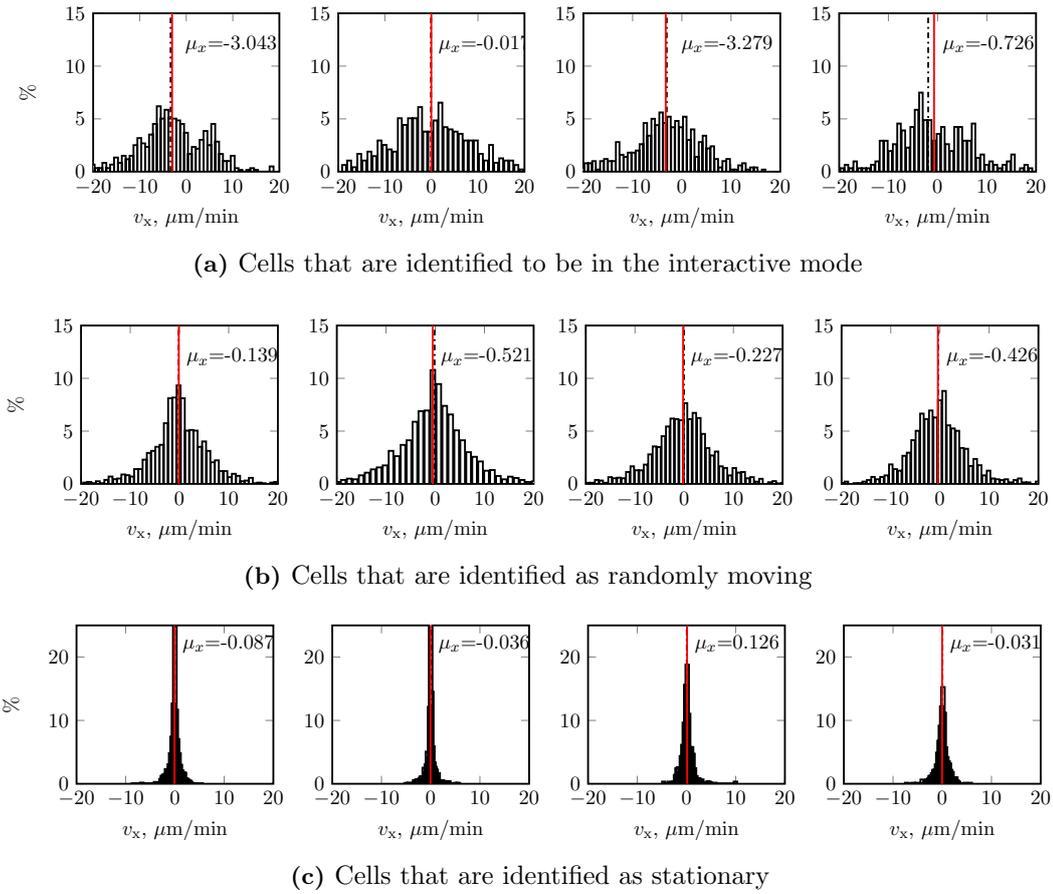


Figure 5.19: Comparison of cell velocity histograms for actively migrating and randomly moving neutrophils during the resolution stage of the inflammatory response. Each column corresponds to an individual larva in the dataset. **(a)** Histograms of the horizontal velocity component corresponding to neutrophils in zebrafish from the reverse migration set that are more likely to be in the active mode. **(b)** Histograms of the horizontal velocity component of neutrophils that are more likely to perform random walk. **(c)** Histograms of the horizontal velocity component of neutrophils that are more probable to be stationary.

the dynamics of an individual cell captures the heterogeneous behaviour of this cell by accounting for several possible modes at each time instance. Thus, the developed framework can be employed to infer the ChA gradients perceived by neutrophils subjected to different inhibiting treatments. It has the potential to be used as an *in silico* tool for preliminary assessment of the effectiveness of novel therapeutic strategies.

5.4.1 Estimation framework

Two methods of MM state estimation have been applied to the recruitment data. The pruning algorithm follows the most probable migratory mode sequence of a cell,

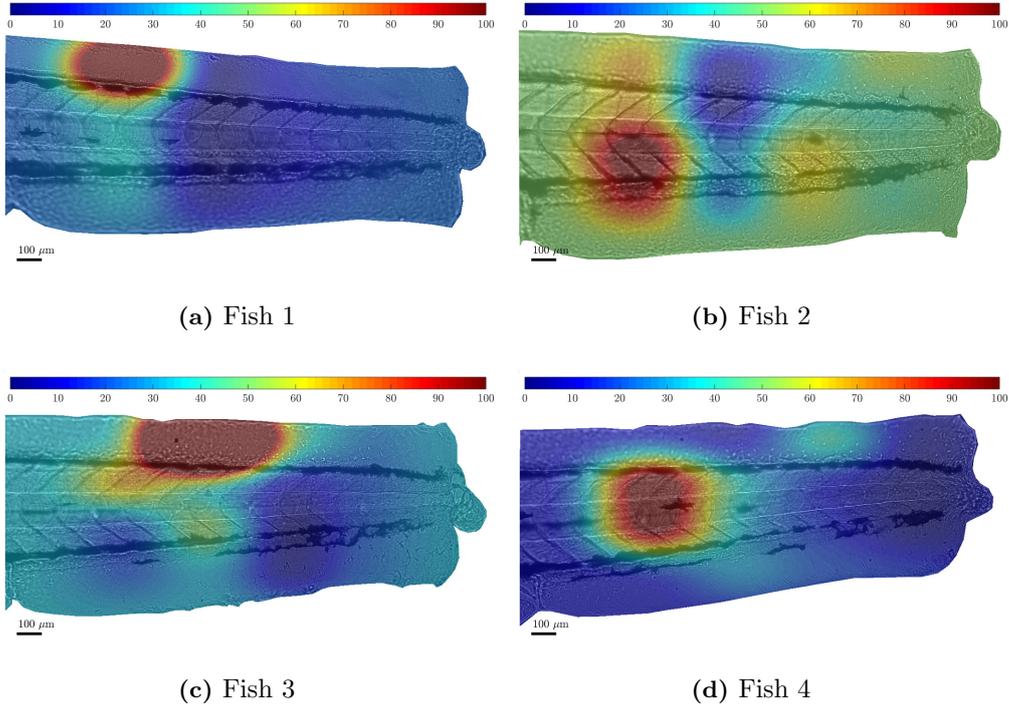


Figure 5.20: The chemoattractant concentration field estimated from the reverse migration data.

Dataset	$V_{N_b=16}$	$\text{Var}(\hat{\Theta})$
Fish 1	5.31×10^{25}	740.14
Fish 2	4.79×10^{28}	294.53
Fish 3	3.65×10^{32}	1.68×10^5
Fish 4	1.02×10^{34}	4.74×10^4

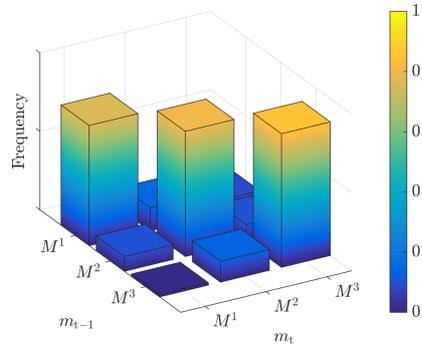


Table 5.7: Volumes of confidence regions and normalised total MLE variance for the reverse migration datasets.

Figure 5.21: Computed mode transition frequencies for the reverse migration datasets.

thus making a definite decision about its behaviour at each time. Only the cells that are identified as responsive contribute to the environment inference. Discriminating the cells with small probabilities of being in the responsive mode leads to a steeper gradient of the estimated chemoattractant field. While the pruning algorithm performance is close to that one of the IMM smoother in this instance, ignoring the less probable modes may have a negative effect on the field inference when the hybrid model describes more types of neutrophil behaviour. The IMM state estimation al-

gorithm outputs the probabilities of each migratory mode of an individual neutrophil at each time. The estimated mode probabilities can be viewed as the importance coefficients in the Gaussian mixture: the higher is the probability of a cell being in a responsive mode, the more information it contributes to the environment inference. This leads to more informed inference results even though the true dynamics of an individual migrating neutrophil is unknown. Hence, the IMM algorithm of hybrid state estimation is considered more appropriate for this application.

The comparison of pruning and merging estimation results for the recruitment data shows that the IMM easily discriminates between the selected modes and assigns very high probability of being in a responsive mode to the cells with clearly defined directional persistence. This is the result of using three easily distinguishable models with random components of different magnitudes. The mode identification quality is expected to worsen if one is to consider several SSMs with similar structure, *e.g.* several models of responding to the environment with varying levels bias and persistence, or pure RWs with varying noise magnitudes to model diffusive behaviour.

Despite the use of a sparse basis function grid, in several cases the ChA field was estimated with local errors, especially in the reverse migration set. This is the effect of using the cardinal basis functions with local support that was previously demonstrated in §4.4.2. Basis with local support only accounts for the high-frequency spatial change, meaning that the tracking data in a local area of the field only informs the estimation of those parameters that correspond to the B-splines that cover that area. A possible way to account for the low-frequency (global) changes is to create a multi-scale decomposition by superimposing a fine grid of smaller basis functions on a sparser grid, as has been done in [128]. However, this solution increases the dimension of the estimation problem and is likely to lead to higher estimation uncertainty. A more promising direction of the future model improvement is to use the hierarchical basis with non-uniform knot grid [134] or using B-spline wavelets as proposed in [191], [192].

5.4.2 Recruitment stage

The analysis of normal wound datasets reveals a common behavioural pattern in recruited neutrophils. Majority of the cells start in the stationary mode deep in the body. Once activated, neutrophils first diffuse away from the starting point in search of the gradient. Upon sensing an uneven concentration, they switch to the responsive mode and rapidly migrate up the gradient. Occasional changes back to the diffusive mode are seen in all fish; they may be the result of sporadic receptor desensitisation to the chemoattractant or obstacles within the tissue. Upon entering the wound site - the area starting at $\sim 100\mu m$ away from the wound, the neutrophils switch back to

the diffusive mode. It is not clear, however, whether the diffusive behaviour at the threat area is the result of neutrophil desensitisation to certain ChA gradients, or chemokinesis in the uniformly high chemoattractant concentration. If cells undergo apoptosis, they are identified as stationary. Thus, the mode M^3 efficiently describe both the cells that near the entry point into the tissue from the circulation in non-activated state and those that have died at the wound site. It can be viewed as a method of outlier detection, because stationary tracks do not contribute to the environment inference.

The estimates of the ChA field for all datasets have much smoother slopes than those obtained in Chapter 3 with the same field parametrisation. This captures the nature of slowly diffusing environments. Although the field inference results are presented in arbitrary units, the estimates for all zebrafish from the same experiment have similar range which demonstrates the consistency of the estimation framework.

While the use of a hybrid model certainly relaxes the assumptions made about the cell dynamics, the analysis carried out in this chapter considers one universal time-invariant chemoattractant field that is expected to drive neutrophils towards the wound. This expectation is contradicted by the results in the tail nick injury dataset, where a small outlier peak in ChA concentration has been inferred from several localised persistent tracks that are directed away from the wound. It has been shown experimentally that throughout the course of the inflammation recruitment neutrophils respond to various, sometimes competing, chemoattractants [193]. Chemoattractant concentration is also expected to diffuse over time [148], so relaxing the assumption about time invariance of the environment will enable us to make more detailed inferences about the process of neutrophil chemotaxis.

5.4.3 Resolution stage

Rather than discriminating between two models of the population dynamics, the framework utilised here accounts for the possibility of an individual cell to change its behaviour. Despite the use of a more flexible model, the estimation results of inflammation resolution data reveal that the neutrophils predominantly diffuse away from the wound, which confirms the findings of earlier works [17], [179].

It has been observed that some neutrophils can arrive to the wound even at 10 hours post injury, and they could be guided by the residual chemoattractant concentration. The framework used in this chapter could either identify these cells as simply not responding to the chemorepellent, or could infer that there is a high chemorepellent concentration in the fish body driving the immune cells outward towards the wound (see Figure 5.20). Thus, the possible extension to consider in the analysis of neutrophil reverse migration is to take into account two conflicting en-

vironments: the attracting environment inducing the chemotactic response in some neutrophils, and the repelling environment driving the cells away from the wound. Considering both chemoattractant and chemorepellent concentrations would require creating two non-matching bases to model each of the environments. Furthermore, a new hybrid model of individual neutrophil dynamics can be proposed with the following modes: cell responding to the attractant concentration, cell responding to the repellent concentration, randomly diffusing cell, and a stationary cell.

5.5 Summary

The analysis carried out in this chapter continues the investigation of neutrophil directed migration. Building on the potential field model of the environment influence and phenomenological model of cell dynamics introduced in Chapter 3, this work further relaxes the assumptions made about the cell-environment interaction prior to the parameter estimation by accounting for the switching nature of neutrophil behaviour highlighted in the experimental literature. The migratory behaviour of an individual neutrophil is described by a hybrid system with three modes: responding to the environment, randomly diffusing, or stationary. It is assumed that an individual cell moves in response to the environment only when it is in the first mode. Based on the observed neutrophil behaviour it has been concluded that each migratory mode can be reached from any other mode, thus the heterogeneous dynamics of a cell are described by the JMS.

The proposed model is calibrated against the tracking data obtained from the zebrafish model of inflammation via the approximate ML framework proposed in Chapter 4. The inference framework employs a combination of several well-known methods for the state estimation of the non-linear hybrid model. Utilising such model is the first attempt to analyse the interactions between the migrating neutrophils and the external environment without restricting the individual cell to a certain type of migratory behaviour. Two multiple model state estimators are compared: the IMM algorithm, that merges the preceding mode histories, and the M-best pruning algorithm, that follows the most probable mode sequence for each cell.

The estimation results for the inflammation recruitment stage demonstrate that majority of the neutrophils actively respond to the ChA gradient on their way to the wound, but upon arrival their behaviour becomes more erratic. The underlying ChA field estimated from the tracking data of responsive cells confirms the idea of the global slowly diffusing concentration, but two fish with tail fin injury demonstrate local competing sources chemoattractant. The estimation results for the reverse migration data support the claim made in [179] about the random nature of neutrophil migration away from the wound during the inflammation resolution stage. The

corresponding field estimates are indecisive and have high variance because they are not informed by the sufficient number of responsive cell tracks.

In summary, this chapter utilises the information contained in localised tracking data of individual neutrophils to infer the shape of the global environment that governs the migration of the entire cell population, thus bringing together two scales of the chemotaxis modelling. The questions that remains open is how does an individual neutrophil translates the perceived ChA concentration gradient into the reorientation of its leading edge or other severe shape changes. The next chapter attempts to answer this question by analysing the subcellular concentrations of the migrating neutrophil observed *in vivo*.

6 | Statistical analysis of neutrophil morphodynamics regulated by the subcellular signalling

Previous chapters of this thesis deal with high-level models that treat an individual cell as a material point that can commit to different types of motion. Extensive biological data is available for the development and validation of these macro-scale models, so they are suitable for prediction and inference of the unobserved aspects of chemotaxis. Whole-cell studies, both theoretical and experimental, provide no insight into the mechanism of translating the environmental signals perceived by the chemotaxing cell to its locomotion. Characterisation of this process is crucial for advancing our understanding of neutrophil function. Moreover, the identification of key players in this translation process may lead to the development of more precise treatment strategies that can target specific proteins within the cell without affecting its other functionality.

This chapter looks into one aspect of the chemoattractant-induced cell locomotion: the relationship between the concentration of subcellular signalling mediators and the changes of the cell membrane. An image-processing framework is developed for extracting the information about the subcellular concentration. Coupled with the active contour tracking of the cell boundary, it relates local shape changes to the concentration of the local activator. The state space model (SSM) of the cell membrane evolutions is tested against neutrophil *in vivo* data via a series of correlation tests. The generic EM algorithm is then employed to estimate the model parameters based on the quantitative data obtained from the image processing modules.

6.1 Background

Mathematical modelling is deservedly recognised as a standard method of studying molecular interactions within the cell [194]. Modelling on the subcellular level often fills in for the missing specific experimental data by providing hypotheses about

the underlying molecular relationships that are considered in the future experiment design. Several types of theoretical models of eukaryotic cell motility that consider subcellular processes are reviewed in [41]. The first type deals exclusively with biochemical characteristics of the cytoskeleton in order to understand the full role of actin in the cell shape change and locomotion. The second type aims to describe the signalling activity and its relation to forces facilitating cell motion. The third type combines the two previous approaches.

The modelling literature is dominated by two hypotheses about the subcellular mechanism driving eukaryotic chemotaxis: the “internal compass” and “pseudopod-centered” models [69]. The internal compass refers to a localised signal inside the cell that represents the direction of the external ChA gradient. Formation of new pseudopods is strongly regulated by the extracellular environment. For example, if the direction of the steepest gradient differs from the current direction of the cell, the new pseudopod is formed to align with the gradient. This paradigm implies direct causal relationship between the environment sensing and the polarisation. This relationship is established mathematically in a series of models employing local excitation, global inhibition (LEGI) mechanism [59]–[61].

The hypothesis of pseudopod-led regulatory mechanism was first advocated by Insall in [69], where he argues that the future direction of cell polarisation is largely defined by the current location of pseudopods or the leading edge. In subsequent works [66], [68] he and colleagues have developed a mathematical model where the external signal processing is implicit and only regulates the rate of subcellular processes which are modelled by the Meinhardt RDS [58] previously discussed in §2.2.3 of the literature review. This means that new pseudopod generation is the result of evolution and splitting of the existing ones, rather than that of the environment sensing. This model can potentially explain the phenomenon of behavioural polarisation discussed in [29]. While pseudopod-led paradigm presumes different mechanism of signal processing, its pattern formation process leading to cell polarisation is similar to the one described in LEGI model. The difficulty of finding biological candidates for the agents of the theoretical RDS hinders validation of these models.

Unlike cell population or even single cell models, the models of intracellular dynamics are rarely phenomenological, and the reasons for it are two-fold. Firstly, the experimental methods are often not sophisticated enough to validate very detailed relationships proposed in these models. On the other hand, some of the aspects of the model may not be connected to real-life phenomena. For example, models based on finite-element decomposition popular in modelling the cell shape and bulk [65], [66] lack the explicit connection between the model parameters and the underlying physical processes. Although they can predict the changes of the shape to a certain degree of accuracy, they do not explain the underlying dynamical processes driving

the severe shape changes of the migrating cell.

In neutrophils, chemoattractant receptors are uniformly distributed along the cell membrane. Gradient sensing is facilitated by rearrangement of the signalling lipids that bind with the receptor-associated G proteins and activate the downstream signalling pathways. As was discussed in §2.1.2 of the literature review, in the presence of the shallow external gradient, neutrophils selectively accumulate lipid products of PI3k near the membrane where there is more receptor activity. One of the key lipids that accumulate at the up-gradient edge are the PI(3,4,5)P₃ – self-catalytic second messengers that initiate activation of multiple downstream signalling cascades. The self-amplifying internal gradient of PIP₃ is much steeper than that of the external ChA. This lipid is considered as a strong candidate for the role of local activator in both LEGI and pseudopod-centred model as it has been proven to trigger the polarisation and cause the bias in random motility of neutrophils [19], [69].

The crucial role of PIP₃ in directed motion has been shown in multiple experimental studies [195], [196]. However, there is no quantitative framework in place that would allow investigating the causal relationship between spatially varying PIP₃ concentration and local changes in the cell shape. The focus of this chapter is to provide a missing link between the theory that delegates the role of a local activator to PIP₃ and the experimental data that shows its distribution along the cell membrane.

A common way to assess the distribution of PIP₃ within the cell is to track the GFP-tagged pleckstrin homology (PH) domain of Akt (PH-Akt-GFP). Akt is accumulated near the membrane and activated via the engagement of its PH domain with PIP₃, so the observed concentration of its molecules is indicative of the underlying concentration of PIP₃ [197]. This chapter analyses the time-lapse images of the PH-Akt-GFP neutrophils observed *in vivo* under the assumption that the local concentration of PIP₃ is proportional to the measurable fluorescence of PH-Akt along the tracked cell boundary.

