Techniques for analysis and integration of histology data and agent-based modelling of T cell-tumour interactions



Róisín Stephens
Hertford College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Michaelmas 2023

For my Grandad.

Acknowledgements

Firstly, I would like to thank my supervisors, Helen and Philip. Thank you for your support, guidance and mentorship. It has been a privilege to work with you both. Lucy, thank you for your encouragement and for hosting me during several trips to Roche, which I thoroughly enjoyed. Thank you to EPSRC and Roche for funding my DPhil studies. Thank you to the SABS DTC for supporting me throughout the DPhil. I would like to thank my transfer and confirmation of status examiners, Ruth Baker, Joe Pitt-Francis and Mark Coles, for your valuable input and advice. I would also like to thank Charlotte and Sandy, for your support and guidance.

To the WCMB family - I feel incredibly lucky to have been part of such a supportive, kind and brilliant research group; there are far too many people to thank individually. Josh, you have been my rock - thank you for the many inspirational discussions through the DPhil which have helped shaped my work, and for being there in both the good and the difficult times. Emma, we started the DPhil together and I am so glad we have been there to support each other. Thanks to all the brilliant people in Office S4.16, past and present - to Maxi and Joanneke for welcoming me to the office, to Duncan, for being by my side (literally) for three years of the DPhil, and to Simon, Carles and Sam for all the laughs, and for keeping me sane during the stressful last stages of my thesis. Thank you Kara, Sofia and Liam for your friendship and all the fun we've had.

Bobby, thank you for your enthusiastic support and the many coffees at Gail's over the years. Gaby, thank you for being there for me, even from the other side of the world. Gina, thank you for your belief in me - we both got our theses written in the end. Azamati, thank you for your friendship and encouraging me to see the positive. Anna and Niamh, thank you for your support, particularly over the lockdowns in Oxford.

A special thanks go to my family. To my grandparents, for all of your love and encouragement. To Kieron, for all of your support. To my sister Katrina, thank you for keeping me sane, and for not allowing me to talk about my research for more than 30 seconds! Finally, thank you to my Mum and Dad. Without your love, faith and unconditional support I would not have got to this stage - you always believed I would do it.

Abstract

Cytotoxic (CD8+) T cells form an important part of the adaptive immune response. They kill infected and abnormal cells, including cancerous cells, via a process known as immunosurveillance. When CD8+ T cells infiltrate tumour tissue, they identify antigens on the surface of tumour cells, and begin tumour cell lysis. Immunotherapies are increasingly being developed to target aspects of the immune system, such as CD8+ T cells, in order to boost its anti-tumour response. In clinical trials of such therapies, patients are monitored in several ways, often including core-needle biopsies in which a small piece of the tumour is removed, stained for specific cell markers, and imaged.

The increasing availability of such spatially-resolved imaging data motivates the development of mathematical models which capture the spatial dynamics of individual CD8+ T cells and tumour cells. Agent-based models (ABMs) allow the simulation of complex biological systems with detailed temporal and spatial resolution. Spatial statistics, for example the pair correlation function (PCF), can quantify the spatial heterogeneity inherent in both biological and synthetic data, and generate insight into the underlying biological system.

The main aims of this DPhil thesis are to explore how mathematical modelling and image analysis techniques can be used in oncology clinical trials to gain insight into the mechanism of action of immunotherapeutics. In the first part of the thesis, we adapt an existing on-lattice ABM of tumour-CD8+ T cell interactions in order to elucidate the key mechanisms driving oscillatory dynamics caused by tumour-T cell interactions, and investigate how specific modelling assumptions influence the spatial dynamics. We then undertake an exploratory analysis of histology data from a Phase I clinical trial of Simlukafusp (FAP-IL2v), a cancer immunotherapy developed by Roche. Here we aim to ascertain whether using the PCF to analyse the imaging data can uncover further insights into the efficacy of the drug, as opposed to the current clinical practice of recording cell counts and visual inspection of CD8+ T cell infiltration in the tumour. In the final part of the thesis we develop a new on-lattice ABM which builds on the findings of our first model, and use it to identify conditions under which spatial statistics enable us to make earlier and more accurate predictions of tumour responses to immunotherapy.

Contents

1 Introduction **1**

1.1 Thesis overview . . . 1

1.2 Review of biological background . . . 2

1.2.1 T cells and the immune system . . . 2

1.2.2 The cancer-immunity cycle . . . 3

1.2.2.1 The 3 Es of cancer immunoediting . . . 4

1.2.3 Cancer immunotherapies targetting T cells . . . 4

1.3 Mathematical models in immuno-oncology . . . 4

1.4 Agent-based modelling . . . 5

1.4.1 ABM implementation . . . 6

1.4.2 Agent-based modelling of tumour-immune interactions . . . 7

1.4.3 Implementation of on-lattice ABMs . . . 11

1.4.3.1 Domain size and boundary conditions . . . 11

1.4.3.2 Neighbourhood choice . . . 12

1.4.3.3 Grid occupancy . . . 12

1.4.3.4 Timestep . . . 12

1.5 Spatial statistics . . . 13

1.5.1 Pair correlation function (PCF) . . . 13

1.5.2 Spatial statistics in image analysis . . . 13

1.6 Integration of image data and mathematical modelling in oncology . . . 15

1.6.1 Scheduling of data sampling . . . 16

1.7 Contributions to biopsy scheduling pipeline and thesis context . . . 17

1.7.1 Image data processing . . . 18

1.7.2 ROI sampling . . . 18

1.7.3 Spatial analysis . . . 19

1.7.4 ABM . . . 20

1.7.5 Hutchinson and Grimm (2022) . . . 21

1.7.6 Summary . . . 21

1.8 Thesis structure and aims 21

2 Emergent dynamics of a cellular automaton model of tumour-immune interactions 24

2.1 Overview 24
2.2 Model development 25
2.2.1 Motivation 25
2.3 Development of simplified model 26
2.3.1 CA structure 26
2.3.2 Initial conditions 27
2.3.3 Tumour cells 28
2.3.4 CTLs 28
2.3.5 Time stepping 30
2.3.6 Key simplifications from the Hutchinson and Grimm (2022) model . 32
2.4 Methods 36
2.4.1 Summary statistics 36
2.4.2 Spatial analysis - on-lattice pair correlation function 36
2.4.3 Parameter values 39
2.4.4 Sensitivity analysis 39
2.5 Results 40
2.5.1 Typical model output 40
2.5.2 Temporal dynamics 43
2.5.3 Spatial dynamics exhibited by the model 43
2.5.4 Pairwise parameter sweeps 49
2.5.5 Optimal value of CTL migration bias for tumour elimination - a trade-off between directed and random CTL motion 51
2.6 Exploring modelling assumptions - choice of time step 55
2.6.1 Overview 55
2.6.2 Effect of time step choice on optimal bias in T cell migration 56
2.7 Discussion 61

3 Spatial analysis of tumour biopsy histology images 65

3.1 Overview 65
3.2 Simlukafusp 66
3.3 The data set 67
3.4 Methodology 69
3.4.1 Previous analysis of the Simlukafusp data set 69

3.4.2 Inclusion of annotation data . . . 70
3.4.3 Data preparation . . . 70
3.4.4 Length-scale verification . . . 73
3.4.5 Summary statistics . . . 74
3.4.6 Image sampling . . . 74
3.4.7 Analysis of ROIs . . . 76
3.4.8 Smoothing of PCFs . . . 80
3.4.9 Average PCF calculation for a set of ROIs . . . 80
3.4.10 Principal component analysis . . . 81
3.5 Results . . . 82
3.5.1 Cell density analysis . . . 82
3.5.2 Identifying responders . . . 83
3.5.3 Dependence of response on the on-treatment biopsy timepoint . . . 86
3.5.4 Clustering using cell density data . . . 87
3.5.4.1 Individual patient analysis . . . 90
3.5.5 Trends in PCF metrics . . . 92
3.5.5.1 Tumour - CD8 cross-PCF . . . 94
3.5.6 Analysis considering primary tumour indication . . . 96
3.6 Robustness of findings to changes in methodology . . . 103
3.6.1 Image-averaged statistics . . . 105
3.7 ROI representativeness of whole image . . . 106
3.8 Addressing some of the issues with the data suitability for these analyses . 107
3.9 Discussion . . . 110

4 Exploring the robustness of spatial analysis methodology using synthetic and real data 112
4.1 Overview . . . 112
4.2 Methods . . . 114
4.2.1 Protocol . . . 114
4.2.2 The model . . . 115
4.2.3 Labelling synthetic tumour outcomes . . . 116
4.3 2-parameter sweep, varying CD8+ T cell proliferation rate and bias in CD8+ T cell motion . . . 118
4.3.1 Variation in PCF metrics over simulation timecourse . . . 121
4.4 Predicting tumour outcome from a single timepoint . . . 125
4.5 Multiple on-treatment biopsy scheduling . . . 129

4.5.1 Predicting tumour outcome using data from two timepoints 130
4.5.1.1 Predictions using PCF metrics 131
4.6 Varying initial condition 134
4.6.1 Converting biopsy data into lattice format 134
4.6.2 Inferring non-proliferating tumour cell population 135
4.6.3 Comparison of spatial distributions of real and synthetic data 136
4.6.4 The effect of ROI size on the spatial evolution of the ABM 138
4.6.5 Dependence of classifier prediction accuracy on initial condition . . . 142
4.7 Clinical relevance 142
4.8 Discussion 145

5 Discussion **147**
5.1 Research summary 148
5.2 Future directions 150
5.2.1 Extension of modelling approach 150
5.2.2 Development of image analysis 151
5.2.3 Integration of mathematical modelling and image data 151
5.3 Concluding remarks 152

A Initial condition **153**
A.1 Initial cell numbers 153
A.2 Procedure of selecting initial condition 154
A.2.1 Burn-in period 155

B Distance transform map for CTL sensing of tumour cells **157**
B.1 Distance transform map 157
B.2 CTL biased random walk 157

C Parameter scaling **162**
C.1 Notation 162
C.2 Non-spatial stochastic model 162
C.2.1 Proliferation probabilities 163
C.2.2 Death probabilities 164
C.2.3 Killing probabilities 164
C.2.4 Migration probabilities 165

D CD8+ biased motion implementation **167**

E Patient data set **171**

E.1 Primary indications and biopsy location 171

F Co-localisation of cells at sub-cell diameter distances **173**

F.1 Overview 173

F.2 Distributions of nearest neighbour distances 173

F.3 Hypotheses for existence of cell-cell pairs separated by distance less than a cell diameter 174

F.4 Testing H1: Mislabelling of segmentation algorithm 176

F.4.1 Relationship between number of overlaps and cell density 176

F.5 Testing H2: Cells in co-localised pairs have formed a complex 178

F.6 Testing H3: Spatial distribution of cell-cell overlaps 178

F.6.1 Plotting location of overlaps 178

F.6.2 Relationship between cell-cell localisation and tissue morphology . . 180

F.6.3 Location of overlaps relative to distance to tissue edge 180

F.7 Comments on H4 and H5 183

F.8 Nearest neighbour distance threshold 183

G Supplementary results for fine-grain parameter sweep over CD8+ T cell biased motion and proliferation probabilities **185**

Bibliography **185**

Chapter 1

Introduction

1.1 Thesis overview

Cancer is a highly complex disease that is caused by many factors, resulting in mutations to normal, healthy cells and leading to uncontrollable cell growth. Many therapies have been developed to treat cancer using different approaches, including radiotherapy, chemotherapy and immunotherapy. The immune system and cancer are significantly intertwined, with many cancers progressing due to mechanisms of immune evasion or hijacking. Immunotherapies, a relatively recent cancer therapy, aim to repair or boost parts of the immune response which may not be working efficiently in order to eradicate the tumour. The efforts of multiple disciplines are required to attempt to analyse the driving mechanisms of cancer growth and make assessments in a patient-centric treatment approach. Digital pathology and mathematical modelling have become increasingly utilised to analyse the tumour microenvironment and predict driving tumoural factors. In particular, mathematical modelling is moving towards the forefront of the scientific process of understanding cancer, with computational advances allowing for *in silico* experiments to be run to inform directions of experimental work, and reduce the parameter space needed to be explored, thus minimising both the time and costs involved in experimental studies. However, mathematical modelling methods require sufficient data to validate and parametrise models. As technology develops, the simulation of complex spatiotemporal models is becoming easier, requiring spatiotemporal data to compare to. Recent advances in digital pathology have allowed for the generation of rich spatiotemporal data sets of, for example, multiplex histology images, which stain for multiple biomarkers. The availability of these data calls for robust methods of analysing the spatial patterns contained in these images. These spatial analyses are also key in bridging between data and models, for both model initialisation, and comparison of model output with data.

The rest of this chapter provides a review of the key literature covering the aspects underpinning the work contained in this thesis. Firstly, the relevant biological background of immunotherapies and T cell infiltration into tumours is covered. Next, we will review existing mathematical models of tumour-immune interactions, with a particular focus on agent-based models. We then discuss existing techniques for histological data analysis and examples of integration of these data with agent-based models. Finally, we will summarise the findings of the literature review, focussing on the parts of the literature and knowledge which motivate the work contained in this thesis. We provide an overview of the rest of the thesis and outline the questions this work aims to address.

1.2 Review of biological background

1.2.1 T cells and the immune system

The immune system has a significant role in the control of cancer in the body. It acts to identify and eliminate abnormal cells, including cancerous cells, from the body, in a process termed immunosurveillance, and is comprised of two subgroups; the innate and adaptive immune response. The innate response is the first-line of defence against foreign bodies, and is composed of cells such as natural killer (NK) cells and macrophages. The adaptive immune response is more specialised and involves T and B lymphocytes, which have a variety of functions. This more specific immune response is responsible for the generation of antibodies and for acquisition of immunity against pathogens.

T lymphocytes (also known as T cells) have varying functions within the immune response, and the most common types are CD8+ (cytotoxic) T cells and CD4+ T cells. All T cells express TCR (T cell receptor) and CD3 receptors. Whether a T cell expresses CD8 or CD4 receptors dictates which types of antigens the cell can detect through MHC (Major Histocompatibility Complex) binding (Sauls et al., 2018). The CD4+ T cell population is divided into conventional T helper (Th) cells, and regulatory T (Treg) cells (Corthay, 2009). Th cells are involved in the recruitment and activation of other immune cells such as cytotoxic T cells, B cells and macrophages (Zhu and Paul, 2010). Treg cells are an essential part of the immune response, suppressing and regulating immune processes, which in healthy individuals is important to prevent autoimmunity (Vignali et al., 2008). However, this function can have negative implications in cancer, for example if tumour cells express "self" antigens which cause the Tregs to downregulate the immune response.

1.2.2 The cancer-immunity cycle

The cancer-immunity cycle, proposed by Chen and Mellman (2013), summarises these interactions, and is shown in Figure 1.1.

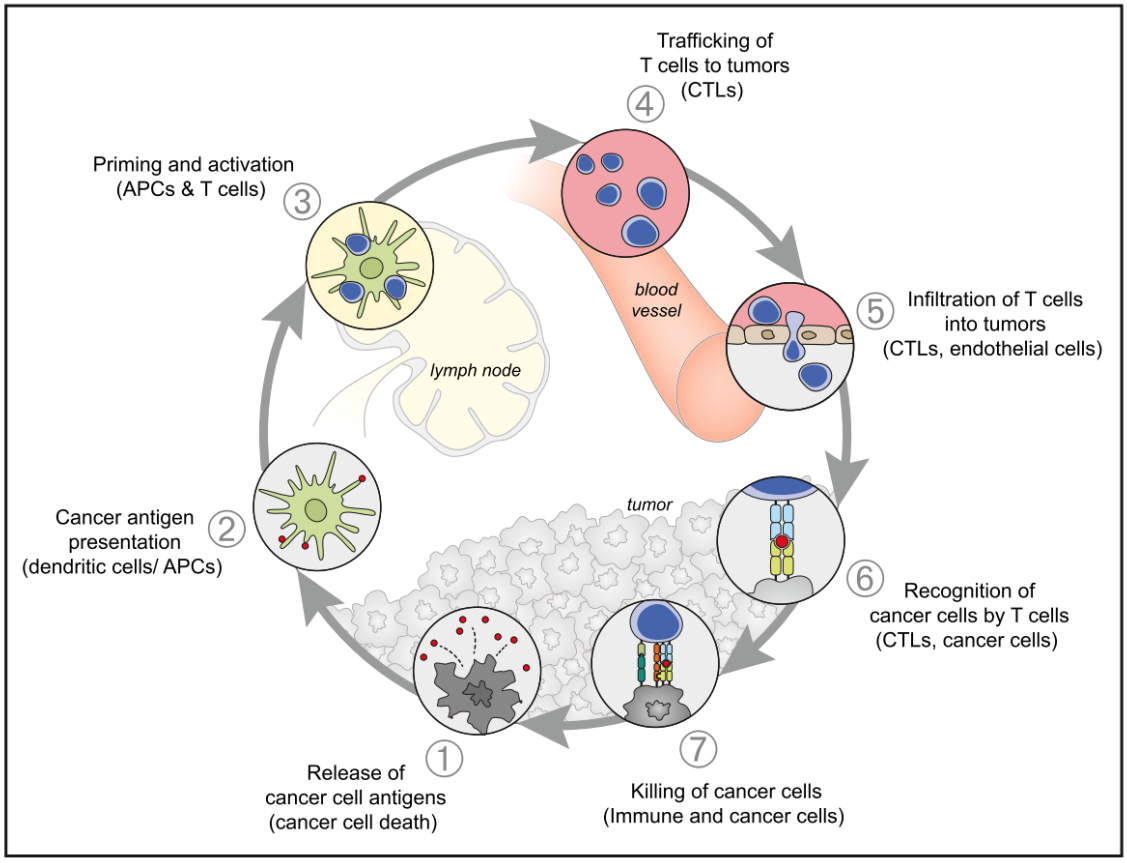


Figure 1.1: The cancer-immunity cycle, reprinted from Chen and Mellman (2013) with permission. APC = Antigen Presenting Cells, CTL = cytotoxic T lymphocyte. See text for details.

The cycle begins with the release of antigens during cancer cell death, which are detected by dendritic cells. Dendritic cells (a type of antigen-presenting cell (APC)) travel to the lymph node, where tumour antigen presentation to CTLs occurs, causing the CTLs to become activated. CTLs are trafficked through the blood vessels, from which they infiltrate the tumour. Upon recognition of the tumour antigen, the killing process is carried out by the CD8+ T cells. This process releases tumour antigens, amplifying the response of the immune system through the cycle. A tumour can escape the immune system response by hijacking one or more of these points in the cycle (Chen and Mellman, 2013).

1.2.2.1 The 3 Es of cancer immunoediting

The tightly regulated interplay between the immune system and tumour cells can result in one of three outcomes, depending on the efficacy of the immune response. The "three Es of cancer immunoediting" refer to the possible outcomes: tumour escape, tumour elimination and tumour-immune equilibrium (Dunn et al., 2004). Tumour elimination occurs when the immune system successfully eradicates the tumour. Equilibrium occurs when this immune-editing response is able to control the tumour, but not strong enough to fully eradicate all tumour cells. Tumours can stay within this equilibrium state for a long period of time, after which tumour escape may occur in which the immune system can no longer contain the tumour. For example, this could result from tumour mutations promoting immune resistance.

1.2.3 Cancer immunotherapies targetting T cells

Immunotherapies target many different aspects of the immune response, and are designed to target specific points in the cancer-immunity cycle which the tumour has evaded, to increase the immune response to the tumour (Chen and Mellman, 2013). There are different types of immunotherapies, including checkpoint inhibitors, bispecific antibodies, immune system modulators, adoptive T cell transfer, and vaccines.

An example of a checkpoint protein is programmed death ligand 1 (PDL1), expressed by healthy cells to prevent autoimmune processes. Anti-PDL1 monoclonal antibodies have been developed to target this checkpoint protein, as PDL1 is also expressed by tumour cells evading the immune system (Weiner et al., 2009). An example of a T cell bispecific antibody is carcinoembryonic antigen TCB (CEA TCB), which binds to CEA-expressing tumour cells and CD-3 expressing T cells (Bacac et al., 2016), increasing T cell infiltration in the tumour and resulting in T cell activation within the tumour. Other immunotherapies, such as immune modulators, act within the tumour microenvironment to boost the immune response. Simlukafusp is an example of such a molecule, which binds to tumour-associated cells and stimulates CD8+ T cell proliferation.

1.3 Mathematical models in immuno-oncology

T cell infiltration has been shown to be a major prognostic factor in cancer patients, and a key indicator as to the outcome of treatments such as immunotherapies. It is therefore critical to the development of immunotherapies to gain a more comprehensive understanding of the intra-tumour distribution of T cells and elucidate the dynamics of these cells, by exploring the mechanisms by which T cells infiltrate the tumour and undergo tumour cell

killing. By understanding the dynamics of T cells within the tumour, the effects of cancer therapies can be better understood and predicted, particularly the effects of immunotherapies, some of which act directly on T cells and perturb their dynamics. In this context, mathematical modelling can play an important role in understanding the biological mechanisms, and provide a method of predicting the response of a tumour with given conditions and application of mono- or combination therapies.

There are many existing mathematical models of tumour-immune interactions, in some cases with the addition of the influence of immunotherapies. Approaches that have been used to model interactions between tumour cells and the immune system can be broadly classified into spatial and spatially well-mixed models. Spatially well-mixed models are typically comprised of systems of ordinary differential equations (ODEs) and assume that spatial effects can be ignored. There are many ODE models of tumour- immune interactions in the literature (de Pillis et al. (2005), Kirschner and Panetta (1998), Kuznetsov et al. (1994), Robertson-Tessi et al. (2012), Wilkie and Hahnfeldt (2013), Wilson and Levy (2012)). ODE models are tractable for mathematical analysis and comparison to experimental data, which often is limited to cell numbers. ODEs can be solved numerically and are computationally inexpensive. However, ODE models may neglect important spatial details. They also neglect stochastic effects, which may be important at low cell numbers. Tumours are spatially heterogeneous structures, and T cell infiltration has been shown to be heterogeneous. Therefore, in many cases more complex and computationally expensive models are required to capture the spatio-temporal dynamics of the system under investigation. Two types of spatial mathematical modelling approaches are systems of partial differential equations (PDEs) (Li et al., 2019, Matzavinos et al., 2004) and agent-based (or individual-based) models (ABM, IBM).

1.4 Agent-based modelling

Mathematical models defined by ODEs and PDEs, such as those discussed in the previous section, summarise the biological system, sometimes with minimal numbers of parameters, and can be analytically tractable. Many of these models are also computationally efficient to simulate, with built-in or customisable solvers. However in each case, some detail is lost or abstracted; in the case of ODEs, we lose spatial information, and in the case of PDEs, modelling behaviours at small population numbers becomes challenging. Complex spatial models which consider the behaviour of individual agents (agent-based models) have become increasingly attainable due to advancements in computing. Although these models remain computationally expensive, packages such as CHASTE (Pitt-Francis et al., 2009), PhysiCell (Ghaffarizadeh et al., 2018) and other software have been developed to optimise

simulation runtime, unit test and provide a universal software platform on which many people can build their models.

There are numerous different types of ABMs, each with a different setup and set of underlying assumptions. These ABMs can be categorised into on-lattice and off-lattice modelling approaches. On-lattice models include cellular automaton models (Alarcón et al. (2003), Kather et al. (2017), Strobl et al. (2022)) and cellular Potts models (Sottoriva et al. (2011), Szabó and Merks (2013)), where the lattice structure is typically composed of squares, although other regular tessellating shapes may be used, such as a regular hexagonal lattice. Cellular automaton (CA) models consider agents to occupy discrete spaces on a lattice structure, and are relatively simple to implement compared to other ABM types. In cellular Potts models, cells can occupy many lattice spaces, such that different cell sizes and shapes can be more realistically modelled, compared to simpler CA models. In these lattice-based models, cells typically interact within a specified neighbourhood of lattice sites. Off-lattice models account for cell size, in addition to forces between cells, and can be considered to be more biologically realistic, as cells are not confined to occupying lattice sites. Off-lattice approaches include overlapping spheres models, where the positions of cell centres are tracked, and cells are modelled as being connected by springs, allowing for mechanical forces, such as cell deformation and exclusion, to be taken into account. Cells interact with other cells within a specified radius from the respective cell centres. Hybrid ABMs combine continuous and discrete approaches (Rejniak and Anderson (2011)), typically using PDEs to model quantities such as chemoattractant and drug concentration in the domain. For example, Mallet and De Pillis (2006) modelled tumour growth supported by a neighbouring blood vessel, in which a PDE was used to model nutrient concentration, and cells were modelled using a CA approach.

1.4.1 ABM implementation

In addition to different types of ABM, there exist several frameworks with which ABMs can be implemented. Examples of these frameworks include CHASTE (Mirams et al. (2013), Pitt-Francis et al. (2009)), PhysiCell (Ghaffarizadeh et al., 2018), CompuCell3D (Swat et al., 2012) and HAL (Bravo et al., 2020). Each framework uses different programming languages and modular structures, so these are important considerations to take into account when formulating the structure and design of an ABM. For simpler ABMs, it may be appropriate to produce bespoke software which allows the author to have complete control over the rules and processes modelled. Next, a choice as to the spatial dimensionality of an ABM needs to be made. Even ABMs modelling 2D space are computationally expensive, so this is one of the main limiting factors when designing algorithms for 3D ABMs. However,

the systems being modelled are three-dimensional, so it could be argued that 3D models allow for more biologically realistic dynamics and model predictions, although it may be challenging to analyse the results they generate. Many ABM models consider 2 spatial dimensions (Christophe et al., 2015, Kather et al., 2017, Macfarlane et al., 2018, 2019, Mallet and De Pillis, 2006), with some 3D tumour spheroid models in the literature (Gong et al., 2017, Kather et al., 2018). The justification for dimensionality choice of an ABM can also depend on the available data with which to calibrate and validate the model. Most data are in the form of scalar, time course data (for example temporal blood counts) or in two spatial dimensions (e.g. histology data), therefore deducing model parameters for 3 spatial dimensions can be challenging. There are some 3D spatial data in the form of experimental tumour spheroid models, which can be compared to in silico tumour spheroids generated by an ABM.

Other considerations in the ABM implementation include the choice of timescale for the simulation, and the time-stepping procedure adopted. There is not a clear procedure for identifying an appropriate time-step size, which is usually a stated choice by the authors of the model, depending on the processes being considered; for example a smaller timestep may be more appropriate for biological processes such as diffusion or migration, whereas a larger timestep may be more appropriate for modelling other processes, such as cell death or proliferation. In the next section, we review existing ABMs of tumour-immune interactions, in the context of the considerations outlined above.

1.4.2 Agent-based modelling of tumour-immune interactions

Examples of ABMs modelling tumour-immune interactions are varied in their focus, biological detail and implementation (Norton et al., 2019), and in this section we will review some of the key models relevant to the work in this thesis. We firstly summarise the main findings of these studies relating to the biology of the model, and then thematically review the aspects of ABM implementation, for example the timestepping and model setup in each case. Many of the works reviewed in this section successively build on previously published ABMs, therefore the structure of this review will go through related publications (those which adopt similar frameworks or model rules) in chronological order.

Ferreira Jr et al. (2002) developed a hybrid CA-PDE model of a growing avascular tumour. Tumour cells were grown on a lattice, with cells permitted to pile up within lattice sites. The effects of the immune system were beyond the scope of this work. They found that various tumour morphologies are produced depending on the nutrient supply and tumour-specific parameters of the model. Mallet and De Pillis (2006) extended the work of Ferreira Jr et al. (2002) by modelling the immune response to a tumour. The "immune"

population is comprised of NK cells and cytotoxic T cells (CTLs), which interact with a growing avascular tumour on a 2D square lattice, with periodic boundary conditions for two opposing sides of the domain, and a source of two nutrients along the other two sides. Boundary conditions apply to both the cell populations and the nutrients. The nutrients control the growth and survival of the cells. Mallet and de Pillis found temporal oscillations in the tumour and immune populations for most immune system parameter choices, and ultimately CTL recruitment and killing parameters were remarked to be significant in tumour removal. Of note, they remark that the CA rules implemented in their model are not "definitive choices" and merit further investigation as to the robustness of the model to the functional forms of the assumed probabilities for cell actions in the model, and thus the appropriateness of the modelling approach. We will refer back to these remarks as part of the motivation of the work presented in Chapter 2 of this thesis.