In summary, this chapter tests the causal relationship between redistribution of PIP₃ induced by the external ChA gradient and local changes of the cell shape. Several computational modules developed for extracting the quantitative data are discussed in Section 6.2. The correlation analysis of raw quantitative data is conducted in Section 6.3 to establish whether there is a strong connection between the measured variables. A simple model of the cell membrane evolution is then presented in a state space form where the state is the local normal velocity of the cell membrane. The introduced model allows to employ the generic EM algorithm for simultaneous smoothing of local velocities and estimation of the model coefficients. The estimation framework is applied to motile neutrophil with multiple protrusions

observed *in vivo* and to the polarised neutrophil-like immortalised cell from [27] observed *in vitro*.

6.2 Methods

6.2.1 Data acquisition

The subcellular activity within recruited neutrophils has been observed *in vivo* using the Tg(lyz:PH-Akt-EGFP) zebrafish line. The injured larvae were observed via time-lapse imaging with an inverted Olympus IX81 microscope, using an Olympus UPlanSApo 60x oil Immersion objective (Germany). The cells were tracked at a specific region near the wound site using the 488 nm laser line that imaged the zebrafish on 20 different depths (along the Z scale). The Z-stack images were converted into the single maximum exposure layer within imageJ software. The compressed 2-D images of several neutrophils are presented in Figure 6.1.

6.2.2 Model of cell boundary evolution

Let us define the cell boundary as an evolving closed curve $\Gamma(t)$, whose motion is governed locally by the interaction of forces acting normal to the boundary. The full model of the boundary evolution is described in detail in §2.2.4. In order to ensure model identifiability, this chapter examines a simplified model and considers only the following forces locally acting on the cell membrane

$$\mathcal{F} = (\mathcal{F}_{\text{pro}} + \mathcal{F}_{\text{visc}} + \mathcal{F}_{\text{ten}} + \mathcal{F}_{\text{vol}})\nu, \quad (6.2.1)$$

where ν is a local outward pointing normal to the segment ds of the curve $\Gamma(t)$, the \mathcal{F}_{pro} is the protrusion force that is proportional to the local concentration of the signalling proteins and lipids

$$\mathcal{F}_{\text{pro}} = \sum_i^{N_a} k_{a,i} a_i. \quad (6.2.2)$$

Other forces in (6.2.1) are defined in §2.2.4. $\mathcal{F}_{\text{visc}}$ is the viscous force caused by the resistance of the environment to the membrane protrusion, and forces \mathcal{F}_{ten} and \mathcal{F}_{vol} are the result of local membrane tension and volume conservation, respectively. The experimental data provides information about only one signalling lipid, PIP₃, so the analysis in this chapter is conducted under the following defining assumption:

Assumption 6.2.1: *PIP₃ is the only activator regulating the cell protrusion, e.g. in equation (6.2.2) $N_a = 1$.*

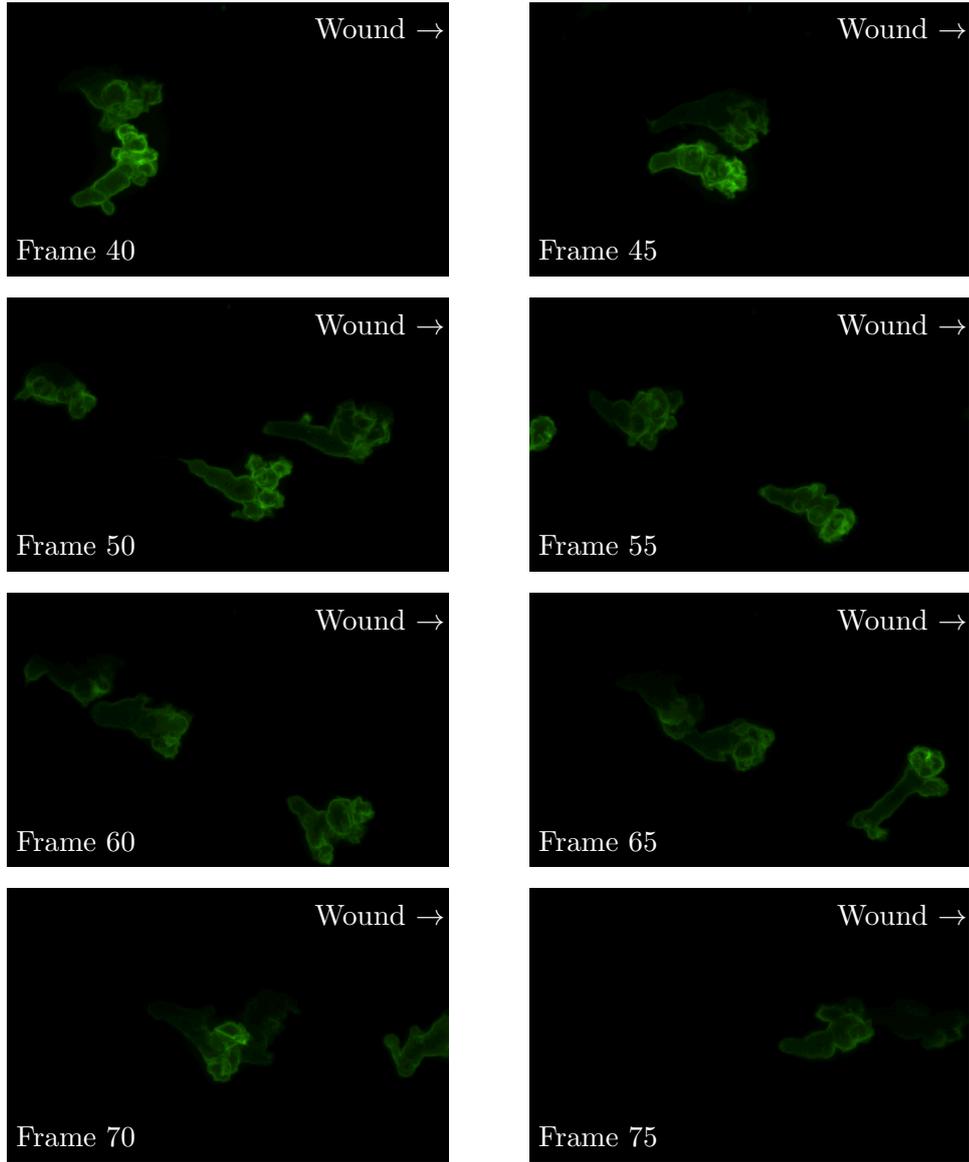


Figure 6.1: Example time-lapse images of the PH-Akt-GFP neutrophils in the zebrafish body migrating towards the tail fin wound.

The computational model proposed in [65] implies that the the force balance holds at each time instance, t , for each segment ds of the curve $\Gamma(t)$

$$\mathcal{F}_t(ds) = 0, \quad (6.2.3)$$

Expanding constituent terms of the above equation renders the expression that analytically links normal velocity of the cell membrane and local concentration of PIP₃:

$$v_\nu(ds) = \left(\frac{(k_a)}{k_{\text{visc}}} a_t(ds) + \frac{\gamma}{k_{\text{visc}}} \kappa_t(ds) + \frac{k_{\text{vol}}}{k_{\text{visc}}} (\mathcal{A}_t - \mathcal{A}_0) \right) \nu, \quad (6.2.4)$$

where $a_t(ds)$ denotes the local concentration of the PIP₃ at the segment $ds \in \Gamma_t$ at a given time. The interpretation of time-lapse cell images deals with the projection of

3-D cell bulk on a 2-D plane, therefore using force balance equation should be viewed as approximate. The discrepancy between the mathematical model of the normally acting force considered here and the actual local force can be represented by random disturbances in local normal velocity change. Instead of obtaining the velocity directly from algebraic the relation, one could formulate the dynamical model that employs the following discrete-time ODE for the normal velocity component

$$\frac{dv_\nu(ds)}{dt} = \frac{1}{m} \mathcal{F}_t(ds) + \mathbf{w}_t(ds), \quad (6.2.5)$$

where m is an arbitrary coefficient corresponding to the material mass of the local membrane area and where the random component $\mathbf{w}_t(ds) \sim \mathcal{N}(0, Q_w)$ accounts for the acceleration caused by the forces not considered in the simplified model. Expanding the constituent terms yields the following expression

$$\frac{dv_{\nu,t}(ds)}{dt} = \frac{1}{m} \left(k_a a_t(ds) + \gamma \kappa_t(ds) + k_{\text{vol}}(\mathcal{A}_t - \mathcal{A}_0) - k_{\text{visc}} v_{\nu,t}(ds) \right) + \mathbf{w}_t(ds). \quad (6.2.6)$$

The tangent component of the local segment motion is not considered in here in accordance to the original model proposed in [65]. Hence, the dynamical model of the cell membrane evolution relies on the following quantities:

- Local normal velocity of the cell boundary, $v_\nu(ds)$;
- Local curvature of the cell boundary, $\kappa(ds)$;
- Local PIP₃ concentration along the cell boundary, $a(ds)$;
- The area enclosed by the observed cell boundary, \mathcal{A} .

The remainder of this section summarises the methods for extracting these quantities from the time-lapse images.

6.2.3 Cell boundary tracking

At first, the time-lapse images are processed using the the active contour tracking tool developed in [198]. Within the tracking framework, cell contours are first identified using the watershed algorithm that highlights the areas of the image. The boundary is then parametrised with a collection of cubic B-splines placed on K equidistant support nodes. These B-splines compose the active contour of the cell that is utilised in constructing the second-order spatio-temporal model of the individual cell migration. In this model, the measurement vector consists of the boundary segments, and the hidden state consists of the positions and velocities of the support nodes. The tracking is then performed by applying the Kalman smoother to the constructed kinetic model with an intermediate procedure of the cell correspondence identification. The cell correspondence is estimated via a common nearest neighbour method [199].

The output of the tracking module is a time series of cell boundaries approximated with a collection of discrete nodes $\mathbf{s}_t^k = [s_x^k, s_y^k]$, $k = 1, \dots, K$. Each node is assigned a velocity vector $\mathbf{v}_t^k = [v_x^k, v_y^k]$ components of which are defined in the absolute reference frame. The model of cell membrane evolution considers only the component of membrane velocity that is normal to the cell boundary. Thus, for the purposes of the analysis carried out in this chapter, velocity vectors of each node must be projected onto the new reference frame relative to the cell boundary. The set of nodes produced by the tracking algorithm $\mathbf{s}_t^k = [s_x^k, s_y^k]$, $k = 1, \dots, K$ can approximate the cell boundary at an individual frame so it can be viewed as a polygon while computing local normals (see Algorithm C.1 in Appendix C). The projection of the velocity vector on the local reference frame produces a new vector $\mathbf{v}_t^k = [v_\nu^k, v_\tau^k]$. The discrete approximation of the boundary also permits straightforward computation of Cartesian curvature, $\kappa(\mathbf{s}_t^k)$, and the area enclosed by the cell membrane, \mathcal{A}_t , using Algorithm C.2 in Appendix C, thus providing all quantities required to construct the discrete-time SSM and test it against the experimental data.

6.2.4 Extracting PIP₃ concentration data from the time-lapse images

It is necessary to extract the information about the local concentration of PIP₃ along the cell membrane and associate it with the tracking data prior to the analysis. Since the imaging data contains fluorescence intensity of the PH domain of Akt that binds to PIP₃ molecules, in stable activation process one can assume proportional relationship between the observed concentration of Akt and the hidden concentration of PIP₃.

Assumption 6.2.2: *The integrated fluorescence intensity obtained from the imaging data is proportional to the local PIP₃ concentration.*

The procedure of extracting local intensity from one frame is implemented as follows. The RGB image of the cell is first converted into the greyscale colour map. Because the GFP-tagged cells are exposed only in the green channel (R = 0, B=0), the resultant image has luminance directly proportional to the original. This means that the relative difference in fluorescence intensity is preserved. Next, the membrane of a cell is identified via mathematical topology method of edge detection [200]. The fluorescence intensity of each pixel along the detected edge is extracted from the image with a threshold of $\eta = 50$ arbitrary units (a.u.) to discriminate the background fluorescence. Finally, the pixels with brightness above the threshold are associated with the nodes available from the active contour tracking module using the nearest neighbour method. The full procedure is summarised in Algorithm 6.1.

Algorithm 6.1 Extraction of local fluorescence intensity along the cell boundary

Input: Sequence of time-lapse images of the migrating neutrophil, Im_t , $t = 1, \dots, T$; sequence of discretised neutrophil boundaries from the tracking module, $s_t^k = \{s_x, s_y\}_t^k$, $t = 1, \dots, T, k = 1, \dots, K$.

Output: Local fluorescence intensity associated with the boundary nodes, $a_t^k, t = 1, \dots, T, k = 1, \dots, K$.

1: **for** $t \leftarrow 1, T$ **do**

2: Convert the image into the greyscale: $\text{Im}_t \rightarrow \text{Gs}_t$;

3: Erode the image: $\text{Gs}_t \ominus \text{Se}$, where Se is a flat structuring element;

4: Obtain the cell boundary image:

$$\text{B}_t = \text{Gs}_t - \text{Gs}_t \ominus \text{Se};$$

5: Identify coordinates of all pixels on the edge with intensity higher than a threshold η :

$$\text{P}_t \{ \text{px} \in \text{B}_t : a(\text{px}) \geq \eta \},$$

where px is an individual pixel and $a(\text{px})$ is the intensity of the pixel;

6: **for** $k \leftarrow 1, K$ **do**

7: Associate the pixels from the set P_t^k with the node s_t^k on the boundary via the NN method;

8: Compute the integrated local intensity of the area

$$a(s_{t-1}^k) = \sum_{\text{P}_t^k} a(\text{px});$$

9: **end for**

10: **end for**

6.2.5 Estimating parameters of the boundary evolution model

Because the mathematical model introduced above considers only the normal velocity components, there is no coupling between the neighbouring nodes and they can be considered as individual points in the population. The discrete-time discrete-space dynamical model that describes the behaviour of an individual point arising from the (6.2.6) is expressed as follows

$$v_\nu(s_t^k) = v_\nu(s_{t-1}^k) - r_v v_\nu(s_{t-1}^k) + r_a a(s_{t-1}^k) + r_\gamma \kappa(s_t^k) + r_{\text{vol}} (\mathcal{A}_t - \mathcal{A}_0), \quad (6.2.7)$$

for $k = 1, \dots, K$ and $t = 1, \dots, T$, where $r_v = k_{\text{visc}} T/m$, $r_a = k_a T/m$, $r_\gamma = \gamma T/m$, and $r_{\text{vol}} = k_{\text{vol}} T/m$ are the unknown model coefficients and where T is the time increment between the frames. The corresponding state-space representation of (6.2.7) is obtained by assuming that the normal velocity of an individual node is a one-dimensional state

$$\mathbf{x}_t^k = \left[v_\nu(s_t^k) \right],$$

and considering the rest of the measured variables as the constituent terms of the deterministic input vector

$$\mathbf{u}_t^k = \left[a(s_t^k), \kappa(s_t^k), (\mathcal{A}_t - \mathcal{A}_0) \right]^\top.$$

Then the model with the familiar structure introduced in §2.3.2 takes a simplified form

$$\mathbf{x}_t = A_t \mathbf{x}_{t-1} + B_t \mathbf{u}_{t-1} + \mathbf{w}_{t-1}, \quad \mathbf{w}_t \sim \mathcal{N}(0, Q_w) \quad (6.2.8a)$$

$$\mathbf{y}_t = C_t \mathbf{x}_t + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(0, R_v) \quad (6.2.8b)$$

with the following dynamical matrices:

$$A = [1 - r_v]; B = [r_a \quad r_\gamma \quad r_{\text{vol}}]; C = [1].$$

The process noise in this case accounts for possible random perturbations of the membrane. Since cell membrane velocities are already smoothed by the tracking module, no additional measurement noise is considered in the model.

Assumption 6.2.3: *The measurement process is noiseless, e.g. $R_v = 0$.*

The parameters of interest in this chapter are the following:

$$\Theta = \{A, B, Q_w\}. \quad (6.2.9)$$

Given the unit measurement matrix C , Assumption 6.2.3 renders the model (6.2.8) a first order autoregressive model (AR(1)) that has no latent variables. The estimate of the unknown parameter vector Θ can thus be straightforwardly obtained via classical Maximum Likelihood (ML) algorithm that is introduced in §2.4.2. The likelihood function for the unknown parameters of the linear Gaussian AR(1) model is defined as follows

$$\begin{aligned} \mathcal{L}(\Theta) &= \prod_{k=1}^K \prod_{t=1}^T \frac{1}{\sqrt{2\pi Q_w}} \times \\ &\times \exp \left\{ -\frac{1}{2} \left(\mathbf{x}_t^k - A \mathbf{x}_{t-1}^k - B \mathbf{u}_{t-1}^k \right)^\top Q_w^{-1} \left(\mathbf{x}_t^k - A \mathbf{x}_{t-1}^k - B \mathbf{u}_{t-1}^k \right) \right\}. \end{aligned} \quad (6.2.10)$$

Taking the logarithm of (6.2.10) renders

$$\begin{aligned} \log \mathcal{L}(\Theta) &= -\frac{1}{2} \left\{ K T \log |Q_w| + Q_w^{-1} \text{tr} (S_{xx} - A S_{xb} - \right. \\ &\quad - S_{xb} A^\top + A S_{bb} A^\top - B S_{ux} - S_{xu} B^\top + \\ &\quad \left. + B S_{uu} B^\top + B S_{uxb} A^\top + A S_{xbu} B^\top) \right\} + c, \end{aligned} \quad (6.2.11)$$

where constant c includes all terms independent of the unknown parameter Θ . The constituent terms of (6.2.11) are defined as follows

$$S_{xx} = \sum_{k=1}^K \sum_{t=1}^T \mathbf{x}_t^k (\mathbf{x}_t^k)^\top. \quad (6.2.12a)$$

$$S_{\mathbf{x}\mathbf{x}\mathbf{b}} = \sum_{k=1}^K \sum_{t=1}^T \mathbf{x}_t^k (\mathbf{x}_{t-1}^k)^\top, \quad S_{\mathbf{x}\mathbf{b}\mathbf{x}} = S_{\mathbf{x}\mathbf{x}\mathbf{b}}^\top. \quad (6.2.12\text{b})$$

$$S_{\mathbf{b}\mathbf{b}} = \sum_{k=1}^K \sum_{t=1}^T \mathbf{x}_{t-1}^k (\mathbf{x}_{t-1}^k)^\top. \quad (6.2.12\text{c})$$

$$S_{\mathbf{u}\mathbf{u}} = \sum_{k=1}^K \sum_{t=1}^T \mathbf{u}_{t-1}^k (\mathbf{u}_{t-1}^k)^\top. \quad (6.2.12\text{d})$$

$$S_{\mathbf{x}\mathbf{u}} = \sum_{k=1}^K \sum_{t=1}^T \mathbf{x}_t^k (\mathbf{u}_{t-1}^k)^\top, \quad S_{\mathbf{u}\mathbf{x}} = S_{\mathbf{x}\mathbf{u}}^\top. \quad (6.2.12\text{e})$$

$$S_{\mathbf{x}\mathbf{b}\mathbf{u}} = \sum_{k=1}^K \sum_{t=1}^T \hat{\mathbf{x}}_{t-1}^k (\mathbf{u}_{t-1}^k)^\top, \quad S_{\mathbf{u}\mathbf{x}\mathbf{b}} = S_{\mathbf{x}\mathbf{b}\mathbf{u}}^\top. \quad (6.2.12\text{f})$$

The parameter MLEs are then straightforwardly by taking partial derivatives of the log-likelihood function and setting them equal to zero. To obtain the dynamics matrices, set

$$\begin{cases} \frac{\partial \log \mathcal{L}(\Theta)}{\partial A} = -S_{\mathbf{x}\mathbf{x}\mathbf{b}} + S_{\mathbf{x}\mathbf{u}} B^\top + S_{\mathbf{b}\mathbf{b}} A^\top = 0; \\ \frac{\partial \log \mathcal{L}(\Theta)}{\partial B} = -S_{\mathbf{x}\mathbf{u}} + A S_{\mathbf{x}\mathbf{b}\mathbf{u}} + B S_{\mathbf{u}\mathbf{u}} = 0. \end{cases} \quad (6.2.13)$$

Then solving the linear system results in the following expression

$$\begin{bmatrix} \hat{A} & \hat{B} \end{bmatrix} = \begin{bmatrix} S_{\mathbf{x}\mathbf{x}\mathbf{b}} & S_{\mathbf{x}\mathbf{u}} \end{bmatrix} \begin{bmatrix} S_{\mathbf{b}\mathbf{b}} & S_{\mathbf{u}\mathbf{x}\mathbf{b}} \\ S_{\mathbf{x}\mathbf{b}\mathbf{u}} & S_{\mathbf{u}\mathbf{u}} \end{bmatrix}^{-1}. \quad (6.2.14)$$

To obtain the MLE of the noise covariance matrix set

$$\frac{\partial \log \mathcal{L}(\Theta)}{\partial Q_w} = K T Q_w^{-1} + Q_w^{-2} \{\zeta\} = 0, \quad (6.2.15)$$

where ζ denotes the quadratic form within the trace operator in (6.2.11). The resultant estimate is then given by

$$\begin{aligned} \hat{Q}_w = \frac{1}{K T} \{ & S_{\mathbf{x}\mathbf{x}} - \hat{A} S_{\mathbf{x}\mathbf{x}\mathbf{b}}^\top - S_{\mathbf{x}\mathbf{x}\mathbf{b}} \hat{A}^\top + \hat{A} S_{\mathbf{b}\mathbf{b}} \hat{A}^\top - \\ & - \hat{B} S_{\mathbf{x}\mathbf{u}}^\top - S_{\mathbf{x}\mathbf{u}} \hat{B}^\top + \hat{B} S_{\mathbf{u}\mathbf{u}}^\top + \hat{B} S_{\mathbf{u}\mathbf{x}\mathbf{b}} \hat{A}^\top + \hat{A} S_{\mathbf{u}\mathbf{x}\mathbf{b}}^\top \hat{B}^\top \}. \end{aligned} \quad (6.2.16)$$

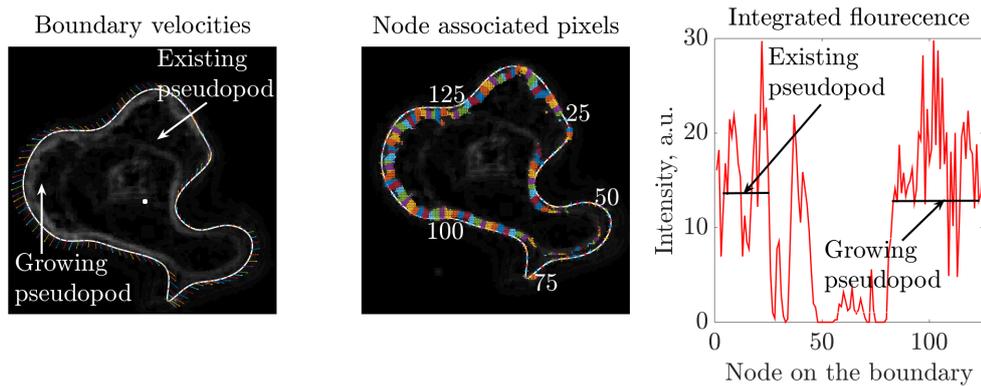
The parameters of interest are estimated in an iterative scheme.

6.3 Results

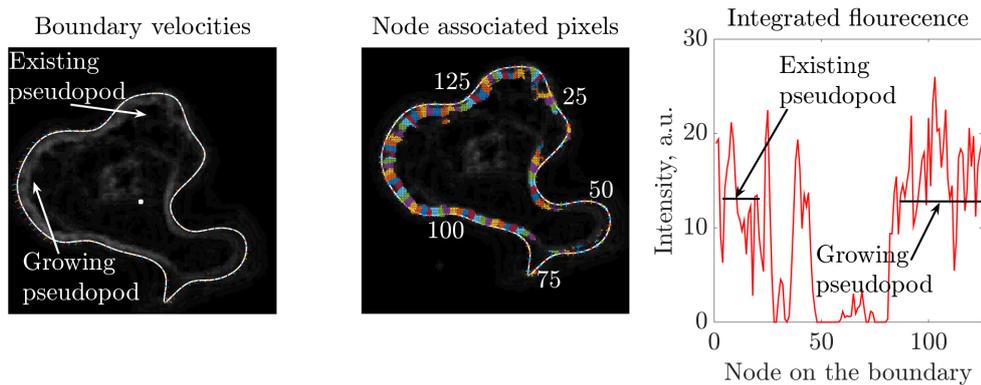
6.3.1 Image processing results

The results of cell boundary tracking and node-association of local fluorescence intensity are presented in Figure 6.2. The cell membrane at each frame is approximated with $K = 125$ control points within the tracking module. Local membrane velocities

and detected pixels with high fluorescence intensity are plotted against the cell edge that has been extracted using image erosion. The pixels on the cell edge that are clustered using Algorithm 6.1 are colour coded according to the node on the boundary with which they are associated. It can be seen from the integrated intensity plot that the brightest fluorescence corresponds to the areas of cell protrusion. However, the intensity peaks locally at the area between nodes 25 and 50 that is retracted.



(a) Image processing results: frame 10



(b) Image processing results: frame 11

Figure 6.2: Processing results of two subsequent frames of a migrating neutrophil. The left image shows node velocities produced by the boundary tracking module. The white point denotes the estimated position of the cell centroid. The middle image shows the pixels along the identified edge associated with each node on the cell boundary. The right image shows the local integrated intensity plotted against the node index.

addition, the tracking module provides the estimates of cell centroid position and velocity. In the shown time frames the centroid velocity is rather small, meaning that the neutrophil at this instance is immobile and assessing the environment by extending an additional pseudopod in the direction opposite from the existing one.

Next, the estimated velocities are projected on the local reference frames at the vertices of the polygon that approximates the cell membrane as presented in Figure 6.3a. The projections of velocity vectors obtained from the tracking module onto

the local reference frame are illustrated in Figure 6.3b. It can be seen on the figure that both the retraction and protrusion of the cell membrane are characterised by large normal velocity component and small tangent component, which allows us to speculate that the local normal velocity is a sufficient measure of the local membrane change.

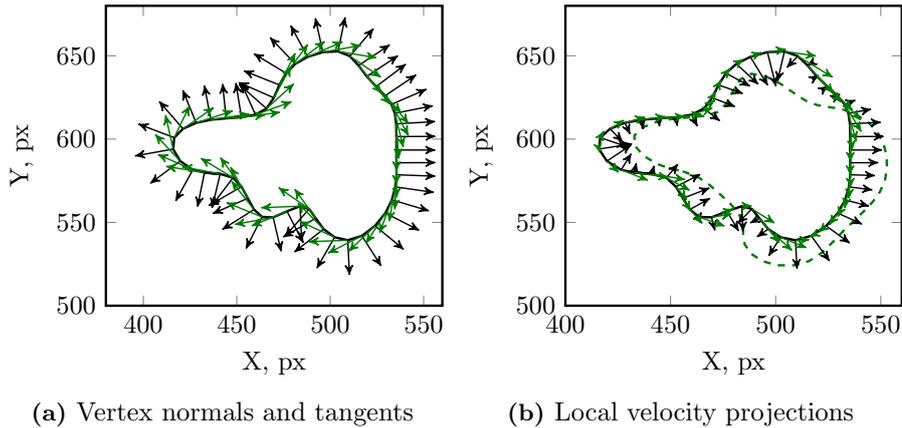


Figure 6.3: Vertex bases and projections of local velocities. **(a)** The outward-facing normal unit vectors with corresponding tangent vectors are computed for each node. The plotted vectors are scaled for visualisation. **(b)** Each velocity vector at time t estimated via the tracking module is projected on the reference frame centred at the corresponding vertex to obtain a normal component (—) and a tangent component (—). The cell membrane at time $t+1$ is defined by the dashed line (---).

As a result of processing two *in vivo* image sequences, 8 datasets have been obtained corresponding to individual neutrophils. Each dataset contains spatio-temporally resolved measurements of the local normal velocity, $v_\nu(s_t^k)$, integrated fluorescence intensity, $a(s_t^k)$, local curvature, $\kappa(s_t^k)$, $k = 1, \dots, K$, and the cell area, \mathcal{A}_t . These quantities are used below to identify the parameters of the proposed dynamical model.

6.3.2 Correlation analysis results

First, the raw quantitative data extracted from the time-lapse images is subjected to the correlation analysis in order to assess whether PIP_3 has strong association with the normal velocity at the following time instance. The correlation plots for 8 cells are presented in Figure 6.4. The values of r_a for each cell are calculated via the standard linear regression. It can be seen that the estimated slope varies for cells 1-5 that all belong to the same imaging sequence, while the results for cells 6-8 are more consistent. However, the scatter plots for all cells shows that the extracted intensity values have rather intermittent nature: even in the nodes with high normal

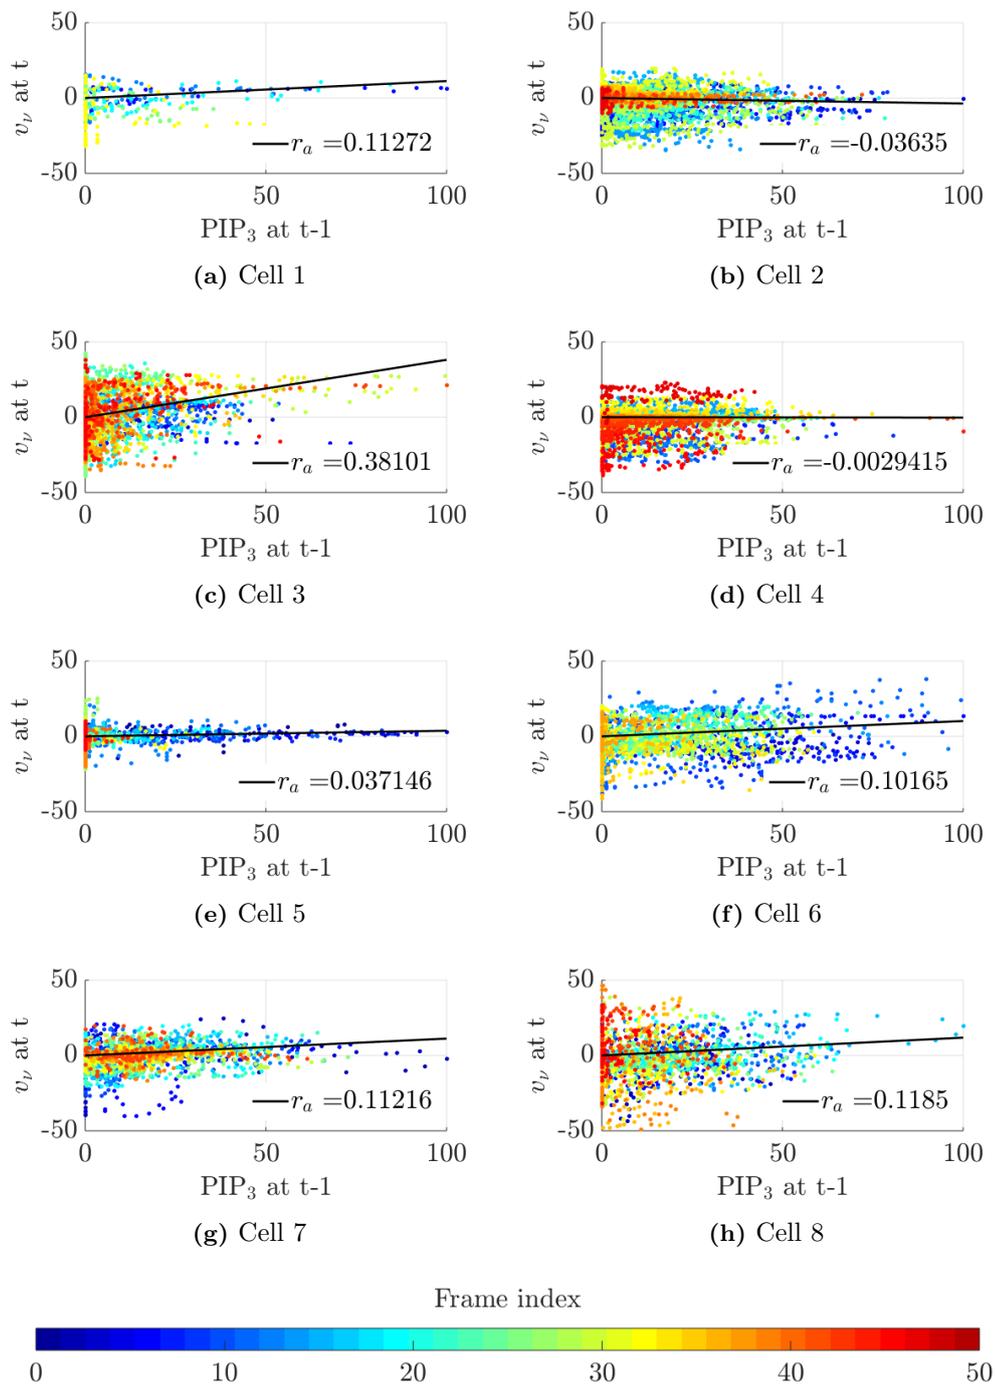


Figure 6.4: Correlation plots of local normal velocity and local concentration of PIP_3 over the entire cell boundary. Scatter plots are colour coded according to the index of the frame from which the values have been extracted. The provided colour bar demonstrates colour-to-frame association. The black line in each plot illustrates the estimated slope of linear regression.

velocity, the local intensity is zero. This is expected to have negative effect on the estimation results.

Since local normal velocities demonstrate rather weak correlation with the extracted intensity distribution along the whole cell boundary, it is interesting to see whether individual segments of the membrane with higher concentration of PIP₃ show more active protrusion. For this analysis, the nodes on the cell membrane boundary have been classified into four groups based on their fluorescence intensity and its first time derivative. First, the areas with high intensity are thresholded at 35 a.u., and then the areas with positive and negative time derivatives are identified. Figure 6.5 demonstrates the distribution of cell velocities in four clusters of nodes for one of the cells: high concentration with positive rate of change (■) and negative rate of change (■), and low concentration with positive rate of change (■) and negative rate of change (■). Observe that the areas with negative derivative have normal velocities predominantly pointing inwards, which is characteristic for the membrane retraction. The negative derivative of the intensity at a particular node is interpreted here as the effect of receding PIP₃ concentration which could coincide with either membrane retraction or the "stalling" area between protrusions. Here, "stalling" refers to the parts of the cell membrane located between conflicting pseudopods and characterised by negative local curvature. An example of the stalling area can be observed near the node 125 in Figure 6.2. This study is concerned only with the areas of PIP₃ activation, *i.e.* with positive time derivative.

Normal velocities in the areas with high (■) and low (■) intensity do not fit a specific distribution type. That is why the shift in velocity magnitude is assessed using the Mann-Whitney test [201] with null hypothesis that it is equally likely that normal velocities in the areas with low fluorescence will be greater or less than the average v_ν in the area of high intensity. The decision about the hypothesis is made based on the test value p that denotes the approximate probability of observing such result. In this study, the alternative hypothesis states that the median of velocity distribution in areas with high PIP₃ concentration is greater than that in the nodes with low a. The results of the hypothesis testing with significance level $\alpha = 0.05$ for all tracked cells are summarised in Table 6.1. In the table, logical variable H denoting whether the null-hypothesis has been rejected ($H = 1$), or accepted ($H = 0$) is shown alongside the probabilities that informed the decision.

It can be seen from the table that for the majority of the cells the null-hypothesis is rejected. Based on the Mann-Whitney test results and computed medians of velocity distributions, we can conclude that in the nodes with higher concentration of PIP₃, on average, velocity of the advancing membrane is higher. This means that, despite weak direct correlation of membrane deformation velocity and the local concentration of PIP₃, there is a median shift in normal velocity distribution in the areas of the cell where higher concentrations of this mediator are observed. From

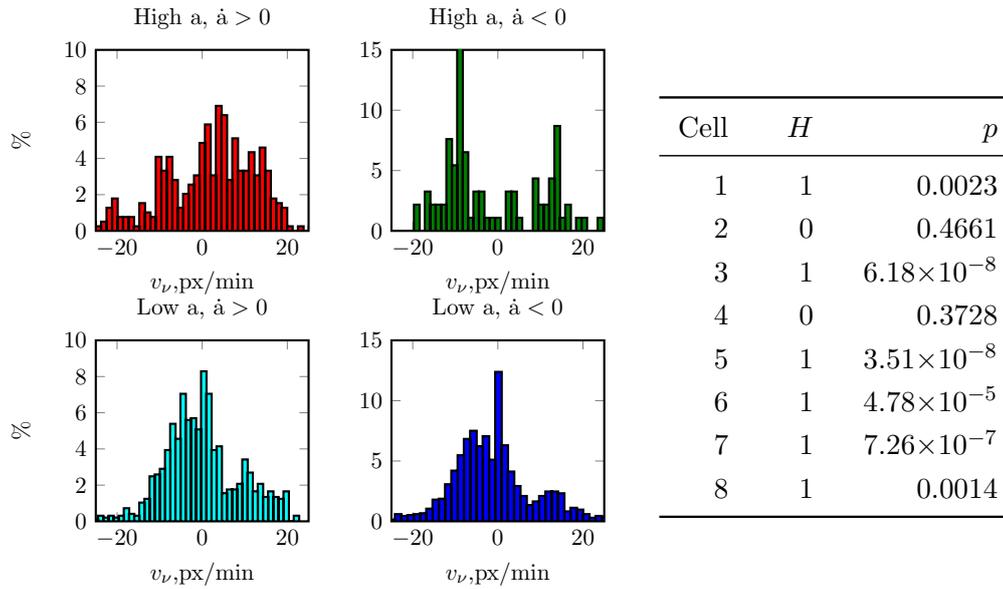


Figure 6.5: Normal velocity histograms for Cell 6.

Nodes are clustered according to concentrations and growth rate of PIP₃ concentration.

Table 6.1: Mann-Whitney U-test results.

the biological point of view, this result can be interpreted as follows: PIP₃ is not essential for pseudopod formation, however its high concentration is associated with accelerated membrane protrusion. These results allow us to make a judgement about the population of nodes as a whole, however they do not provide information about the dynamics of the cell membrane. Given that an association has been detected, it is interesting to examine the dynamical relationship between cell membrane velocity and the local PIP₃ concentration.

6.3.3 Example of model identification from *in vivo* data

In order to evaluate the influence of PIP₃ concentration on local membrane evolution, the dynamical model introduced in §6.2.5 is calibrated against the extracted quantitative data via the EM algorithm. The algorithm is initialised by the LS-type step that computes the initial estimate of the parameter vector, $\hat{\Theta}^0$, using the results of cell boundary tracking and the integrated intensity smoothed over the cell boundary.

The non-normalised integrated fluorescence intensity is smoothed over the cell boundary domain prior to the estimation. Two smoothing algorithms have been tested: the moving average algorithm and the B-spline smoothing. The smoothing results around the cell boundary for one time instance are compared in Figure 6.6. The MA method is selected for this analysis because it is more sensitive to sharp changes in the magnitude and its output remains positive, unlike that of the B-spline

smoother. The smoothed intensity and local curvature are shown in Figure 6.7 as

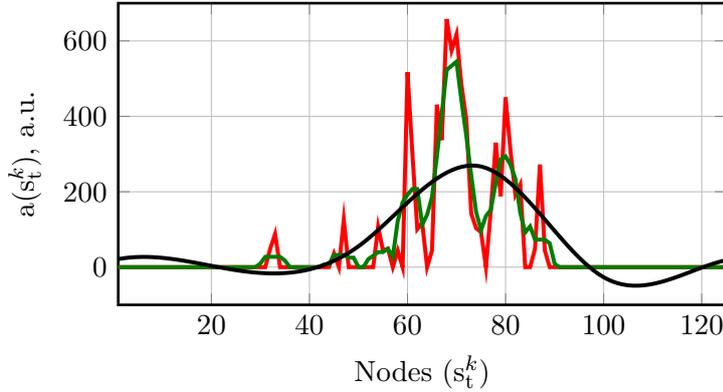


Figure 6.6: Smoothing of the extracted fluorescence intensity over the cell boundary at time t . Measured fluorescence intensity (—) smoothed with MA (—) and B-spline (—) methods.

colour maps where the Y-axis corresponds to the node index on the unwrapped cell boundary, and the X-axis to time instance. The smoothed intensity captures the diffusive nature of the local activator, with highest concentration between nodes 100 and 20. Local curvature is the highest and the cell area computed for each frame are used to estimate the model parameters via the EM algorithm.