López et al. (2017*a*) build on the work of Mallet and de Pillis, developing a model which focussed on investigating the lysis of tumour cells by CTLs. Tumour growth (without the immune system) was simulated, with tumour cells permitted to pile up within lattice sites. The 2D projection of this tumour (i.e. one cell per lattice site) was then surrounded by CTLs and the simulation run forward for 4 hours to observe how tumours with different morphologies are lysed. The authors found in their model that tumour killing was limited by crowding of cells at the tumour perimeter, and highlighted how the tumour is killed "layer-by-layer". López et al. (2017*b*) extend the model from López et al. (2017*a*) to explore the long-time dynamics of the system. The model is simulated on the order of 1000s of days, and produces behaviour representing the three Es of immunoediting. Again, the model adopts a similar framework to that of Mallet and De Pillis (2006), but the authors note that the relative timescales of CTL killing of tumour cells and tumour cell proliferation are taken into consideration, by separating the tumour and CTL updates, i.e. updating all tumour cells on a large timestep, and then updating the CTL cells on smaller, intermediate timesteps. Lopez et al. found long periods of tumour dormancy, which was unstable and depended on a balance between the tumour cell proliferation rate and the immune response.

Christophe et al. (2015) developed an on-lattice ABM of CTL interactions with melanoma tumour nodules. Within this model, they examined two CTL displacement models, considering a symmetric random walk, and a self-interacting reinforced random walk, where CTL migration is biased towards "scout" CTLs which have hit the tumour nodule boundary. CTLs are subject to progressive exhaustion as they kill, and tumour cells are stochastically able to acquire immune resistance as they proliferate. This resistance takes the form of requiring a longer killing time by the CTLs. They found that the CTL population size (which was kept constant throughout the simulation), had a great effect on the initial collision

frequency with tumour cells, and thus on the efficacy of tumour removal. The time for the tumour cell killing process by CTLs (and thus the corresponding killing rate) also affected the efficacy of tumour removal - increasing the duration of the killing time from its default (measured) value intuitively resulted in a longer time for tumour eradication, however decreasing the killing time did not significantly speed up tumour eradication. To explain this, Christophe et al. (2015) refer to the difference in timescales between the tumour killing process (on the order of minutes) and tumour proliferation (on the order of hours). Thus, if the tumour proliferation rate was much faster, then a decrease in the CTL killing time may have a larger effect as these processes would be happening on similar timescales. The CTL efficacy was found to be further boosted by incorporating CTL chemotactic movement towards scout CTLs. Christophe et al. (2015) concluded that increasing the exhaustion limit and increasing chemotaxis work in synergy to eradicate the tumour.

Macfarlane et al. (2018) studied the movement of activated and unactivated cytotoxic T lymphocytes (CTLs) on a 2D spatial grid, and interactions with tumour cells and dendritic cells (DCs). A similar model setup to Christophe et al. (2015) was used, with some parameter values adopted from this earlier work, however the focus of Macfarlane et al. (2018) and the mechanism of CTL movement used was different. Macfarlane et al. (2018) produced qualitatively similar CTL migration tracks to those found in previous experimental work (Boissonnas et al., 2007), which had observed a change in motion of CTLs upon exposure to tumour antigen. Brownian motion was assumed for antigen-activated CTLs, and inactive CTLs moved according to a Lévy walk, reproducing experimental cell tracking observations (Harris et al., 2012). DCs are activated by interacting with tumour cells. Activated DCs can subsequently activate inactive CTLs. Tumour cells can be killed by activated CTLs, and these CTLs will continue to search for tumour cells, i.e. deactivation and exhaustion is not modelled. Initial conditions of the model are a circle of densely packed tumour cells (400 cells) in the centre of the domain and inactive CTLs and DCs are randomly placed within the domain. Macfarlane et al. (2018) found that increasing the number of both CTLs and DCs causes crowding effects and reduces the access to central tumour cells, and thus results in reduced tumour cell killing. Increasing just the number of CTLs resulted in faster tumour elimination, whereas increasing the number of DCs was found to cause crowding effects and lead to reduced tumour killing. It was also found that increasing the activation rates of CTLs and DCs had little effect on tumour clearance.

Macfarlane et al. (2019) extended their model from Macfarlane et al. (2018) to incorporate different tumour antigen expression and specific recognition of these antigens by the CTLs. Tumour cells are modelled to change their antigen profile through epimutations (at any time) or mutations (during mitosis). Each CTL can only be activated against one of the

tumour antigens. They found that the tumour outcome was highly dependent on the initial tumour antigen distribution, with homogeneous tumours eradicated compared to initially very heterogeneous tumours, which the authors state points to the merit of having a wide repertoire of CTL specificity.

Kather et al. (2017) developed a 2D CA model of a growing avascular colorectal tumour, considering infiltrating immune cells from the microenvironment. They incorporated immunotherapy implicitly in their model, referred to as an "immune boost", adjusting the immune cell influx at the time of treatment. Tumours with different tumour-immune phenotypes were generated by varying the model parameters, and were visually compared to histology data. Kather et al. (2018) subsequently extended this model to 3D, including more biological detail, for example naive and pro-tumour macrophages, and clonal evolution of the tumour caused by immunogenic mutations. Treatment combinations of chemotherapy and immunotherapy were also modelled by adjusting the model parameters corresponding to the drug mechanism of action. Model parameters were estimated by fitting model output (in particular, ratios of cell population numbers) to patient data. Spatial information of the cell distributions was not considered quantitatively, however a qualitative comparison was made between the cross-section of the synthetic tumour and histology data.

The 2D model developed by Kather et al. (2017) was further extended by Hutchinson and Grimm (2022) to form part of a pipeline to inform on-treatment biopsy scheduling using digital pathology and mathematical modelling. The model was extended to incorporate multiple-hit killing, whereby tumour cells require multiple attacks from cytotoxic T cells to be killed. The boundary conditions and domain were also adapted to model a region of the tumour, rather than the whole tumour as in Kather et al. (2017). This model is the motivation for the modelling work in Chapter 2 of this thesis, and will be discussed in more detail there.

Gong et al. (2017) developed a 3D multiscale model of tumour-immune interactions, focussing on the effect of anti-PDL1 immunotherapy. At the tissue scale, agents are modelled on-lattice, and at the molecular scale, PDEs are used to model cytokine processes. Tumour cells are modelled as two populations with different PDL1 expression levels, and T cells are modelled as effector, cytotoxic or suppressed. Immune evasion is modelled by PDL1+ tumour cells deactivating T cells, which can only kill PDL1- tumour cells. The parameters controlling the deactivation of T cells by PDL1+ tumour cells were adjusted to model the anti-PDL1 therapy. IL-2 cytokines are included in the model as they are secreted by activated T cells, and promote T cell proliferation.

Ghaffarizadeh et al. (2018) studied tumour growth under an immune response using a model developed in PhysiCell (physics-based multicellular simulator), a framework in

which agents are modelled off-lattice with force-governed dynamics (Ozik et al., 2019). Ghaffarizadeh et al. (2018) identified immune migration as an important factor in the degree of mixing of tumour and immune agents, and the resulting effectiveness of the immune response.

1.4.3 Implementation of on-lattice ABMs

In this thesis, we will focus on the development of on-lattice ABMs of tumour-T cell interactions, for several reasons. Firstly, our collaborators at Roche have developed an on-lattice ABM which we use as a motivation for the first research chapter. Secondly, we reason that as a starting point for modelling a new system (in this case, a new treatment), a simple model which can be clearly understood and broken down into its constituent parts is ideal. Complexity and detail can be added into a model once the fundamentals are established. Finally, these are arguably the most computationally efficient ABMs, and generally many simulations are required in order to gain a deep understanding of the emergent dynamics of an ABM.

As previously noted, ABM implementations can greatly differ in the literature, and as there are many aspects of the setup of an ABM to consider, we will now consider these thematically and review the previously discussed studies in these contexts.

1.4.3.1 Domain size and boundary conditions

In most studies concerning tumour-immune interactions, a solid growing avascular tumour is modelled within empty space or host tissue. In some select models, a region of tumour is modelled, such that the whole domain represents the tumour microenvironment (Hutchinson and Grimm, 2022, Lai et al., 2019). This modelling choice is of importance as it impacts on the choice of boundary condition for the model. In models of whole tumours, the tumour is not typically grown to the domain boundary, so boundary conditions are less significant, only affecting the dynamics of other populations, such as immune cells, which may move through the domain boundary periodically (Alarcón et al., 2003, López et al., 2017*b*, Mallet and De Pillis, 2006) or reflect off the boundary (Macfarlane et al., 2018). The boundary conditions may have secondary effects on the tumour dynamics through interactions with these other cell populations. There may be additional boundary conditions at the tumour interface in ABMs, modelling surface tension and cell adhesion. These boundary conditions are usually present in models which explicitly model forces, and tend to be incorporated into off-lattice models.

1.4.3.2 Neighbourhood choice

In on-lattice models, agents typically interact within a defined "neighbourhood", usually chosen as the von Neumann neighbourhood (direct four neighbours, up/down and left/right) (Christophe et al., 2015) or the Moore neighbourhood (direct eight neighbours, the von Neumann neighbourhood and the four diagonal neighbours) (Hutchinson and Grimm, 2022, Kather et al., 2017, Macfarlane et al., 2018). Tzedakis et al. (2015) explore the effect of neighbourhood choice on tumour growth in a 2D hybrid on-lattice model. Tumour cells undergo oxygen-dependent growth and the authors investigate chemotaxis and haptotaxis separately, and the effect of neighbourhood choice under each process of cell migration. They found the tumour attained a more circular morphology in the Moore neighbourhood scheme in comparison to the von Neumann scheme, and, as the number of possible sites to move/proliferate into doubles between von Neumann and Moore, Tzedakis et al. (2015) found significant differences in the migration of cells under single cell occupancy rules. They concluded that a Moore neighbourhood better approximates circular tumour morphology compared to using a von Neumann neighbourhood.

1.4.3.3 Grid occupancy

On-lattice models restrict cells to occupy discrete sites within a domain. These models typically restrict lattice sites to one cell per site to model spatial exclusion. This setup assumes that all cells occupy the same amount of space, which in the case of a single cell population is a reasonable assumption. Recently proliferating or dying cells may not take up the same amount of space in reality, however this is a simplifying assumption of such models.

1.4.3.4 Timestep

As previously noted, there is not a unified procedure for ascertaining a suitable timestep with which to simulate an ABM, and this choice is widely model-specific, depending on the timescales of the biological processes being modelled. Numerical convergence analysis can be used to determine the largest timestep at which the output converges. However, this convergence analysis would be for a specific set of parameter values, and typically the parameter space of an ABM is of large dimensions. Typically in ABMs with fixed timesteps, the timescale of the shortest process is adopted as the timestep size, for example Macfarlane et al. (2018) simulate their model using timesteps of 1.44 min based on the time for a CTL to migrate one lattice space. Other models separate the timescales of biological processes, by incorporating smaller, intermediate timesteps on which faster processes occur, and larger timesteps for slower processes. For example, Kather et al. (2017) adopt a 12 hour timestep,

within which all tumour cells act, followed by the immune cell actions. Immune migration occurs on smaller timesteps to proliferation and death. López et al. (2017*b*) use a similar approach of simulating the tumour cell actions on a different timescale to the immune cell actions, maintaining an ordering of tumour, then immune, actions in their model. Another approach to the timestep choice issue is to use a scheme such as the Gillespie algorithm, which uses a variable timestep with which to simulate cell actions.

1.5 Spatial statistics

Spatial statistics are widely used in many different fields, for example, ecology, to describe the spatial distribution of species and population dynamics. Spatial statistics have previously been used to characterise the spatial information contained within biopsy images, particularly the spatial patterns of T cell infiltration. Quantification of the spatial distribution of cells in the biopsy tissue also enables direct comparison with simulation results from an ABM (Agnew et al., 2014, Dini et al., 2018, Fozard et al., 2011, Gavagnin et al., 2018). In this thesis, we focus on using the pair correlation function (PCF) to characterise both histology data and results from our agent-based modelling.

1.5.1 Pair correlation function (PCF)

The Pair Correlation Function (PCF) is a spatial statistic that identifies length scales on which clustering (and dispersal) occur within a point pattern (Chiu et al., 2013). The PCF, $g(r)$ indicates the length scale on which clustering occurs within a point pattern, compared to the case where the points are randomly distributed. The PCF is calculated by taking each point in turn, calculating the number of points around the point of interest in concentric annuli of increasing radius, r, and normalising by the area of these disks (see Figure 1.2). The distribution is then averaged over the total number of points in the region of interest, and divided by the average cell density.

If $g(r) > 1$, then there are more cells separated by distance r than would be expected if the cells were randomly distributed. If $g(r) < 1$, then there are fewer cells separated by distance r than in a random distribution.

1.5.2 Spatial statistics in image analysis

Histology images span many types of samples, from whole tumour sections, in the case of surgical tumour treatment, to smaller biopsy samples, collected through procedures such as core needle biopsies. Analysis of histology images firstly involves methods to locate cells within an image and classify them. In this thesis, we concentrate on the methods used to

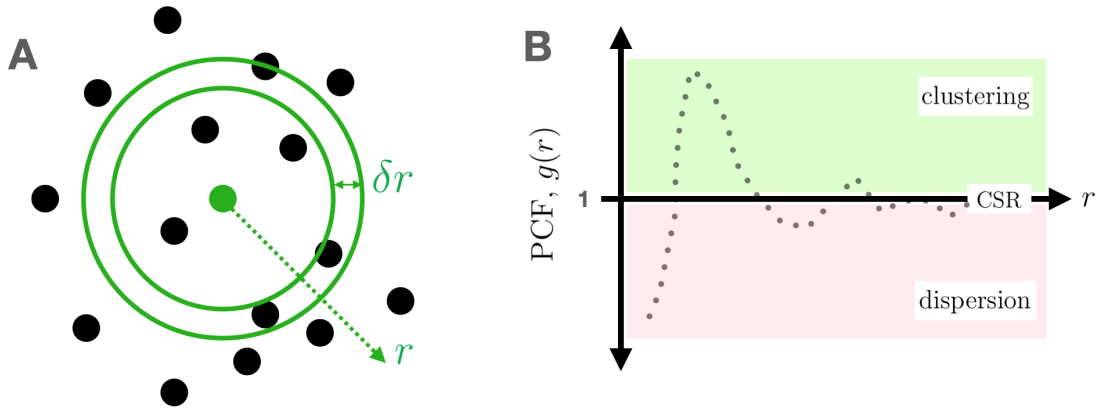


Figure 1.2: (A) Schematic illustrating how the PCF is calculated. An annulus of radius r and thickness δr is considered about each point (in this example, the point shown in green is considered). This annulus moves out up to radius r_{max}, indicated by the green dotted arrow. The density of points in each annulus is recorded and normalised by the average point density over the region. Note that points are shown zoomed in for clarity of the schematic; as they are cell centre coordinates, points would either be inside or outside the annulus. (B) The PCF, $g(r)$, takes positive values, with $g(r) = 1$ corresponding to a point cloud with a comparable spatial distribution to Complete Spatial Randomness (CSR). Clustering of points at radius r is indicated by $g(r) > 1$ (shaded in green) and dispersal of points at radius r is indicated by $g(r) < 1$ (shaded in pink). An example PCF is illustrated by the dotted line.

analyse the point cloud data generated by the application of such algorithms, so we do not cover this area of image analysis in this review.

The PCF has been applied in many studies, from quantifying immune infiltration in tumours (Bull et al., 2020) to detecting necrotic regions of tumour spheroids (Dini et al., 2016). Carstens et al. (2017) analysed tumour tissue cores from patients with pancreatic ductal adenocarcinoma to investigate the immune composition of the disease. Multiplex imaging was used to mark eight T cell sub-populations, and the percentage composition of the total T cell population was calculated. Patient survival was found to correlate with total T cell infiltration, and, in particular, cytotoxic CD8+ T cells and CD4 effector (helper) T cells were correlated with an increase in survival. They then conducted a spatial analysis of the T cell sub-populations relative to tumour cells, focussing on the area within a $20\mu m$ radius from the tumour cells, as this was noted as a region of "enhanced cell-cell contact". The L-function was computed (a spatial statistic which measures point clustering), from which they found that only the presence of CD8+ T cells within this radius of tumour cells showed a significant correlation with prolonged patient survival.

Gong et al. (2019) investigated intra-tumoral clustering of CD8+ T cells in whole slide images, using spatial point pattern statistics such as Ripley's K function. They used a sliding

window in which to calculate the spatial statistics, and examine the spatial heterogeneity within a given sample. They also used a clustering algorithm which used the CD8+ T cell density data to identify clusters. The clusters were then enclosed by an alpha-shape, to define the cluster morphology. They compared the calculated spatial metrics between responders and non-responders to an immunotherapy, and found increased numbers of high density CD8+ T cell clusters in responders. Response was not found to correlate with average CD8+ T cell density within the tissue.

There are several different spatial statistics which we could employ for our analysis in this thesis. We choose to use the PCF in this work, as it is one of the most commonly used spatial metrics, and is straight forward to explain and understand. The latter point is crucial if we are trying to develop a methodology to be used within a clinical setting; explainability of the models and metrics used is very important. Furthermore, an analogue for the PCF for application to on-lattice models is available in the literature (Gavagnin et al., 2018) and the simplicity of the metric allows for the versatility of application to off-lattice and on-lattice data. Other spatial statistics to analyse point patterns could have been chosen (such as the L function), in addition to topological data analysis (TDA) methods.

1.6 Integration of image data and mathematical modelling in oncology

Up to this point we have reviewed the key literature pertaining to mathematical modelling and histology analysis in immuno-oncology. We now shift focus to literature in which these two concepts have been linked or integrated. Validating ABMs with data is challenging, as they often have many parameters whose values may not be known or may be difficult to measure. Furthermore, for ABMs which consider spatial interactions, validation methods should include a measure of the spatial characteristics of the model output. In the last section, it was discussed how Kather et al. (2017, 2018) quantitatively compared model output to data using cell counts, but only qualitatively validated the spatial structure of the tumours modelled. One of the aims of our work is to develop such validation methods, which are important when developing an ABM (Cosgrove et al., 2015).

Approximate Bayesian Computation (ABC) has been used to infer parameters of mathematical models, in which summary statistics from model and data are compared. The distance between the data and the simulated model output is calculated, and the model parameter set is accepted if this distance is below a specified threshold (Lambert et al. (2018), Sunnåker et al. (2013)). Johnston et al. (2014) used ABC to estimate the parameters (cell diffusivity and cell proliferation rate) of an on-lattice model of collective cell spreading using data from scratch assays. The PCF and cell number were considered in

the summary statistic. Browning et al. (2018) also used ABC to estimate the parameters of an off-lattice model of this system. These methods of parameter inference typically require data at multiple timepoints, and the summary statistics must be found to adequately distinguish different behaviours.

1.6.1 Scheduling of data sampling

One of the motivations of this thesis is the question of when the on-treatment biopsy should be scheduled in an immunotherapy clinical trial, and whether mathematical modelling could help in personalising this decision for each patient and aid in clinical decision making. Hutchinson and Grimm (2022) aim to answer this question by optimising the parameters of their ABM (initialised using pre-treatment images, and developed from the model by Kather et al. (2017)) for each patient using the corresponding on-treatment images. Only one on-treatment image was available per patient, however each image was sampled by several regions of interest (ROIs). The distributions of the CD8-CD8 PCFs calculated for the simulated data and real on-treatment ROI data were compared and 12 parameter sets identified as being the best at replicating the spatial distribution of CD8+ T cells at the on-treatment timepoint, with a reported mean accuracy of 77%. However, the authors note the limited availability of data, namely the small biopsy samples. This study also highlights the question of how useful a quantitative approach can be in the early stages of a clinical trial, where sufficient on-treatment data are unavailable to be able to optimise the model parameters and use the model predictions to inform decision making. Furthermore, there is a question in particular as to how informative a single on-treatment timepoint can be for optimising the high-dimensional parameter space of an ABM. As found in a synthetic study by Bull and Byrne (2023), tumours may exhibit a wide range of spatial cell distributions, indicating different immune responses, at different timepoints before reaching a steady state. Furthermore, tumours are known to be highly heterogeneous, so there is the added complication of how representative a single biopsy is of the tumour at any given time.

Lai et al. (2019) present a similar pipeline to Hutchinson and Grimm (2022), integrating MRI and histology data (collected using needle biopsies) with an on-lattice hybrid 2D CA model, to investigate chemotherapy scheduling. Their hybrid ABM models cells and blood vessels as discrete agents, and uses PDEs to model oxygen and drug concentrations. The model is validated using tumour volumes and apparent diffusion coefficients from MRI imaging.

1.7 Contributions to biopsy scheduling pipeline and thesis context

The work contained in this thesis builds on work carried out during the first year of my DPhil, in collaboration with Lucy Hutchinson (LH) and Oliver Grimm (OG) at Roche, based on ideas that they had previously developed. After that time, LH and OG continued to develop the project without my involvement, publishing their findings in Hutchinson and Grimm (2022). In this paper, biopsy data from a clinical trial of a novel immunotherapeutic is combined with an agent-based model (ABM) to generate a simulated time course of biopsies to support clinicians in biopsy scheduling decisions. Prior to my involvement, all aspects of the pipeline already existed, except for the sensitivity analysis and parameter optimisation steps (see Figure 1.3 for summary of the pipeline when I joined the project).

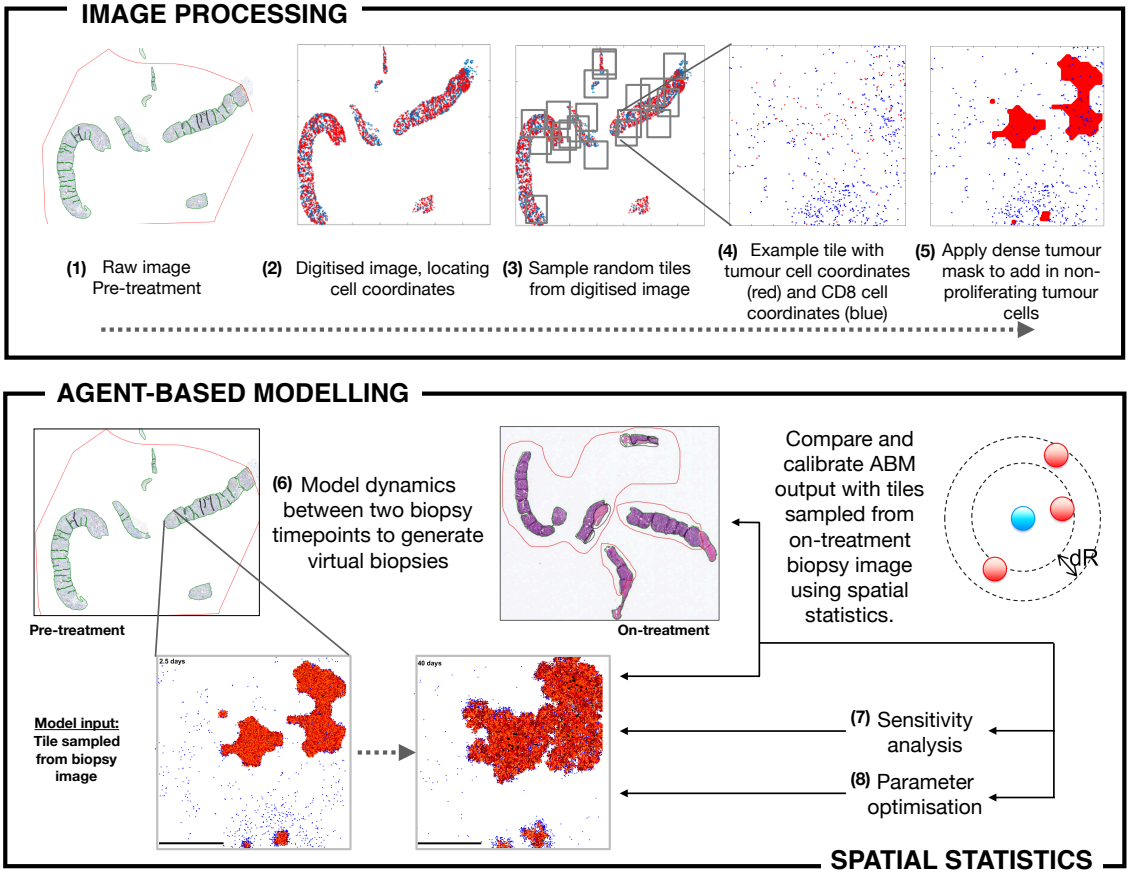


Figure 1.3: Visual summary of pipeline at point of me joining the project. I put the figure together, but the inset figures were produced by LH and OG.

My contribution to this project involved further developing the methodologies underpinning the pipeline, spanning methods for the image processing, image sampling, spatial analysis and the agent-based model. My contributions to each of these aspects of the

pipeline are summarised as follows. I use "I" to refer to parts of the pipeline which I personally developed and wrote code for, or decisions which I made. During this year of being on the project, myself, LH and OG had frequent collaborative meetings to discuss my findings and directions going forward; i.e. none of the developments were carried out or accepted into the pipeline without discussion with LH and OG. Points marked with (*) in the sections below were decisions and ideas which were retained within the pipeline after my involvement in the project, and were further developed by LH and OG.

1.7.1 Image data processing

- ***Automated and sped up pre-processing of images using pathologist annotations to replace manual deletion of control tissue** - OG provided pathologist annotations for the images and we all agreed that these were important to incorporate into the pipeline. I wrote the code to incorporate these annotation data into the image processing portion of the pipeline and wrote functions to automatically trim each image using the tumour annotation, such that just the patient tissue was remaining in the image and the control tissue was removed. This sped up the pipeline as user input was not required to trim the image.

- ***Organised image data into data structure which automatically separated data from tissue fragments within an image** - Tissue and tumour annotation overlaps were detected and stored separately in a data structure which I designed and implemented to store data (including cell coordinates) for each tissue fragment separately. I used the exclude annotations to intersect with the tumour-tissue intersection regions to reflect the updated tissue morphology and border, and deleted cells within the exclude regions, as these may have been regions of artefacts, tears or folds in the tissue and may affect the spatial analysis. LH and OG state that final algorithms were reworked after my involvement in the project to optimise them for use in new samples.

1.7.2 ROI sampling

- ***Reduced ROI size from 400 to 100 cell widths** - I found that a large fraction of empty space was being sampled and the biopsy tissue was not being uniformly sampled. I reduced the ROI dimension to 100 cell widths such that tiles could fit within most of the biopsy tissue pieces in the images across the patient cohort.

- ***Only sampled ROIs within a single tissue fragment** - As in some of the images samples were broken up into several pieces of tissue, I changed the protocol to

not allow sampling of ROIs which contained areas from multiple pieces of tissue, as their positions on the slide are arbitrary, and so tissue pieces should be sampled from independently.

- **Implemented new ROI sampling algorithm to maximise number of sampled ROIs without replacement** - sampled ROIs from the image without replacement, to ensure that one area of the image was not over-sampled. After we discussed possible approaches, I implemented an ROI sampling algorithm which scans through all possible square ROIs (dimension 100 cell widths) within the image at its default orientation, and records the coordinates of ROIs that satisfy the modified conditions detailed in the previous paragraph. I use this ROI sampling algorithm later in the thesis, in Chapter 3. The minimum cell threshold for ROI acceptance was set to a 1% empty space tolerance. This methodology was reworked iteratively by OG and LH after my involvement to allow for ROI overlap and a larger empty space tolerance to optimise the algorithms for new samples in the dataset.

1.7.3 Spatial analysis

- ***Replaced spatial analysis method with on-lattice PCF given by Gavagnin et al. (2018)** - The spatial analysis method was changed after I concluded from the spatial analysis of the images that it was sensitive to the tissue morphology and edge effects and was not normalised to reflect the lattice data. A parameter of a bi-exponential fit of the previous spatial statistic was very sensitive to empty space in the ROI, which I found through analysis of the ABM simulations. LH and OG performed further analysis after my involvement in the project on how best to use the published method in Gavagnin et al. (2018), which led to the development of the method in Hutchinson and Grimm (2022) to compare distributions of RDFs from different samples.

- ***Decided to use CD8 PCF instead of CD8-TU PCF** - I identified the CD8 to CD8 T cell PCF as the most reliable statistic for characterising the images and comparing to the ABM data due to the missing non-proliferating tumour cell data, by calculating the PCFs and cross PCFs for different tumour densities, and found variation in the resulting CD8-TU PCFs, which was discussed with LH and OG. As we sample ROIs only from the tumour area, this gives information on the level of CD8+ T cell infiltration into the tumour, and the spatial distribution of CD8+ T cells within the tumour.

1.7.4 ABM

- **Investigated burn-in period of ABM with changing initial tumour cell density to inform value chosen for this density** - I found after an initial burn-in period the dynamics converged at timepoints greater than approximately 20 days. These results showed the extent of the burn-in period for each density which could inform the decision of what initial tumour density to adopt.

- ***Model simplification** - I removed necrosis and fibrosis from the model as we did not know the oxygen distribution within the sample or have data regarding fibrosis.