The estimation results for all cells are summarised in Table 6.2. It can be seen that most of the parameter MLEs are inconsistent for different cells, and the estimated noise covariances suggest that the estimation algorithm attributes the dominant effect on the membrane dynamics to the random disturbance. An example

Cell	\hat{r}_v	\hat{r}_a	\hat{r}_γ	\hat{r}_{vol}	\hat{Q}_w	\hat{P}_0^k
1	0.3080	0.2289	4.4175	-0.0013	69	38.07
2	0.0672	-0.0896	-41.1960	-0.0003	51	2.28
3	0.5059	0.1419	14.1492	-0.0003	124	106.03
4	0.2203	-0.0826	29.2137	-0.0026	60	10.22
5	0.2532	0.1243	2.1701	-0.0001	198	6.67
6	0.4714	0.0247	-16.0153	-0.0001	47	48.13
7	0.5875	0.0096	12.6756	-0.0000	44	58.14
8	0.1076	0.2217	-0.3254	-0.0004	58	15.09

Table 6.2: Estimation results for PH-Akt-GFP neutrophils.

of the local velocity predicted by the model with estimated parameters from the data shown above is compared to the EM output in Figure 6.8. The model is simulated under the assumption of noiseless process to evaluate the mean value behaviour of the evolving membrane. It can be seen that magnitudes of the predicted velocity

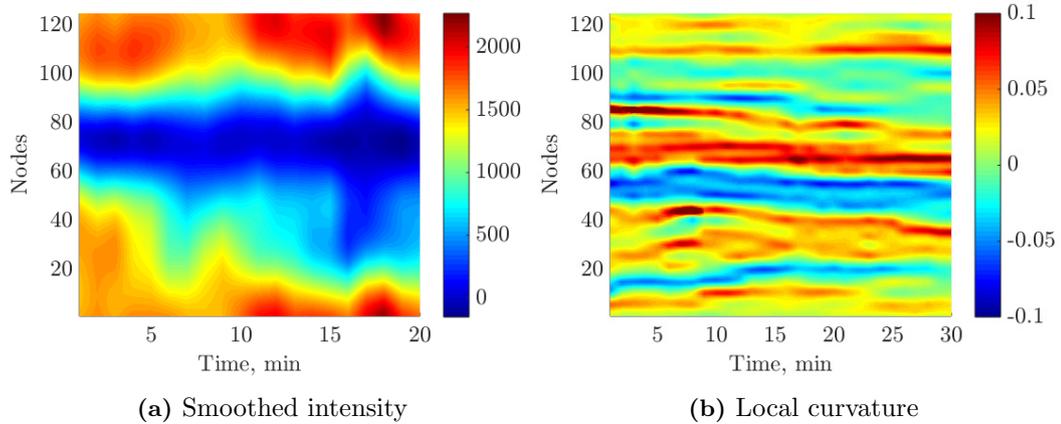


Figure 6.7: Cell properties extracted from the imaging data. **(a)** Smoothed local fluorescent intensity of PH-Akt-GFP is assumed to be proportional to the PIP_3 concentration. **(b)** Local curvature of the cell membrane.

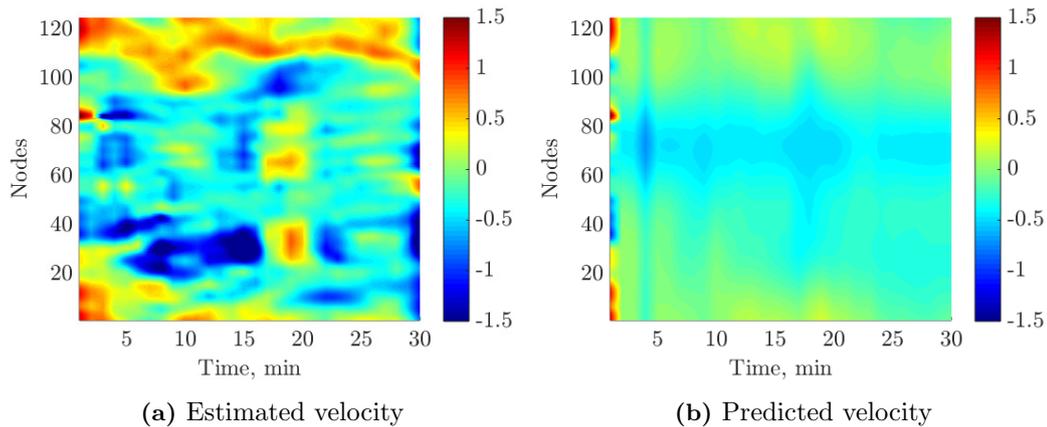


Figure 6.8: Estimation results for a motile neutrophil (Cell 6). **(a)** The local normal velocity (in px/min) filtered during the expectation step of the EM algorithm. **(b)** The local normal velocity (in px/min) is predicted using the dynamical model with the estimated coefficients.

are smaller than the estimated ones. Besides, the model produces values that are directly proportional to the intensity magnitudes shown in Figure 6.7. This shows a significant mismatch between the extracted cell velocities and their predictions. An additional estimation procedure has been conducted under the assumption that the noise covariance is known ($\hat{Q}_w = 1$), however the predicting capabilities of the estimated dynamical model saw no improvement.

Poor estimation results can be interpreted as the result of describing non-persistent polarity with a simple linear model. No meaningful predictions can be made about the evolution of multiple protrusions from the available PIP_3 distribu-

tion along the cell membrane. This leads to the conclusion that the hypothesis of direct influence of PIP_3 on pseudopod protrusion as specified by the linear SMM used here is falsified for neutrophils performing chemotaxis. Alternative mathematical representations of cell morphodynamics should be sought in future analyses.

6.3.4 Example of model identification from *in vitro* data

While the examined model poorly describes the morphological changes of highly motile cells such as neutrophils, it is interesting to consider whether it can fit the data from polarised cells. This analysis is performed with a different cell line which is often used to model neutrophil behaviour – immortalised mammalian leukocytes, HL-60. A set HL-60 cell with the the same GFP labelling as the neutrophils investigated in this chapter is taken from [27]. The morphodynamics of one cell in the set that maintains polarity for 14 frames is analysed here. The image processing results presented in Figures 6.9b- 6.9c show that the highest fluorescence intensity corresponds to the leading edge of the cell.

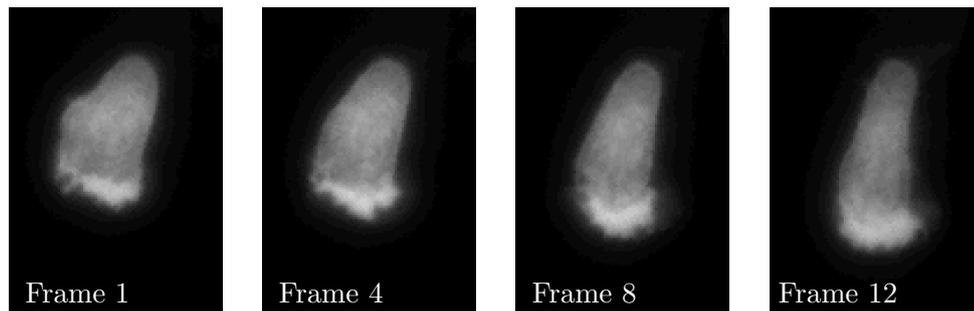
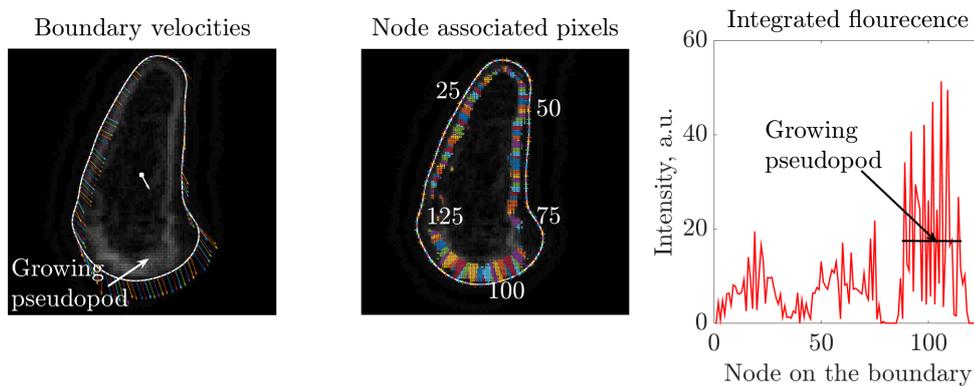
The node-vs-time colour maps presented in Figure 6.10 show that the highest intensity is maintained between nodes 75 and 120, which correspond to the leading edge. The second, although much smaller, increase in value corresponds to the trailing edge of the cell. The highest curvature is observed at the trailing edge and bordering nodes of the advancing pseudopod.

The results of applying the EM algorithm to the polarised cell data are shown in Table 6.3. Because only one cell is analysed, it is not clear whether the MLEs of \hat{r}_v and \hat{r}_a will be more consistent for polarised cells. Similarly with the results obtained from *in vivo* data, estimated parameter values suggest that the cell area change has little effect on the boundary evolution. It could be argued that the effect of volume conservation can only be observed in 3-D, so considering it as an aspect in analysing 2-D data can only be approximate. On the other hand, the estimated noise covariance is much smaller compared to that for motile neutrophils.

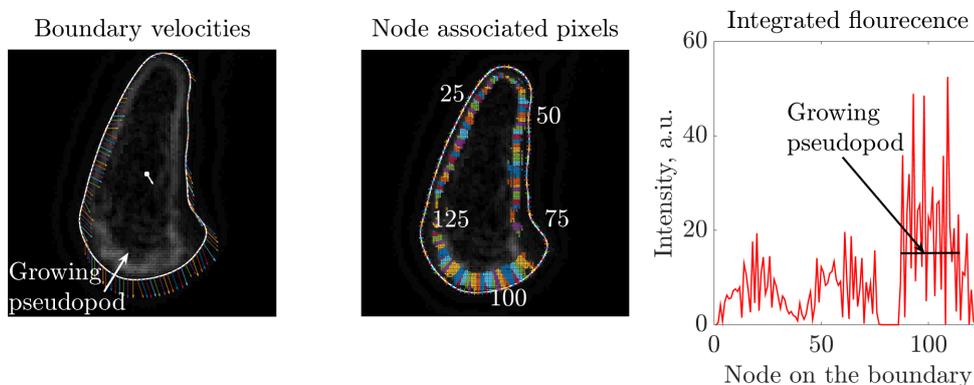
Parameter	\hat{r}_v	\hat{r}_a	\hat{r}_γ	\hat{r}_{vol}	\hat{Q}_w	\hat{P}_0^k
Value	0.4872	0.1503	-102.03	-0.0004	1.3637	1.0007

Table 6.3: Estimation results for a polarised cell.

The estimates from Table 6.3 have been substituted into the dynamical model to predict the local normal velocity given the concentration of PIP_3 , local curvature, and the cell area. Because of the linear relationship between the components presumed by the model, higher concentration at the trailing edge contributes to the increase in normal velocity, however this effect is compensated by large curvature

(a) Examples of time-lapse *in vitro* images

(b) Image processing results: frame 9



(c) Image processing results: frame 10

Figure 6.9: Time-lapse images of a polarised HL-60 and selected processing results. (a) Selected *in vitro* frames of the immortalised mammalian neutrophil with GFP-tagged PH domain of Akt performing the chemotactic run. (b)-(c) Processing results of two subsequent frames of the cell. The left image shows node velocities produced by the boundary tracking module. The middle image shows the pixels along the cell boundary associated with each node on the cell boundary. The right image shows the local integrated intensity plotted against the node index.

values. This signifies that the model captures the regulatory role of the local membrane tension in the polarised cell. The noiseless prediction results are compared

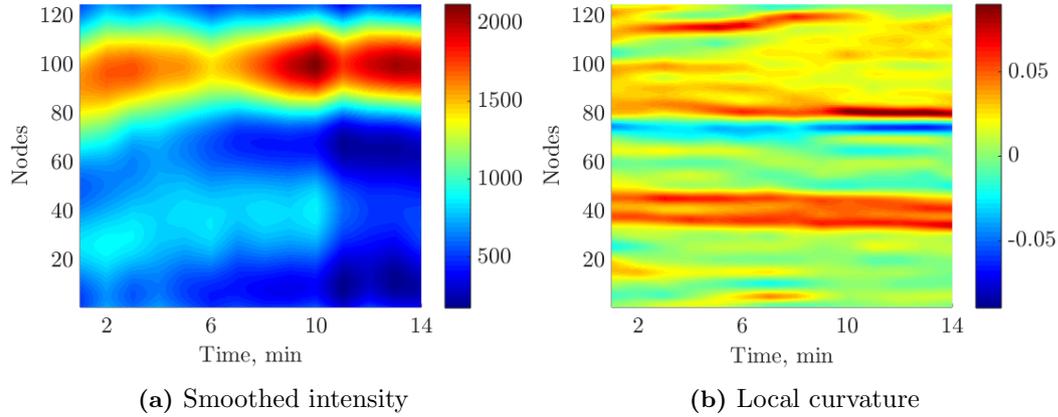


Figure 6.10: Quantitative data extracted from the images of a polarised HL-60 cell. (a) Smoothed local fluorescent intensity of PH-Akt-GFP is assumed to be proportional to the PIP_3 concentration (b) Local curvature of the polarised cell. The highest curvature correspond to the cell rear.

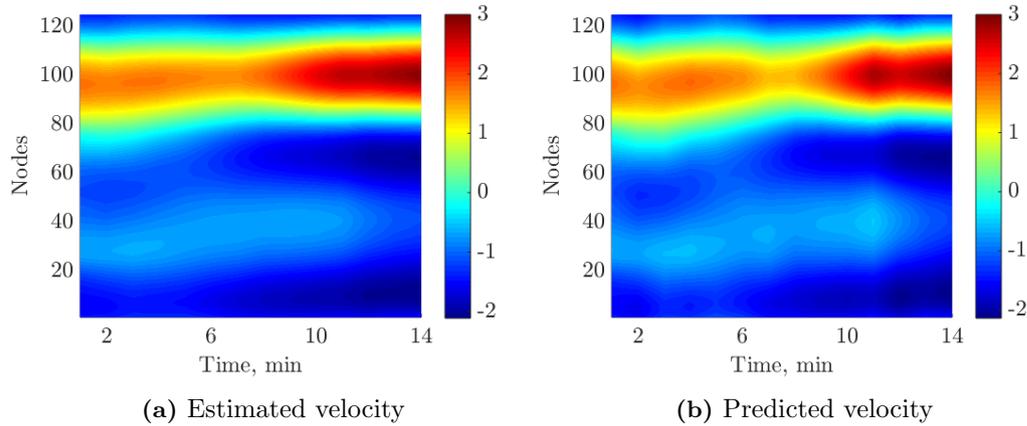


Figure 6.11: Estimation results for a polarised HL-60 cell. (a) Local normal velocity (in px/min) filtered during the expectation step of the EM algorithm. (b) Normal velocity (in px/min) is predicted using the dynamical model with the estimated coefficients.

to the values estimated by the EM algorithm in Figure 6.11. It can be seen that the predicted velocity map aligns with the estimated values. However, the estimated process noise covariance still indicates rather large noise-to-signal ratio in the resultant model.

This is the only cell for which the parameter estimation resulted into a good prediction, suggesting that PIP_3 is a participant of the signalling pathways in an established leading edge, rather than the activator of new protrusions. Combined with the results from the *in vivo* data, this indicates that the linear dynamical model

that allocates the role of the “internal compass” to the internal gradient of PIP_3 cannot be considered as a candidate for the role of the local activator that induces membrane protrusion in the polarisation process.

6.4 Discussion

The results demonstrate that, while on average higher PIP_3 concentration corresponds to the faster protrusion of the cell membrane, the examined model does not support the hypothesis that PIP_3 is the direct contributor to the pseudopod growth. These results can be interpreted from two perspectives: characterisation of the biological process and adequacy of the mathematical description. On one hand, the estimation outcomes agree with several experimental studies which demonstrate on various knock-out cell mutants that the PI3k- PIP_3 signalling pathway is not essential for accurate chemotaxis, although it accelerates protrusion growth [202], [203]. The analysis conducted in this chapter provides quantitative evidence obtained from cells with normal PIP_3 generation rate that complements the experimental findings from knock-out mutant cells. On the other hand, the results indicate that the commonly used linear dynamical model with a single activator cannot sufficiently describe the severe morphological changes of a migrating neutrophil.

The fact that the tangent movement of the membrane segments is not considered in the model is also rather limiting because in some protrusions only the front part moves forward, while the rest of the pseudopod has a larger tangent component of the velocity (see, for example, Figure 6.12). It can be seen in the figure that the upper pseudopod, where the high concentration of PIP_3 is observed, is characterised by a “slide” in the membrane. This leads to the mismatch between the selected model of boundary evolution and the observed protrusion, where, despite the high concentration of PIP_3 , the dominant component of local velocity is tangent to the membrane. This means that the dynamical model needs to be extended to the local tangent velocity evolution, however such representation may imply spatial dependence between the nodes. The state space model must then consider the state of the entire curve rather than an individual node, which may complicate the model parameter estimation.

Another possible extension of the SSM is to consider the spatial aspect of the local activator dynamics using one of the RD models discussed in §2.2.3. Because the nodes are equally spaced on the discretised curve representing the cell boundary, the PDE describing the spatio-temporal dynamics of the activator can be straightforwardly discretised using finite differences method. The difficulties of adapting such models arise from the fact that the evolution of the activator is coupled with extracellular environment and often a global inhibitor. The former cannot be observed

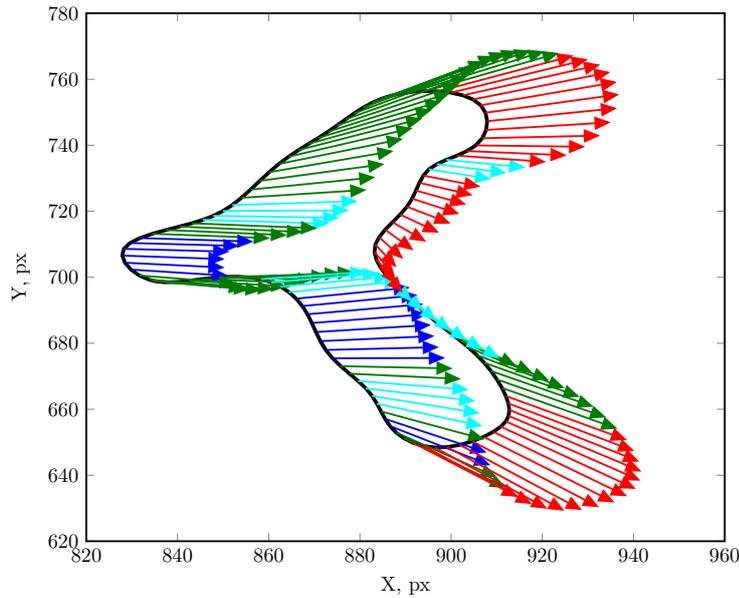


Figure 6.12: Example of the cell membrane classification based on the local curvature and local velocity components. Node velocities are colour coded as follows. Advancing membrane: $v_\nu \gg \epsilon$, $|v_\nu| \gg |v_\tau|$ (—); retracting membrane: $\kappa > 0$, $v_\nu \ll -\epsilon$, $|v_\nu| \gg |v_\tau|$ (—); sliding membrane: $v_\tau \gg \epsilon$, $|v_\tau| \gg |v_\nu|$ (—); stalling membrane: $\kappa < 0$, $|v_\nu| \ll \epsilon$, $|v_\tau| \ll \epsilon$ (—). The threshold ϵ is arbitrarily set to 10 px/min for the given cell. The threshold depends on the time increment between frames and the resolution of time-lapse images.

experimentally which leads to the problem of simultaneous estimation of multiple aspects of the highly non-linear process, while the latter is an artificial parameter that is not linked to any physical aspects of the subcellular signalling. Developing a simplified model of the local activator redistribution is therefore deemed more feasible.