- ***Changed T cell biased random walk implementation** - I changed the distance metric used in the mapping function to ensure consistency with the Moore neighbourhood used in the ABM, and the PCF on-lattice normalisation used.

- ***Changed boundary conditions** - As the cells would now be in contact with all domain edges in the updated ROI sampling regime, I changed the no flux boundary conditions to implement a doubly periodic boundary condition to the domain and the underlying maps to reflect our assumption that the areas outside of the ROI had the same spatial cell distribution as within the ROI.

- **Sensitivity analysis (SA)** - I ran pairwise parameter sweeps of parameters governing CD8 cell behaviour (note that this SA is independent and separate to that presented in Hutchinson and Grimm (2022)). One of my main findings was that increasing CD8 proliferation results in a decrease in CD8 cells, and was also found independently by LH and OG in their paper. I use this SA approach later in the thesis, in Chapter 2, where we simulate the effect of hypothetical immunotherapeutics by conducting these pairwise parameter sweeps.

- **Consistency analysis using exponential fit parameters of PCF** - Analysis using Spartan package in R to check robustness of number of cells and PCF parameters to the number of stochastic repeats of the ABM simulation.

The above bullet points outline my contributions to the code which were implemented into the pipeline at the point of my involvement in the project. LH and OG refined and reworked the code to further develop the ABM after this point.

1.7.5 Hutchinson and Grimm (2022)

The methods I outlined above were coded by me and were implemented into the pipeline at the stage at which I left the project, and LH states that LH and OG since developed and reworked the code after this point to incorporate further decisions made by them regarding the workflow.

I had no involvement in the methodologies except for the development of the image processing and ROI sampling, choice and implementation of spatial statistic, and adaptations to the ABM. I was not involved in developing the method of comparing PCFs between image and ABM (the spatial agreement metric), and I did not conduct the sensitivity analysis or parameter optimisation presented in the paper.

1.7.6 Summary

The work contained in this thesis was motivated by questions which arose during the development of the pipeline described above in collaboration with LH and OG, and uses some of the methods I developed and implemented while undertaking this preliminary work. All work presented in the rest of this thesis is my own, unless clearly stated in the text.

1.8 Thesis structure and aims

In this chapter, we have reviewed the literature covering the themes underpinning the work presented in this thesis. In this thesis, we aim build on the work to date in the literature, by presenting three studies centred around analysis and modelling of tumour biopsy samples. This thesis combines agent-based modelling with image data, using spatial analysis, with the aim to explore how mathematical modelling and image analysis techniques can be used in oncology clinical trials to gain insight into the mechanism of action of immunotherapeutics and extend the analysis and insights that can be obtained. This work has been motivated by biopsy data from a clinical trial of FAP-IL2v, an immunotherapy which was under development at Roche. During the first year of my DPhil I worked on a project with my collaborators at Roche aiming to optimise the scheduling of on-treatment biopsies in oncology clinical trials. The work presented in this thesis stems from questions that arose during this preliminary work. We now summarise the content and aims of each chapter as follows.

In Chapter 2, we start by simplifying the Hutchinson and Grimm (2022) model, to investigate the driving factors behind oscillatory dynamics seen in their model. We aim to understand if these effects are biologically driven or an artefact of the modelling assumptions. We then simulate the effects of hypothetical immunotherapies by varying the CD8+

T cell properties in a pairwise manner, extending the local sensitivity analysis in Hutchinson and Grimm (2022), to find combinations of CD8+ T cell properties which minimise tumour elimination time for a given tumour type. In our survey of the literature of ABM models in tumour-immune interactions, the existence of processes occurring on different timescales was a common point of discussion and motivation for different implementations of the models. In Chapter 2, we investigate how these results dependent on the underlying modelling assumptions and time-scaling of biological processes.

In Chapter 3, we aim to investigate whether there are any differences between the pre- and on-treatment biopsy data, and to see if the drug affects the spatial structure of the tumour. Biopsy data are typically assessed by a pathologist, examining the cell counts and the level of T cell infiltration within the tumour, and the presence of stroma. One of the aims of this chapter is to assess what further information can be gained by analysing the data using spatial statistics, and to establish which spatial biomarkers are prognostic for patient outcome. Alongside the biologically motivated questions, we address methodological questions in this chapter. These questions concern the parameters used to calculate the pair correlation function (i.e. the annulus width, and minimum number of cells with which to calculate the PCF), and how to sample the data.

In Chapter 4, we use the model developed in Chapter 2 with the spatial analysis methodology from Chapter 3 to investigate whether mathematical modelling can aid in the scheduling of on-treatment biopsies, and if spatial analysis can separate the synthetic tumour outcomes sooner than simply analysing cell counts. We explore the robustness of the spatial analysis to changes in model parameters. We then integrate our model with samples from the Simlukafusp data set, and investigate how the ROI size sampled affects the separation of tumour outcomes.

In Chapter 5, we summarise and discuss the findings of this thesis in the context of the wider literature, and point to future directions of research to further develop this work.

In this thesis, we present several novel findings and contributions to the field. In Chapter 2, we explore in detail the effects of changing modelling assumptions on the emergent behaviour that the model produces. We find an emergent feature in the model, with CD8+ T cells forming ring-like structures, which can result in non-linear effects on the tumour elimination times. We highlight how this feature is the result of the modelling assumptions inherent in our on-lattice ABM, and thus our finding has wider relevancy for similar models. In Chapter 3, we present a methodology for analysing core-needle biopsy samples, outlining the considerations which need to be taken with these often small and fragmented tissue pieces. We also provide hypotheses as to why patient response was non-uniform across the cohort, informed by the results of the spatial analysis. We propose a novel methodology

in Chapter 4 to investigate the accuracy of predicting tumour outcome based on limited timepoint data. This method leverages the ability to generate vast amounts of synthetic data from an ABM, which allows us to estimate the upper bound of tumour outcome prediction accuracy using a single on-treatment timepoint.

Chapter 2

Emergent dynamics of a cellular automaton model of tumour-immune interactions

2.1 Overview

In Chapter 1, baseline T cell infiltration was highlighted as an indicator of cancer patient survival. Furthermore, the spatial distribution of T cells within the tumour has been found to be a significant indicator of patient outcomes (Carstens et al., 2017). In Chapter 1 we also reviewed existing mathematical models of tumour-immune interactions, many of which are ABMs. These models capture behaviours and interactions between distinct cells, building a timecourse of the spatial evolution of the system.

Our collaborators at Roche have developed a pipeline (Hutchinson and Grimm, 2022) integrating biopsy imaging data and mathematical modelling to optimise the timing of the collection of the on-treatment biopsy, which is examined by pathologists to assess drug efficacy. The biopsy data were taken from a Phase I clinical trial of Simlukafusp. This molecule binds to FAP-expressing tumour-associated cells and stimulates the IL2β/γ receptors on T cells, promoting T cell proliferation (Waldhauer et al., 2021). An ABM was used to simulate the dynamics between the pre-treatment and on-treatment biopsies for each patient in the cohort, adapted from a model by Kather et al. (2017). Analysis of the Hutchinson and Grimm (2022) model revealed emergent dynamics which were difficult to interpret due to the model complexity. Agent-based models by definition are often complex, making it difficult to pinpoint the origin of their emergent behaviours. Hutchinson and Grimm (2022) note that an extensive sensitivity analysis of their model was beyond the scope of their paper, and that co-dependencies between model parameters had not been considered. Computational constraints are a major bottleneck in the pursuit of complete analysis of these models. Model reduction involves identifying key mechanisms and components of the

detailed model, in order to retain the essence of the dynamics whilst developing a model that is easier to analyse.

To perform a more detailed sensitivity analysis, in this chapter we propose a new cellular automaton (CA) model of interactions between tumour cells and CD8+ cytotoxic T cells (CTLs), based on the models developed by Kather et al. (2017) and Hutchinson and Grimm (2022). The aim of this work is to develop a minimal model that captures the key biological features of this system of tumour-CTL interactions, such that the effect on the tumour response of varying the CTL properties can be elucidated. Our analysis of the model has the potential to identify targets for immunotherapy.

In this chapter we aim to answer the following research questions:

- How can we simplify the Hutchinson and Grimm (2022) model in order to obtain a reduced model that exhibits the same qualitative behaviours, and thus understand their origins?
- For a given tumour type, what does the simplified model predict to be the optimal CTL properties to minimise tumour elimination time?
- How do assumptions of how to implement the biological processes and the timescales considered affect the model results?
- What modifications or adjustments to the modelling assumptions could be made to improve the reliability of the model to qualitatively predict the tumour-CTL dynamics?

This chapter is structured as follows. In Section 2.2 we review the findings of Kather et al. (2017) and Hutchinson and Grimm (2022) on which the work in this chapter is based. In Section 2.3 we introduce a new, simplified version of the Hutchinson and Grimm (2022) model, and methods of analysing the simulated data in Section 2.4. In Section 2.5 we present typical simulations of the model and a sensitivity analysis considering the properties of CTLs. In Section 2.6 we explore the effect of changing the timescales simulated on the model outputs. Finally, Section 2.7 provides a summary of findings of this chapter.

2.2 Model development

2.2.1 Motivation

When constructing our CA, it is important to consider the motivation behind our modelling choices. In particular, the main biological mechanisms involved in the tumour-CTL system,

and the kinds of data that could be used to initialise and calibrate such models, should inform the modelling approach that is taken.

In Chapter 1, we discussed the importance of the spatial distribution of CTLs within the tumour as an indicator of patient prognosis. This, paired with the increasing availability of spatially-resolved imaging data of individual cells within tumour tissue, motivates our choice to develop a spatial model, with the extended aim of integrating image data with our mathematical model. Our decision to use an ABM rather than other spatial modelling methods is motivated by the availability of data with individually-labelled cells and the ability to readily enforce rules governing individual cell behaviour.

Many existing ABMs model tumour-immune interactions using a CA framework. There are many advantages to using CAs, from their relatively simple design and implementation through to their computational efficiency, compared to other types of ABMs. These reasons motivate us to start with a CA framework and use as simple a model as possible, providing ease of comparison and fitting to data. Kather et al. (2017) developed a CA model to study lymphocyte-tumour-stroma interactions and understand why immunotherapies have shown little benefit in colorectal cancer. The authors of this study used an ABM framework to capture the spatial patterns observed in histology samples of colorectal cancer. Our collaborators at Roche have adapted the model by Kather et al. (2017) in a project to optimise the scheduling of biopsy sampling in oncology clinical trials (Hutchinson and Grimm, 2022), which provides further motivation for adopting this framework. Their model inspires the model presented in this chapter, such that our findings are relevant to ongoing pharmaceutical questions and studies.

2.3 Development of simplified model

2.3.1 CA structure

The aim of the simple model developed in this chapter is to retain the key features of the Hutchinson and Grimm (2022) model and use our model to compare tumour responses to different immunotherapies. We chose to implement our ABM from scratch in Python, due to the relatively simple rules and to easily implement rules consistent with Hutchinson and Grimm (2022). The model retains the basic biological mechanisms of the model by Kather et al. (2017). The dynamics of two interacting cell populations are simulated. Tumour cells and CTLs move and interact on a two-dimensional spatial lattice. Only one cell may occupy a lattice site, and all cells are assumed to occupy the same area in the grid, regardless of their type. Each lattice site is 14.9 μm in length, corresponding to the diameter of a tumour cell (Kather et al., 2017). All actions take place within the Moore neighbourhood

of a cell (its eight nearest neighbours). Periodic boundary conditions are imposed along the perimeter of the domain, so that cells move in, and out, of the ROI. By adopting these boundary conditions we are assuming that the tissue surrounding the ROI has the same composition and structure as the tissue inside the domain. The choice to model an ROI of the tumour rather than the whole tumour (as in Kather et al. (2017)) is motivated by the biopsy data; the same approach was used by Hutchinson and Grimm (2022). During the collection and preparation of tissue samples, the edges of the tissue can become damaged. Therefore to exclude such damaged tissue, attention focusses on ROIs that do not include edges. The ABM domain size has dimension 100 by 100 lattice sites.

2.3.2 Initial conditions

To initialise the model, synthetic tiles are generated in which $n_T(t = 0)$ tumour cells and $n_E(t = 0)$ CTLs are randomly placed at lattice sites on a regular 2D grid of dimension 100 by 100 cell diameters (see Figure 2.1).

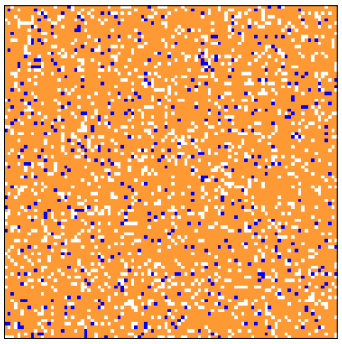


Figure 2.1: Sample initial condition: 100 x 100 cell diameter square lattice, randomly seeded with 8000 tumour cells (orange squares) and 500 CTLs (blue squares). Empty space is shown in white.

The random cell distribution shown in Figure 2.1 was used to initialise all simulations in this chapter for simplicity and to see how when running the model, the system dynamics deviate from an initially well-mixed condition. In what follows, each simulation is initialised with 8000 tumour cells and 500 CTLs. The initial condition is identical for all stochastic realisations of the model. In Appendix A we consider how representative it is to take this approach, compared to generating a new random distribution of cells for each stochastic realisation. We find that, with the exception of one outlier realisation, the standard deviation of the cell population dynamics of 100 realisations using the single condition overlaps with that of considering a different initial configuration for each realisation. We also find

that the PCFs were consistent in the two cases. Thus we justify proceeding with the usage of this initial condition throughout the rest of this chapter.

2.3.3 Tumour cells

During each iteration (12 hours), each tumour cell, labelled i, is randomly selected, and a random number, $\xi^i_{TU} \in [0,1]$, drawn between 0 and 1. This number is compared to the probabilities with which the tumour cell will proliferate, migrate, die or remain idle as described below. These probabilities must sum to one. The rules for each tumour action are detailed as follows.

Proliferation A tumour cell will proliferate if $0 \leq \xi^i_{TU} \leq r_T$, where r_T is the tumour proliferation probability. A proliferation event will only be enacted if there is at least one empty lattice site in the cell's neighbourhood; one of the empty sites is chosen at random to place the daughter cell.

Migration A tumour cell will migrate if $r_T < \xi^i_{TU} \leq r_T + m_T$, where m_T is the tumour migration probability. A migration event is enacted if there is at least one empty lattice site in the cell's neighbourhood; one of the empty sites is chosen at random for the tumour cell to move to. Tumour cells marked for migration will move only one lattice space during a time step.

Death A tumour cell will die if $r_T + m_T < \xi^i_{TU} \leq r_T + m_T + d_T$, where d_T is the tumour death probability. Natural cell death is distinct from CTL-induced killing of tumour cells, according to the tumour death probability parameter. If a tumour cell is marked to die, it is instantaneously removed from the grid.

If $r_T + m_T + d_T < \xi^i_{TU} \leq 1$ then the tumour cell will remain idle for that time step.

2.3.4 CTLs

As for the tumour cells, the CTL processes (proliferation, migration, natural death) are carried out in a random order. A random number ξ^j_{CTL} is generated for each CTL (labelled j) and used to determine its action. Influx of CTLs into the domain is neglected; we assume that any influx of CTLs is incorporated into the proliferation term in the model, which implicitly assumes that influx is proportional to the number of CTLs in the domain.

Proliferation A CTL will proliferate if $0 \leq \xi^j_{CTL} \leq r_E$, where r_E is the CTL proliferation probability. A proliferation event is implemented if there is at least one empty

lattice site in the cell's neighbourhood; one of the empty sites is chosen at random and the daughter cell placed in that site.

Migration A CTL will migrate if $r_E < \xi^j_{CTL} \leq r_E + m_E$, where m_E is the CTL migration probability. A migration event is permitted if there is at least one empty lattice site in the cell's neighbourhood. CTLs move via a biased random walk which is determined by a bias parameter, β. This parameter takes values between zero and one, with a value of zero corresponding to random movement and a value of one corresponding to movement which is biased towards the closest tumour cell. When a CTL is marked for migration, a random number ξ_B is generated before each migration step and compared to the value of the bias parameter which determines whether movement is random or biased towards the nearest tumour cell, whose location is determined using a distance map - a simple example is shown in Figure 2.2. If $\xi_B \leq \beta$ then the CTL will move to the neighbouring space that is closest to the nearest tumour cell. If $\xi_B \geq \beta$ then the CTL will move to an empty neighbouring space chosen at random. It is assumed that if a CTL is selected to migrate and there is at least one empty neighbour site, then it will migrate, irrespective of the value of the tumour distance map at its original site. A similar distance map approach was used in the Kather et al. (2017) model, except that they used the Euclidean distance when calculating the distance to the closest tumour cell from a given cell site. We use the Chebychev distance to mirror the Moore neighbourhood scheme adopted in the model rules. Thus we consider diagonal neighbours in the same way as orthogonal ones. Appendix B details the implementation of this tumour distance map and quantifies the effect of different assumptions when generating this map.

The Chebychev distance, $D_{chebychev}$, between cell sites A and B is defined as follows:

$$D_{chebychev}(A, B) := max(|x_A - x_B|, |y_A - y_B|) \tag{2.1}$$

where $x_{A,B}$ and $y_{A,B}$ are the x and y coordinates of cell sites A and B.

CTLs marked to migrate move up to η lattice spaces during the time step if there is space to migrate. CTL migration is terminated if a tumour cell is detected within its Moore neighbourhood (see Tumour killing section below).

Tumour killing During each 12hr timestep, each CTL will scan its Moore neighbourhood for tumour cells to kill. If a tumour cell is detected in a CTL's neighbourhood, then a random number $\xi_k \in [0, 1]$ is generated and compared to the tumour kill probability parameter, κ. If $\xi_k \leq \kappa$ then the tumour cell will be killed and instantaneously removed from the grid. If more than one tumour cell is in the T cell's neighbourhood, then only one tumour cell will be randomly selected for killing. If $\xi_k > \kappa$ then a killing event does not happen.

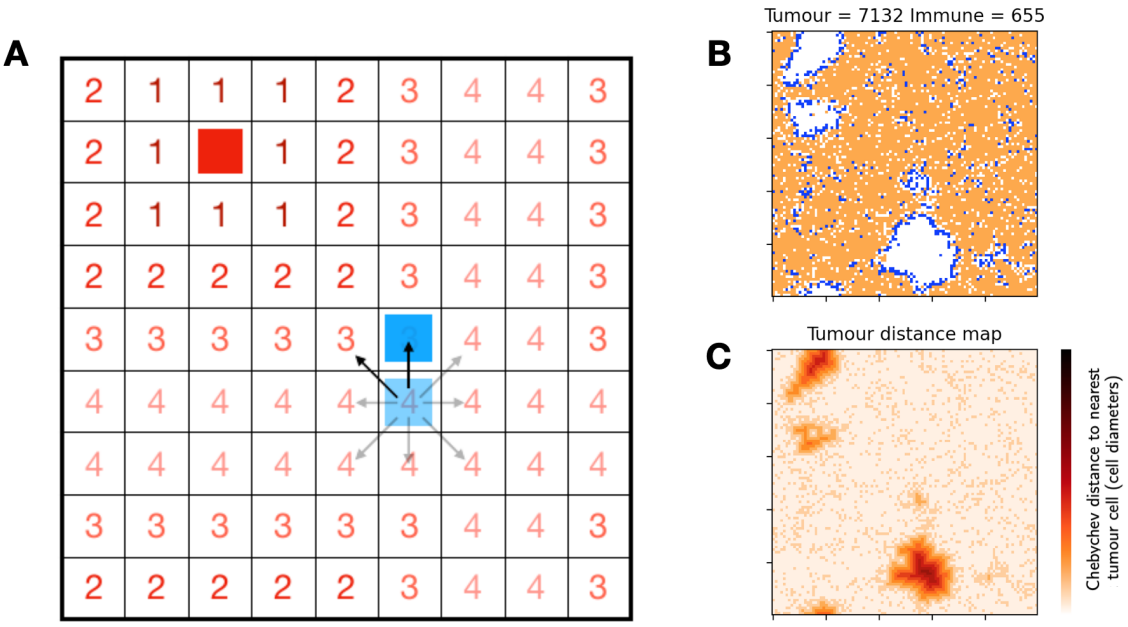


Figure 2.2: Example of how CTLs (blue) sense, and migrate towards, tumour cells (red). After all tumour cells have acted on a particular timestep, a distance transform is used to generate a map which gives the Chebychev distance from each grid site to the closest tumour cell in the domain. The numbers in each of the lattice sites indicate the Chebychev distance between that lattice site and the closest tumour cell. If a CTL is selected to migrate, then a random number is generated and compared to the value of the random walk bias parameter to determine whether the CTL will migrate towards the nearest tumour cell (the bold black arrows in panel (A) indicate the CTL's possible movement directions when motion is biased towards the tumour cells). If biased motion is not chosen, then the CTL moves into one of the eight neighbouring sites at random (dependent on the neighbouring sites being empty). Panel (C) shows the distance transform map generated for the simulation snapshot shown in panel (B).

Death A CTL will die if $r_E + m_E < \xi^j_{CTL} \leq r_E + m_E + d_E$, where d_E is the CTL death probability. If a CTL is marked to die, then it is instantaneously removed from the grid.

Algorithm 1 outlines the algorithm used to simulate this model, which was implemented in Python 3.8.10.

2.3.5 Time stepping

The length of a time step (iteration) in the model is fixed at 12 hours. Tumour cells act first; the distance map (used to account for the attraction of CTLs to tumour cells) is updated, and then the CTLs act (see Figure 2.3). This time step and algorithm updating scheme were chosen for consistency with the Kather et al. (2017) and Hutchinson and Grimm (2022) models and also because the biopsy data are sampled at timepoints that are tens of days

Algorithm 1: Simulation algorithm

```
Domain initial condition for tumour and T cell agents
while t < t_final do
    // act on all tumour cells first
    Shuffle list of tumour cells
    for all tumour cells do
        Generate random number ξ^i_TU
        if ξ^i_TU ≤ r_T then
            if empty neighbour then
                Cell proliferates → add new tumour cell into random empty neighbour
            end
        else
            if r_T < ξ^i_TU ≤ r_T + m_T then
                if empty neighbour then
                    Cell migrates → move tumour cell into random empty neighbour
                end
            else
                if r_T + m_T < ξ^i_TU ≤ r_T + m_T + d_T then
                    Cell dies, remove from grid
                else
                    Cell remains idle
                end
            end
        end
    end
    // after all tumour cells act, update distance transform map
    Calculate distance transform on tumour population
    Shuffle list of T cells
    for all T cells do
        Generate random number ξ^j_CTL
        if ξ^j_CTL ≤ r_E then
            if empty neighbour then
                Cell proliferates → add new T cell into random empty neighbour
            end
        else
            if r_E < ξ^j_CTL ≤ r_E + m_E then
                for imspeed intermediate timesteps do
                    if next to tumour cell and CTL has moved at least one space then
                        // migration stops if T cell encounters a tumour cell in its
                         neighbourhood
                        return
                    else
                        if empty neighbour then
                            Generate random number ξ_B
                            if ξ_B ≤ β then
                                Move T cell to empty neighbour with lowest value in tumour distance
                                  transform map
                            else
                                Move T cell into random empty neighbour
                            end
                        end
                    end
                end
            else
                if r_E + m_E < ξ^j_CTL ≤ r_E + m_E + d_E then
                    Cell dies, remove from grid
                else
                    Cell remains idle
                end
            end
        end
        // T cell killing tumour cells
        if next to tumour cell then
            Generate random number ξ_k
            if ξ_k ≤ κ then
                T cell kills tumour cell → tumour cell removed from grid
            end
        end
    end
    t ← t + dt
end
```

apart so, for computational efficiency, the time step cannot be very small.

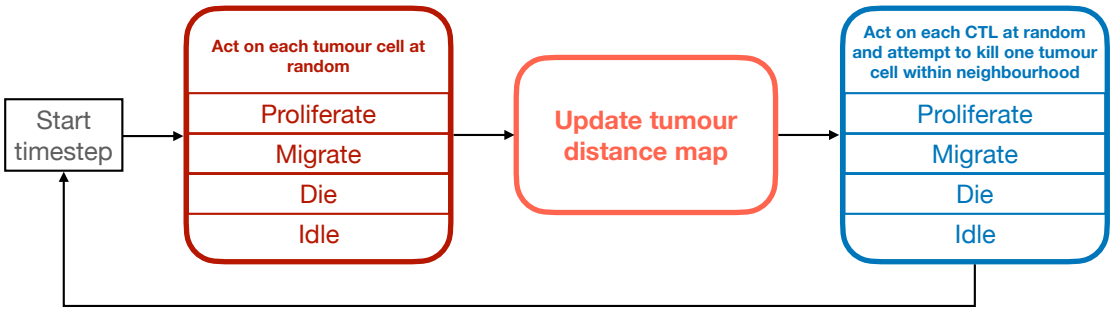


Figure 2.3: Flowchart showing how our CA model is updated at each time step.

Although the structure of the algorithm updating scheme is the same as that used by Hutchinson and Grimm (2022) (i.e. updating the tumour cells, then updating the CTLs), our algorithm contains several simplifications; these are detailed below (see Section 2.3.6).

2.3.6 Key simplifications from the Hutchinson and Grimm (2022) model

We have made a number of simplifications to the Hutchinson and Grimm (2022) model. These simplifications can be classified into biological simplifications and algorithmic simplifications.

In Hutchinson and Grimm (2022), non-stem tumour cells and CTLs have a limited proliferative capacity, such that there is a maximum number of proliferation events a cell could make before dying. We remove the cell proliferation capacity, so that the tumour cells form a homogeneous mass in which the number of proliferation events that a single tumour cell or CTL can make is unlimited. Thus, tumour stem cells are not considered; all tumour cells are assumed to be identical. This simplification reduces the number of model parameters. We assume further that cell death due to reaching maximum number of proliferation events can be implicitly incorporated into the natural cell death probability d_T, as Kather et al. (2017) found the value of this parameter by apoptosis markers in histology samples, which would not distinguish between the two death modes.

We also removed the killing capacity for each CTL, thereby neglecting CTL exhaustion. Under this simplification the effect of exhaustion is implicitly incorporated into the tumour killing probability parameter, κ. Engagement of CTLs when killing tumour cells is also neglected in our model, however when a killing event is permitted, the CTL involved will not migrate for the rest of the timestep, which implicitly models this engagement period, and also reflects findings in the biological literature that CTLs arrest their migration when they encounter an antigen-expressing tumour cell (Boissonnas et al., 2007). In contrast, in

Hutchinson and Grimm (2022), killing and migration events are simulated on intermediate timesteps, so the engagement period can be explicitly modelled; a CTL that has just killed is unable to kill again for a given number of intermediate timesteps, but is able to migrate.

Hutchinson and Grimm (2022) also assume a tumour cell needs to undergo two killing events before being removed from the grid, reflecting findings that tumour cells require multiple hits to be killed (Weigelin et al., 2021). We assume that only one killing event is required to remove a tumour cell, and that the multiple hit hypothesis may be implicitly considered in the choice of the value of the kill probability κ, an assumption which Kather et al. (2017) made in their model.

In our model, we neglect the influx of CTLs into the tissue. This assumption reduces model complexity. Further, to the best of our knowledge, there are no data regarding the influx rate of CTLs. By making this simplification, we assume that CTL influx is implicitly incorporated into the CTL proliferation term, and that CTL influx is proportional to the number of CTLs within the tissue.

We also simplified the model algorithm. The most significant difference between the implementation of our model and the Hutchinson et al. (2022) model is that each CTL is only allowed to perform one of the following actions per timestep: migrating x grid spaces, proliferation, death, or remaining idle. Figure 2.4 shows the algorithm used in the model by Hutchinson and Grimm (2022).

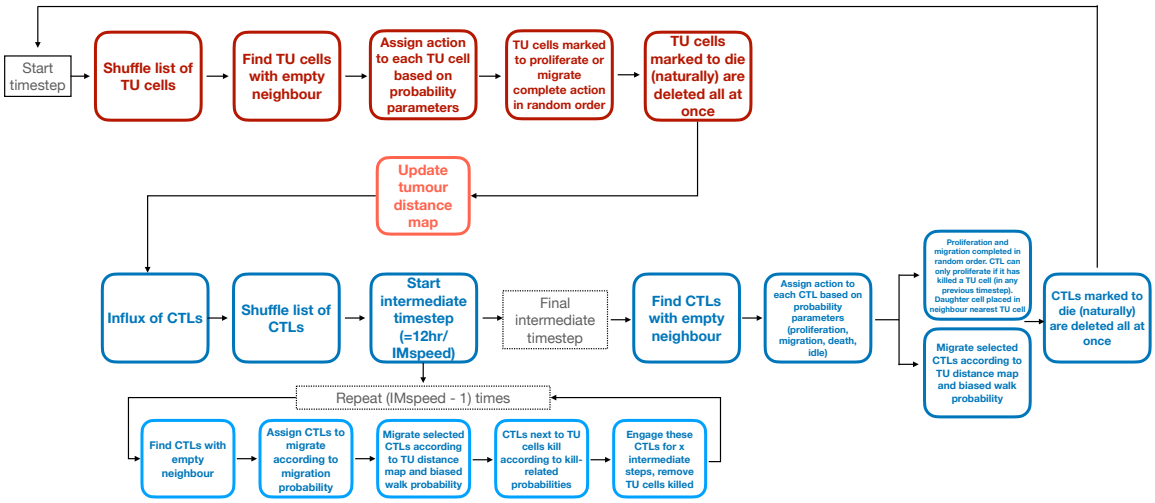


Figure 2.4: Flowchart showing how the Hutchinson and Grimm (2022) model is updated at each timestep. Red and blue boxes contain actions pertaining to the tumour cell population and CTL population, respectively.

Each CTL is chosen randomly and will carry out one of these actions. After its action, if a CTL has not died in that timestep, then it attempts to kill a neighbouring tumour cell. In

the Hutchinson and Grimm (2022) model, the CTL migration steps happen first - for each intermediate timestep, CTLs are randomly selected to migrate one grid space, and search the neighbourhood for a tumour cell to kill (if killing is successful they remain engaged in killing for m intermediate timesteps). CTLs are then randomly chosen to proliferate, and the same for death. Thus in their model, migration, proliferation and death happen sequentially, rather than in a random order. As a result, all cells marked to proliferate attempt proliferation first, then all cells marked for natural death die together. This can lead to some cell proliferation events being blocked due to crowding; randomising the ordering of the proliferation and death processes would reduce the impact of this phenomenon. In the Hutchinson and Grimm (2022) algorithm a CTL cell may carry out multiple actions (i.e. migrate and proliferate) in the same timestep. In our model we consider each CTL at random and follow through its actions for that timestep (e.g. move x spaces then kill) and then pick the next CTL at random to perform its actions. In contrast, the Hutchinson and Grimm (2022) model mixes the CTLs during the migration phase. Our algorithm is simpler to conceptualise and ensures that any given CTL does not perform more than one of the migration, proliferation or death events in a single timestep. Our algorithm also assumes that tumour killing and CTL migration are independent. By contrast, in the Hutchinson and Grimm (2022) algorithm, killing is enacted within the migration phase, so if the CTL migration probability m_E is reduced, this also reduces the number of killing events.

In Hutchinson and Grimm (2022), only tumour cells with an empty neighbouring grid space are considered for proliferation, migration and death events. Algorithmically this was an efficient approach, however to only consider tumour cells that have an empty neighbour for the death event does not have an obvious biological motivation - on the contrary, cells in a crowded environment may have a greater probability of cell death due, for example, to lack of nutrients. In order not to overcomplicate our model, all cells are considered when selecting for the different cell events. Their neighbourhood status is only checked after the cell action has been selected. In Hutchinson and Grimm (2022) the mechanism of action of the immunotherapy used in the clinical trial was modelled by requiring that a given CTL must have killed a tumour cell in any previous time step before being permitted to proliferate. This was not included in our model.

Finally, we note that data extracted from the biopsy images in the FAP-IL2v data set do not contain information about proliferation capacity, killing capacity, exhaustion, CTL influx, CTL engagement period, stem or non-stem tumour cells. This is another reason for neglecting these biological features in our simplified model.

The main differences between our model and those by Kather et al. (2017) and Hutchinson and Grimm (2022) are detailed in Table 2.1.

Model feature	**Kather et al. (2017)**	**Hutchinson and Grimm (2022)**	**Our model**
Model context	Avascular tumour growth from single stem cell	Modelling a region of interest (ROI) within a tumour, with avascular growth	see Hutchinson and Grimm (2022)
Initial conditions	Single tumour stem cell in domain centre	ROI from biopsy tissue (stem and non-stem tumour cells, CD8+ T cells)	Synthetic ROI with randomly distributed tumour and CD8+ T cells
Boundary conditions	No flux	Doubly periodic	Doubly periodic
Agents	Stem and non-stem tumour cells, CTLs	see Kather et al. (2017)	Tumour cells and CTLs
Other biological features	Fibrosis, necrosis	Fibrosis and necrosis removed	see Hutchinson and Grimm (2022)
Cell properties	Tumour cells: proliferation capacity, stemness. CTLs: proliferation capacity, killing capacity, engagement status	see Kather et al. (2017)	All properties in Kather et al. (2017) were removed during model simplification.
Biological processes modelled	Cell proliferation, migration, death, CTL killing of tumour cells, formation of fibrotic and necrotic regions	Cell proliferation, migration, death, CTL killing of tumour cells	Cell proliferation, migration, death, CTL killing of tumour cells
Timescales	12 hour timestep, with shorter timesteps when modelling CTL migration.	see Kather et al. (2017)	12 hour timestep, each CTL completes full migration path (i.e. moves up to η spaces) before the next cell chosen to act.
Tumour killing	Migrating CTLs check neighbourhood for tumour cell to kill at each intermediate timestep, will kill and become engaged for x intermediate steps	Migrating CTLs check neighbourhood for tumour cell to kill at each intermediate timestep, will kill and become engaged for x intermediate step	CTLs check neighbourhood after each migration step and if a tumour cell is detected in the neighbourhood the CTL migration path terminates and the CTL attempts to kill the tumour cell

Table 2.1: Comparison of our model with the models of Kather et al. (2017) and Hutchinson and Grimm (2022).

2.4 Methods

2.4.1 Summary statistics

We now define the summary statistics which will be used to characterise the output from the ABM. Firstly, we are interested in the temporal evolution of cell numbers of the two cell populations; the number of tumour cells at time t days, $n_T(t)$, and the number of CTLs at time t days, $n_E(t)$; which are extracted after each timestep. We are also interested in extracting spatial information from our model, in particular the spatial distribution of the CTL population as this is known to be a significant indicator of patient outcomes (Carstens et al., 2017). To characterise the model output, we will use the pair correlation function between CTLs, corrected for on-lattice data (Gavagnin et al., 2018).

The summary statistics used to analyse the model output are listed in Table 2.2.

Summary statistic	**Definition**
$n_T(t)$	Number of tumour cells at time t days
$n_E(t)$	Number of CTLs at time t days
t_{elim}	Time (days) for the tumour population to be eliminated.
$g_{EE}(r)$	Pair correlation function for CTLs to CTLs at radial distance r (units of cell diameters).
$g_{TE}(r)$	Pair correlation function for tumour cells to CTLs at radial distance r (units of cell diameters)

Table 2.2: Summary statistics used to analyse output of our CA model.

2.4.2 Spatial analysis - on-lattice pair correlation function

Gavagnin et al. (2018) discussed how PCF calculations should be modified for on-lattice data. Such modifications are needed because, in contrast to off-lattice models (or biological data), cells can only occupy a limited number of sites (see Figure 2.5). They considered different neighbourhoods, for example the von Neumann neighbourhood and Moore neighbourhood ; and extended annuli of these neighbourhoods for greater distances. They also computed PCFs, assuming periodic and non-periodic boundary conditions. In all cases, they assumed that at most one agent can occupy a single lattice site.

In our model we consider Moore neighbourhoods and impose periodic boundary conditions. Therefore we use Gavagnin et al.'s results for the PCF calculation of this particular geometry, referred to as the *square uniform metric* in their paper, to analyse our image data and ABM output.

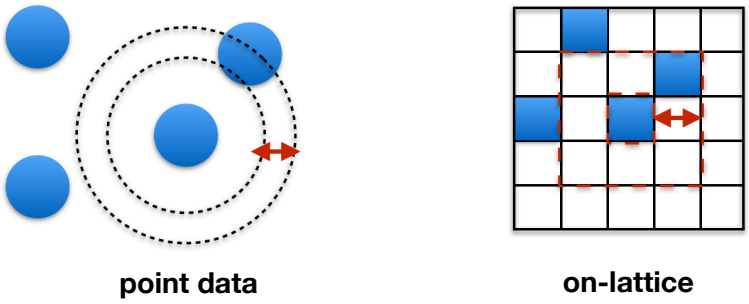


Figure 2.5: Schematic illustrating the difference when calculating spatial statistics on point data and on-lattice data. The blue points and filled lattice spaces represent cells. The dashed lines show the hollow disk in the off-lattice and on-lattice frameworks that is used to calculate the density of points at increasing distance from the point of interest. This disk is circular in the off-lattice framework, and square in the on-lattice framework (for a Moore neighbourhood assumption, where diagonal neighbours are considered as being the same distance from the point of interest as the neighbours in the x and y directions) (Gavagnin et al., 2018).

The PCF, $g(m)$, is calculated by first counting the number of pairs of agents, $c(m)$, which are separated by the discrete distance $m = 1, 2, 3, ...$, and then dividing by a normalisation factor, E, so that:

$$g(m) = \frac{c(m)}{E}. \tag{2.2}$$

The normalisation factor E is calculated as follows. Let $q(m)$ be the number of lattice sites at distance m from a specific site. Then, assuming a Moore neighbourhood, we have

$$q(m) = 8m. \tag{2.3}$$

The number of distinct pairs of sites $q(m)$ at distance m from each other must be multiplied by the total number of lattice sites and normalised by 2, to yield the number of *pairs* of lattice sites, $s(m)$ at distance m:

$$s(m) = \frac{q(m)L_xL_y}{2} = 4mL_xL_y \tag{2.4}$$

where L_x and L_y are the integer x and y dimensions of the lattice, respectively, with $L_x, L_y > 1$.

The normalisation factor used in Gavagnin et al. (2018) was the expected value for the number of pairs of agents separated by distance m when N agents are randomly distributed on the lattice (Binder and Simpson, 2013). In particular, they assumed that

$$E(m) = \left(\frac{N}{L_x L_y}\right)\left(\frac{N-1}{L_x L_y - 1}\right)s(m) \tag{2.5}$$

Substituting for $s(m)$ from Equation 2.4, we have that

$$E(m) = \frac{4mN(N-1)}{L_x L_y - 1}. \tag{2.6}$$

A limitation of this PCF calculation is that the pairwise distance m cannot be greater than half the dimension of the lattice (Gavagnin et al., 2018).

For PCFs where two populations of cells are considered (i.e. the PCF for CTLs around tumour cells) where N and M are the number of CTLs and tumour cells respectively, the normalisation factor is:

$$E(m) = \frac{8mNM}{L_x L_y - 1}. \tag{2.7}$$

The factor of 2 between the two normalisation factors originates from the double counting of pairs of cells which happens when the two cell types in the PCF calculation are the same (i.e. calculating the PCF of CTLs around CTLs).

We note that the PCF used here assumes that diagonal distances are equivalent to orthogonal distances, which may confound the results. For example the value of the PCF at a distance of 10 cell diameters will be calculated taking into account cells that are at a Euclidean distance of $\sqrt{200} \approx 14.1$ cell diameters (on the diagonal). We motivate our use of the square uniform metric by considering the neighbourhood we use in our model, which is the same approach taken by Owen et al. (2020), in which they analyse on-lattice simulation data using the square uniform PCF to reflect their choice to use the Moore neighbourhood for cell-cell interactions. Finally, we justify the use of this metric by highlighting that short-range interactions are of most interest in our work, i.e. co-localisation of tumour cells and CTLs (as this is what results in tumour cell killing) and so we want our spatial metric to reflect the distance metric we use in the model on the short length scale in particular. We adapt the Matlab code by Gavagnin et al. (2018) in Python, and use this for the PCF calculations.

2.4.3 Parameter values

In their model, Kather et al. (2017) consider tumour cells and tumour infiltrating lymphocytes, referred to as the "immune" population. In Hutchinson and Grimm (2022) and our work, the immune population is assumed to comprise CTLs only, due to the data available in the biopsy images. This is a reasonable assumption, as in the Kather et al. (2017) model the immune cells are assumed to be purely cytotoxic. Kather et al. (2017) parametrised their model using histology data, where lymphocytes were identified through their morphology and were stained for Ki67 and aCasp3 to extract the proliferative and dying cell fractions. Assuming all lymphocytes in the data are cytotoxic may overestimate some CTL model parameters (proliferation and death probabilities). Hutchinson and Grimm (2022) adopted some of the parameter values from the Kather et al. (2017) model, however the default parameter values of the Hutchinson and Grimm (2022) model cannot be directly applied to our model as many of the biological details were removed. For example, the proliferation capacities and T cell exhaustion feature were removed, which impacts the growth rates of both cell populations. We have developed a simplified version of the Hutchinson and Grimm (2022) model in order to encompass a minimal set of mechanisms to model this system, with the aim of identifying T cell characteristics driving tumour elimination, which could form possible targets for immunotherapies, whilst maintaining a simple model with few parameters such that it could be used to fit to histology data.

Our simplified model contains nine input parameters; Table 2.3 details each parameter, the range of its values, and the default values used by Hutchinson and Grimm (2022).

2.4.4 Sensitivity analysis

As acknowledged previously, we have made several simplifications to the Hutchinson and Grimm (2022) model when developing our model, relating, in particular, to the CTL population. Therefore it may not be appropriate to adopt the same parameter values from their paper in our model. In order to establish a region of parameter space within which to explore our simplified model, in Section 2.5.3 we simulate the model as we vary four of the CTL parameters; CTL migration probability m_E, CTL killing probability κ, CTL proliferation probability r_E and CTL random walk bias β, in a coarse-grain parameter sweep. We consider a specific set of tumour parameters which will remain fixed (in effect considering a specific tumour type). These parameter values are the same as those used by Hutchinson and Grimm (2022). This approach is adopted for simplicity and to focus on the effects of varying the CTL response in the model, however future work could consider inter-patient heterogeneity, and the effect of varying the tumour properties as well as the CTL properties.

Parameter	Description	Possible value range	Default value used by Hutchinson et al. (2022)
r_T	Tumour cell proliferation probability	[0, 1]	0.5
m_T	Tumour cell migration probability	[0, 1]	0.35
d_T	Spontaneous tumour cell death probability	[0, 1]	0.12
r_E	CTL proliferation probability	[0, 1]	0.049
m_E	CTL migration probability	[0, 1]	0.2
d_E	Spontaneous CTL death probability	[0, 1]	0.015
κ	Probability of CTL killing of a tumour cell	[0, 1]	0.1
β	Probability of bias in CTL random walk towards closest tumour cell	[0, 1]	0.2
η	Speed of CTL movement, corresponding to the number of lattice spaces moved across on a given time step	$[0, \infty]$	100

Table 2.3: Default parameter values for the CA model, from Hutchinson and Grimm (2022), and their possible value ranges.

We identify behaviours exhibited by our model and by examining the resultant dynamics, we establish a region of parameter space to explore further.

Further sensitivity analysis presented in Section 2.5.4 will investigate the effects of varying CTL parameters in a pairwise manner. As drugs may target more than one biological property of the immune system, it is important to consider the compounding effects of changing multiple parameter values. Our parameter sensitivity analysis goes beyond that presented in Hutchinson and Grimm (2022), where a local sensitivity analysis was performed.

2.5 Results

2.5.1 Typical model output

Our model exhibits a range of qualitative behaviours, broadly classified into tumour elimination, tumour escape, and coexistence of tumour cells and CTLs. Depending on the parameter values, the model exhibits one of these behaviours, or multiple behaviours over many realisations of the model. Similar qualitative behaviours were reported by Hutchinson

and Grimm (2022). Figure 2.6 shows the different qualitative behaviours that our model exhibits when it is initialised using identical initial conditions, identical parameter values (as in Figure 2.6), and different values of CTL-tumour killing probability, κ. Timecourses are shown for one realisation of the model, sampled from 100 stochastic realisations run for each parameter set. The data for the other realisations, including the tumour and CTL population dynamics, are presented in Figures 2.7 and 2.8.

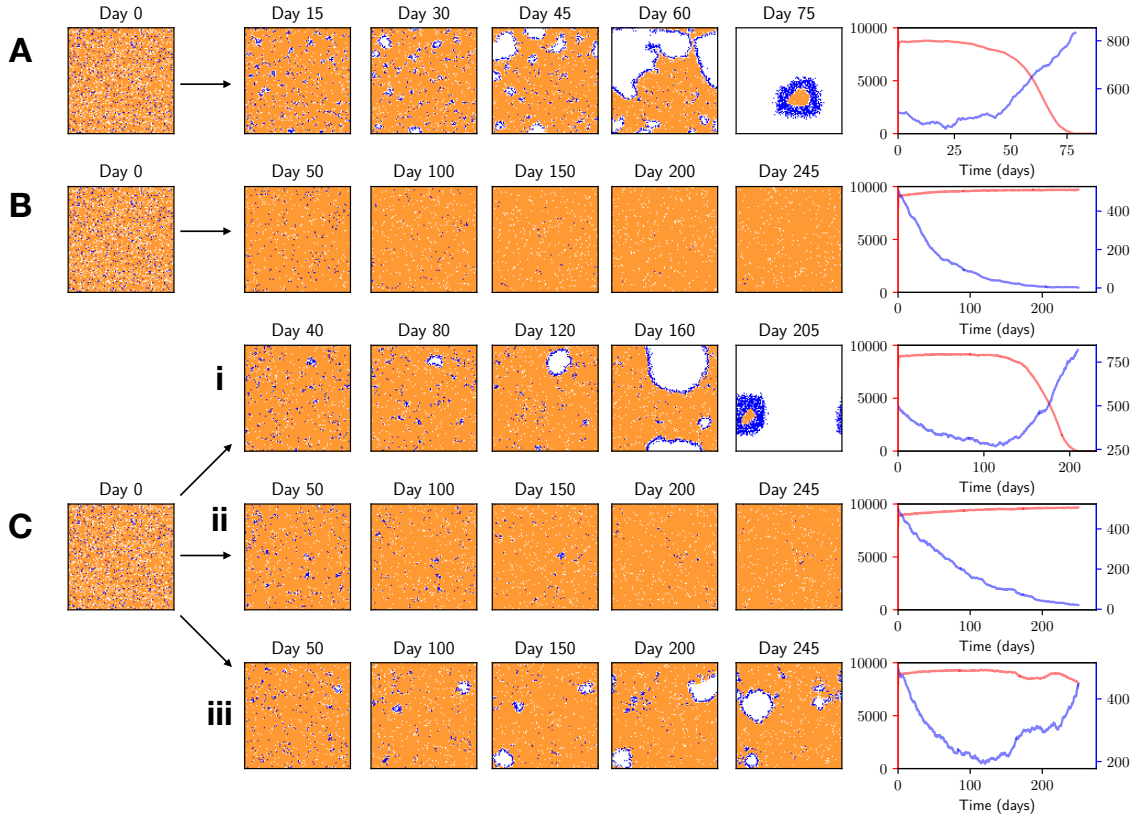


Figure 2.6: Typical behaviours exhibited by the model. Snapshots of the ROI are taken at the indicated times from a single realisation of the model, simulated up to 250 days. Each realisation is initialised with the same initial condition, shown in the LHS column and labelled "Day 0". The RHS column shows the evolution of the tumour population (LHS y-axis, red line) and the CTL population (RHS y-axis, blue line). (A) Tumour elimination - the CTL population eradicates the tumour population. Parameter values: $\{r_T = 0.5, m_T = 0.35, d_T = 0.12, r_E = 0.025, m_E = 0.8, d_E = 0.015, \kappa = 0.9, \beta = 0.5, \eta = 100\}$. (B) Tumour escape - the CTL population does not eradicate the tumour and goes extinct. Parameter values as in (A) with $\kappa = 0.1$. (C) Three different behaviours from the same initial condition and parameter values (same as in (A) with $\kappa = 0.5$). (i) Tumour elimination, (ii) tumour escape and (iii) tumour-CTL coexistence state.

Figure 2.6A shows snapshots from a single realisation of the simulation where each of the 100 realisations resulted in tumour elimination. The parameter values are $\{r_T = 0.5, m_T = 0.35, d_T = 0.12, r_E = 0.025, m_E = 0.8, d_E = 0.015, \kappa = 0.9, \beta = 0.5, \eta = 100\}$. Small

clusters of CTLs form at early times; through CTL killing of tumour cells these clusters then expand to form rings of varying sizes. The CTL rings expand further and increase in thickness through CTL proliferation (which is facilitated by the empty space contained within the CTL ring) while progressively eliminating the tumour population. Finally, the CTL population surrounds the remaining tumour cells and eliminates them by killing from the tumour margin inwards. Figure 2.6A (Day 75) shows that there is a clear separation between the tumour mass and the surrounding CTLs - there is no CTL infiltration into the tumour and the killing rate is dependent on the tumour margin accessible to the CTLs. We note that in the elimination cases shown in Figure 2.6, the CTL population experiences a decrease from the initial population number, before recovering and then expanding. In future work it would be interesting to investigate the dependence of this CTL population recovery period on the model parameters.

Figure 2.6B shows an example of a simulated timecourse for which all realisations of the model resulted in tumour escape, i.e. the CTL population became extinct. The parameter values are $\{r_T = 0.5, m_T = 0.35, d_T = 0.12, r_E = 0.025, m_E = 0.8, d_E = 0.015, \kappa = 0.1, \beta = 0.5, \eta = 100\}$. In this case, the CTL population does not develop clusters and expand into the tumour mass. Instead, the CTLs decrease in number over time. The tumour population expands to fill the space left by the CTLs and eventually reaches a steady state within the domain. It is likely that the initial condition we employ may be a contributing factor to the size of the region of parameter space where tumour escape is observed, as one could envision a higher probability of CTL survival if the CTLs were initialised in clusters, rather than distributed randomly within the tumour tissue and at such a density that many "lone" CTLs are surrounded by tumour cells and unable to move or proliferate.

Figure 2.6C shows an example of a parameter set that can generate three different qualitative behaviours. The parameter values are $\{r_T = 0.5, m_T = 0.35, d_T = 0.12, r_E = 0.025, m_E = 0.8, d_E = 0.015, \kappa = 0.5, \beta = 0.5, \eta = 100\}$. Within 100 realisations of the model, cases of (i) tumour elimination, (ii) tumour escape and (iii) tumour-CTL coexistence are observed. The tumour elimination and escape cases follow similar patterns as in (A) and (B). The coexistence case in (C) shows similar behaviour to the case of tumour elimination, with CTLs forming ring-like structures, which grow in size as CTLs kill tumour cells and proliferate. However, in case (C)(iii), the expansion of the rings is slower than in the elimination case shown in (C)(i), and breaks in the CTL rings can be observed (e.g. on Day 200), due to either CTL death or migration, or the CTLs killing faster and expanding the ring faster than they can replenish and reconstitute the ring through proliferation. The tumour cell population may recover if there is a break in a CTL ring and the tumour cells infiltrate the void at the centre of the ring. This void enables rapid tumour cell expansion,

because the proliferation rate of the tumour cells is much greater than that of the CTLs. The update rules of the CA may amplify this effect, as tumour cells and CTLs are updated sequentially. This is a mechanism by which the total number of tumour cells in the domain oscillates over time; transient oscillations in the tumour population are caused by the birth and death of CTL ring structures. We note that in this case the model was simulated up to 250 days, and that if it was run for longer, the tumour may eventually be eliminated. However, as the latest on-treatment biopsy taken in the data set considered in Hutchinson and Grimm (2022) was taken at 113 days since first dose, our coexistence result is relevant on this timescale.

2.5.2 Temporal dynamics

We examine the model behaviour further by varying four of the CTL parameters; CTL proliferation probability r_E, CTL random walk bias β, CTL migration probability m_E and CTL killing probability κ. The CTL death probability d_E and CTL speed η were fixed at the values given in Table 2.3. Figures 2.7 and 2.8 show the evolution of the tumour and CTL populations respectively, as the four CTL parameters are varied. Each line shows one of the 100 realisations run for each parameter set. Simulations were run up to 250 days.

The model exhibits tumour escape for all stochastic realisations (and CTL elimination) for $\kappa = 0.1$ and $r_E = 0.025$, and the tumour population approaches a steady state as the ROI is filled with tumour cells. Tumour elimination is found in all stochastic realisations of many parameter regimes shown in Figure 2.7. Tumour elimination is observed for all parameter sets with $\kappa = 0.9$. As κ decreases, higher r_E values are required for all realisations to result in tumour elimination (parameter sets where this occurs are $\{\kappa = 0.5,\, r_E = [0.05, 0.1]\}$ and $\{\kappa = 0.1,\, r_E = 0.1\}$).

Bifurcating dynamics are found in the parameter sets with $\{\kappa = 0.5, r_E = 0.025\}$ and $\{\kappa = 0.1, r_E = 0.05\}$, in which some realisations result in tumour escape, and others result in tumour elimination (or are heading towards tumour elimination and the simulations were not run long enough to see this behaviour; in this case we see a "coexistence" state on the timescales of interest). In some of the regions of parameter space with bifurcating dynamics, oscillations in the tumour population are present, for example for the parameter set $\{\kappa = 0.5,\, m_E = 0.8,\, r_E = 0.025,\, \beta = 1\}$. These oscillations are more prominent as β increases. Oscillations are not seen in the CTL population (see Figure 2.8).

2.5.3 Spatial dynamics exhibited by the model

Although the model has been simplified, it exhibits a wide range of qualitative behaviours. The dimensions of the parameter space of ABMs are normally large, and so this poses chal-

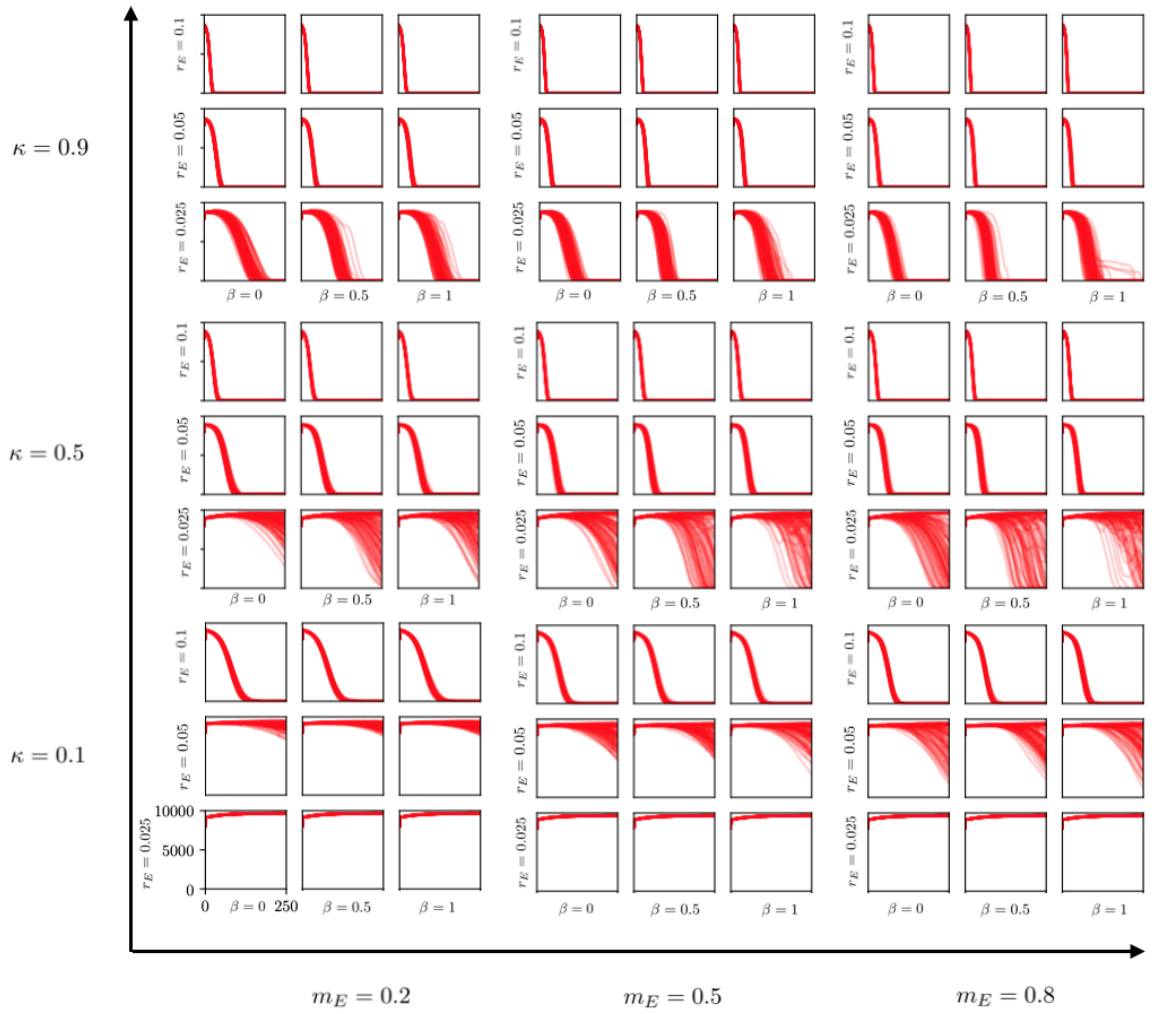


Figure 2.7: Tumour population dynamics over time, varying four of the CTL parameters; CTL migration probability m_E, CTL killing probability κ, CTL proliferation probability r_E and CTL random walk bias β. Each line corresponds to a single realisation of the model, which was simulated 100 times for each parameter set. For each individual plot, axis labels have been omitted for plotting clarity - vertical axis corresponds to cell number, and horizontal axis corresponds to time (days). The axes scalings are indicated in the bottom left plot.

lenges as to how to identify and catalogue the qualitative behaviours that a model exhibits. Bull and Byrne (2023) sweep over two parameters of their model of macrophage-tumour interactions to show the range of qualitative behaviours exhibited. We present the results of a similar parameter sweep in Figure 2.9. We show snapshots of single simulations for a coarse-grained parameter sweep, varying four of the CTL parameters; CTL proliferation probability, CTL random walk bias, CTL migration probability and CTL killing probability. The CTL death probability d_E and CTL speed η were fixed at the values given in Table 2.3. The snapshots were taken from the data set presented in Figures 2.7 and 2.8.