In polarised cells, the relationship between the PIP_3 redistribution and the direction of the membrane protrusion is more evident. This is demonstrated in the HL-60 cell observed *in vitro*. The model with estimated parameters can predict the magnitude of normal velocity at the leading edge based on the local concentration of PIP_3 and the curvature of the membrane. Indeed, the majority of modelling works focus exclusively on persistently polarised shape [41], [60], [66] which is only suitable for cells with stable leading edges. It can be argued that the single leading edge is not maintained throughout the process of chemotaxis, especially in *in vivo* neutrophils whose motion is characterised by multiple protrusions (recall the neutrophil shown in Figure 6.1). At the same time, both neutrophils and HL-60 cells demonstrate an inherent propensity to polarise in absence external gradient [204]. Thus, a possible extension to this analysis would be to introduce a switching model that reflects different modes of neutrophil membrane evolution: periods of clear polarity can be

described by the model considered in this chapter, while the periods of multiple protrusions require a different description with a higher degree of randomness. The switching process could be conditioned on the number of protrusions of the cell. Development of such a model is beyond the scope of this thesis.

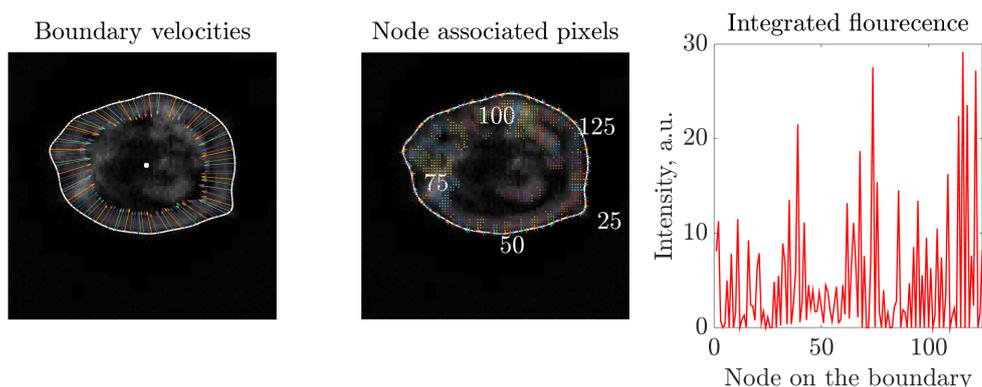
The lack of correlation between the messenger concentration and the local velocity may be the result of various imaging artefacts that are caused by observing a 3-D object in two dimensions. In case of neutrophil population considered in previous chapters, the cell movement along the Z-axis (away or towards the camera) can be neglected because the "depth" of the zebrafish tail fin is rather small compared to the distances the neutrophils must cover over the course of the inflammatory response. After zooming in to the level of an individual cell the Z-component of the local velocity becomes comparable in size with the other two components. Thus, the high concentration of PIP₃ may correspond to the membrane advancing towards the camera that is obscured by compressing the Z-stack images into one layer (see Figure 6.13). For more conclusive results, individual Z-layers can be processed using the tools developed in this chapter.

Because of the limited size of the experimental dataset, the biological significance of the analysis results is disputable. It is also not clear whether there is a stronger link between PIP₃ and the protrusion formation in other types of eukaryotic cells as only neutrophils and neutrophil-like cells are analysed. The estimation procedure developed in this chapter can be applied to the different cell types or intracellular messengers such as Rac that has been successfully imaged in [205], [206], PTEN, and other species that may regulate local shape changes in a migrating cell.

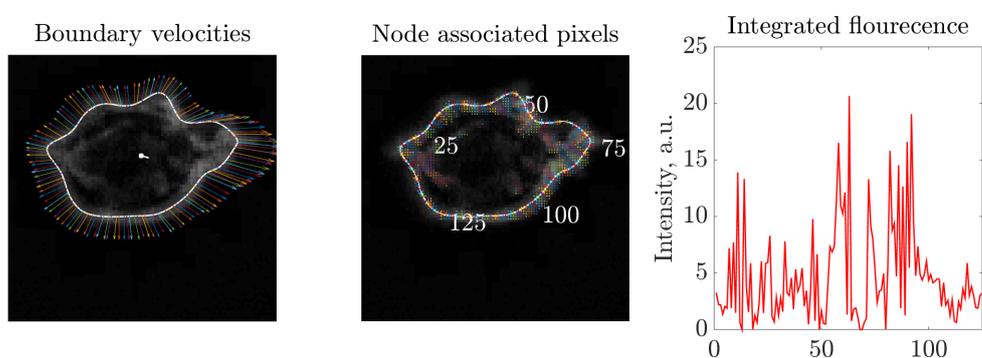
6.5 Summary

This chapter introduces a computational module for processing the time-lapse images of neutrophil subcellular activity and utilises the processing results to assess the assumption commonly made in the modelling literature about the relationship between the environment sensing mechanism and shape changes in migrating eukaryotic cells. Processed are the images of migrating neutrophils tracked *in vivo* with GFP-tagged PH-Akt domain whose fluorescence corresponds to the local concentration of PIP₃ – a crucial second messenger that is considered as a candidate for the role of local activator in several models of the subcellular signalling mechanism driving chemotaxis.

The mathematical model of cell shape evolution that is examined in this chapter is tested against the *in vivo* images of zebrafish neutrophils obtained by the collaborators on the project, as well as *in vitro* images of immortalised mammalian neutrophil (HL-60 line) with persistent polarisation from [27]. The correlation analysis



(a) Cell moving away from the camera



(b) Cell moving towards the camera

Figure 6.13: The effect of 2-D microscopy on the single-cell scale. Since the Z-stack time-lapse images are compressed into one layer, the motion along the Z-axis is not accounted for in the cell boundary tracking. (a) The neutrophil migrating away from the camera. (b) The neutrophil migrating towards the camera. Both cells are characterised by similar distribution of Akt intensity.

of the data extracted from *in vivo* imaging results reveals that there is no strong connection between the distribution of PIP_3 and the speed of advancing or retracting membrane. This could be attributed to imaging artefacts or tracking errors. However, on average, areas with the highest local concentration and the positive time derivative of the second messenger correspond to the accelerated membrane protrusion which conforms to the hypothesis that PIP_3 accelerates the pseudopod extension. The processing results of the polarised HL-60 cell are used to estimate the parameters of the evolution model and predict local velocities in the leading edge based on the PIP_3 concentration and the local curvature of the cell membrane. This indicates that the simplified model of the membrane evolution heavily relies on the assumption of persistent polarisation which is rarely the case for chemotaxing neutrophils. The main focus of the future work should therefore be on the development of more sophisticated data-driven models of pseudopod formation and retraction.

The progress in this direction may be impeded by slow advances in the imaging of subcellular activity *in vivo*.

In summary, this chapter demonstrates a set of methods that can be used to evaluate the relationship between various subcellular species and deformations of the cell membrane. This analytical framework can be instrumental in future studies of cell morphodynamics for processing of time-lapse imaging data that focuses on internal concentrations of signalling proteins.

7 | Conclusion

7.1 Summary of contributions

The ultimate goal of researching the neutrophil migration is to increase the precision of the therapeutic treatment for various inflammation-mediated diseases while minimising potentially harmful side effects. This is impossible to achieve without thorough knowledge of the key processes regulating the responsiveness of neutrophils to the external environment which drives their migration towards and away from the area of the infectious threat. A colossal leap towards clearer understanding of neutrophil function during the inflammatory response has been made with the introduction of the *in vivo* models of inflammation which allow researchers to observe cells in naturally occurring inflammatory event. However, *in vivo* microscopy often obscures the important aspects of the process, such as the underlying environments that induce neutrophil chemotaxis. Along with lengthy and expensive experiment design, this shortcoming slows down the process of treatment development and increases its costs. Data-driven computational models of neutrophil migration have the potential to simplify the process of testing various therapeutic strategies which can reduce the costs of medical trials. In order to make meaningful predictions of cell behaviour, these models must first be calibrated against the experimental data. This thesis provides an account of developing statistical inference tools that can be used at various stages of the model calibration process.

The background chapter provides an overview of the inflammatory response from the neutrophil-centred perspective, discusses mathematical models of cell chemotaxis on varying scale, and concludes with a brief but comprehensive review of the statistical inference methods that are utilised throughout the thesis. The state space approach to modelling the dynamical processes is advocated as a suitable way to interpret the problem posed in the thesis.

Chapter 3 introduces a novel problem of simultaneous estimation of neutrophil velocities and the hidden chemoattractant environment that drives their migration during the recruitment stage of the inflammation. A simple solution is sought based on the assumption that the hidden environment acts on the moving cell as a potential

field. A state space model of individual neutrophil dynamics is formulated to explicitly include the parametrised model of the global attractant concentration. The hidden environment is parametrised via the basis function decomposition where the scaling coefficient corresponding to individual basis functions are the model parameters. The resultant state space model is linear with respect to the unknown parameters and non-linear with respect to the cell state. Cell states and field model parameters are estimated via the Expectation Maximisation (EM) algorithm that employs an unscented smoother during the expectation step. An additional step approximates cell positions in order to ensure the analytical solution of the maximisation step. The inference results are in agreement with the widely accepted hypothesis that the neutrophils are actively driven to the wound area by the external chemoattractant concentration, however the estimation is constricted by the assumption that all neutrophils uniformly respond to the chemoattractant.

Chapter 4 considers a generic problem of the environment inference based on the observed migration of multiple models with hybrid dynamics. A jump Markov model of individual object dynamics is combined with the environment model proposed in Chapter 3. The potential field method of modelling the object-environment interaction leads to a hybrid model that is linear in terms of the unknown parameters. An approximate EM solution is derived for the simultaneous estimation of object full states, mode probabilities, and the environment model parameters. Because one of the state space models within the hybrid system is non-linear with respect to the state, an additional approximation step is introduced that computes the forecasts of object positions that are further utilised to solve the maximisation problem in closed form. The algorithm also employs the IMM-URTS smoother on the expectation step that provides the mode probabilities and mode-conditioned states of each object at reduced computational cost. The simulation examples reveal three crucial factors that affect the accuracy of the environment estimation: the tracking data localisation, the choice of initial conditions in the inference framework, and the closeness of the utilised hybrid model to the true object dynamics. One of the simulation examples demonstrates that a good estimate of the underlying environment can be achieved irrespective of the knowledge of the true environment model. Although the estimation framework is developed in the interest of processing neutrophil tracking data, some of the presented simulation examples speculate on other potential applications where the inference of the global environment may be required.

The hybrid model of the individual neutrophil dynamics considered in Chapter 5 relaxes the assumption of uniform behaviour in the population of migrating cells. The proposed model accounts for three types of behaviour observed experimentally:

- responsive mode, in which the neutrophil moves in response to the spatial changes in external chemoattractant (ChA) concentration;

- diffusing mode, in which the neutrophil is performing a mean-reversing random walk;
- stationary mode, which describes apoptotic or non-activated cells.

Flexibility of the exploited JMS makes it possible to utilise the model in the novel analysis of both recruitment and resolution stages of inflammation. In the analysis, the hybrid system is embedded into the estimation framework proposed in Chapter 4 to infer chemoattractant and chemorepellent concentrations from the zebrafish tracking data. Two alternative methods of the approximate state estimation of a hybrid model are considered: IMM algorithm that merges the mode histories, and the pruning algorithm that favours the most probable migratory mode of a neutrophil at each time. The environment inference for all recruitment datasets recovers the smooth global environment with higher concentration at the wound site, however in two instances for the tail fin nick injury the outlier peaks are identified from the persistent migration of several neutrophils in the direction opposite from the wound. Mode identification results show that majority of the neutrophils remain responsive on their route to the injury site, yet upon arrival to the site they switch to random movement. This change of behaviour could be caused by either loss of sensitivity to the gradient or by chemokinesis in uniformly high ChA concentration.

To author's best knowledge, Chapter 5 describes the first attempt to analyse the reverse migration process using the modelling framework that incorporates the influence of the environment in the dynamical model of the cell. Moreover, the hybrid model utilised for this analysis takes into account the fact that a neutrophil may switch between various migratory modes throughout the process of reverse migration. The estimation results conform with previously proposed hypothesis that neutrophils predominantly diffuse away from the wound area [17], [179].

The level of detail in the data-driven model is usually defined by the experimental data in hand. While a large part of this thesis examines the process of directed cell migration on the level of population and the individual whole cell, Chapter 6 moves on to a finer scale and analyses the subcellular mechanisms of the migrating neutrophil. The hypothesis that the distribution of PIP_3 along the cell boundary acts as a local mediator of protrusion growth is tested against the imaging data of individual neutrophils. A simple model of local cell membrane evolution is formulated in the state space domain to linearly link local PIP_3 concentration, local curvature, and the cell area to the local normal velocity of the boundary. The state-parameter estimation of the developed model is performed via the ordinary Least Squares algorithm to recover the proportionality coefficients and the process noise variance. The preliminary analysis of raw quantitative data obtained from the *in vivo* time-lapse images of migrating neutrophils reveals that there is no strong correlation between the candidate mediator and local advances of the cell membrane, although

the areas with higher concentration of PIP_3 are characterised by the median shift of local velocity distributions. The dynamical model with estimated parameter values could not reproduce the observed membrane dynamic of motile cells, leading to the conclusion that the redistribution of PIP_3 alone may not be sufficient to mediate multiple protrusions. However, the estimation results of a polarised cell observed *in vitro* suggest that the examined model may sufficiently describe the dynamics of a cell with a well defined leading edge.

In summary, this thesis makes an important contribution to the effort of closing the gap between the experimental data and the mathematical modelling of directed cell migration during the inflammatory response. The environment inference framework developed in Chapters 3-5 is flexible, amendable to various types of wounds and migration patterns, and can be utilised as a decision-support tool in the design process of various inflammation treatments. On the other hand, the image processing tools and the estimation algorithm developed in Chapter 6 can be of assistance in identifying the subcellular messengers that regulate the deformation of the cell shape.

7.2 Future Work

While the work conducted in this thesis significantly extends the scope of the investigation of neutrophil migration during the inflammatory response, several limitations of the proposed models and estimation methods have been identified throughout the thesis. Overcoming these limitations is the basis of future work outlined below.

Implementation of the estimation framework

Several aspects of the estimation framework proposed in Chapter 4 can be improved. First, it is necessary to assess how various approximations of the non-linear dynamical model may improve the estimation performance compared to the forecasting object positions that has been employed in this work. A numerical optimisation of the expected log-likelihood can also be considered in case of a sparse basis function grid. For example, a multi-start search will be able to identify multiple stationary points of the log-likelihood function, though at the cost of increased computational complexity. The proposed algorithm is developed in the interest of application where no prior information about the field is available. As a Maximum Likelihood solution, it provides only a point estimate of the model parameters. Extending the EM algorithm to produce the MAP estimate of the underlying environment will improve its applicability to situations where the prior knowledge about the environment parametrised model can be utilised. Furthermore, employing variational Bayes methods

to approximate the posterior *pdf* of model parameters will extend the framework to a fully Bayesian estimation [207].

The mode transition probability matrix is assumed to be known to the estimation framework. Its estimation can be included by extending the developed algorithm to the expectation conditional maximisation scheme, in which one of the conditional M-steps will estimate the underlying environment, and the other will estimate the transition frequency matrix as described in [155].

Several simulation examples in Chapter 4 demonstrate the disadvantage of using a simple tensor B-spline grid with uniform knot sequence. Local support of the individual basis functions combined with the tracking data confined to a particular area of the field lead to poor accuracy of individual parameter estimates. This effect can be observed in Chapter 5 in the analysis of reverse migration data. Several alternatives to the simple B-spline decomposition can be considered to account for the spatial changes of the environment on varying scale: the hierarchical knot placement, multi-resolution approximation with B-spline wavelets, or using a combination of fine and coarse grids. The multi-resolution model of the environment will be able to capture small details in the areas with high density of object tracks and "smooth" the areas where limited data is available. This extension can serve as foundation for further improvements in the application to neutrophil migration that are discussed below.

Modelling and estimation of cell-environment interaction

The analysis of neutrophil chemotaxis in Chapter 5 is performed under the assumption that the underlying chemoattractant concentration field is time-invariant. This implies a single slowly-diffusing agent acting on the recruited neutrophils, which contradicts recent experimental studies that observe multiple chemoattractants throughout the course of inflammation recruitment. For example, the rapidly diffusing hydrogen peroxide (H_2O_2) has been identified as a candidate for the initial wound-to-cell signalling [208]. The experiments reveal that the peroxide concentration peaks near the wound at 20 min post injury and diffuses until 120 min to be replaced with intermediate attractants. The future analysis should account for the spatio-temporal dynamics of the ChA concentration. Because the inherent batch-based structure of the developed framework requires collecting the tracking data prior to the field inference, the simplest solution would be to divide the neutrophil tracks into hourly sets and perform the estimation for each hour separately using the algorithm proposed in Chapter 4. However, the sparsity of the hourly tracking data may lead to poor accuracy of field estimation, especially if the field model is parametrised using the basis functions with local support. A more sophisticated approach would involve recursive environment estimation: at each time interval,

the concentration field will be inferred from the neutrophil tracks that have been accumulated by the given time.

Another restricting assumption made in this analysis is that there is only one chemoattractant source that promotes neutrophil recruitment. Multiple experimental studies have demonstrated that neutrophils navigating in a complex ChA field have to select between various local gradients [209]. The estimation results for the tail nick injury in Chapter 5 support this statement: the shape of the inferred global environment indicates that there may be two or more threat areas located next to each other. Neutrophil chemotactic behaviour in case of competing ChA sources has been investigated in [177], where the gradient model is assumed to be known. Adding distinct models of neutrophil interaction with each of the competing gradients to the framework proposed in this thesis will enable the inference of these non-trivial environments based only on the tracking data.

The natural direction for future analysis of reverse neutrophil migration is to consider the conflicting chemoattractant and chemorepellent environments. This can be implemented by introducing an additional migratory mode to the hybrid model of neutrophil dynamics that will describe the recruitment process. Two unique basis function decompositions of the competing fields are required for this type of analysis.

The estimates of chemoattractant concentrations are presented in arbitrary units that only describe the relative spatial change of the environment. In order to relate the inference results to the physical chemoattractant concentration, the estimation framework should be calibrated on *in vitro* experiments where the underlying ChA field can be manipulated and measured.

Modelling and estimation of cell shape evolution

The main direction of the future research, however, should focus on the study of subcellular processes driving neutrophil migration. Several outstanding problems that arise from the analysis in Chapter 6 need to be addressed. Firstly, a more complex dynamical model that accounts for the spatio-temporal dynamics of the activator can be used to describe the evolution of the cell membrane. Multiple models for that account for diffusive nature of the activator have been reviewed in §2.2.3, however most of them couple local mediators with a global inhibitor. Approximating the partial differential equation that describes spatio-temporal dynamics of the local activator with the finite differences method will lead to coupling between neighbouring nodes. The state of the system will then have to include the velocities and concentrations along the entire cell boundary:

$$\mathbf{x}_t = \left[\{v_\nu(s_t^k)\}_1^K, \{a(s_t^k)\}_1^K \right]^\top,$$

where K denotes the number of equidistant nodes used to approximate the boundary. Because the distance between nodes may be changing between the time-lapse frames, the state transition dynamics obtained from the finite-differences algorithms will become time-varying. The expectation step of the algorithm will have to include a time-varying smoother for the state estimation.

Secondly, the model used in Chapter 6 also presumes that the local activator concentration is directly proportional to the local external gradient sensed by the cell. A further extension would be to include the external field along the cell membrane as a hidden state that can be estimated simultaneously with the model parameters. This extension would provide a link between the subcellular processes and the external environment. The anticipated challenge of this extension is spatio-temporal modelling of the extracellular domain that would ensure identifiability of the system.

Finally, the most feasible extension of this dynamical modelling framework is to account for the switching behaviour of migrating neutrophils on the subcellular level. Since the estimated model parameters obtained from the polarised cell allow us to predict the local velocity from the observed concentration of PIP_3 , it may be sufficient to describe morphodynamics of polarised cells. An alternative model can be introduced to describe the motile behaviour of neutrophils, and the switching from polarised to motile state can be defined by the presence of multiple protrusions of the neutrophil membrane. An example quantitative method that combines boundary tracking and detecting protrusions in a migrating cell is presented in [210]. The estimation results suggest that the PIP_3 cannot be assumed as the only significant contributor to multiple pseudopod extension, so additional mediators should be considered when modelling membrane evolution of motile cells. Data-driven development of such models is still in the early stages and relies on improvements in the *in vivo* microscopy of subcellular concentrations. Future work in this direction should seek ways to quantify multiple subcellular environments simultaneously.