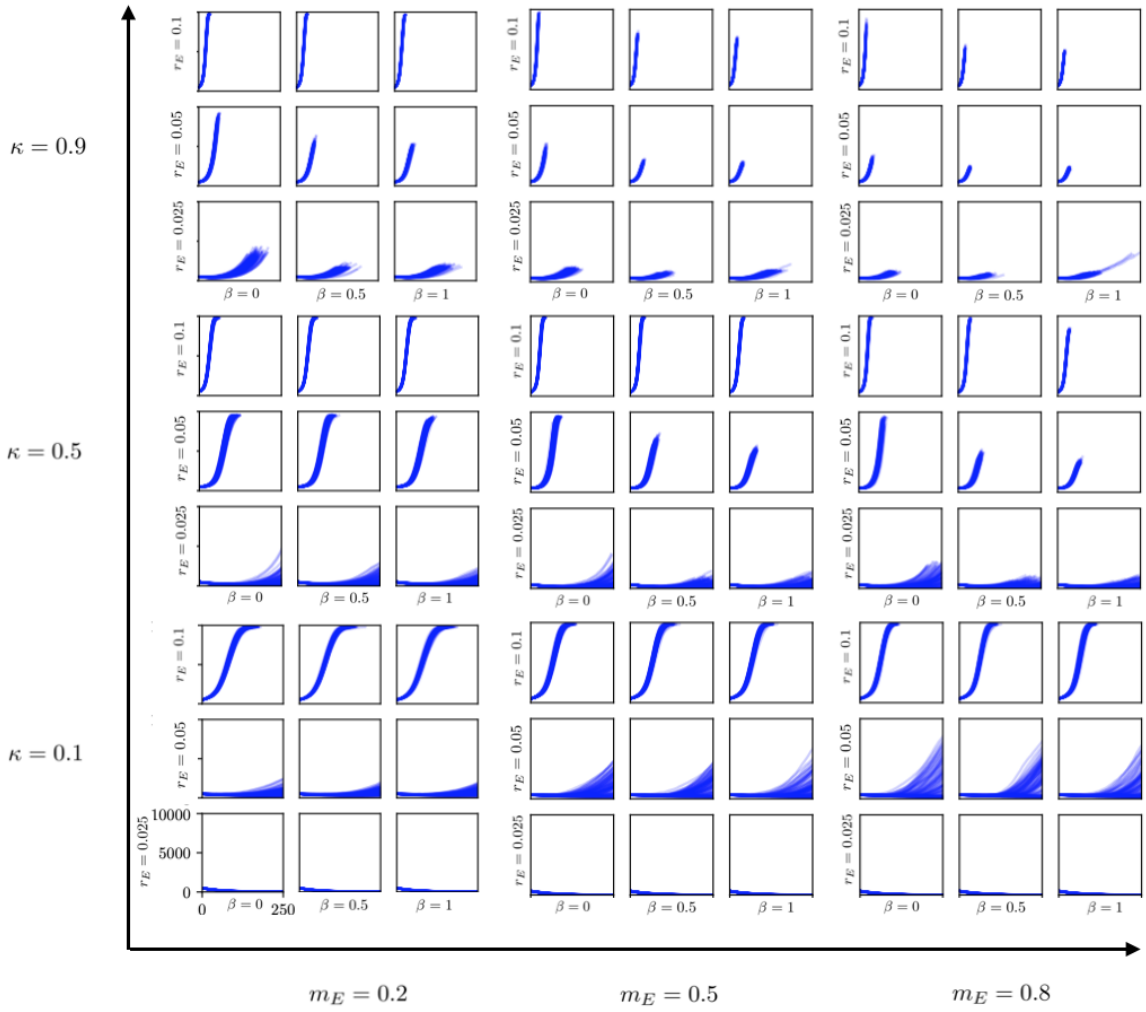


Figure 2.8: CTL population dynamics over time, varying four of the CTL parameters; CTL migration probability m_E, CTL killing probability κ, CTL proliferation probability r_E and CTL random walk bias β. Each line corresponds to a single realisation of the model, which was simulated 100 times for each parameter set. For each individual plot, axis labels have been omitted for plotting clarity - vertical axis corresponds to cell number, and horizontal axis corresponds to time (days). The axes scalings are indicated in the bottom left plot.

Along the large axes in Figure 2.9 we vary the CTL migration probability m_E (along the x axis) and CTL killing probability κ (along the y axis). In each inset parameter set we vary the CTL random walk bias β (along the x axis) and CTL proliferation probability r_E (along the y axis). Each snapshot is shown at the time when the tumour burden (number of tumour cells) has reduced to 50% of its initial level; if the tumour was not reduced to this extent, then the snapshot at $t = 150$ days is shown. For parameter sets where tumour elimination and escape were found in different stochastic realisations, a snapshot is shown from a realisation where the tumour burden reached 50%. In this analysis the

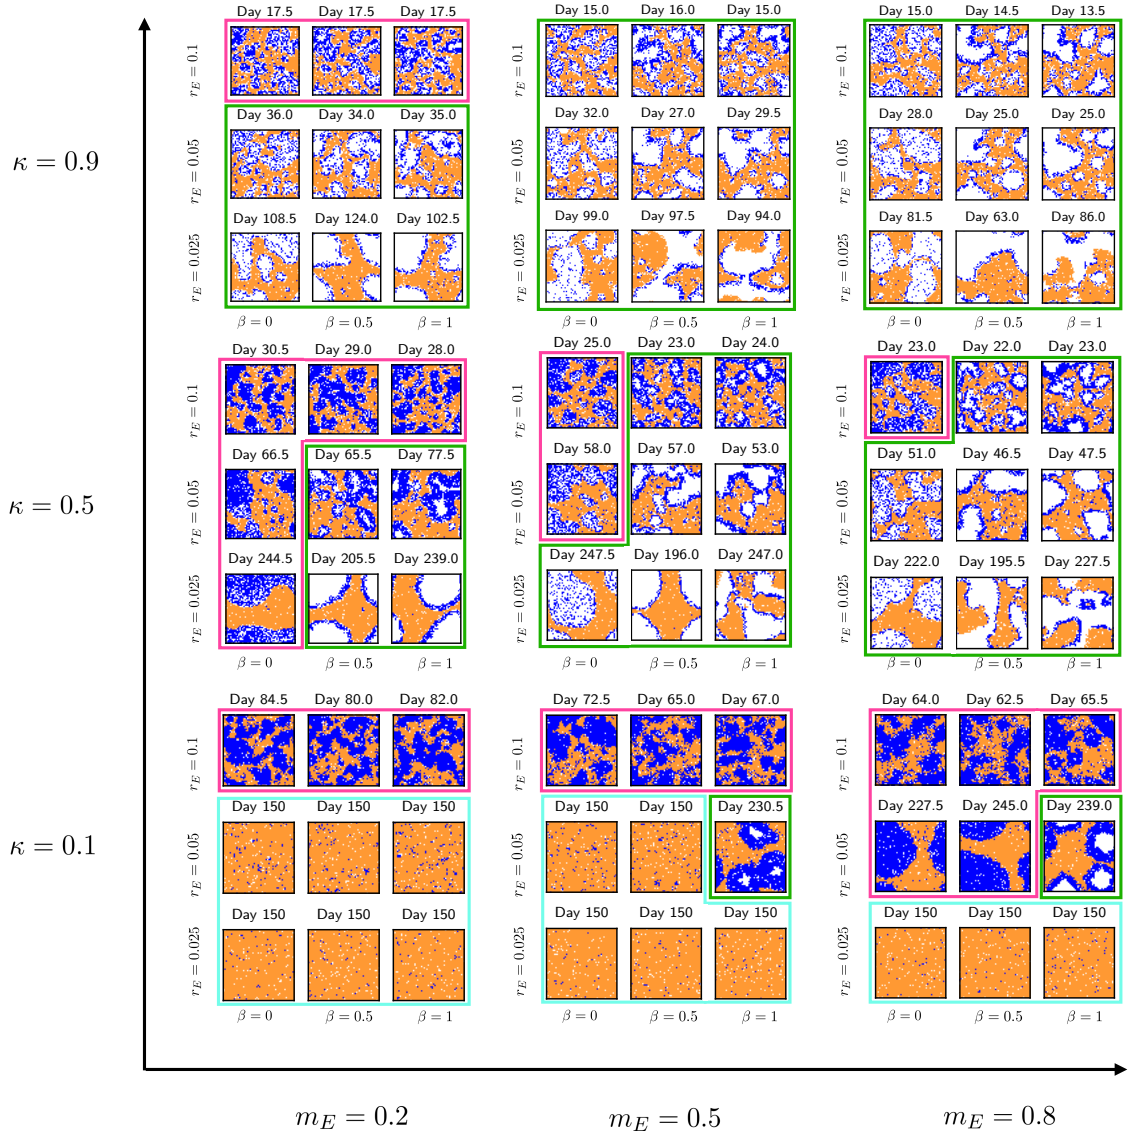


Figure 2.9: Snapshots of the distribution of tumour cells (orange), CTLs (blue) and empty space (white) for a single simulation, varying four of the CTL parameters; CTL migration probability m_E, CTL killing probability κ, CTL proliferation probability r_E and CTL random walk bias β. Snapshots are taken at 50% tumour burden, or, in the case that none of the 100 stochastic repeat simulations resulted in 50% reduction of the tumour population, a snapshot at 150 days is shown. Each snapshot shown was visually classified according to the presence of CTL clusters (pink boxes), CTL rings (green boxes) and CTL dispersion within the tumour tissue (cyan boxes).

aim is to examine qualitative differences in how the CTL population attempts to eliminate the tumour in different regions of parameter space, rather than comparing the dynamics at the same time points. As the timescale of the dynamics changes greatly as we sweep over parameter space, if spatial patterns were compared at the same time point it would not

be possible to visualise the full extent of the different behaviours that the model exhibits, i.e. at a fixed time point the tumour may have already been eliminated in some regions of parameter space, but in other regions the CTL population may not yet have expanded sufficiently to be able to start eliminating the tumour population.

The model shows three qualitative behaviours in the emergent CTL dynamics; rings, clusters and dispersion, marked in Figure 2.9. For low r_E and low κ, the CTLs are dispersed throughout the tissue and eventually die, resulting in tumour escape. For high r_E and low κ, the CTLs do not clear as much space in the domain through tumour killing; they dominate the grid through proliferation and form dense clusters. For high values of κ, ring structures are formed by the CTLs, with the thickness of the rings and CTL occupancy within the void regions depending on the values of r_E and β.

Due to the differences between our model and that of Hutchinson and Grimm (2022), it is difficult to directly compare the results from our parameter sweep with theirs. Our parameter sweep was conducted in order to identify regions of parameter space which merit further exploration. To this end, we identify regions of parameter space with $0.025 \leq r_E \leq 0.05$ and $\kappa \geq 0.1$ for further analysis due to the variety of qualitative behaviours observed. Furthermore, this region of parameter space excludes cases with very dense clustering of CTLs such that the tumour cells are only eliminated when the CTLs occupy the entire domain (for example when $\kappa = 0.1, m_E = 0.8, r_E = 0.05, \beta = 0$) which may not be biologically realistic.

We extended the spatial analysis of this parameter sweep by calculating the CTL-CTL and the TU-CTL PCFs, and the results are presented in Figures 2.10 and 2.11, respectively.

Both PCFs exhibit similar behaviour across the parameter space, with the CTL-CTL PCF peaking at short distances (corresponding to short-range clustering due to proliferation) and decreasing as the radial distance increases. Most of the CTL-CTL PCFs tend towards 1 as the radius increases, corresponding to the distribution expected under complete spatial randomness. By contrast, when there are large clusters of CTLs (for example $\kappa = 0.1, m_E = 0.8, r_E = 0.05, \beta = 0, 0.5$), $PCF = 1$ occurs at larger distances, and the PCF does not tend towards 1 as r increases. Snapshots with CTL rings show up in the CTL-CTL PCF as the PCF oscillating around the value 1. Oscillations in the PCF are most notable in parameter sets for which $r_E = 0.025$ and $\beta > 0$, where there is less CTL proliferation, reduced localisation of CTLs at the tumour margin, and therefore more empty space.

Like the CTL-CTL PCF, the TU-CTL PCF shows a similar signature across all parameter sets. Dispersion is found at short distances. As the tumour proliferation probability is much higher than the CTL proliferation probability, there is a higher chance of tumour cells co-localising with each other. The PCF increases as radial distance increases, and passes

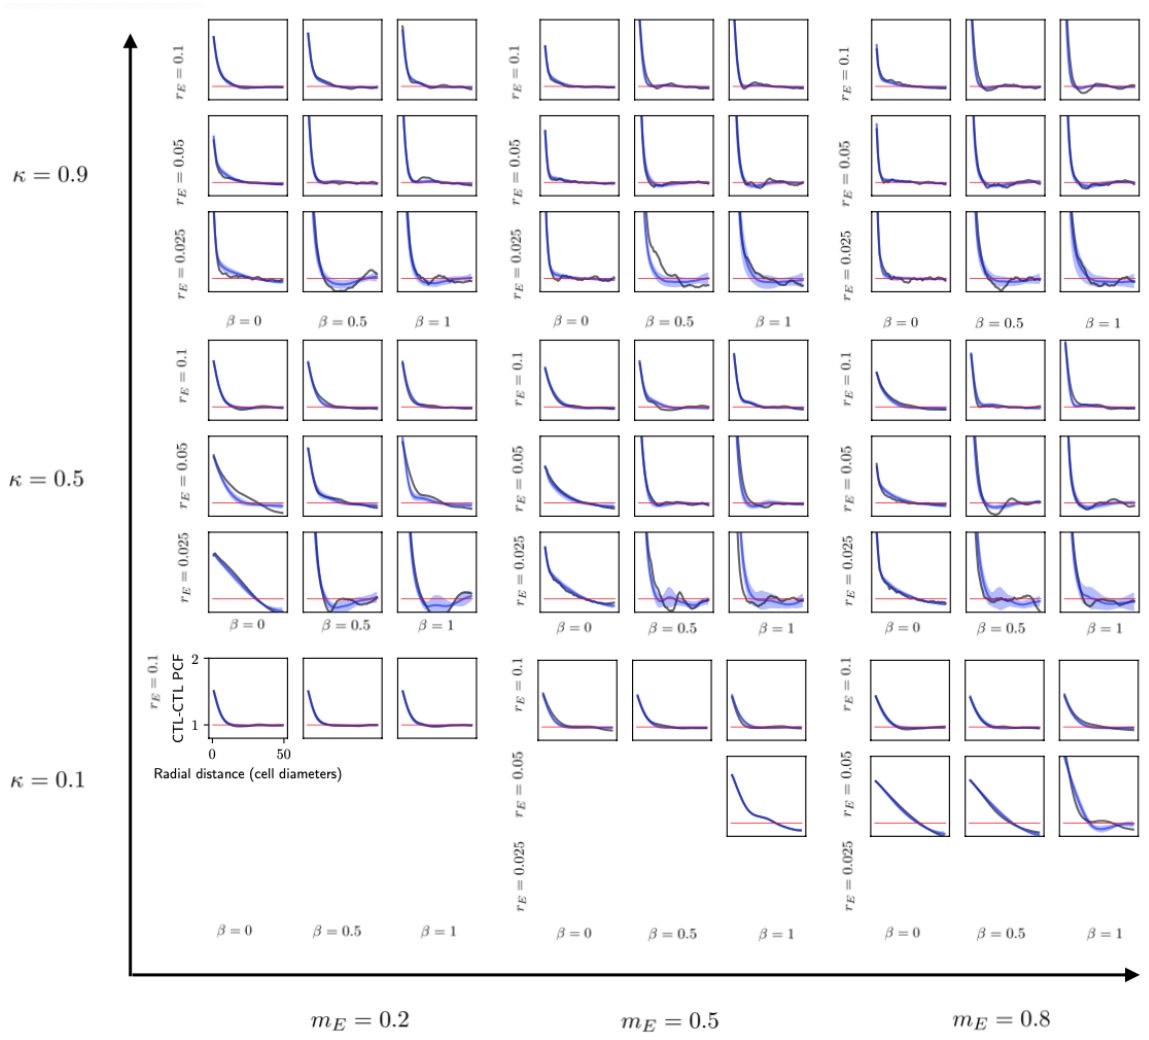


Figure 2.10: CTL-CTL PCF at 50% tumour burden, varying four of the CTL parameters; CTL migration probability m_E, CTL killing probability κ, CTL proliferation probability r_E and CTL random walk bias β. Averaged PCF over up to 100 stochastic realisations shown in blue, with standard deviation corresponding to the shaded region (for some parameter sets not all 100 realisations reduced tumour burden by 50% so were not analysed here). Red line indicates PCF=1. Black line shows individual PCF corresponding to the snapshot shown in Figure 2.9. For each individual plot, axis labels have been omitted for plotting clarity - vertical axis corresponds to the CTL-CTL PCF, and horizontal axis corresponds to Chebychev distance (cell diameters). The axes scalings are indicated in the bottom left plot. Regions of the parameter sweep with no plots correspond to the parameter sets which did not result in any of the model realisations exhibiting a 50% reduction of the tumour burden.

through the value of 1, indicating the average distance we start to see CTLs clustering. Generally this distance is greater for higher CTL proliferation probability and low CTL migration bias.

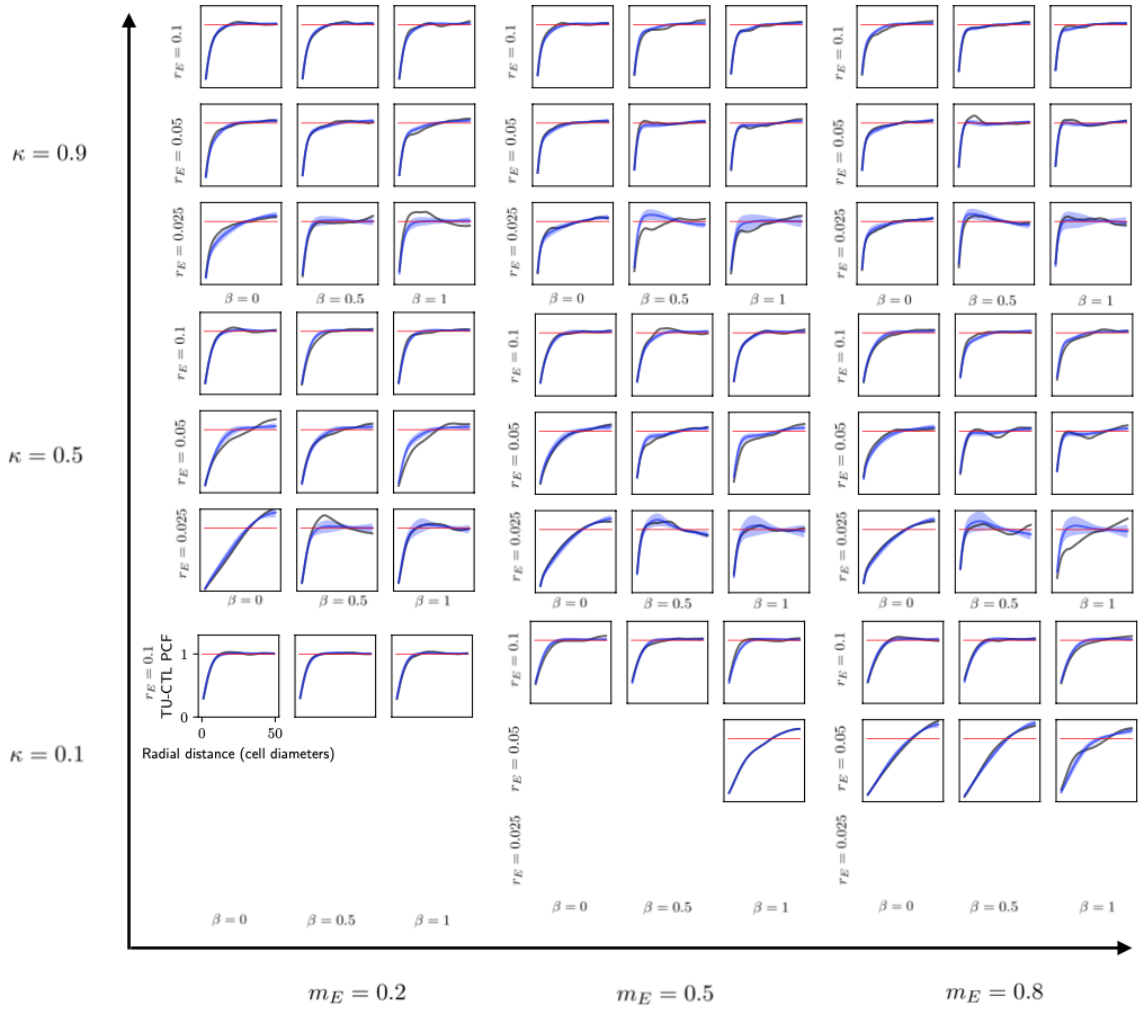


Figure 2.11: TU-CTL PCF at 50% tumour burden, varying four of the CTL parameters; CTL migration probability m_E, CTL killing probability κ, CTL proliferation probability r_E and CTL random walk bias β. Averaged PCF over up to 100 stochastic realisations shown in blue, with standard deviation corresponding to the shaded region (for some parameter sets not all 100 realisations reduced tumour burden by 50% so were not analysed here). Red line indicates PCF=1. Black line shows individual PCF corresponding to the snapshot shown in Figure 2.9. For each individual plot, axis labels have been omitted for plotting clarity - vertical axis corresponds to the TU-CTL PCF, and horizontal axis corresponds to Chebychev distance (cell diameters). The axes scalings are indicated in the bottom left plot. Regions of the parameter sweep with no plots correspond to the parameter sets which did not result in any of the model realisations exhibiting a 50% reduction of the tumour burden.

2.5.4 Pairwise parameter sweeps

From the analysis in Section 2.5.3, we have identified baseline parameter values to use in further analysis of the model - these values are given in Table 2.4.

Parameter	Description	Baseline value	Values considered in parameter sweep
r_T	Tumour cell proliferation probability	0.5	0.5
m_T	Tumour cell migration probability	0.35	0.35
d_T	Spontaneous tumour cell death probability	0.12	0.12
r_E	CTL proliferation probability	0.025	{0.02, 0.025, 0.03, 0.035, 0.04}
m_E	CTL migration probability	0.8	{0.1, 0.3, 0.5, 0.7, 0.9}
d_E	Spontaneous CTL death probability	0.015	0.015
κ	Probability of CTL killing of a tumour cell	0.9	{0.1, 0.3, 0.5, 0.7, 0.9}
β	Probability of bias in CTL random walk towards closest tumour cell	0.5	{0, 0.2, 0.4, 0.6, 0.8, 1}
η	Speed of CTL movement, corresponding to the number of lattice spaces moved across on a given time step	100	100

Table 2.4: Baseline parameter values and the range of values considered in the pairwise parameter sweep analysis.

Figure 2.12 shows the mean tumour elimination time t_{elim}, varying two CTL parameters at a time. These plots show that increasing CTL proliferation, migration and tumour killing decreases t_{elim}, as might be expected. There is, however, an apparent saturation effect as these parameters increase, for example as the CTL proliferation probability is increased, the incremental decrease in t_{elim} becomes smaller, suggesting that at higher values of r_E, there is limited benefit to increasing CTL proliferation further. This points to crowding effects in the domain which will be discussed later in this chapter.

In the region of parameter space of the baseline parameter values we have chosen, t_{elim} is very sensitive to the CTL kill probability κ, with tumour escape at low κ - Figure 2.13 shows the eliminated fraction of tumours over 100 stochastic realisations for each parameter set. Increasing each parameter results in a monotonic decrease in t_{elim}, except when increasing the CTL migration bias β, where non-monotonicity is seen in each of the three heatmaps showing varying β. In the next section we explore in more detail this optimal value β^* for minimising t_{elim}.

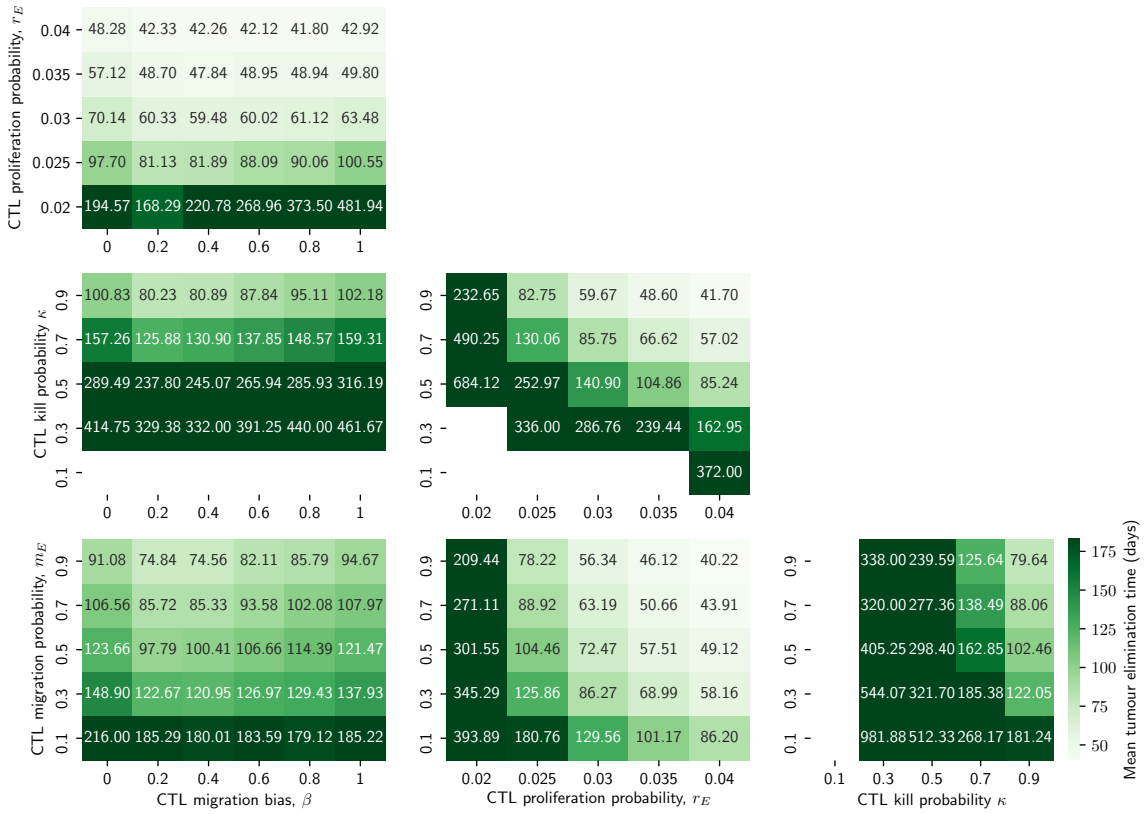


Figure 2.12: Mean tumour elimination time (days) for pairwise parameter sweeps. For each parameter combination, the model was run for 100 stochastic realisations, up to 1500 days. Heatmap colour indicates t_{elim} with an upper cut-off value of half a year. The mean t_{elim} was calculated for those stochastic realisations that resulted in tumour elimination. Empty spaces indicate parameter sets which did not result in tumour elimination for any realisations of the model.