Appendices

A | Derivation of the maximisation step for the EM algorithm

Chapters 3 and 4 solve the problem of the environment inference with the EM algorithm. Since the hidden environment is presented by the same model, the maximisation steps in both chapters are derived using the same principle. The proof below is presented for Lemma 4.3.2 that considers a more complex hybrid model, but can be easily amended to prove Lemma 3.3.1, if one is to consider that the hybrid model has only one discrete state. The notation in this appendix is consistent with the notation adopted throughout the thesis.

Recall the expression for the \mathcal{Q} -function for the jump Markov system from Lemma 4.3.1

$$\mathcal{Q}(\Theta, \hat{\Theta}^i) = \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \mathbb{E}[\log p(\mathbf{x}_t^k | \mathbf{x}_{t-1}^k, \mathbf{s}_{t-1}^k, \mathbf{m}_t^k, \Theta) | \mathbf{m}_t^k, \mathbf{y}_t^k, \hat{\Theta}^i] + c, \quad (\text{A.1})$$

where the constant c denotes all terms independent of Θ , \mathbf{s}_{t-1}^k is the augmented state and where

$$\mu_t^{k,j} \triangleq \mathbb{P}(\mathbf{m}_t^k = M^j | \mathbf{y}^k, \Theta) \geq 0. \quad (\text{A.2})$$

Then the expression (4.3.19) can be achieved as follows.

Proof. Expand (A.1) using the definition of the state update probability

$$\begin{aligned} \mathcal{Q}(\Theta, \hat{\Theta}^i) = & c_1 - \frac{1}{2} \mathbb{E} \left[\sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \left\{ (Q_\omega^j)^{-1} \times \right. \right. \\ & \left. \left. \times \left((G^j)^\dagger (\mathbf{x}_{t|t}^k - A^j \mathbf{x}_{t-1|t}^k - B^j \phi(\mathbf{s}_{t-1}^k) \Theta) \right)^2 \right\} \right], \end{aligned} \quad (\text{A.3})$$

where the expectation is taken over states $\mathbf{x}_{t-1}^k, \mathbf{x}_t^k$, and where constant c_1 includes the terms that will be discarded during the maximisation procedure. Recall that the term $\phi(\mathbf{s}_{t-1}^k)$ is approximated using the computed \mathbf{s}_{t-1}^k prior to the maximisation step and thereafter is referred to as ϕ_{t-1}^k . Noting (4.3.20), expand constituent terms

of (A.3) further as,

$$\begin{aligned}
\mathcal{Q}(\Theta, \hat{\Theta}^i) = c_1 - \frac{1}{2} \mathbb{E} \left[\sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \left\{ (\mathbf{x}_t^k)^\top \Sigma_\omega^j \mathbf{x}_t^k - \right. \right. \\
- 2(\mathbf{x}_t^k)^\top \Sigma_\omega^j A^j \mathbf{x}_{t-1}^k + \\
+ 2\Theta^\top (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j A^j \mathbf{x}_{t-1}^k - \\
- 2\Theta^\top (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j \mathbf{x}_t^k + \\
+ (\mathbf{x}_{t-1}^k)^\top (A^j)^\top \Sigma_\omega^j A^j \mathbf{x}_{t-1}^k + \\
\left. \left. + \Theta^\top (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j B^j \phi_{t-1}^k \Theta \right\} \right]. \tag{A.4}
\end{aligned}$$

We can obtain parameter estimates corresponding to the maximum of the log-likelihood function by taking the partial derivative of (A.4) with regard to Θ and setting it equal to zero

$$\begin{aligned}
\frac{\partial \mathcal{Q}(\Theta, \hat{\Theta}^i)}{\partial \Theta} = \mathbb{E} \left[\sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \left\{ (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j \mathbf{x}_t^k - \right. \right. \\
\left. \left. - (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j A^j \mathbf{x}_{t-1}^k - (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j B^j \phi_{t-1}^k \Theta \right\} \right]. \tag{A.5}
\end{aligned}$$

Using properties of the expectation listed in the appendix B, we can rewrite the above equation to produce

$$\begin{aligned}
\frac{\partial \mathcal{Q}(\Theta, \hat{\Theta}^i)}{\partial \Theta} = \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \left\{ (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j \mathbb{E}_{\mathbf{m}_t^k = M^j} [\mathbf{x}_t^k] - \right. \\
\left. - (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j A^j \mathbb{E}_{\mathbf{m}_t^k = M^j} [\mathbf{x}_{t-1}^k] - (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j B^j \phi_{t-1}^k \Theta \right\}. \tag{A.6}
\end{aligned}$$

Solving (A.6) for Θ leads directly to (4.3.19). Furthermore, the second partial derivative of the \mathcal{Q} -function defined by

$$\frac{\partial^2 \mathcal{Q}(\Theta, \hat{\Theta}^i)}{\partial \Theta^2} = - \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{j=1}^{N_m} \mu_t^{k,j} \left\{ (\phi_{t-1}^k)^\top (B^j)^\top \Sigma_\omega^j B^j \phi_{t-1}^k \right\} \tag{A.7}$$

is clearly negative definite owing to its auto-product structure, which verifies that the new parameter estimate $\hat{\Theta}^{i+1}$ is located at a maximum, global or local, completing the proof. \blacksquare

Remark A.1: *If one is to assume the single SSM for the object dynamics as has been done in Chapter 3, then the following substitutions must be made in the proof above:*

- *There is only one candidate SSM $N_m = 1$;*

- The probability of the considered SSM is $\mu_t^{k,j} = 1$ at all times for all objects;
- All mode-conditioned expectations are replaced by the expectations conditioned on the parameter vector $\mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_{t-1}^k] \triangleq \mathbb{E}[\mathbf{x}_{t-1}^k | \mathbf{y}^k, \hat{\Theta}^i]$ and $\mathbb{E}_{\hat{\Theta}^i} [\mathbf{x}_t^k] \triangleq \mathbb{E}[\mathbf{x}_t^k | \mathbf{y}^k, \hat{\Theta}^i]$.

Then solving the simplified equation (A.6) for Θ leads directly to (3.3.14), proving Lemma 3.3.1.

B | Properties of the expectation

Definition B.1: Let $g(x)$ be a function of a random variable with pdf $p(x)$. The expected value of $g(x)$ is defined as follows

$$\mathbb{E}[g(x)] = \int_x g(x)p(x)dx. \quad (\text{B.1})$$

Definition B.2 (from [74]): Let $g(x, y)$ be a function of two random variables with joint pdf $p(x, y)$. The expected value of $g(x, y)$ is defined as follows

$$\mathbb{E}[g(x, y)] = \int_x \int_y g(x, y)p(x, y)dxdy. \quad (\text{B.2})$$

Definition B.3 (Conditional expectation): For two random variables with a continuous joint pdf $p(x, y)$, the marginal pdf of y is given by

$$p_y(y) = \int_x p(x, y)dx. \quad (\text{B.3})$$

Then the conditional expected value of x is computed relative to the conditional distribution $p(x | y) = p(x, y)/p_y(y)$

$$\mathbb{E}[x | y = Y] = \int_y xp(x | y)dy, \quad (\text{B.4})$$

with the randomness inherited from randomness in y .

The following properties are used in derivation of the EM algorithms.

Property B.1 (Linearity of expectation): Let $f(x)$ and $g(x)$ be functions of x and a and b be constants, then, if the expectation exists, it satisfies the following

$$\mathbb{E}[af(x) + bg(x)] = a\mathbb{E}[f(x)] + b\mathbb{E}[g(x)]. \quad (\text{B.5})$$

Property B.2 (Jensen's Inequality): Let $f(x)$ be a function of x , then

$$f(\mathbb{E}[x]) \leq \mathbb{E}[f(x)]. \quad (\text{B.6})$$

Proof. see [211]. ■

Property B.3 (Expectation of quadratic form): *Let $x \in \mathbb{R}^{n \times 1}$ be a vector and $A \in \mathbb{R}^{n \times n}$ be a square matrix, then*

$$\mathbb{E} [x^\top A x] = \text{tr} (\mathbb{E} [x^\top A x]). \quad (\text{B.7})$$

Theorem B.1 (Law of iterated expectations): *Suppose that a random variable x has the expectations $\mathbb{E} [x]$, then for any continuous random variable y defined in the same probability space it holds that*

$$\mathbb{E} [x] = \mathbb{E} [\mathbb{E} [x | y]] \quad (\text{B.8})$$

Proof. The conditional expectation $\mathbb{E} [x | y]$ can itself be treated as a random variable, then by definition B.1 its expected value is given by

$$\mathbb{E} [\mathbb{E} [x | y]] = \int_y \mathbb{E} [x | y = Y] p_y(y) dy, \quad (\text{B.9})$$

where $p_y(y)$ is the marginal *pdf* of y defined by (B.3). Noting the definition B.3, the above equation can be expanded as follows

$$\begin{aligned} \mathbb{E} [\mathbb{E} [x | y]] &= \int_y \int_x x p(x | y) dx p_y(y) dy \\ &= \int_y \int_x x p(x | y) p_y(y) dy dx \\ &= \int_x \int_y x p(x, y) dy dx, \end{aligned} \quad (\text{B.10})$$

where x can be moved outside the internal integral and where $p(x, y)$ is the joint *pdf* of two random variables, marginalisation of which renders

$$\mathbb{E} [\mathbb{E} [x | y]] = \int_x x \int_y p(x, y) dy dx = \int_x x p_x(x) dx. \quad (\text{B.11})$$

The integration over the marginal *pdf* $p_x(x)$ by definition B.1 gives the following

$$\mathbb{E} [\mathbb{E} [x | y]] = \int_x x p_x(x) dx = \mathbb{E} [x]. \quad (\text{B.12})$$

■

Corollary B.1: *Suppose that a random variable x has the expectation $\mathbb{E} [x]$, then for a discrete random variable y that can take a finite number of values $\{Y^j\}$ it holds that*

$$\mathbb{E} [x] = \sum_j \mathbb{E} [x | y] P(y = Y^j). \quad (\text{B.13})$$

C | Geometry of the cell boundary

This appendix presents definitions and methods for calculating auxiliary quantities that were necessary for the computation of the cell boundary evolution model. Note that previously adopted notation is abandoned in this appendix.

Definition C.1 (Local Curvature [212]): *Suppose that $\Gamma(t)$ is a twice continuously differentiable curve, then its local curvature at a node $s(t)$ is the rate of change of the tangential angle $\phi(t)$ and is defined as*

$$\kappa(s) = \frac{\phi'(t)}{\|l'(t)\|}, \quad (\text{C.1})$$

where $l(t)$ is the local arc length, the prime denotes the first order derivative with respect to the variable t , $\frac{d}{dt}$, and $\|\cdot\|$ denotes Euclidean norm.

Property C.1 (Local curvature in Cartesian coordinates): *For a curve defined by Cartesian coordinates $\Gamma(t) = [x(t); y(t)]$, the local curvature satisfies*

$$\kappa(s) = \frac{x'(t)y''(t) - y'(t)x''(t)}{[(x'(t))^2 + (y'(t))^2]^{3/2}}. \quad (\text{C.2})$$

Proof. Denote the tangent vector as $s'(t) = [x'(t); y'(t)]$ and tangent angle as $\phi(t)$. By definition of the Euclidean norm:

$$\kappa(s) = \frac{\phi'(t)}{\sqrt{(x'(t))^2 + (y'(t))^2}}. \quad (\text{C.3})$$

The change of tangent angle can be obtained using the definition of tangent

$$\tan\phi = \frac{dx(t)}{dy(t)}, \quad (\text{C.4})$$

which is the equivalent of

$$\tan\phi = \frac{dx(t)/dt}{dy(t)/dt} = \frac{x'(t)}{y'(t)}. \quad (\text{C.5})$$

The derivative of the tangent is defined as follows

$$\frac{d}{dt} \tan\phi = \sec^2\phi \frac{d\phi}{dt}. \quad (\text{C.6})$$

On the other hand,

$$\frac{d}{dt} \tan \phi = \frac{d}{dt} \left(\frac{x'(t)}{y'(t)} \right) = \frac{x'(t)y''(t) - y'(t)x''(t)}{(x'(t))^2}. \quad (\text{C.7})$$

After combining equations (C.6), (C.7) and simple manipulations we get

$$\frac{d\phi}{dt} = \frac{x'(t)y''(t) - y'(t)x''(t)}{(x'(t))^2 + (y'(t))^2}, \quad (\text{C.8})$$

Substituting numerator in (C.3) by (C.8) leads to the final expression for local curvature in the form (C.2), completing the proof. \blacksquare

Local outward-pointing normal. Since the cell boundary curve is approximated with a discrete set of nodes, it is represented by a polygon. Then for each vertex s_k the local normal $\bar{\nu}(s_k)$ can be computed as a polygon vertex normal using algorithm C.1. The result of running one iteration of the algorithm is demonstrated in Figure C.1.

Algorithm C.1 Computation of polygon vertex bases

Input: Sequence of nodes defining the polygon, $s_k = [x_k, y_k]$, $k = 1, \dots, K$, ordered in the counter-clockwise direction with $s_1 = s_K$.

Output: A set of outward-facing normals, $\bar{\nu}(s_k)$, $k = 1, \dots, K-1$; a set of corresponding tangent vectors, $\bar{\tau}(s_k)$, $k = 1, \dots, K-1$.

- 1: **for** $k \leftarrow 1, K-1$ **do**
- 2: Form a segment vector $\bar{AB} = s_k - s_{k-1}$;
- 3: Form a segment vector $\bar{BC} = s_{k+1} - s_k$;
- 4: Compute the outward face normals

$$\begin{aligned} \bar{\nu}_{AB} &= \text{null}(\bar{AB}), \\ \bar{\nu}_{BC} &= \text{null}(\bar{BC}), \end{aligned}$$

such that $\bar{\nu}_{AB} \times \bar{AB} > 0$ and $\bar{\nu}_{BC} \times \bar{BC} > 0$;

- 5: Set the origin of both face normals at the vertex s_k ;
- 6: Obtain the vertex normal as a normalised geometric sum of two pointing normals:

$$\bar{\nu}(s_k) = \frac{\bar{\nu}_{AB} + \bar{\nu}_{BC}}{\|\bar{\nu}_{AB} + \bar{\nu}_{BC}\|};$$

- 7: Obtain the corresponding tangent vector $\bar{\tau}(s_k) = \text{null}(\bar{\nu}(s_k))$ such that $\bar{\nu}(s_k) \times \bar{\tau}(s_k) > 0$;
 - 8: **end for**
-

Area enclosed by the cell boundary. The cell area can be computed as the area of the non-convex polygon via triangulation. One of the most popular techniques is Delaunay triangulation that maximises the smallest angle each triangle (see Figure C.2a). This method requires solving an optimisation problem on for each cell at each frame that proves to be computationally expensive. A simpler method for computing the area with a standard triangulation is summarised in algorithm C.2. The advantage of representing the triangle vertexes as a matrix is that it will automatically detect the triangles that are outside the polygon (coloured grey in

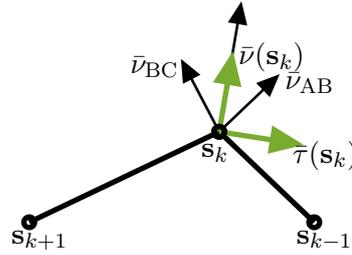


Figure C.1: An example of the basis located at a vertex of the polygon.

Algorithm C.2 Computation of the polygon area with standard triangulation

Input: Sequence of nodes defining the polygon, $s_k = [x_k, y_k]$, $k = 1, \dots, K$, ordered in the counter-clockwise direction with $s_1 = s_K$.

Output: The area of the polygon, \mathcal{A} .

- 1: Initialise the algorithm with $\mathcal{A} = 0$;
- 2: **for** $k \leftarrow 3, K - 2$ **do**
- 3: Draw diagonal from s_k to s_1 ;
- 4: **end for**
- 5: **for** $k \leftarrow 2, K - 2$ **do**
- 6: Calculate the area of the triangle with vertexes s_k , s_{k+1} , and s_1 as a determinant of the matrix:

$$\mathcal{A}_\Delta = \frac{1}{2} \det \begin{bmatrix} x_1 & y_1 & 1 \\ x_k & y_k & 1 \\ x_{k+1} & y_{k+1} & 1 \end{bmatrix};$$

- 7: $\mathcal{A} = \mathcal{A} + \mathcal{A}_\Delta$;
 - 8: **end for**
-

Figure C.2b) as the ones with negative area. Both methods have the same output as demonstrated in Figure C.2.

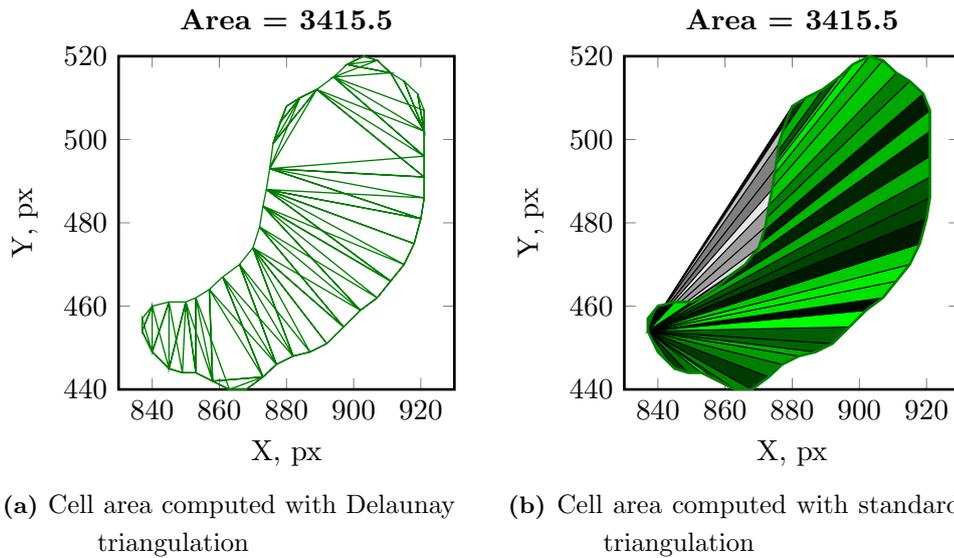


Figure C.2: Approaches to calculating the area of non-convex polygons.

Bibliography

- [1] E. A. Harvie & A. Huttenlocher, “Neutrophils in host defense: New insights from zebrafish”, *J. Leukocyte Biol.*, vol. 98, no. 4, pp. 523–537, 2015.
- [2] T. W. Starnes & A. Huttenlocher, “Neutrophil reverse migration becomes transparent with zebrafish”, in *Advances in Hematology*, 2012, pp. 1–11.
- [3] S. Nourshargh, S. A. Renshaw & B. A. Imhof, “Reverse migration of neutrophils: Where, when, how, and why?”, *Trends Immunol.*, vol. 37, no. 5, pp. 273–286, 2016.
- [4] E. Kolaczowska & P. Kubes, “Neutrophil recruitment and function in health and inflammation”, *Nat. Rev. Immunol.*, vol. 13, pp. 159–175, 2013.
- [5] W. L. Lee, R. E. Harrison & S. Grinstein, “Phagocytosis by neutrophils”, *Microbes Infect.*, vol. 5, no. 14, pp. 1299–1306, 2003, Forum in Immunology on Neutrophils.
- [6] C. D. Sadik, N. D. Kim & A. D. Luster, “Neutrophils cascading their way to inflammation”, *Trends Immunol.*, vol. 32, no. 10, pp. 452–460, 2011.
- [7] M. J. Tindall, P. K. Maini, S. L. Porter & J. P. Armitage, “Overview of mathematical approaches used to model bacterial chemotaxis ii: Bacterial populations”, *Bull. Math. Biol.*, vol. 70, no. 6, pp. 1570–1607, 2008.
- [8] M. U. Ehrenguber, D. A. Deranleau & T. D. Coates, “Shape oscillations of human neutrophil leukocytes: Characterization and relationship to cell motility”, *J. Exp. Biol.*, vol. 199, pp. 741–747, 1996.
- [9] W. L. Lee, R. E. Harrison & S. Grinstein, “Phagocytosis by neutrophils”, *Microbes Infect.*, vol. 5, no. 14, pp. 1299–1306, 2003.
- [10] P. Lacy, “Mechanisms of degranulation in neutrophils”, *Allergy, asthma, and clinical immunology : official journal of the Canadian Society of Allergy and Clinical Immunology*, vol. 2, no. 3, pp. 98–108, 2006.
- [11] V. Brinkmann *et al.*, “Neutrophil extracellular traps kill bacteria”, *Science*, vol. 303, no. 5663, pp. 1532–1535, 2004.
- [12] A. Mócsai, B. Walzog & C. A. Lowell, “Intracellular signalling during neutrophil recruitment”, *Cardiovasc. Res.*, vol. 107, no. 3, pp. 373–385, 2015.
- [13] A. D. Kennedy & F. R. DeLeo, “Neutrophil apoptosis and the resolution of infection”, *Immunol. Res.*, vol. 43, pp. 25–61, 2009.
- [14] C. Gray *et al.*, “Simultaneous intravital imaging of macrophage and neutrophil behaviour during inflammation using a novel transgenic zebrafish”, *Thromb. Haemost.*, vol. 105, no. 5, pp. 811–819, 2011.
- [15] B. Hamza *et al.*, “Retrotaxis of human neutrophils during mechanical confinement inside microfluidic channels”, *Integrative biology : quantitative biosciences from nano to macro*, vol. 6, no. 2, pp. 175–183, 2014.
- [16] F. Vianello, I. T. Olszak & M. C. Poznansky, “Fugetaxis: Active movement of leukocytes away from a chemokinetic agent”, *J. Mol. Med.*, vol. 83, no. 10, pp. 752–763, 2005.