2.5.5 Optimal value of CTL migration bias for tumour elimination - a trade-off between directed and random CTL motion

The histology data motivating the development of our model are from biopsies sampled in a Phase 1 clinical trial of an immunomodulator which promotes the localisation of CTLs within the tumour microenvironment, and stimulates CTL proliferation (Hutchinson and Grimm, 2022). In our model, these behavioural changes can be modelled by increasing the bias in the CTL migration (represented by the parameter β) and the CTL proliferation probability, r_E. Therefore we conduct a sweep across values of these parameters. Figure 2.14 shows how the mean tumour elimination time (days) changes as r_E and β vary.

As r_E is increased the tumour elimination time decreases, because as more CTLs are produced they will be more effective at killing tumour cells and occupy space within the domain. However, as β is increased, there is a non-monotonic effect on the tumour elimination

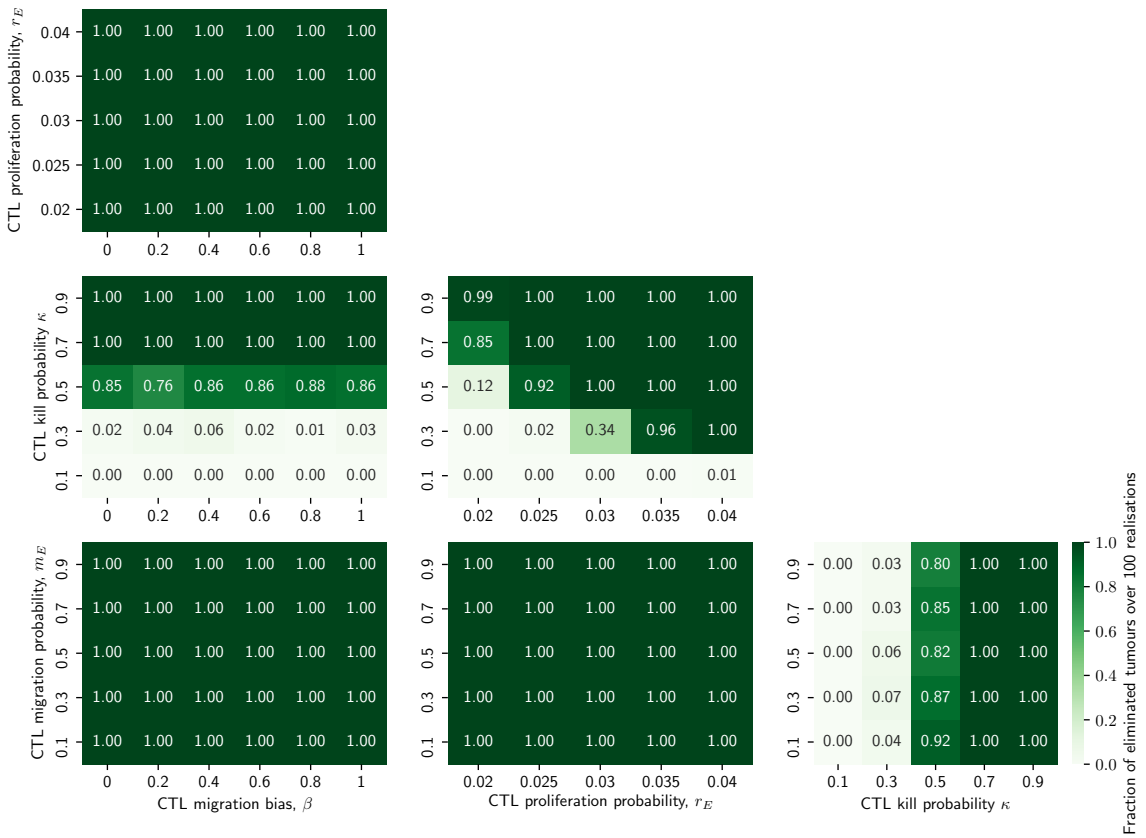


Figure 2.13: Fraction of eliminated tumours out of 100 stochastic realisations of the model, for pairwise parameter sweeps. The simulations were run up to 1500 days, and elimination is defined as tumour eradication within 1500 days.

time. For small values of β (i.e. as the bias increases from zero) the tumour elimination time decreases. However, at a critical value, β^*, the tumour elimination time reaches a minimum, increasing the bias above this value results in longer elimination times. The optimal value of β increases as the proliferation probability r_E increases. We note that the standard deviations shown by the error bars in Figure 2.14 overlap significantly, however the mean behaviour shows this non-monotonic dependence. From the cell counts alone it is unclear what causes this behaviour.

Figure 2.15 shows a timecourse of a single simulation for three values of the bias parameter β. In these simulations $r_E = 0.025$.

Visual inspection of the spatial distribution of the cells reveals that, for small values of r_E, fully directed motion (i.e. $\beta = 1$) causes CTLs to get stuck in high density regions at the tumour margin. Their proliferation rate also decreases due to crowding. If a break in the CTL ring appears, due to cell death, then the tumour can invade the empty space within the ring, resulting in transient oscillations and longer tumour elimination times. For

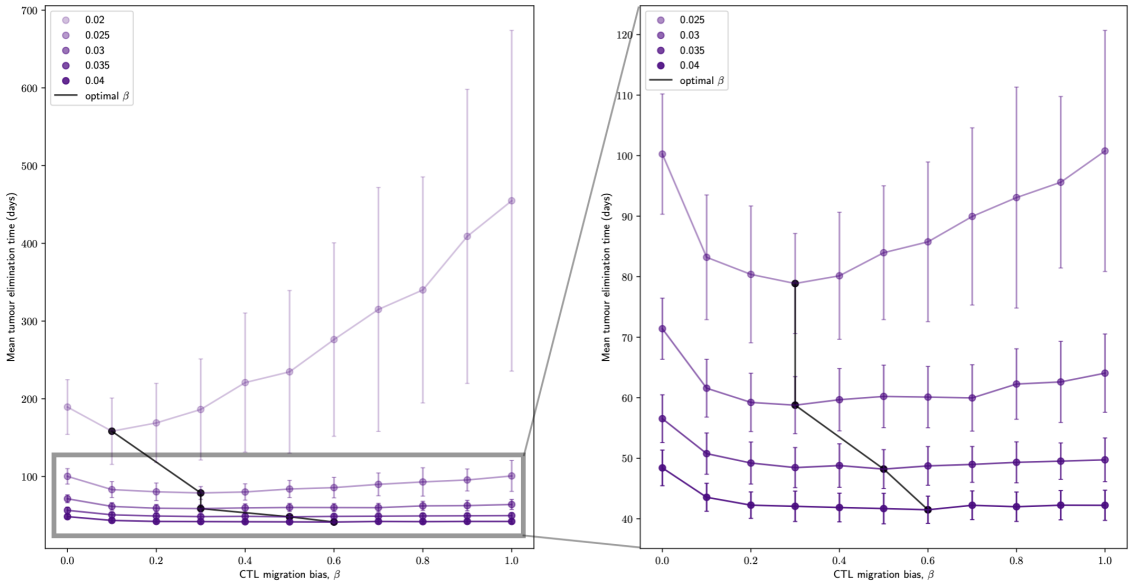


Figure 2.14: Mean tumour elimination time (days) as the bias in the CTL biased random walk varies. Error bars indicate the standard deviation over 100 repeat simulations for each parameter set. Baseline parameter values given in Table 2.4 are used. RHS plot is zoomed in version of LHS plot.

fully random CTL migration, CTLs that surround the tumour boundary may encounter a tumour cell by chance and halt their migration. CTLs may also move randomly within the voids, ready to attack any tumour cells which try to invade the empty space. However, when CTL movement is not biased, the CTL rings expand less rapidly than when they move via a combination of biased and random motion. In the latter case, a balance between directed and random movement is needed for the CTLs to minimise the tumour elimination time such that tumour cells that invade empty space are controlled and CTLs have free space in their Moore neighbourhood, into which they can proliferate.

At high r_E, CTL proliferation dominates migration and clusters of CTLs form (instead of the ring structures shown in Fig. 2.15). Increasing the migration bias at high values of r_E does not significantly alter the tumour elimination time.

In summary, our model predicts that, depending on the proliferation rate of the CTLs, a balance needs to be struck between directed motion towards the tumour cells, and random motion in order to minimise tumour elimination time. In particular, if the co-localisation of CTLs to tumour cells is too large, this may result in a worse outcome (i.e. longer time to tumour elimination).

On further examination of the case when $r_E = 0.025$ and $\beta = 1$, we observe some outliers in the tumour cell count data which have long elimination times. Figure 2.16 shows the timecourse for one of these outliers. We observe that the fully biased movement of the

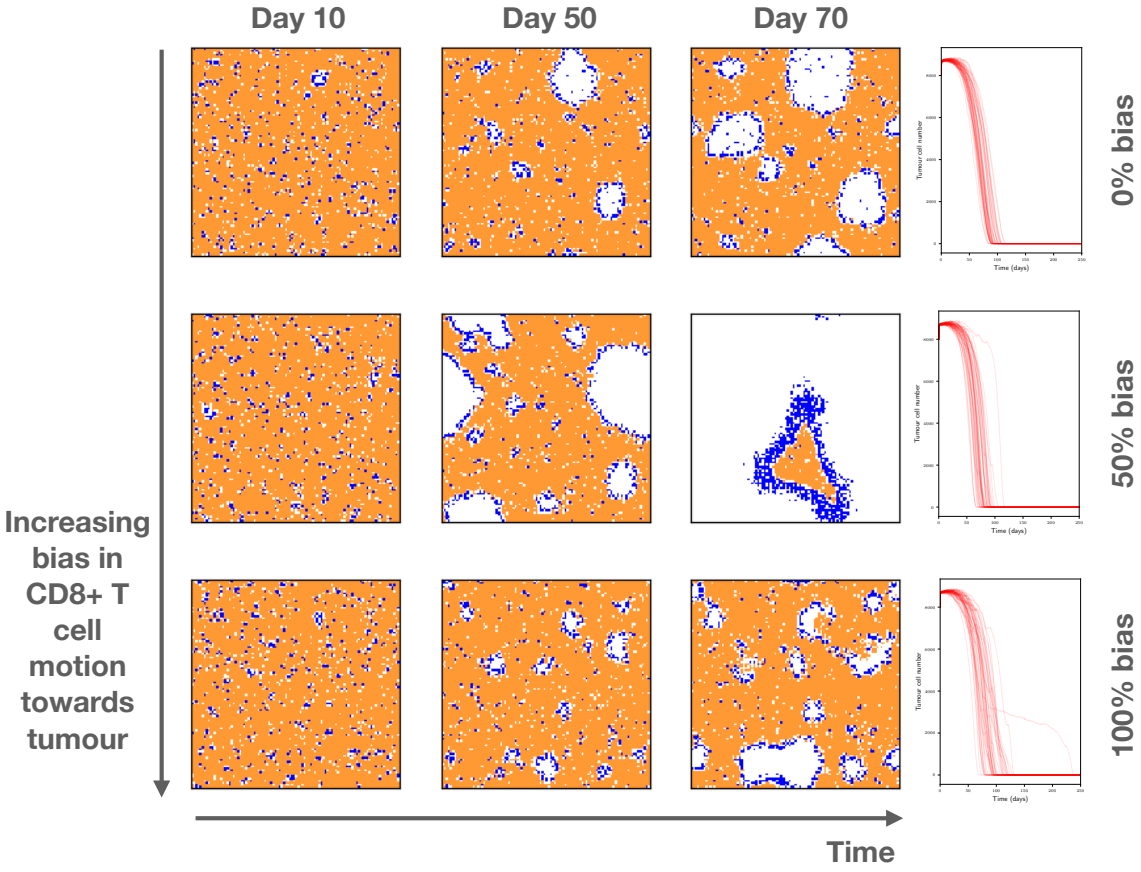


Figure 2.15: Series of snapshots showing how the cell distributions at times t = 10, 50 and 70 days change as the bias in CTL migration, β, varies. The RHS panels show the tumour trajectories for 50 repeat simulations for each of the three parameter sets.

CTLs, combined with the periodic boundary conditions, causes the tumour cells to evolve to a travelling wave, followed by the CTLs (in the simulation the tumour cells move to the right of the domain, and then appear on the left due to the boundary conditions). This causes a transient period of slow tumour killing as the CTLs are only able to access part of the tumour margin. This behaviour is not biologically realistic, and so in future analysis and integration with data, these outliers could be identified through the PCF analysis and omitted from the results.

To verify that the optimal bias result still holds under different boundary conditions, we performed the same analysis on the model under non-periodic (reflective) boundary conditions. Figure 2.17 shows that the mean tumour elimination time depends non-monotonically on the bias parameter when reflective boundary conditions are imposed.

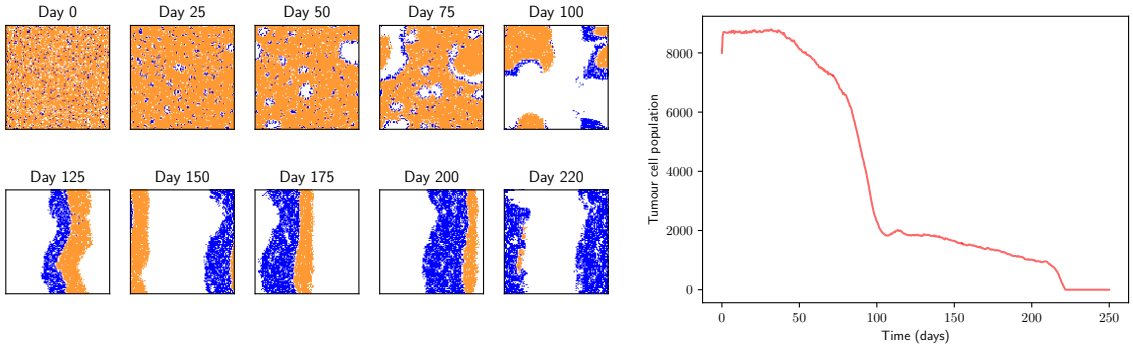


Figure 2.16: Example of dynamics exhibited in some stochastic realisations of the model with parameter set $\{\kappa = 0.9, m_E = 0.8, r_E = 0.025, \beta = 1\}$.

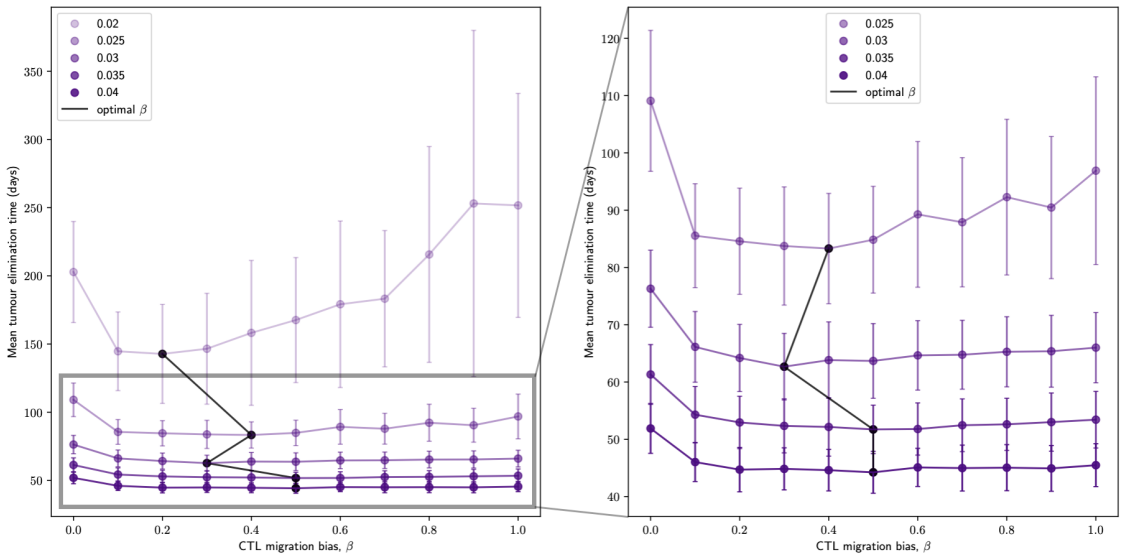


Figure 2.17: Mean tumour elimination time (days) as a function of the bias in the CTL biased random walk, with non-periodic (reflective) boundary conditions imposed at the domain edges. Error bars indicate the standard deviation over 100 repeat simulations for each parameter set. Baseline parameter values given in Table 2.4 are used. RHS plot is zoomed in version of LHS plot.

2.6 Exploring modelling assumptions - choice of time step

2.6.1 Overview

In our model, when β, the bias in the T cell motion, is increased, there is non-monotonic behaviour in the mean tumour elimination time, with the optimal value depending on the choice of the other model parameters. We also observed dynamic interactions between the tumour cells and CTLs, with the formation of ring structures in the T cell population in some regions of parameter space, and the invasion of tumour cells into these areas of empty space inside the rings, resulting in oscillatory dynamics in the tumour population.

When analysing the results from a model, it is important to consider the extent to which they are due to the biological mechanisms included in the model, or the modelling approach itself. In particular, the counter-intuitive results generated from this model could be due to implementation artefacts, since there are many ways of implementing an ABM. Therefore it is important to consider whether our results are robust to perturbations in the model implementation.

In this section we investigate whether the size of the timestep affects the model dynamics. The timestep procedure used to implement the ABM was taken from that used by Hutchinson and Grimm (2022) and originally proposed by Kather et al. (2017) model, in order to be consistent with their models, however we did not use the intermediate timesteps for CTL movement for simplicity, which we discussed earlier in Section 2.3.6.

2.6.2 Effect of time step choice on optimal bias in T cell migration

We expanded our analysis to investigate whether the value of the bias parameter which minimises time to tumour elimination is robust to changing the time step when the CTL proliferation probability r_E and CTL migration bias parameter β are varied.

The parameter set used in these simulations is given in Table 2.5. The same initial condition as in the previous analyses was used. For each timestep, the parameter values were scaled appropriately, e.g. for a timestep of 0.12hr, migrating CTLs moved one grid space per timestep, and the probabilities for proliferation for the 12hr timestep model were re-scaled, with the exception of the immune migration probability, which was not re-scaled as the distance the CTLs could move per timestep was scaled. The cell death rates were scaled by considering the probability that the cell would not die over a given period, shown in Table 2.5. We used a different scaling for the death probabilities as, compared to proliferation, death is a process that a cell undergoes only once. Appendix C contains more detailed reasoning for making these parameter scalings.

Parameter	Description	12hr	1.2hr	0.12hr
r_T	Tumour cell proliferation probability	0.5*	$(1+0.5)^{1/10}-1$	$(1+0.5)^{1/100}-1$
m_T	Tumour cell migration probability	0.35*	0.035	0.0035
d_T	Spontaneous tumour cell death probability	0.12*	$1-(1-0.12)^{1/10}$	$1-(1-0.12)^{1/100}$
r_E	CTL proliferation probability	0.025	$(1+0.025)^{1/10}-1$	$(1+0.025)^{1/100}-1$
m_E	CTL migration probability	0.8	0.8	0.8
d_E	Spontaneous CTL death probability	0.015*	$1-(1-0.015)^{1/10}$	$1-(1-0.015)^{1/100}$
κ	Probability of CTL killing of a tumour cell	0.9	0.09	0.009
β	Probability of bias in CTL random walk towards closest tumour cell	1	1	1
η	Speed of CTL movement, corresponding to the number of lattice spaces moved across on a given time step	100*	10	1

Table 2.5: Default parameter values for the CA model. Values marked * are taken from Kather et al. (2017). The value for η was chosen to be 100 instead of 97 as used by Kather et al. (2017), to aid scaling down to smaller timesteps. Values for the parameters r_E, m_E, κ and β can each be set within the range [0,1] for the 12hr default timestep case.

Figure 2.18 shows how the mean tumour elimination time changes as the bias β varies for two values of the CTL proliferation probability. The error bars indicate the standard deviation across 10 model realisations for three values of the timestep; 12 hours, 1.2 hours and 0.12 hours.

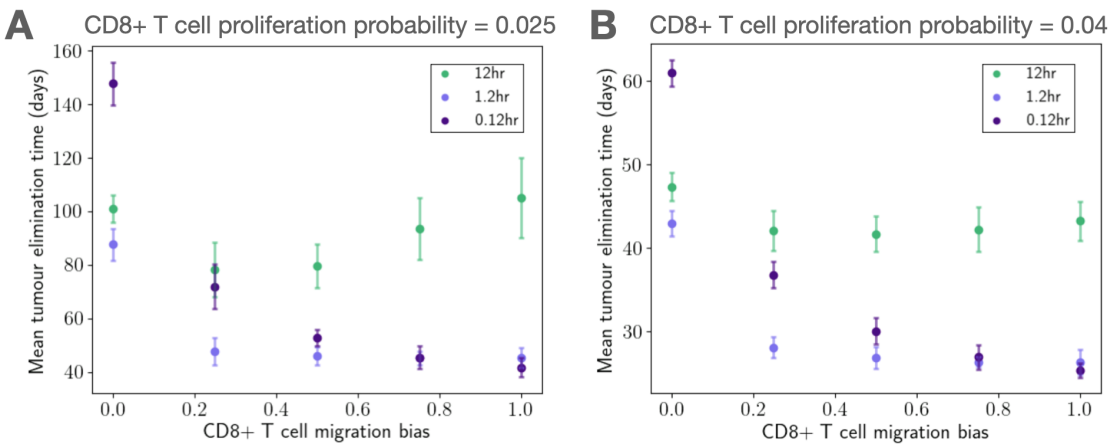


Figure 2.18: Sensitivity analysis, varying CTL proliferation probability and CTL migration bias towards tumour cells. Each plot shows how the mean tumour elimination time changes while increasing the bias parameter, for a fixed CTL proliferation probability. The error bars indicate the standard deviation across 10 realisations of the model. Simulations using the 12hr, 1.2hr and 0.12hr timesteps are shown in green, light purple and indigo respectively.

The same qualitative behaviour is seen across the timestep regimes, i.e. the tumour population is eliminated in all realisations of the ABM. However, the timescale on which the tumour is eliminated varies between time step regimes, and depends on the parameter set. At low CTL migration bias, the 1.2 hr timestep results in the fastest tumour elimination time. As the bias parameter is increased, a cross-over in the trajectories for the different timesteps is observed and at high bias values the small timesteps (1.2hr and 0.12hr) predict faster elimination times than the 12hr timestep.

Figure 2.18 shows that the 12hr and 1.2hr timesteps exhibit similar qualitative behaviours in the relationship between elimination time and bias, i.e.

for $t_{steps} \geq 1.2hr$, $\ni \beta^* \in [0,1] \quad s.t. \quad T_E(\beta^*) = min_{\beta \in [0,1]}\, T_E(\beta)$.

For the default 12hr timestep the mean tumour elimination time depends non-monotonically on the bias parameter. However, for the smallest timestep ($\Delta t = 0.12hr$), the qualitative behaviour changes; the tumour elimination time decreases with increasing bias. We conclude that the non-monotonic behaviour is not robust to changes in the timestep for the current model setup. We note that for a timestep of 1.2hr the elimination time asymptotes as $\beta \to 1$. This limiting value decreases as the CTL proliferation probability r_E increases, reflecting the fact that as r_E increases, CTLs form dense clumps, and the importance of migration for the co-localisation of CTLs and tumour cells becomes less important in tumour cells killing. In summary, from the point of view of drug design, both the 12hr and 1.2hr

timesteps predict that increasing the T cell biased movement above a certain threshold is not beneficial for tumour elimination. This threshold is larger for the 1.2hr timestep than for the 12hr timestep, and above this threshold the 12hr timestep predicts an increasing tumour elimination time, compared to no difference in the 1.2hr timestep. However, for $t_{step} = 0.12hr$, $T_E(\beta)$ decreases as β increases; behaviour we might expect to see when the biased movement of CTLs towards tumour cells increases. This suggests that the default time step of 12hr is too long, and that the simulation should be updated on the scale of the shortest time process (one CTL migration step), to recover expected behaviour in the tumour population when the CTL migration bias is increased.

However, before concluding that the time step should be adjusted, it is imperative to understand why the results change as the timestep varies. When there is no bias ($\beta = 0$), CTLs migrate via a random walk, with no biased motion towards tumour cells. In the 12hr case, a CTL may move up to 100 grid spaces in a timestep (or until it finds a tumour cell in its Moore neighbourhood). In the 1.2hr case, CTLs move up to 10 grid spaces, and in the 0.12hr case CTLs move one space. This means that in a given timestep a CTL covers more space in the 12hr step compared to the 1.2hr step. With the 1.2hr steps added up to 12 hours, there is a probability that a given CTL will be selected to migrate in all of the ten 1.2hr steps, given by $(P_{CTLmigration})^{10}$, which for this parameter set is $0.8^{10} = 0.107$. So, even though with the scaling down of parameter values the same proportion of CTLs will migrate compared to the 12hr timestep, the proportion of CTLs that migrate continuously for a given 12hr period is much less in the 1.2hr timestep. This suggests that in the absence of biased CTL motion, in this model it is advantageous to have a smaller proportion of cells that migrate a longer distance, rather than having cells able to migrate a shorter distance. Since the size of the clusters of CTLs are limited by the area of the domain they can explore, a larger timestep, with more directed migratory CTLs, allows the CTL rings to expand more rapidly, and to coalesce with other clusters sooner. In this way the CTLs are able to gain control over the tumour cells.

When biased motion is active (i.e. $\beta > 0$), a qualitative change occurs and the tumour is eliminated more rapidly on the shorter timestep. As the bias increases, CTL ring structures start to form, rather than clusters which emerge when $\beta = 0$. As CTL migration becomes more efficient, and directed towards the tumour, the elimination time reduces significantly. In this case, having a shorter timestep means more CTLs move over the course of 12hrs for timesteps of 0.12hr and 1.2hr than for the default 12hr timestep, and so tumour elimination occurs more rapidly.

In summary, we have found that by decreasing the time step used to implement the ABM we can recover different dependencies of the tumour elimination time on the CTL

migration bias parameter. For a short timestep, on the scale of one CTL migration step, we find that the tumour elimination time decreases monotonically with increasing bias. By contrast, for longer timesteps, CTL migration halts when a CTL reaches a tumour cell and attempts killing. Thus for a timestep of 12hrs most CTLs do not complete all of the 100 possible migration steps during the 12hr period. In contrast, for the 0.12hr timestep, a CTL will move one space, attempt to kill if there is a tumour cell in its neighbourhood, and then migrate again in the next 0.12hr timestep. By decreasing the time step we are not considering this "latent" time which effectively accounts for the engagement time for CTL killing of a tumour cell, which can be up to 6 hours in duration (Breart et al., 2008, Eisenhauer et al., 2009). Further work is needed to investigate whether such a large discrepancy in the results could arise between the different timesteps if the engagement period is accounted for in the 0.12hr timestep case.

Finally, we investigate if the implementation of the map used to determine the CTL biased motion affects the qualitative behaviour of the optimal bias when scaling down the timestep. We note that with the distance-based map, CTLs move in a biased way towards the nearest tumour cell, and the concentration of tumour cells is therefore not considered. In other models, a chemokine may be modelled using a PDE, in which case the density of the chemokine-expressing tumour cells would be taken into account in the CTL biased motion. However, to keep the model simple, we implement a Gaussian kernel over the tumour cell population, to produce an interaction map to implement the biased CTL motion. This implementation is detailed further in Appendix D. Figure 2.19 shows the corresponding results using the Gaussian implementation of the biased CTL motion. Here we see that the qualitative behaviour of an optimal CTL bias is conserved as the timestep is scaled down. Consequently, we find that decreasing the timestep causes a change in tumour outcomes for some stochastic realisations, with increasing number of tumour escape outcomes as the timestep is reduced. An explanation for this could be that this implementation draws CTLs into the densest regions of tumour cells, due to the increased signal in regions of high densities of tumour cells, thus the CTLs do not have space to proliferate. Whereas, in the distance map implementation, the CTLs simply move to the nearest tumour cell, which may not be in the most dense region of tumour cells.