- [17] G. R. Holmes *et al.*, “Repelled from the wound, or randomly dispersed? reverse migration behaviour of neutrophils characterized by dynamic modelling”, *J. R. Soc. Interface*, vol. 9, no. 77, pp. 3229–3239, 2012.
- [18] S. D. Kobayashi & F. R. DeLeo, “Role of neutrophils in innate immunity: A systems biology-level approach”, *Wiley Interdiscip. Rev. Syst. Biol. Med.*, vol. 1, no. 3, pp. 309–333, 2009.
- [19] P. A. Iglesias & P. N. Devreotes, “Navigating through models of chemotaxis”, *Curr. Opin. Cell Biol.*, vol. 20, no. 1, pp. 35–40, 2008.
- [20] T. P. Stossel, “The machinery of blood cell movements”, *Blood*, vol. 84, no. 2, pp. 367–379, 1994.
- [21] S. H. Zigmond, H. I. Levitsky & B. J. Kreel, “Cell polarity: An examination of its behavioural expression and its consequences for polymorphonuclear leukocyte chemotaxis”, *J. Cell Biol.*, vol. 89, no. 3, pp. 585–592, 1981.
- [22] S. H. Zigmond, “Ability of polymorphonuclear leukocytes to orient in gradients of chemotactic factors”, *J. Cell Biol.*, vol. 75, pp. 606–616, 1977.
- [23] R. Snyderman & E. Goetzl, “Molecular and cellular mechanisms of leukocyte chemotaxis”, *Science*, vol. 213, no. 4510, pp. 830–837, 1981.
- [24] K. Futosi, S. Fodor & A. Mócsai, “Neutrophil cell surface receptors and their intracellular signal transduction pathways”, *Int. Immunopharmacol.*, vol. 17, no. 3, pp. 638–650, 2013.
- [25] R. T. Tranquillo, D. A. Lauffenburger & S. H. Zigmond, “A stochastic model for leukocyte random motility and chemotaxis based on receptor binding fluctuations”, *J. Cell Biol.*, vol. 106, pp. 303–309, 1988.
- [26] S. Yoo *et al.*, “Differential regulation of protrusion and polarity by PI(3)K during neutrophil motility in live zebrafish”, *Dev. Cell*, vol. 21, 2011.
- [27] G. Servant *et al.*, “Polarization of chemoattractant receptor signaling during neutrophil chemotaxis”, *Science*, vol. 287, no. 5455, pp. 1037–1040, 2000.
- [28] Y. Xiong, C.-H. Huang, P. A. Iglesias & P. N. Devreotes, “Cells navigate with a local-excitation, global-inhibition-biased excitable network”, *Proceedings of the National Academy of Sciences*, vol. 107, no. 40, pp. 17079–17086, 2010.
- [29] M. P. Wymann *et al.*, “Oscillatory motion in human neutrophils responding to chemotactic stimuli”, *Biochem. Biophys. Res. Commun.*, vol. 147, no. 1, pp. 361–368, 1987.
- [30] D. Irimia, G. Balazsi, N. Agrawal & M. Toner, “Adaptive-control model for neutrophil orientation in the direction of chemical gradients”, *Biophys. J.*, vol. 96, no. 10, pp. 3897–3916, 2009.
- [31] K. M. Henry, C. A. Loynes, M. K. B. Whyte & S. A. Renshaw, “Zebrafish as a model for the study of neutrophil biology”, *J. Leukocyte Biol.*, vol. 94, no. 4, pp. 633–642, 2013.
- [32] B. Stramer *et al.*, “Live imaging of wound inflammation in drosophila embryos reveals key roles for small GTPases during *in vivo* cell migration”, *J. Cell Biol.*, vol. 168, no. 4, pp. 567–573, 2005.
- [33] S. A. Renshaw & C. A. Loynes, “A transgenic zebrafish model of neutrophilic inflammation”, *Blood*, vol. 108, no. 13, pp. 3976–3978, 2006.
- [34] C. Schwab, A. Klegeris & P. L. McGeer, “Inflammation in transgenic mouse models of neurodegenerative disorders”, *Biochimica et Biophysica Acta - Molecular Basis of Disease*, vol. 1802, no. 10, pp. 889–902, 2010.
- [35] J. Seok, H. S. Warren *et al.*, “Genomic responses in mouse models poorly mimic human inflammatory diseases”, *Proceedings of the National Academy of Sciences*, vol. 110, no. 9, pp. 3507–3512, 2013.
- [36] W. Razzell, W. Wood & P. Martin, “Swatting flies: Modelling wound healing and inflammation in drosophila”, *Disease Models & Mechanisms*, vol. 4, no. 5, pp. 569–574, 2011.
- [37] G. Lieschke & P. Currie, “Animal models of human disease: Zebrafish swim into view”, *Nat. Rev. Genet.*, vol. 8, no. 5, pp. 353–367, 2007.
- [38] M. Tomura *et al.*, “Monitoring cellular movement *in vivo* with photoconvertible fluorescence protein “kaede” transgenic mice”, *Proceedings of the National Academy of Sciences*, vol. 105, no. 31, pp. 10871–10876, 2008.

- [39] A. L. Robertson *et al.*, “A zebrafish compound screen reveals modulation of neutrophil reverse migration as an anti-inflammatory mechanism”, *Sci. Transl. Med.*, vol. 6, no. 225, pp. 225–229, 2014.
- [40] J. Chen *et al.*, “Akt isoforms differentially regulate neutrophil functions”, *Blood*, vol. 115, no. 21, pp. 4237–4246, 2010.
- [41] W. R. Holmes & L. Edelstein-Keshet, “A comparison of computational models for eukaryotic cell shape and motility”, *PLoS Comput. Biol.*, vol. 8, no. 12, e1002793, 2012.
- [42] E. F. Keller & L. A. Segel, “Model for chemotaxis”, *J. Theor. Biol.*, vol. 30, no. 2, pp. 225–234, 1971.
- [43] G. A. Dunn & A. F. Brown, “A unified approach to analysing cell motility”, *J. Cell Sci.*, vol. 1987, no. Supplement 8, pp. 81–102, 1987.
- [44] C. S. Patlak, “Random walk with persistence and external bias”, *Bull. Math. Biol.*, vol. 15, pp. 311–338, 1953.
- [45] T. Hillen & K. J. Painter, “A user’s guide to PDE models for chemotaxis”, *J. Math. Biol.*, vol. 58, no. 1, pp. 183–217, 2008.
- [46] D. Lauffenburger, R. Aris & K. Keller, “Effects of cell motility and chemotaxis on microbial population growth”, *Biophys. J.*, vol. 40, no. 3, pp. 209–219, 1982.
- [47] W. Alt, “Biased random walk models for chemotaxis and related diffusion approximations”, *J. Math. Biol.*, vol. 9, pp. 147–177, 1980.
- [48] K. J. Painter, “Mathematical models for chemotaxis and their applications in self-organisation phenomena”, *J. Theor. Biol.*, 2018.
- [49] G. E. Uhlenbeck & L. S. Ornstein, “On the theory of the Brownian motion”, *Phys. Rev.*, vol. 36, pp. 823–841, 5 1930.
- [50] J. Doob, “Brownian movement and stochastic equations”, *Ann. Math.*, vol. 43, no. 2, pp. 351–369, 1942.
- [51] C. L. Stokes, D. A. Lauffenburger & S. K. Williams, “Migration of individual microvessel endothelial cells: Stochastic model and parameter measurement”, *J. Cell Sci.*, vol. 99, pp. 419–430, 1991.
- [52] R. J. Petrie, A. D. Doyle & K. M. Yamada, “Random versus directionally persistent cell migration”, *Nat. Rev. Mol. Cell Biol.*, vol. 10, no. 8, pp. 538–549, 2009.
- [53] R. T. Tranquillo & D. A. Lauffenburger, “Stochastic model of leukocyte chemosensory movement”, *J. Math. Biol.*, vol. 25, pp. 229–262, 1987.
- [54] P. J. M. Jones *et al.*, “Inference of random walk models to describe leukocyte migration”, *Phys. Biol.*, vol. 12, no. 6, pp. 1–12, 2015.
- [55] R. S. Hartman, K. Lau, W. Chou & T. D. Coates, “The fundamental motor of the human neutrophil is not random: Evidence for local non-Markov movement in neutrophils”, *Biophys. J.*, vol. 67, no. 6, pp. 2535–2545, 1994.
- [56] A. Jilkine & L. Edelstein-Keshet, “A comparison of mathematical models for polarization of single eukaryotic cells in response to guided cues”, *PLoS Comput. Biol.*, vol. 7, no. 4, pp. 1–15, 2011.
- [57] A. Gierer & H. Meinhardt, “A theory of biological pattern formation”, *Kybernetik*, vol. 12, no. 1, pp. 30–39, 1972.
- [58] H. Meinhardt, “Orientation of chemotactic cells and growth cones: Models and mechanisms”, *J. Cell Sci.*, vol. 112, pp. 2867–2874, 1999.
- [59] A. Levchenko & P. A. Iglesias, “Models of eukaryotic gradient sensing: Application to chemotaxis of amoebae and neutrophils”, *Biophys. J.*, vol. 82, no. 1, pp. 50–63, 2002.
- [60] C. Shi, C.-H. Huang, P. N. Devreotes & P. A. Iglesias, “Interaction of motility, directional sensing, and polarity modules recreates the behaviors of chemotaxing cells”, *PLoS Comput. Biol.*, vol. 9, no. 7, pp. 1–17, 2013.
- [61] L. Ma *et al.*, “Two complementary, local excitation, global inhibition mechanisms acting in parallel can explain the chemoattractant-induced regulation of PI(3,4,5)P3 response in dictyostelium cells.”, *Biophys. J.*, vol. 87, no. 6, pp. 3764–3774, 2004.

- [62] Y. Xiong, C.-H. Huang, P. A. Iglesias & P. N. Devreotes, “Cells navigate with a local-excitation, global-inhibition-biased excitable network”, *Proceedings of the National Academy of Sciences*, vol. 107, no. 40, pp. 17 079–17 086, 2010.
- [63] A. Mogilner, “Mathematics of cell motility: Have we got its number?”, *J. Math. Biol.*, vol. 58, no. 1, pp. 105–134, 2008.
- [64] M. P. Neilson *et al.*, “Chemotaxis: A feedback-based computational model robustly predicts multiple aspects of real cell behaviour”, *PLoS Biol.*, vol. 9, no. 5, pp. 1–11, 2011.
- [65] C. M. Elliott, B. Stinner & C. Venkataraman, “Modelling cell motility and chemotaxis with evolving surface finite elements”, *J. R. Soc. Interface*, vol. 9, no. 76, pp. 3027–3044, 2012.
- [66] G. MacDonald, J. A. Mackenzie, M. R. Nolan & R. H. Insall, “A computational method for the coupled solution of reaction–diffusion equations on evolving domains and manifolds: Application to a model of cell migration and chemotaxis”, *J. Comput. Phys.*, vol. 309, pp. 207–226, 2016.
- [67] D. Shao, W.-J. Rappel & H. Levine, “Computational model for cell morphodynamics”, *Biophys. Rev. Lett.*, vol. 105, p. 108 104, 10 2010.
- [68] M. P. Neilson, J. A. Mackenzie, S. D. Webb & R. H. Insall, “Modeling Cell Movement and Chemotaxis Using Pseudopod-Based Feedback”, Tech. Rep., 2011, pp. 1–21.
- [69] R. H. Insall, “Understanding eukaryotic chemotaxis: A pseudopod-centred view”, *Nat. Rev. Mol. Cell Biol.*, vol. 11, no. 6, pp. 453–458, 2010.
- [70] J. Durbin & S. Koopman, *Time Series Analysis by State Space Methods*, ser. Oxford Statistical Science Series. Clarendon Press, 2001.
- [71] M. Aoki, *State Space Modelling of Time Series*. Berlin, Heidelberg: Springer-Verlag, 1986.
- [72] Y. Bar-Shalom & X. R. Li, *Estimation and Tracking: Principles, Techniques, and Software*. Norwood, MA: Artech House, Inc, 1993.
- [73] D. Sorensen & D. Gianola, *Likelihood, Bayesian, and MCMC Methods in Quantitative Genetics*, ser. Statistics for Biology and Health. Springer New York, 2007.
- [74] M. H. DeGroot & M. J. Schervish, *Probability and Statistics*, 4th ed. Pearson Education Ltd., 2012.
- [75] R. T. Cox, “Probability, frequency and reasonable expectation”, *Am. J. Phys.*, vol. 14, no. 1, pp. 1–13, 1946.
- [76] M. G. Kendall, “On the reconciliation of theories of probability”, *Biometrika*, vol. 36, no. 1-2, pp. 101–116, 1949.
- [77] M. J. Bayarri & J. O. Berger, “The interplay of Bayesian and frequentist analysis”, *Statistical Science*, vol. 19, no. 1, pp. 58–80, 2004.
- [78] M. J. Zyphur & F. L. Oswald, “Bayesian estimation and inference: A user’s guide”, *J. Manage.*, vol. 41, no. 2, pp. 390–420, 2015.
- [79] J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer, 1985.
- [80] M. A. Girshick & L. J. Savage, “Bayes and minimax estimates for quadratic loss functions”, in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, 1951, pp. 53–73.
- [81] S. Särkka, *Bayesian Filtering and Smoothing*, ser. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2013.
- [82] S. Brooks, A. Gelman, G. Jones & X.-L. Meng, *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.
- [83] K. Csillery, M. G. B. Blum, O. E. Gaggiotti & O. François, “Approximate Bayesian computation (ABC) in practice”, *Trends in Ecology and Evolution*, vol. 25, no. 7, pp. 410–418, 2010.
- [84] G. McLachlan & T. Krishnan, *The EM Algorithm and Extensions*, 2nd ed. Wiley, 2008.
- [85] E. L. Lehman & G. Casella, *Theory of Point Estimation*, 2nd ed. Springer, 1998.
- [86] B. Efron & D. V. Hinkley, “Assessing the accuracy of the Maximum Likelihood estimator: Observed versus expected Fisher information”, *Biometrika*, vol. 65, no. 3, pp. 457–483, 1978.

- [87] N. R. Draper & I. Guttman, “Confidence intervals versus regions”, *J. R. Stat. Soc. Ser. D (Statistician)*, vol. 44, no. 3, pp. 399–403, 1995.
- [88] Y Bar-Shalom, X Rong Li & T Kirubajan, *Estimation with Applications to Tracking and Navigation. Theory, Algorithms, and Software*. 2001.
- [89] R. Bassett & J. Deride, “Maximum a posteriori estimators as a limit of Bayes estimators”, *Math. Program.*, vol. 174, no. 1, pp. 129–144, 2019.
- [90] H. W. Sorenson, “Least Squares estimation: From Gauss to Kalman”, *IEEE Spectr.*, vol. 7, no. 7, pp. 63–68, 1970.
- [91] A. Aitken, “On Least Squares and linear combination of observations”, *Proceedings of the Royal Society of Edinburgh*, vol. 55, pp. 42–48, 1936.
- [92] C. M. Bishop, *Pattern Recognition and Machine Learning: Christopher M. Bishop*. Springer-Verlag, 2006.
- [93] A. P. Dempster, N. M. Laird & D. B. Rubin, “Maximum Likelihood from incomplete data via the EM algorithm”, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, vol. 39, no. 1, pp. 1–38, 1977.
- [94] J. C. F. Wu, “On the convergence properties of the EM algorithm”, *Ann. Stat.*, vol. 11, no. 1, pp. 95–103, 1983.
- [95] M. R. Gupta & Y. Chen, “Theory and use of the EM algorithm”, *Foundations and Trends in Signal Processing*, vol. 4, no. 3, pp. 223–296, 2011.
- [96] R. H. Shumway & D. S. Stoffer, “An approach to time series smoothing and forecasting using the EM algorithm”, *J. Time Ser. Anal.*, vol. 3, no. 4, pp. 253–264, 1982.
- [97] Z. Ghahramani & G. E. Hinton, “Parameter estimation for linear dynamical systems”, University of Toronto, Tech. Rep., 1996.
- [98] R. Lee, *Optimal Estimation, Identification, and Control*, ser. M.I.T. Press research monograph, no. 28. M.I.T. Press, 1964.
- [99] D. Fraser & J. Potter, “The optimum linear smoother as a combination of two optimum linear filters”, *IEEE Trans. Autom. Control*, vol. 14, no. 4, pp. 387–390, 1969.
- [100] G. Kitagawa, “Non-Gaussian state-space modeling of nonstationary time series”, *J. Am. Stat. Assoc.*, vol. 82, no. 400, pp. 1032–1041, 1987.
- [101] R. E. Kalman, “A new approach to linear filtering and prediction problems”, *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.
- [102] B. D. O. Anderson & J. B. Moore, *Optimal Filtering*. Prentice-Hall, Inc., 1979.
- [103] M. S. Grewal & A. P. Andrews, “Applications of Kalman filtering in aerospace 1960 to the present [historical perspectives]”, *IEEE Control Syst. Mag.*, vol. 30, no. 3, pp. 69–78, 2010.
- [104] C. R. Rao, “A note on Kalman filter”, *Proceedings of the National Academy of Sciences*, vol. 98, no. 19, pp. 10 557–10 559, 2001.
- [105] G. Verghese & T. Kailath, “A further note on backwards Markovian models (corresp.)”, *IEEE Trans. Inf. Theory*, vol. 25, no. 1, pp. 121–124, 1979.
- [106] H. E. Rauch, F. Tung & C. T. Striebel, “Maximum Likelihood estimates of linear dynamic systems”, *AIAA J.*, vol. 3, no. 8, pp. 1445–1450, 1965.
- [107] R. H. Shumway & D. S. Stoffer, *Time Series Analysis and Its Applications (Springer Texts in Statistics)*. Springer-Verlag, 2005.
- [108] A. Jazwinski, *Stochastic Processes and Filtering Theory*, ser. Mathematics in Science and Engineering. Elsevier Science, 1970.
- [109] S. J. Julier & J. K. Uhlmann, “New extension of the Kalman filter to nonlinear systems”, vol. 3068, 1997.
- [110] E. A. Wan & R. V. D. Merwe, “The unscented Kalman filter for nonlinear estimation”, in *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000, pp. 153–158.