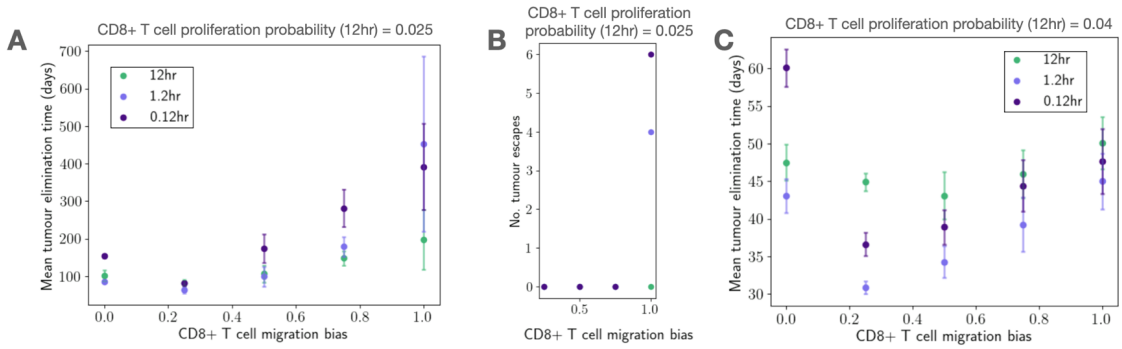


Figure 2.19: (A,C) Mean tumour elimination time with increasing bias parameter, for a fixed CTL proliferation probability (0.025 and 0.04 for the 12 hr timestep case). The error bars indicate the standard deviation across 10 realisations of the model. Simulations using the 12hr, 1.2hr and 0.12hr timesteps are shown in green, light purple and indigo respectively. (B) Number of tumour escapes for each timestep.

2.7 Discussion

In this chapter, we developed an on-lattice ABM of tumour-CTL interactions, based on models developed by Kather et al. (2017) and Hutchinson and Grimm (2022). The motivation behind this work was to develop a simple model, which captures the underlying mechanisms of the existing models. By simplifying the model, we could explore the effects of changing more than one parameter value, thus extending the work of Hutchinson and Grimm (2022).

We showed that our model exhibits the same qualitative behaviours (the 3 E's of immunoediting; tumour elimination, tumour escape and tumour-CTL coexistence) as those found by Hutchinson and Grimm (2022). By examining the spatial distribution of cells we identified the key mechanism for CTL elimination of the tumour in this model, namely the clustering and subsequent formation and expansion of rings of CTLs which kill tumour whilst expanding outwards, and eventually coalesce to eradicate the tumour population. In some regions of parameter space, the tumour was not eliminated within the simulated time, with the CTL population either becoming extinct (referred to here as tumour escape), or a CTL-tumour coexistence state forming. In some of these coexistence states, oscillatory dynamics in the tumour population were observed, which through spatial inspection of the model were identified as the product of the formation and breakage of the CTL rings.

We subsequently identified a baseline set of parameter values for our model from inspection of the model dynamics over a sweep of four CTL parameters. As the model was greatly simplified from the Hutchinson and Grimm (2022) model, we acknowledged that we could not directly relate the parameters of their model to our own. One approach that

could have been taken to parametrise our model would be to fit the results of our model to those of their model, for example using Bayesian methods. However, as the aim of the work in this chapter was to reproduce the qualitative behaviour seen in the Hutchinson and Grimm (2022) model and identify the range of dynamics that these ABMs exhibit, it was sufficient to manually select a region of parameter space to explore. Further work would involve expanding this analysis, for example conducting a global sensitivity analysis of the model (Marino et al., 2008).

A pairwise parameter sweep of four CTL parameters was conducted, to investigate how changing the properties of the CTL population affected the elimination of the tumour, for a specific (fixed) set of tumour parameters, emulating possible immunotherapies. The main finding of this analysis was the identification of a value of β, the bias in CTL migration, that minimises the tumour elimination time. We used our model to show further how β^*, the value of β at which the tumour elimination time was minimised, depended on other parameter values. We found that β^* increased as the CTL proliferation rate increased. This was because the crowding effects of CTLs at the tumour margin at high β values result in reduced CTL proliferation due to a lack of space. Randomly moving "surveillance" CTLs located within the empty space in the rings were found to be important in preventing tumour invasion of these empty spaces. This emergent behaviour was interesting as it seems to have both a biological and methodological underpinning. Biologically, having a trade-off between directed and random motion within the CTL population would be beneficial in terms of the surveillance strategy discussed above. The combination of random and chemotactic cell migration has also been investigated in other systems (Byrne et al., 1998, Byrne and Owen, 2004). However, it is also imperative to consider whether the behaviour we see in our model may be driven partly by the modelling framework and assumptions made. In our model we assume CTLs and tumour cells occupy the same amount of space, which is not biologically realistic; CTLs are smaller than tumour cells. We also neglect differences in the motion of the two cell types, for example CTLs are known to deform their shape and have the ability to "squeeze" between tumour cells. This is not possible in a single occupancy model. As a result the CTLs build up around the tumour margin and do not infiltrate into the tumour mass (unless a tumour cell dies and space becomes available).

The spatial patterns produced by the model, with CTLs building up around the tumour margin, are similar to those seen in other models in the literature, for example models developed by López et al. (2017*a*). They also are reminiscent of histology images of immune-excluded tumour regions (Tiwari et al., 2023). We conclude that we may not be able to reproduce all immune profiles using the model in its current form. It could be argued

that CAs may be appropriate for modelling *in vitro* setups, for example a cell monolayer, however the model still does not capture squeezing of CTLs past the tumour cells.

The observation of CTLs building up around the tumour margin points to the relaxation of the grid occupancy rules to allow more than one cell to occupy a lattice site. It would also be interesting to investigate whether the qualitative behaviours exhibited by our model persist if we use an off-lattice approach. This also means that in our current state of the model the killing rate is limited to the CTL coverage of the tumour margin rather than the number of CTLs, which you might expect.

We must acknowledge that although our model was inspired by Hutchinson and Grimm (2022), and produces the oscillatory and bistable dynamics that they noted (personal communication), further work is needed to ascertain whether this is a faithful simplification, such that it can be concluded whether these results are relevant to the current models in the literature.

In particular, the separation of timescales between CTL migration (short timescale) and tumour cell proliferation (long timescale) is treated differently in each model. In Hutchinson and Grimm (2022), intermediate CTL migration steps are used, where, since each CTL moves a single step, they are more "well-mixed" in their migration. By contrast in our model a CTL will move its full migration path before the next CTL migrates. We chose this method so that the proliferation and death events would be randomly implemented between migration events, instead of being conducted en masse after the CTL migration events. We also made this decision to account for the fact that during one tumour proliferation event, a CTL may move more than one step due to the difference in timescales of these two processes. It could be argued that the timestep should be chosen to match the timescale of the most rapid process in the model; in this case one CTL migration step. This was explored in Section 2.6 where we rescaled the probability parameters and reduced the timestep by a factor of 100. This change resulted in a monotonic decrease in the mean tumour elimination time when increasing β, compared to the non-monotonic dependence found using the default 12hr timestep. Although this result may suggest that the 12hr timestep is too large to achieve convergence of the numerical scheme, we also must consider that reducing the timestep is not only algorithmic but is also effectively modelling different biological behaviours. When we use the 0.12hr timestep, more cells move shorter distances (1 space) compared to the 12hr timestep, where very few cells move the maximum 100 spaces per timestep. Furthermore, as proliferation and death events are assumed to be instantaneous in the model (cells appear or are deleted instantaneously), the timescale of proliferation, death and killing events is not considered. This could be considered in the 0.12hr timestep model by incorporating an engagement time (as in Iarosz et al. (2011)), or

allowing a probability that cells can occupy the same lattice site as a dying cell, with this probability becoming increasingly larger as the time since death increases, until the time period of cell death is complete.

In future work, we could reintroduce some of the biological details that we omitted for simplicity and due to lack of data. For example, CTL influx is an important process but was omitted to simplify the model and because no suitable data are available to guide parameter estimation. Nevertheless, our preliminary simulations suggest that this process could seed more rings/clusters which may increase the probability of a tumour elimination event for those regions of parameter space that were mainly/all tumour escape. Crowding effects are also important here as an influxing CTL would only be able to influx into an empty space in the domain - in some regions of parameter space there would be high probability that this would be in the region of parameter space where the model generates rings/empty voids. We hypothesise that in those regions of parameter space, CTL influx may not contribute significantly to the time to tumour elimination, and may have the same effect as simply increasing the CTL proliferation probability. Further studies could investigate the effect on tumour outcomes of changing the initial spatial distribution of cells in the domain, for example seeding the initial condition with clusters of CTLs. We could also explore heterogeniety within the CTL population, for example CTL sub-populations with different killing rates (Vasconcelos et al., 2015), or how T cell exhaustion affects the spatial dynamics.

In summary, this chapter has presented a simple CA model of tumour-CTL interactions, inspired by Hutchinson and Grimm (2022). We were able to reproduce behaviours seen in their model, and explain the mechanisms behind those behaviours using spatial analysis. We further highlighted the need to identify how the modelling assumptions, as well as the biological mechanisms can contribute to emergent behaviours, and propose suggestions of model adaptations which may allow us to capture the biology more realistically.

Chapter 3

Spatial analysis of tumour biopsy histology images

3.1 Overview

Cytotoxic (CD8+) T cells form an important part of the adaptive immune response. They kill infected and abnormal cells, including cancer cells, via a process known as immunosurveillance. When CD8+ T cells infiltrate tumour tissue, they identify antigens on the surface of tumour cells, and begin tumour cell lysis. Immunotherapies are increasingly being developed to target aspects of the immune system, such as CD8+ T cells, in order to boost its anti-tumour response. Cancer patients in clinical trials are monitored in different ways, including biopsy sampling. Core needle biopsies are invasive, and therefore this procedure is usually limited to the collection of two samples; one biopsy is taken before treatment starts, and the other during treatment (termed in this report as "pre-treatment" and "on-treatment", respectively). The on-treatment biopsy is typically taken a few weeks into treatment to determine whether the drug is having an effect within the tumour. As only one on-treatment biopsy can be taken from a patient, it is imperative that as much information as possible is extracted from these data. At present, clinical studies typically base drug scheduling decisions and conclusions about drug efficacy on quantities such as cell counts extracted from images of the biopsy samples. As such, they do not exploit all of the information available in the data. This discrepancy motivates the work presented in this chapter, in which we analyse a dataset of paired biopsy images taken during Phase I clinical trials of Simlukafusp (FAP-IL2v), an immunomodulatory molecule which localises within the tumour microenvironment, and stimulates IL2β/γ receptors on NK cells and CD8+ T cells. There are a range of primary indications and biopsy locations within the data set, and treatment protocols (for which the data are not available) varied between patients.

In this chapter, we apply existing spatial methods (Bull and Byrne, 2023) to the core-needle biopsy dataset in order to assess their suitability for analysing the data.

Through our analysis, we aim to address the following research questions:

- Do we see any trends in the pre-treatment and on-treatment biopsies?
- What information can we gain from clustering the patients using spatial statistics, as opposed to cell counts?
- Do the spatial statistics shed light on what effect the drug is having?
- What spatial biomarkers or features are most significant in distinguishing patients?
- What limitations are there in applying this method to the data?
- Can we predict how a patient might respond to the treatment by analysing metrics extracted from the pre-treatment biopsy?
- Do patients with different primary indications respond differently to the treatment?

The chapter is structured as follows. We summarise the biological background relating to Simlukafusp in Section 3.2. In Section 3.3 we present the data set, and in Section 3.4 we introduce the methods used to process and analyse it. We present our results in Section 3.5, and our discussion and conclusions in Section 3.9.

3.2 Simlukafusp

Simlukafusp is an immunomodulatory molecule that consists of an antibody which binds to FAP (fibroblast activation protein), and an IL-2 variant which stimulates IL-2β/γ receptors on immune cells (Hutchinson and Grimm, 2022). FAP is highly expressed by cancer-associated fibroblasts (CAFs) (Bughda et al., 2021), which are spindle-shaped cells found within or surrounding the tumour. They are abundant in tumour stroma, and provide physical support to the tumour cells (Liu et al., 2019). FAP may also be expressed by tumour cells themselves, depending on the cancer type, and in these cancers this FAP-expression by tumour cells is associated with worse prognosis (increased metastasis and tumour invasion), compared to cancers where FAP is only expressed within the stroma (Liu et al., 2015, Zhang et al., 2023). The antibody binds to FAP-expressing cells within the tumour micro-environment, which localises IL2 in these regions. The IL-2 variant preferentially stimulates natural killer (NK) and CD8+ T cells over regulatory CD4+ T cells (Siebert et al., 2022), causing the NK and CD8+ cells to proliferate (Waldhauer et al., 2021).

Our analysis aims to use the biopsy data to determine if the treatment performs in practise as it is theorised to, thus suggesting a positive patient response. Hence we seek to find evidence for:

- An increase in the density of proliferating CD8+ T cells between the pre- and on-treatment biopsies, and therefore an increase in the total number of CD8+ T cells (proliferating and non-proliferating).
- An increase in local clustering of proliferating CD8+ T cells due to drug localisation.
- A decrease in the density of proliferating tumour cells, due to cytotoxic events initiated by NK cells and CD8+ T cells.

3.3 The data set

The data set was provided by our collaborators at Roche, comprising spatial coordinates and cell types contained in paired pre- and on-treatment histology images of biopsies sampled from patients from Phase 1 clinical trials for Simlukafusp. The raw images had already been processed by our collaborators using Roche image processing algorithms to produce $\{x, y\}$ for $x, y \in \mathbb{Z}$ coordinates and cell type classification for each cell. Primary indications and locations sampled for the pre- and on-treatment biopsies, with the time between biopsies, are given in Appendix E, where available. Information regarding the dosing schedule for each patient was not available, however each patient received several drug infusions before the on-treatment biopsy sample was taken (the number of infusions may have varied patient-to-patient). Information about what treatments patients received before the trial were not available. The primary indications, biopsy locations and on-treatment biopsy sampling times vary across the patient cohort.

We were provided with data from 71 patients. Biopsy samples were collected via a core needle biopsy, sliced and mounted on plates. The biopsy tissue was stained using Ki67 and CD8 markers, to quantify proliferation and CD8 expression respectively, and locate the positions of the following cells:

- Proliferating tumour cells (Ki67+ CD8-)
- Proliferating CD8+ T cells (Ki67+ CD8+)
- Non-proliferating CD8+ T cells (Ki67- CD8+).

The Ki67 protein is present in cell cycle phases G1, S, G2 and M (i.e. cells that are actively progressing through the cell cycle), but not in quiescent cells in phase G0 (Li et al., 2015, Scholzen and Gerdes, 2000). The biopsy images from the IRIS data set at Roche were digitised and cell types and spatial coordinates provided to us. The data set also contains pathologist annotations of the biopsy images, denoting the tumour, tissue, control,

necrosis and exclude regions. These annotations are discussed in Section 3.4.2 and can be summarised as follows:

- *Tissue* annotations mark the boundary of the tissue extracted in the biopsy procedure.
- *Tumour* annotations mark regions of this tissue that have been classified as tumour by the pathologist.
- *Control* annotations denote the position of the control tissue on the slide; this control tissue is used to verify that the stains are working.
- *Exclude* annotations denote tissue areas that have been folded or distorted during sampling and plating. As such they are not reliable and excluded from the analysis.
- *Necrosis* annotations mark the boundary of necrotic regions within the tissue.

An example of processed images of two paired biopsies from a single patient is shown in Figure 3.1; cell centre positions and pathologist annotations are also included. The tumour annotations are typically more broad-brushed than the other annotations, indicating the general region of the tumour, without detailed boundaries, and may contain stromal regions.

Four patients in the data set had two on-treatment biopsy samples. It is unknown if these samples were taken at the same time, or from the same location. In this chapter, we pool the on-treatment data from the two on-treatment samples. Further work could be undertaken to examine if there are differences in the two on-treatment samples.

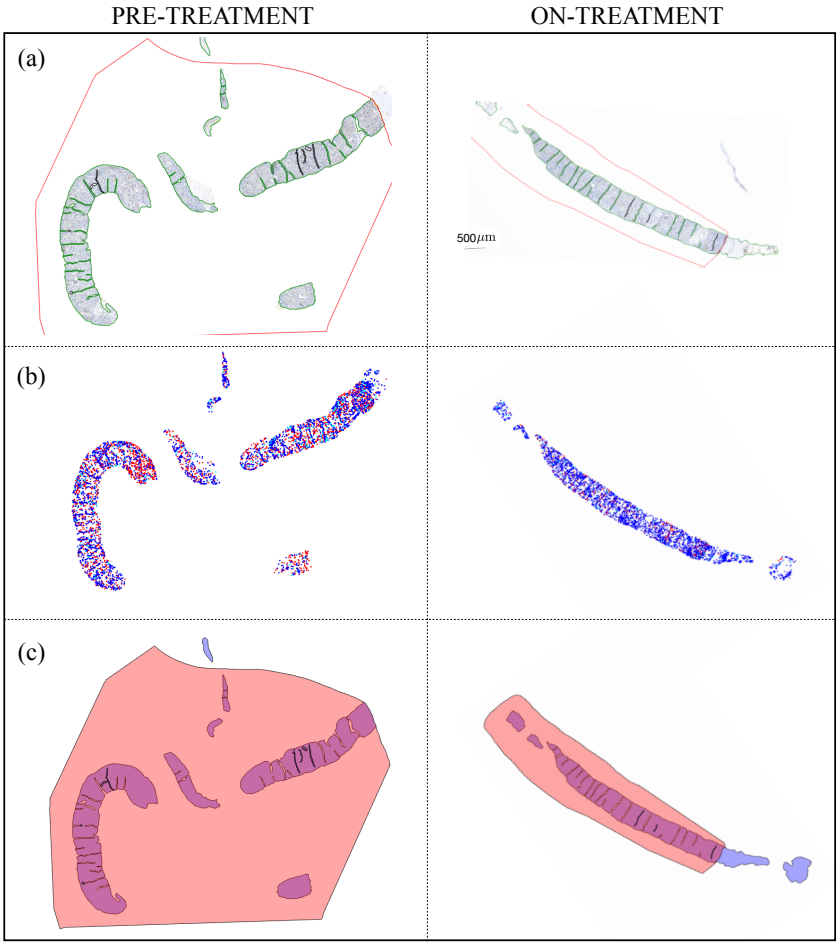


Figure 3.1: Paired data from a typical patient showing pre-treatment and on-treatment biopsies. (a) Stained biopsy tissue samples. Red outline denotes the tumour annotation, green line denotes the tissue annotations and black lines denote the exclude annotations. White space in the images are empty areas of the slide and do not contain any tissue. (b) Digitised images, showing cell centres identified from the histology stains. Key: proliferating (Ki67 stained) tumour cells in red; non-proliferating CD8+ T cells in dark blue; proliferating CD8+ T cells in cyan. (c) Pathologist annotation data: tumour in red shaded areas; tissue in blue shaded areas; exclude in black shaded areas. The tumour annotations describe regions of the biopsy tissue that have been classified as tumour by the pathologist. These are the only regions considered in the analysis of the biopsy data. Raw images in (a) provided by L. Hutchinson and O. Grimm, from the IRIS data set at Roche.

3.4 Methodology

3.4.1 Previous analysis of the Simlukafusp data set

Hutchinson and Grimm (2022) also analysed the data set used in this chapter. They mapped the data onto a lattice and sampled the images using square tiles. The mean and standard deviation of the CD8+ T cell count over the tiles were recorded for pre- and on-treatment

biopsies for each patient. The number of tiles sampled from each image varied depending on the size of the tissue sample. Tiles were sampled with a specified tolerance for empty space. The tiles sampled from on-treatment images were also analysed spatially using the on-lattice correction for the PCF (Gavagnin et al., 2018), for comparison with the output of an ABM used to simulate the tumour's growth dynamics during the time between the two biopsies in order to predict an optimal scheduling of the on-treatment biopsy. In this chapter we analyse the raw point data of the cell locations (pre-scaling onto a lattice), as some data can be lost during the scaling onto lattice process with single cell occupancy rules, and we are interested in whether we can identify any trends or patterns in the data without using modelling.

3.4.2 Inclusion of annotation data

Pathologist annotation data were incorporated into the image processing pipeline to increase automation. The annotations denote the tumour, tissue, exclude and control boundaries. Control tissue present in the image was used to corroborate effective staining of the markers, and was not relevant to the analysis of the biopsy tissue. The tissue annotations were extremely detailed and traced the boundary of the tissue in the biopsy. The tumour annotations provided a coarser description of the outlines of the tumour margin. The exclude annotations denoted areas of the biopsy tissue that had been deformed or damaged during collection and plating of the sample. For example, the tissue edges are often deformed during this process, and folds in the tissue can also emerge. The annotations were used to develop the image processing pipeline as follows. First, the coordinates of the tumour annotations were used to automate the removal of the control tissue from the image. As the tumour annotations were less detailed than the tissue annotations, the area of tissue occupied by the tumour was defined as the intersection of the tumour and tissue annotations. It was calculated using Matlab polyshape functions, and used to delete all cells not contained within the tumour-tissue region. Cells within the exclude and necrotic regions were also deleted and not considered in subsequent sampling and analysis (see Figure 3.2).

3.4.3 Data preparation

The images in the data set vary in terms of the tissue sample size, the degree of tissue fragmentation, and the cell density. The data were analysed to filter out any samples which would not be suitable for analysis. One of the main reasons for removal from the data set was low cell density in the ROIs sampled from a given image. Other reasons for patient removal from the data set included: lack of pre- and on-treatment samples, no tumour annotation data, and a tissue sample being too small to contain an ROI without inclusion

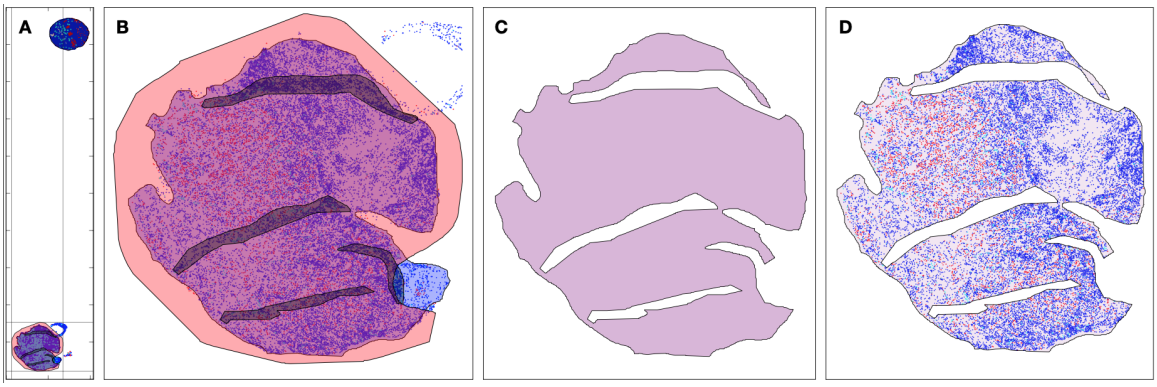


Figure 3.2: Annotation data were integrated into the image processing pipeline to identify regions of interest. (A) Digitised image of the biopsy tissue and control tissue for a single patient. Biopsy tissue is located in the bottom left of the image and control tissue in the top right of the image. The tumour annotation is used to trim the image (removing the control tissue which is not required). Trim lines are plotted in grey. (B) Trimmed image of biopsy tissue (zoomed in on lower tissue piece in (A)), showing cell locations as points and annotation data. Key: Proliferating tumour cells shown as red points, proliferating CD8+ T cells shown as cyan points, non-proliferating CD8+ T cells shown as dark blue points, tumour annotation shown in red shading, tissue annotation in blue shading, exclude annotations in grey shading. (C) Intersection of tumour, tissue and exclude annotations. (D) Cells located outside of the intersection annotation shown in (C) are deleted.

of tissue edges. Figure 3.3 shows the distribution of total cell number (sum of tumour cells and CD8+ cells within the extended ROI).

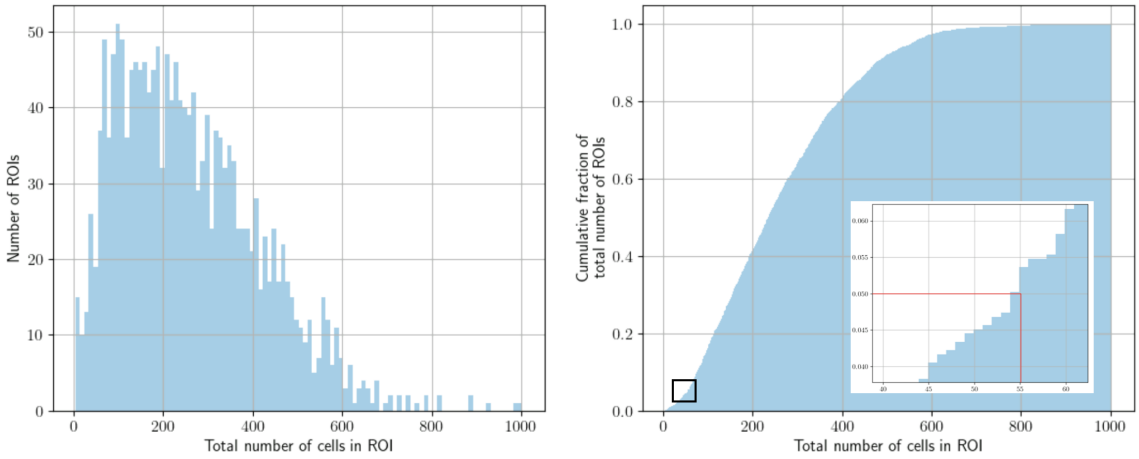


Figure 3.3: (A) Histogram of total number of cells in each ROI. (B) Cumulative fraction of total number of ROIs as a function of total number of cells in each ROI.

A minimum total cell threshold was set at 55 cells, thereby discarding approximately 5% of ROIs (cumulative fraction less than 0.05 in Figure 3.3B). This decision was taken due to lack of necrotic annotations in some of the images. It is unknown whether these

annotations were not available because there were no necrotic regions identified in these images, or whether they were missing from the data set. Figure 3.4 highlights an example of an image which did not have a corresponding necrosis annotation, but which contained regions of no/low cell numbers.

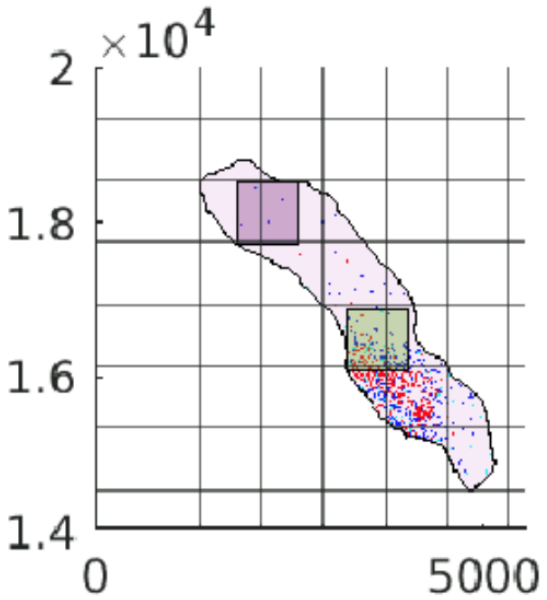


Figure 3.4: Example of low density ROI sampled from an image (top left ROI) which is removed during the data processing (patient 36, on-treatment). Axes denote image distance units, the coordinates used to define the cell locations.

Maintaining the ROIs sampled from these regions, whilst sampling as much of the image area as possible, may not be practical for analysis of the treatment action, and may have the potential to confound results, as one could reasonably assume that these were in fact necrotic regions that had not been marked. Application of these methods on future data sets would require clear annotations and complete data sets such that these assumptions may be rendered obsolete.

The number of patients filtered out at this stage, and the reason for their removal, are given in Table 3.1.

Table 3.1: Quality control of dataset - reasons for removal and number of patients.

Category	**Number of patients**
Pre- or on-treatment data unavailable	4
No ROIs sampled from pre- and/or on-treatment image due to small tissue sample size	8

Table 3.1 – continued from previous page

Category	Number of patients
All ROIs from an image excluded due to low cell number	2
Patients kept in dataset to analyse	57

3.4.4 Length-scale verification

We were not provided with the pixel-distance conversion factor for these images. In order to compare statistics from different patients we first confirmed that the images used the same lengthscale. We calculated the minimum distance between proliferating tumour cell centres, illustrated in Figure 3.5. Reasoning that these cells would be tightly packed, we estimated that the minimum distance between these cells should be approximately one cell diameter.

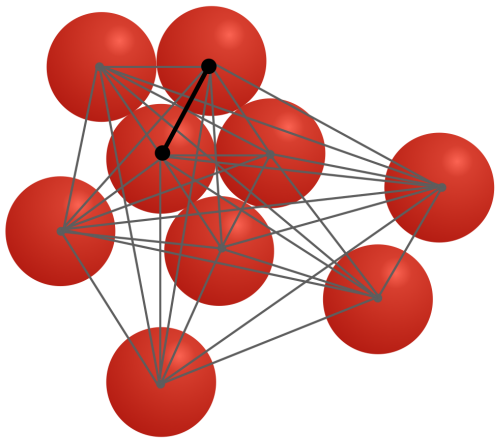


Figure 3.5: Schematic of lengthscale identification process. The nodes represent the cell coordinates provided in the data, and an illustration of the Ki67+ stained tumour cells shown by red spheres. In dense regions of the tumour, we may expect tight packing of tumour cells. Distances between pairs of tumour cells are recorded (shown by grey lines) and the minimum of these distances (black line) assumed to be equal to one tumour cell diameter.