- [111] S. J. Julier & J. K. Uhlmann, “Unscented filtering and nonlinear estimation”, *Proc. IEEE*, vol. 92, no. 3, pp. 401–422, 2004.
- [112] S. Sarkka, “Unscented Rauch–Tung–Striebel smoother”, *IEEE Trans. Autom. Control*, vol. 53, no. 3, pp. 845–849, 2008.
- [113] A. Gelb, *Applied Optimal Estimation*, ser. Applied Optimal Estimation. MIT Press, 1974.
- [114] I. Arasaratnam & S. Haykin, “Cubature Kalman filters”, *IEEE Trans. Autom. Control*, vol. 54, no. 6, pp. 1254–1269, 2009.
- [115] K. Ito & K. Xiong, “Gaussian filters for nonlinear filtering problems”, *IEEE Trans. Autom. Control*, vol. 45, no. 5, pp. 910–927, 2000.
- [116] A. Doucet & A. M. Johansen, *A tutorial on particle filtering and smoothing: Fifteen years later*, 2011.
- [117] Z. Ghahramani & G. E. Hinton, “Variational learning for switching state-space models”, *Neural Comput.*, vol. 12, no. 4, pp. 831–864, 2000.
- [118] O. Costa, M. Fragoso & R. Marques, *Discrete-Time Markov Jump Linear Systems*, ser. Probability and Its Applications. Springer London, 2006.
- [119] F. Gustafsson, *Adaptive Filtering and Change Detection*. Wiley, 2000.
- [120] H. A. P. Blom, “An efficient filter for abruptly changing systems”, in *The 23rd IEEE Conference on Decision and Control*, 1984, pp. 656–658.
- [121] H. A. P. Blom & Y. Bar-Shalom, “The interacting multiple model algorithm for systems with Markovian switching coefficients”, *IEEE Trans. Autom. Control*, vol. 33, no. 8, pp. 780–783, 1988.
- [122] —, “Time-reversion of a hybrid state stochastic difference system with a jump-linear smoothing application”, *IEEE Trans. Inf. Theory*, vol. 36, no. 4, pp. 836–847, 1990.
- [123] R. E. Helmick, W. D. Blair & S. A. Hoffman, “Fixed-interval smoothing for Markovian switching systems”, *IEEE Trans. Inf. Theory*, vol. 41, no. 6, pp. 1845–1855, 1995.
- [124] B. Chen & J. K. Tugnait, “Interacting multiple model fixed-lag smoothing algorithm for Markovian switching systems”, *IEEE Trans. Aerosp. Electron. Syst.*, vol. 36, no. 1, pp. 243–250, 2000.
- [125] W. Koch, “Fixed-interval retrodiction approach to Bayesian IMM-MHT for maneuvering multiple targets”, *IEEE Transactions Aerospace and Electronic Systems*, vol. 36, pp. 2–14, 2000.
- [126] N. Nadarajah, R. Tharmarasa, M. McDonald & T. Kirubarajan, “IMM forward filtering and backward smoothing for maneuvering target tracking”, *IEEE Trans. Aerosp. Electron. Syst.*, vol. 48, no. 3, pp. 2673–2678, 2012.
- [127] A. Kadochnikova, H. M. Isles, S. A. Renshaw & V. Kadiramanathan, “Estimation of hidden chemoattractant field from observed cell migration patterns”, *IFAC-PapersOnLine*, vol. 55, no. 15, pp. 766–771, 2018.
- [128] V. Kadiramanathan *et al.*, “The neutrophil’s eye-view: Inference and visualisation of the chemoattractant field driving cell chemotaxis *in vivo*”, *PLoS One*, vol. 7, no. 4, e35182, 2012.
- [129] J. Liepe *et al.*, “Calibrating spatio-temporal models of leukocyte dynamics against *in vivo* live-imaging data using approximate Bayesian computation”, *Integr. Biol.*, vol. 4, pp. 335–345, 2012.
- [130] M. Zhang, Y. Shen, Q. Wang & Y. Wang, “Dynamic artificial potential field based multi-robot formation control”, in *2010 IEEE Instrumentation Measurement Technology Conference Proceedings*, 2010, pp. 1530–1534.
- [131] A. Wallar & E. Plaku, “Path planning for swarms in dynamic environments by combining probabilistic roadmaps and potential fields”, in *2014 IEEE Symposium on Swarm Intelligence*, 2014, pp. 1–8.
- [132] H. Siu, M. Figueroa, J. Holt & S. Biaz, “A dynamic swarm approach to artificial potential field collision avoidance”, Tech. Rep., 2012, pp. 1–13.
- [133] C. Nusslein-Volhard & R. Dahm, *Zebrafish: A Practical Approach*, 1st ed. Oxford University Press, 2002.

- [134] K. Hollig & J. Horner, *Approximation and Modeling with B-Splines*. Society for Industrial and Applied Mathematics, 2013.
- [135] S. Gibson & B. Ninness, “Robust Maximum Likelihood estimation of multivariable dynamic systems”, *Automatica*, vol. 41, no. 10, pp. 1667–1682, 2005.
- [136] M. Dewar & V. Kadiramanathan, “A canonical space-time state space model: State and parameter estimation”, *IEEE Trans. Signal Process.*, vol. 55, no. 10, pp. 4862–4870, 2007.
- [137] Z. Ghahramani & S. T. Roweis, “Learning nonlinear dynamical systems using an EM algorithm”, in *Advances in Neural Information Processing Systems 11*, MIT Press, 1999, pp. 431–437.
- [138] X.-L. Meng & D. B. Rubin, “Maximum Likelihood estimation via the ECM algorithm: A general framework”, *Biometrika*, vol. 80, no. 2, pp. 267–278, 1993.
- [139] X.-L. Meng, “On the rate of convergence of the ECM algorithm”, *Ann. Stat.*, vol. 22, no. 1, pp. 326–339, 1994.
- [140] J. Yun, F. Yang & Y. Chen, “Augmented particle filters”, *J. Am. Stat. Assoc.*, vol. 112, no. 517, pp. 300–313, 2017.
- [141] J. P. Gleeson, S. Melnik, J. A. Ward, M. A. Porter & P. J. Mucha, “Accuracy of mean-field theory for dynamics on real-world networks.”, *Phys. Rev. E, Statistical, nonlinear, and soft matter physics*, vol. 85, p. 026 106, 2 2012.
- [142] M. Massar & S. Massar, “Mean-field theory of echo state networks”, *Phys. Rev. E*, vol. 87, p. 042 809, 4 2013.
- [143] T. Orchard & M. A. Woodbury, “A missing information principle: Theory and applications”, in *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Theory of Statistics*, University of California Press, 1972, pp. 697–715.
- [144] T. A. Louis, “Finding the observed information matrix when using the EM algorithm”, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, vol. 44, no. 2, pp. 226–233, 1982.
- [145] E. Walter, *Identifiability of Parametric Models*. Elsevier Science, 2014.
- [146] T. Lammermann *et al.*, “Neutrophil swarms require LTB4 and integrins at sites of cell death *in vivo*.”, *Nature*, vol. 498, no. 5, pp. 371–375, 2013.
- [147] B. Holmquist, “Moments and cumulants of the multivariate normal distribution”, *Stochastic Analysis and Applications*, vol. 6, no. 3, pp. 273–278, 1988.
- [148] M. A. Qasaimeh *et al.*, “Neutrophil chemotaxis in moving gradients”, *Advanced Biosystems*, vol. 2, no. 7, p. 1 700 243, 2018.
- [149] T. Kirubarajan, Y. Bar-Shalom, K. R. Pattipati & I. Kadar, “Ground target tracking with variable structure IMM estimator”, *IEEE Trans. Aerosp. Electron. Syst.*, vol. 36, no. 1, pp. 26–46, 2000.
- [150] O. Khatib, “Real-time obstacle avoidance for manipulators and mobile robots”, *Int. J. Rob. Res.*, vol. 5, no. 1, pp. 90–98, 1986.
- [151] S. S. Ge & Y. J. Cui, “Dynamic motion planning for mobile robots using potential field method”, *Autonomous Robots*, vol. 13, no. 3, pp. 207–222, 2002.
- [152] J. Murphy & S. Godsill, “Simultaneous localization and mapping for non-parametric potential field environments”, in *2012 Workshop on Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, 2012, pp. 1–6.
- [153] I. Rapoport & Y. Oshman, “Efficient fault tolerant estimation using the IMM methodology”, *IEEE Trans. Aerosp. Electron. Syst.*, vol. 43, no. 2, pp. 492–508, 2007.
- [154] J. W. Armond *et al.*, “A stochastic model dissects cell states in biological transition processes.”, *Sci. Rep.*, vol. 4, p. 3692, 2014.
- [155] J. Tugnait, “Adaptive estimation and identification for discrete systems with Markov jump parameters”, *IEEE Trans. Autom. Control*, vol. 27, no. 5, pp. 1054–1065, 1982.
- [156] A. Svensson, T. B. Schön & F. Lindsten, “Identification of jump Markov linear models using particle filters”, in *53rd IEEE Conference on Decision and Control*, 2014, pp. 6504–6509.

- [157] T. T. Ashley & S. B. Andersson, “A sequential Monte Carlo framework for the system identification of jump Markov state space models”, in *2014 American Control Conference*, 2014, pp. 1144–1149.
- [158] S. Yildirim, S. S. Singh & A. Doucet, “An online Expectation–Maximization algorithm for change-point models”, *J. Comput. Graph. Stat.*, vol. 22, no. 4, pp. 906–926, 2013.
- [159] —, “Interacting multiple model methods in target tracking: A survey”, *IEEE Trans. Aerosp. Electron. Syst.*, vol. 34, no. 1, pp. 103–123, 1998.
- [160] L. Blackmore, S. Gil, S. Chung & B. Williams, “Model learning for switching linear systems with autonomous mode transitions”, in *2007 46th IEEE Conference on Decision and Control*, 2007, pp. 4648–4655.
- [161] S. Yildirim, S. S. Singh & A. Doucet, “An online Expectation–Maximization algorithm for change-point models”, *J. Comput. Graph. Stat.*, vol. 22, no. 4, pp. 906–926, 2013.
- [162] A. B. Chan, Z.-S. J. Liang & N. Vasconcelos, “Privacy preserving crowd monitoring: Counting people without people models or tracking”, in *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–7.
- [163] V. Rabaud & S. Belongie, “Counting crowded moving objects”, in *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, vol. 1, 2006, pp. 705–711.
- [164] B. Leibe, K. Schindler & L. V. Gool, “Coupled detection and trajectory estimation for multi-object tracking”, in *2007 IEEE 11th International Conference on Computer Vision*, 2007, pp. 1–8.
- [165] A. Logothetis & V. Krishnamurthy, “Expectation Maximization algorithms for MAP estimation of jump Markov linear systems”, *IEEE Trans. Signal Process.*, vol. 47, no. 8, pp. 2139–2156, 1999.
- [166] L. A. Johnston & V. Krishnamurthy, “An improvement to the interacting multiple model (IMM) algorithm”, *IEEE Trans. Signal Process.*, vol. 49, no. 12, pp. 2909–2923, 2001.
- [167] E. Özkan, F. Lindsten, C. Fritsche & F. Gustafsson, “Recursive Maximum Likelihood Identification of jump Markov nonlinear systems”, *IEEE Transactions on Signal Processing*, vol. 63, no. 3, pp. 754–765, 2015.
- [168] D. M. Blei, A. Kucukelbir & J. D. McAuliffe, “Variational inference: A review for statisticians”, *Journal of the American Statistical Association*, vol. 112, no. 518, pp. 859–877, 2017.
- [169] M. Opper & G. Sanguinetti, “Variational inference for Markov jump processes”, in *Advances in Neural Information Processing Systems*, J. C. Platt, D. Koller, Y. Singer & S. T. Roweis, Eds., vol. 20, Curran Associates, Inc., 2008, pp. 1105–1112.
- [170] M. Opper, A. Ruttner & G. Sanguinetti, “Approximate inference in continuous time Gaussian-Jump processes”, in *Advances in Neural Information Processing Systems*, J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel & A. Culotta, Eds., vol. 23, Curran Associates, Inc., 2010, pp. 1831–1839.
- [171] M. J. Beal & Z. Ghahramani, “The variational bayesian em algorithm for incomplete data: With application to scoring graphical model structures”, in *Bayesian Statistics*, J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith & M. West, Eds., vol. 7, 2003, pp. 453–464.
- [172] A. Zammit-Mangion, G. Sanguinetti & V. Kadiramanathan, “Variational estimation in spatiotemporal systems from continuous and point-process observations”, *IEEE Transactions on Signal Processing*, vol. 60, no. 7, pp. 3449–3459, 2012.
- [173] J. Cao, Y. Liang & L. Liu, “Variational bayesian inference for jump markov linear systems with unknown transition probabilities”, in *2018 21st International Conference on Information Fusion (FUSION)*, 2018, pp. 2065–2071.
- [174] V. Pavlovic & J. M. Rehg, “Impact of dynamic model learning on classification of human motion”, *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, 788–795 vol.1, 2000.
- [175] E. Cinquemani, R. Porreca, G. Ferrari-Trecate & J. Lygeros, “A general framework for the identification of jump Markov linear systems”, in *2007 46th IEEE Conference on Decision and Control*, 2007, pp. 5737–5742.

- [176] A. Boyarsky & P. Noble, "A Markov chain characterization of human neutrophil locomotion under neutral and chemotactic conditions.", *Can. J. Physiol. Pharmacol.*, vol. 55, no. 1, pp. 1–6, 1977.
- [177] H. Weavers *et al.*, "Systems analysis of the dynamic inflammatory response to tissue damage reveals spatiotemporal properties of the wound attractant gradient", *Curr. Biol.*, vol. 26, no. 15, pp. 1975–1989, 2016.
- [178] J. R. Mathias *et al.*, "Resolution of inflammation by retrograde chemotaxis of neutrophils in transgenic zebrafish", *J. Leukocyte Biol.*, vol. 80, no. 6, pp. 1281–1288, 2006.
- [179] G. R. Holmes *et al.*, "Drift-diffusion analysis of neutrophil migration during inflammation resolution in a zebrafish model", *Advances in Hematology*, vol. 2012, pp. 1–8, 2012.
- [180] G. R. Holmes, "The spatiotemporal dynamics of inflammatory neutrophil populations", 2013.
- [181] F. Ellett *et al.*, "Defining the phenotype of neutrophils following reverse migration in zebrafish", *J. Leukocyte Biol.*, vol. 98, pp. 975–981, 2015.
- [182] G. Dixon, P. Elks, C. Loynes, M. Whyte & S. Renshaw, "A method for the *in vivo* measurement of zebrafish tissue neutrophil lifespan.", *ISRN Hematology*, vol. 2012, p. 915 868, 2012.
- [183] J. K. Tugnait & A. H. Haddad, "A detection-estimation scheme for state estimation in switching environments", *Automatica*, vol. 15, no. 4, pp. 477–481, 1979.
- [184] P. Andersson, "Adaptive forgetting in recursive identification through multiple models", *Int. J. Control*, vol. 42, no. 5, pp. 1175–1193, 1985.
- [185] A. S. Rahmathullah, L. Svensson & D. Svensson, "Two-filter Gaussian mixture smoothing with posterior pruning", in *17th International Conference on Information Fusion (FUSION)*, 2014, pp. 1–8.
- [186] C. A. Loynes *et al.*, "Pivotal advance: Pharmacological manipulation of inflammation resolution during spontaneously resolving tissue neutrophilia in the zebrafish", *J. Leukocyte Biol.*, vol. 87, no. 2, pp. 203–212, 2010.
- [187] S. R. McColl & I. Clark-Lewis, "Inhibition of murine neutrophil recruitment *in vivo* by CXC chemokine receptor antagonists", *J. Immunol.*, vol. 163, no. 5, pp. 2829–2835, 1999.
- [188] A. H. Lopes *et al.*, "DF2755A, a novel non-competitive allosteric inhibitor of CXCR1/2, reduces inflammatory and post-operative pain", *Pharmacol. Res.*, vol. 103, pp. 69–79, 2016.
- [189] G. F. Baxter, "The neutrophil as a mediator of myocardial ischemia-reperfusion injury: Time to move on", *Basic Res. Cardiol.*, vol. 97, no. 4, pp. 268–275, 2002.
- [190] A. Robertson, "A novel anti-inflammatory mechanism identified by an *in vivo* chemical genetic screen", 2013.
- [191] M. J. Johnson, Z. Shen & Y. Xu, "Scattered data reconstruction by regularization in B-spline and associated wavelet spaces", *J. Approx. Theory*, vol. 159, no. 8, pp. 197–223, 2009.
- [192] P. Aram *et al.*, "Spatiotemporal multi-resolution approximation of the Amari type neural field model", *Neuroimage*, vol. 66, pp. 88–102, 2013.
- [193] J. T. O'Flaherty *et al.*, "Selective desensitization of neutrophils: Further studies with 1-o-alkyl-sn-glycero-3-phosphocholine analogues.", *J. Immunol.*, vol. 127, no. 2, pp. 731–737, 1981.
- [194] J. C. Sible & J. J. Tyson, "Mathematical modelling as a tool for investigating cell cycle control networks", *Methods*, vol. 41, no. 2, pp. 238–247, 2007.
- [195] R. H. Insall & O. D. Weiner, "PIP3, PIP2, and cell movement – similar messages, different meanings?", *Dev. Cell*, vol. 1, no. 6, pp. 743–747, 2001.
- [196] K. A. Hinchliffe, "Cellular signalling: Stressing the importance of PIP3", *Curr. Biol.*, vol. 11, no. 9, R371–R373, 2001.
- [197] F. Wang *et al.*, "Lipid products of PI(3)Ks maintain persistent cell polarity and directed motility in neutrophils", *Nat. Cell Biol.*, vol. 4, pp. 513–518, 2002.
- [198] Y. Zhang, "A Dynamical Systems Modelling Framework for Breast Cancer Cell Motility and Morphology Analysis", 2017.

- [199] X Rong Li & Y. Bar-shalom, "Tracking in clutter with nearest neighbor filters: Analysis and performance", *IEEE Trans. Aerosp. Electron. Syst.*, vol. 32, pp. 995–1010, 1996.
- [200] W. S. AlAzawee, I. Abdel-Qader & J. Abdel-Qader, "Using morphological operations – erosion based algorithm for edge detection", in *2015 IEEE International Conference on Electro/Information Technology (EIT)*, 2015, pp. 521–525.
- [201] J. D. Gibbons & S. Chakraborti, *Nonparametric Statistical Inference*. Marcel Dekker, Inc., 2011.
- [202] O. Hoeller & R. R. Kay, "Chemotaxis in the absence of PIP3 gradients", *Curr. Biol.*, vol. 17, no. 9, pp. 813–817, 2007.
- [203] B. Heit *et al.*, "Pi3k accelerates, but is not required for, neutrophil chemotaxis to fMLP", *J. Cell Sci.*, vol. 121, no. 2, pp. 205–214, 2008.
- [204] K. Wong, O. Pertz, K. Hahn & H. Bourne, "Neutrophil polarization: Spatiotemporal dynamics of RhoA activity support a self-organizing mechanism", *Proceedings of the National Academy of Sciences*, vol. 103, no. 10, pp. 3639–3644, 2006.
- [205] H. Zhang, C. Sun, M. Glogauer & G. M. Bokoch, "Human neutrophils coordinate chemotaxis by differential activation of Rac1 and Rac2", *J. Immunol.*, vol. 183, no. 4, pp. 2718–2728, 2009.
- [206] C. C. Campa *et al.*, "Rac signal adaptation controls neutrophil mobilization from the bone marrow", *Sci. Signal.*, vol. 9, no. 459, ra124, 2016.
- [207] J. Bernardo *et al.*, "The variational Bayesian EM algorithm for incomplete data: With application to scoring graphical model structures", *Bayesian statistics*, vol. 7, pp. 453–464, 2003.
- [208] P. Niethammer, C. Grabher, A. T. Look & T. J. Mitchison, "A tissue-scale gradient of hydrogen peroxide mediates rapid wound detection in zebrafish", *Nature*, vol. 459, no. 7249, 996–999, 2009.
- [209] D. Wu & F. Lin, "Modeling cell gradient sensing and migration in competing chemoattractant fields", *PLoS One*, vol. 6, no. 4, pp. 1–12, 2011.
- [210] Y. Xiong & P. A. Iglesias, "Tools for analyzing cell shape changes during chemotaxis", *Integr. Biol.*, vol. 2, pp. 561–567, 11-12 2010.
- [211] J. Yeh, *Real Analysis: Theory of Measure and Integration*. World Scientific, 2006.
- [212] E. Kreyszig, *Differential Geometry*, ser. Differential Geometry. Dover Publications, 1991.