The minimum distance between proliferating tumour cells is approximately 6 distance units across the set of images. A few images had a smaller minimum distance which, upon further inspection, revealed them to be likely outliers, caused by a classification error, cell morphology deviances, or actively dividing cells which have not completed mitosis. Based on this study, we can proceed with the analysis over the patient cohort, assuming that

lengthscales are consistent between images, and that each unit distance is approximately equivalent to one sixth of a tumour cell diameter.

3.4.5 Summary statistics

We use scalar and spatial metrics to analyse the image data. The scalar metrics are the densities of each cell type. We use the pair correlation function (PCF) to characterise the spatial information contained within the biopsy images, particularly the patterns of T cell infiltration. When calculating the PCF, boundary effects may be important. Periodic boundary conditions may be employed, or the normalisation area corrected for cells at the edges by extrapolating the expected number of cells in an annulus using the number of cells contained in the fraction of the annulus that lies within the tissue region - this will be discussed in Section 3.4.6. When calculating the PCF for each image, several considerations need to be made as some images consisted of a single piece of tissue, and others contained several tissue fragments. These are discussed in Section 3.4.6.

3.4.6 Image sampling

After deciding which metrics we will use, we must next decide how to apply them to the data. The core needle biopsy samples are typically irregular, long, thin pieces of tissue. Many images comprise multiple small, fragmented tissue pieces. Many images also contain "exclude" regions. As many samples are small, we wish to maximise the amount of data we can analyse, whilst minimising the effects of artefacts. Figure 3.6 shows different approaches for analysing the data.

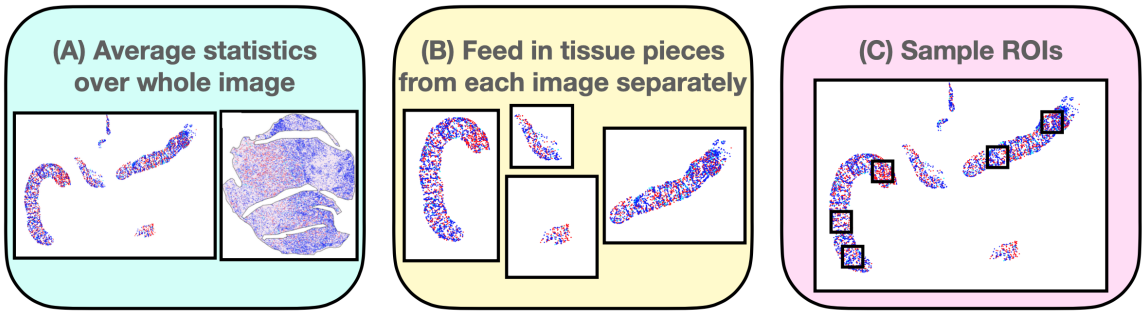


Figure 3.6: Schematic of three possible approaches for analysing the histology data set. ROI = Region of Interest.

Approach A involves calculating statistics on all cells contained within the tissue-tumour-exclude intersection. For images containing more than one piece of tissue, the average cell density over all pieces is calculated. The PCF is calculated for each tissue piece, and weighted according to the number of cells in the tissue pieces. Approach A uses

all the available data, which is important for small samples. However, this approach requires consideration of the tissue boundaries. While edge effects can be corrected for by extrapolating the number of expected points at a given radius (see Figure 3.7), this assumes that the spatial distribution of the cells in the surrounding tissue is radially symmetric. Consequently this correction to the PCF depends on the length of the perimeter of the tissue, which will vary between images. We are also not capturing the intra-patient heterogeneity within the sample, for example, regions with different levels of CD8+ T cell infiltration, as any intra-sample heterogeneity is averaged out in the PCF calculation.

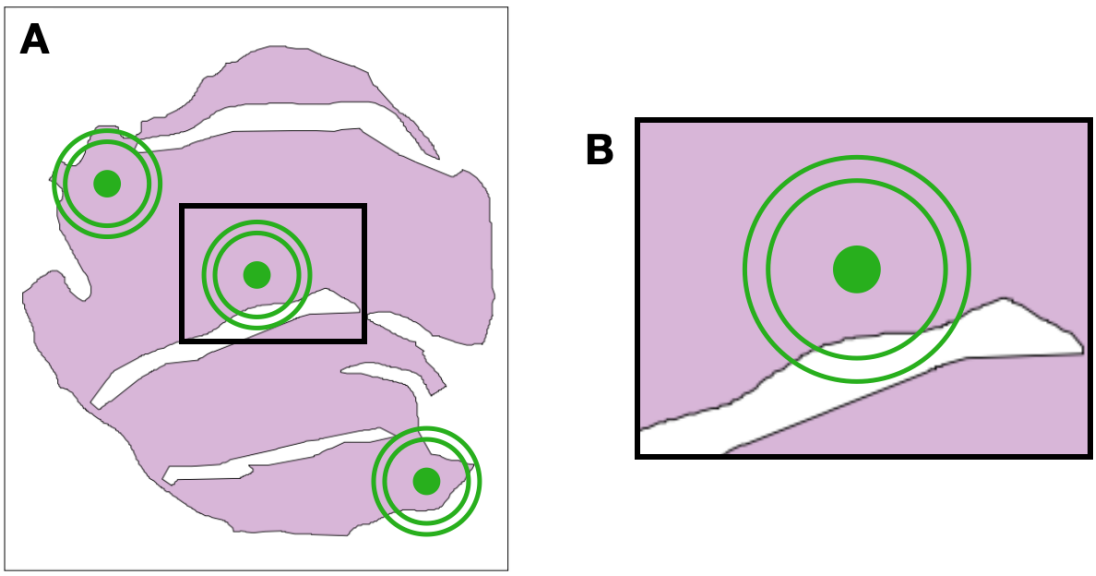


Figure 3.7: Schematic illustrating how edge effects may be accounted for in the PCF calculations. The PCF value at the radius indicated by the green annulus may be approximated by weighting it by the area of the green annulus that lies within the purple tissue annotation.

Approach B in Figure 3.6 treats each tissue piece separately. As such, it captures some of the heterogeneity of the image and uses all of the available data. However, images are treated differently depending on how many pieces of tissue they contain, i.e. some heterogeneity is only captured in images which have more than one piece of tissue.

Approach C involves sampling square regions of interest (ROIs) from within the tissue-tumour-exclude intersection. This method captures intra-patient heterogeneity by sampling different areas of tissue, and provides a means of treating each image in the same manner by analysing regions of tissue of the same size across the cohort. ROIs can be chosen with a specified empty space tolerance (i.e. sample regions outside of the tissue area) to minimise edge effects, and can be sampled with or without replacement. Although not all of the data are sampled, this method is widely used in the literature (Bull et al., 2020, Hutchinson and

Grimm, 2022), and we will use this method of analysing the images going forward. ROIs are a useful sampling method as localised regions of tissue responding to treatment can be identified, even if most of the tissue is not responding and averaged statistics may obscure these responding regions.

ROIs are sampled without replacement (hence no overlapping tiles), with an algorithm we developed that was optimised to maximise the number of ROIs sampled, without inclusion of the tissue boundaries or holes. ROIs are sampled by scanning across the image. Each possible ROI is checked to ensure it falls completely within the tissue area, and does not overlap with previous ROIs, and is accepted if these criteria are met. Other tools for generating ROIs are available, such as PyHIST (Muñoz-Aguirre et al., 2020), which divides a tissue sample image using a fixed grid, with options for ROI overlap and tissue edge inclusion. However, using a grid for small samples, such as those in our dataset, is not optimal to maximise the tissue area sampled, as shown in Figure 3.4 - if a grid was used to sample the ROIs, such as the grid shown in the background of the plot, no ROIs would be accepted, due to tissue edge inclusion. Our tiling algorithm is particularly useful for sampling ROIs without edge inclusion when the ROI dimension is comparable to the width of the tissue sample. Unlike PYHIST, our algorithm also does not require rotation of the image to align the major axis of the tissue sample with a grid.

3.4.7 Analysis of ROIs

The ROIs are spatially analysed using PCFs (between cells of the same type, e.g. tumour-tumour) and cross-PCFs (between cells of different types, e.g. tumour-CD8). The PCF is calculated using the following method. The "source" cells (cells on which the annuli are centred) are the cells contained within the inner region of the ROI (see Figure 3.8). The "target" cells (cells whose density is recorded at a given radius from the source cells) are detected within the outer region (i.e. inner region and a border). The cell density measurements for each ROI are calculated over the whole area of the ROI (inner and outer region).

The cross-PCF, $g_{ST}(r)$, for source cells of cell type S with locations $\mathbf{x}_i = (x_i, y_i)$, surrounded by target cells of type T with locations $\mathbf{x}_j = (x_j, y_j)$, is calculated by using Equation 3.1:

$$g_{ST}(r) = \frac{1}{N_S} \frac{1}{N_T/A} \sum_{i=1}^{N_S} \sum_{j=1}^{N_T} \frac{I(|\mathbf{x}_i - \mathbf{x}_j|)}{A_r(\mathbf{x}_i),} \tag{3.1}$$

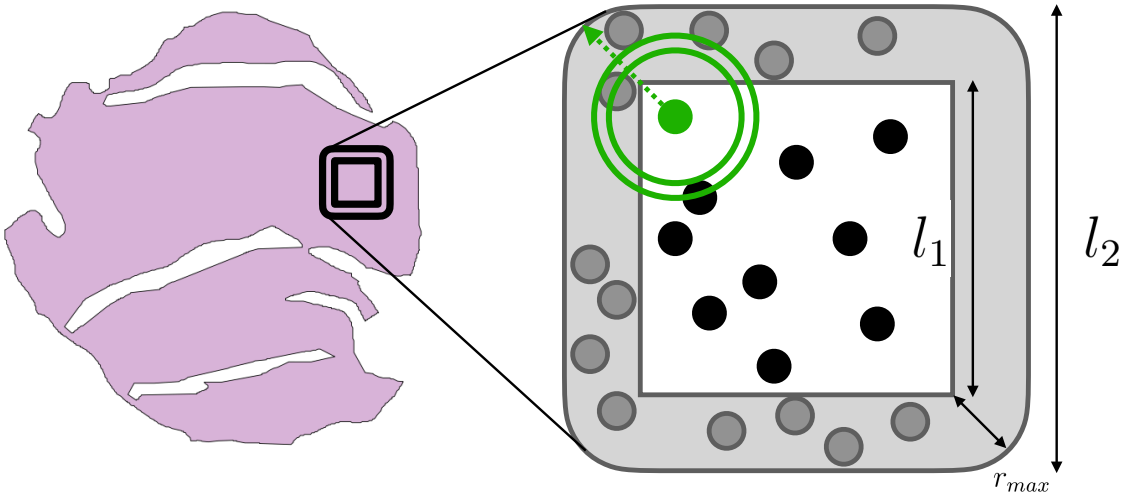


Figure 3.8: Schematic illustrating the sampling of an ROI from a tissue sample. The ROIs are composed of an inner square region of width l_1 (shown in white) and an outer region (shown in grey) of width l_2, with rounded corners of radius r_{max} from the inner region. Cells in black are located in the inner region, and are the "source" cells in the PCF calculations. These "source" cells can detect cells in the outer region (grey cells). This is illustrated by the green annulus representing the calculation of the PCF centred on the green cell in the inner region.

where N_S is the number of source cells in the ROI, N_T is the number of target cells in the ROI. The total area of the ROI, A, is defined as

$$A = \pi r_{max}^2 + l_1^2 + 4 l_1 r_{max} \tag{3.2}$$

where l_1 is the width of the square inner region and r_{max} is the thickness of the outer region (see Figure 3.8).
$I(|\mathbf{x}_i - \mathbf{x}_j|)$ is an indicator function defined by:

$$I(|\mathbf{x}_i - \mathbf{x}_j|) = \begin{cases} 1 & \text{if } r - \delta r \leq |\mathbf{x}_i - \mathbf{x}_j| < r \\ 0 & \text{otherwise} \end{cases}$$

$A_r(\mathbf{x}_i)$ is the area of the annulus with inner radius $r - \delta r$ and outer radius r around a source cell in position $\mathbf{x}_i$, where $\mathbf{x}_i = (x_i, y_i)$ is the position vector of the source cell, and is defined by:

$$A_r(\mathbf{x}_i) = \pi(r^2 - (r - \delta r)^2) \tag{3.3}$$

The PCF between cells of the same cell type uses Equation 3.1 with $S \equiv T$. The PCFs are calculated up to a radius of 100 distance units (approximately 15 tumour cell widths). This maximum radius was chosen to focus on the local cell distributions, and enable sampling the largest possible tiles from the images without inclusion of empty space, as increasing the maximum radius of PCF calculations means that the outer region (border) of the ROI has to be increased, which means that the inner region has to be made smaller, resulting in fewer source cells considered in the PCF calculations.

The mean number of proliferating CD8+ T cells in each ROI is small (typically less than 20 cells). Therefore, when calculating the PCFs, the total CD8+ T cell population is considered, comprising proliferating and non-proliferating CD8+ T cells. Three PCFs are calculated on each ROI; tumour to tumour, tumour to CD8, and CD8 to CD8. Figure 3.9 shows an example of an ROI sampled from an image, and the corresponding PCFs. In this example, the tumour-tumour PCF tends to a value of 1 as r is increased, whereas the CD8-CD8 PCF remains above a value of 1 for large r, and this is reflected visually in the ROI plot, where tumour cells (red) appear more randomly distributed than the CD8+ T cells (blue). The tumour-CD8 cross PCF exhibits co-localisation at small distances and dispersion at intermediate distances, asymptoting to one at larger distances.

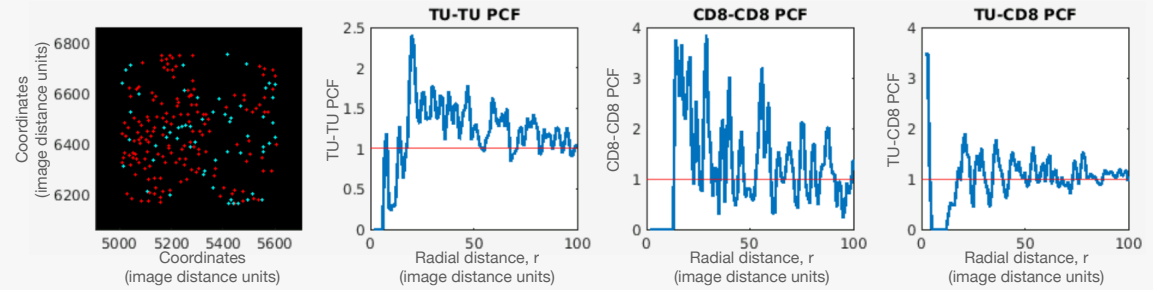


Figure 3.9: Example of a ROI sampled from an image, with the corresponding PCFs and cross-PCFs. TU = proliferating tumour cells, CD8 = total CD8+ T cells (proliferating and non-proliferating). The image of the ROI shows proliferating tumour cells in red, proliferating CD8+ T cells in cyan and non-proliferating CD8+ T cells in blue. The PCFs are calculated using $\delta r = 2$.

To determine a suitable choice for the annulus width δr, PCFs were calculated for the example tile with $\delta r \in \{2, 12\}$. A sliding window annulus is considered, by which the annulus moves radially outwards one distance unit at a time, and the PCF calculated for each value of δr. Figure 3.10 shows how the PCFs depend on δr.

In order to capture the features of the data with a sufficiently fine resolution, whilst reducing the amount of noise in the PCFs, we set $\delta r = 3$. These choices reflect the findings of a previous spectral analysis of the choice of δr, in which the optimal value was found

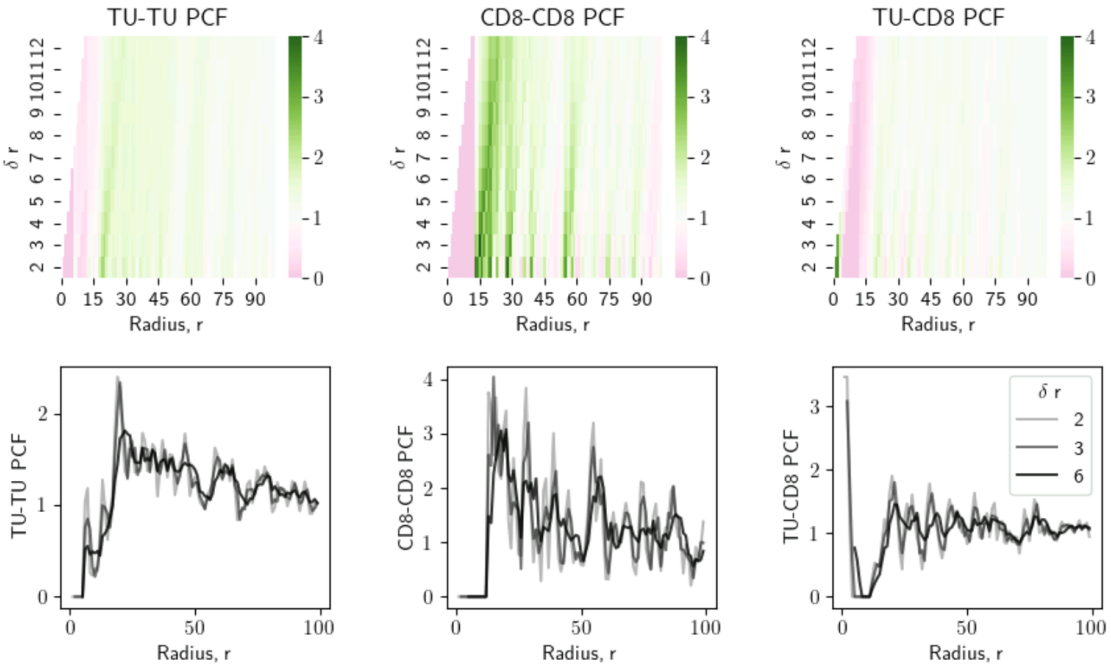


Figure 3.10: Sensitivity of the PCF calculations to the choice of value for the annulus thickness, δr. Heatmaps show the PCFs along the x axis for varying δr along the y axis. Green, white and pink regions correspond to clustering, CSR and dispersion respectively. TU = proliferating tumour cells, CD8 = total CD8+ T cells (proliferating and non-proliferating).

to be approximately equal to the cell diameter of the population which is being analysed (Binder and Simpson, 2015). As CD8 T cells are typically smaller than tumour cells, we consider a value of $\delta r = 3$, which is half the approximate width of a tumour cell.

Features seen at finer resolutions, for example fluctuations around $g_{ST}(r) = 1$, combine into single features as δr increases. Interestingly, short-range clustering of tumour and CD8 cells can only be detected from the cross-PCFs when $\delta r < 6$. Earlier in this chapter we found that the closest tumour cells were located 6 distance units apart, and concluded that the diameter of a tumour cell in these images must be approximately equal to 6 distance units. We interpret co-localisation of CD8+ T cells and tumour cells at distances smaller than the width of a tumour cell as a signature of actively engaged CD8+ T cells and tumour cells (i.e. tumour cells which are being killed by CD8+ T cells). This is an important feature of the data which could be overlooked if a suitable value for δr was not chosen.

During the lengthscale verification process, we noted that there were a few outlier cases of co-located proliferating tumour cells at less than 6 distance units, which we discussed earlier in the chapter may be mislabelled or actively proliferating cells, i.e. two daughter cells which have not fully separated yet. All image segmentation algorithms have an associated

level of accuracy, so this co-localisation may be due to mislabelling. Without the raw images available, we aim to determine whether these co-localised cells are artefacts, or biological features. We summarise and test the hypotheses as to the origin of the co-localised cells in Appendix F. In summary, we found that most CD8-TU overlaps were found in the same patient pre- and on-treatment, suggesting it could be a biomarker. 6 out of 64 patients had overlapping cells in both pre- and on-treatment samples. Furthermore, there was not a significant difference between the number of tumour-nonproliferating CD8 pairs (32) and tumour-proliferating CD8 pairs (26), suggesting it is not a labelling error. A weak correlation between frequency of overlapping cells and number of CD8 cells in the tissue was found, however, this does not disprove that there is an artefact, as with increased T cells we may expect the immune system to be functioning more efficiently and hence it could be causing increased co-localisation of tumour and CD8 cells.

3.4.8 Smoothing of PCFs

The PCFs calculated from the ROIs have varying levels of noise depending on the number of cells within each ROI. To mitigate the effects of noise on the metrics extracted from the PCFs, each PCF is smoothed using a Savitsky-Golay filter (Savitzky and Golay, 1964), which has previously been used to smooth PCFs (Bull, 2018). As some of the raw PCF vectors contain zero values, the smoothing filter occasionally produces negative values, which are physically unrealistic, so any negative values of the smoothed PCF were set to zero. Figure 3.11 shows examples of the TU-CD8 PCF which have been smoothed.

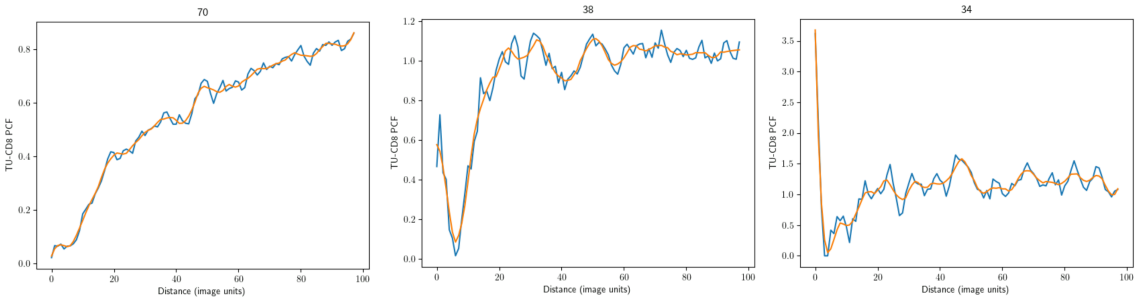


Figure 3.11: Examples of TU-CD8 PCFs for three patients, averaged over ROIs sampled from on-treatment images. The raw mean PCFs are plotted in blue, and the Savistsky-Golay filter with a window size of 15 and 3rd order polynomial is plotted in orange, to smooth the PCF signal.

3.4.9 Average PCF calculation for a set of ROIs

The average PCFs for each patient are calculated as follows. The PCF g_i for each ROI, i, is calculated, and weighted by the number of tumour cells within the inner region, rel-

ative to the total number of tumour cells summed over the inner regions of all the ROIs, $N_i^{TU} / \sum_{i=1}^{n} N_i^{TU}$, where n is the number of ROIs. The weighted PCFs are then summed to calculate the mean PCF for the patient.

3.4.10 Principal component analysis

Principal component analysis (PCA) reduces the dimensionality of a dataset into a user-specified number of principal components (PCs) (Jolliffe and Cadima, 2016). This technique is commonly applied before using a clustering algorithm (for example K-means clustering) to categorise the data. The optimal number of clusters in K-means clustering can be determined by producing an elbow plot. To do this, the squared distance of all points within a cluster to the cluster centre is calculated and summed over the cluster, and then over all clusters to produce the elbow score. This elbow score is computed as the number of clusters, k, varies. The "elbow" in the resulting plot defines the optimal number of clusters $k = k_{optimal}$ and is where the elbow score does not decrease significantly for $k > k_{optimal}$. In many cases, the elbow may not be clear, and the silhouette score can be used to determine the number of clusters (Rousseeuw, 1987). The silhouette score measures how well each point belongs to its cluster. This score is calculated for each point and averaged over all points to produce the average silhouette score for each value of k. The silhouette score ranges between -1 and 1, and is maximised when $k = k_{optimal}$. The Python library scikit-learn was used to apply these methods to the data.

Bull and Byrne (2023) calculated spatial statistics on synthetic data generated from an off-lattice ABM of tumour-macrophage interactions. These statistics were concatenated to produce a feature vector for each frame of the ABM. By applying PCA to these feature vectors, it was possible to characterise each frame as *tumour elimination*, *tumour escape* or *tumour-immune equilibrium* phenotypes. We apply this method of analysis to our clinical dataset to test if it can pick up patterns between pre- and on-treatment samples.

The statistics for each ROI are concatenated to produce a feature vector $\mathbf{a}_F$ of dimension (1 x 297), as indicated below:

$$\mathbf{a}_F = [\rho_T, \rho_E^{np}, \rho_E^{p}, g_{TE}(\mathbf{r}), g_{EE}(\mathbf{r}), g_{TT}(\mathbf{r})],$$

where ρ_T, ρ_E^{np} and ρ_E^{p} are the densities of proliferating tumour cells, non-proliferating CD8+ T cells and proliferating CD8+ T cells, respectively, while $g_{AB}(\mathbf{r}) = [g(r_0), g(r_1), ..., g(r_{max})]$ represents the PCF in the case of $A \equiv B$ or cross-PCF otherwise, calculated on source cell type A and target cell type B. PCFs and cross-PCFs are calculated starting at $r = \delta r = 3$, so in our analysis $r_0 = 3$ image units, and $r_{max} = 100$ image units. The radius is incremented by 1 image unit such that $r_{i+1} - r_i = 1$.

3.5 Results

3.5.1 Cell density analysis

We start our analysis of the patient cohort by computing the densities of proliferating tumour cells (Ki67+ TU cells), proliferating CD8+ T cells (Ki67+ CD8+ cells) and non-proliferating CD8+ T cells (Ki67- CD8+ cells) in the samples. The aim of this analysis is to ascertain how much we can learn about the effect of the treatment from non-spatial metrics. We sample square ROIs of dimension 800 distance units, using the procedure described in Section 3.4.6. We calculate the cell densities in these ROIs sampled from pre- and on-treatment histology images, and calculate the mean cell densities over the ROIs for each image.

Figure 3.12 shows how the percentage change in tumour cell density between pre- and on-treatment samples varies with mean CD8 cell density pre-treatment. A weak positive correlation is found ($R^2 = 0.016$), and the residuals show a high variance in the data. This suggests that the patient tumour response is independent of the number of CD8 cells pre-treatment.

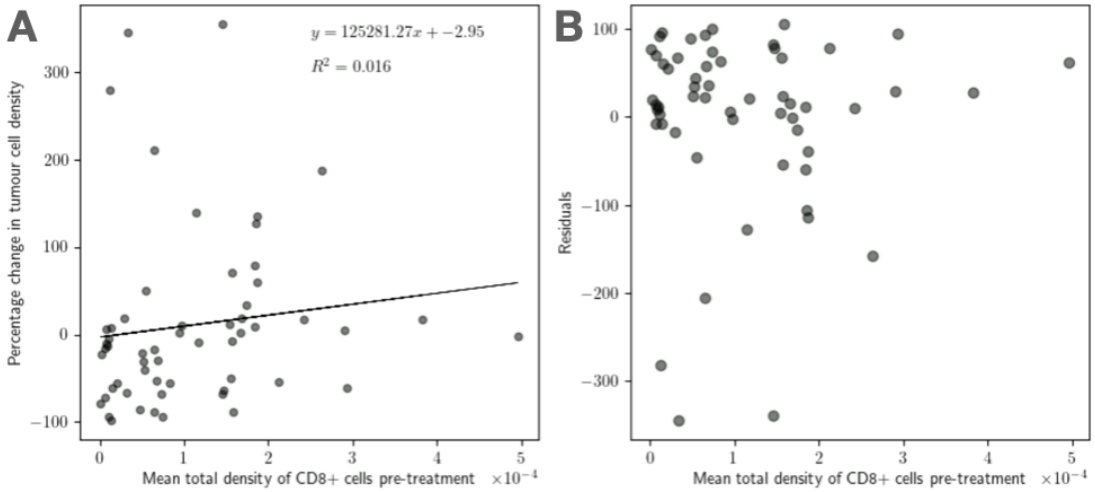


Figure 3.12: (A) Percentage change in tumour cell density plotted against CD8 cell density pre-treatment. Each data point represents an individual patient. (B) Residuals of the linear fit to the data in (A).

3.5.2 Identifying responders

Firstly, we use the data to quantify changes in tumour burden. For each patient we calculate the percentage change in tumour cell density relative to the baseline pre-treatment tumour cell density. We use this relative measure, as tumour cell density will vary between patients due to the different primary indications.

Figure 3.13 shows the percentage change in tumour cell density and total CD8 cell density (sum of proliferating and non-proliferating cells) over the patient cohort, with patients ordered by this percentage change.

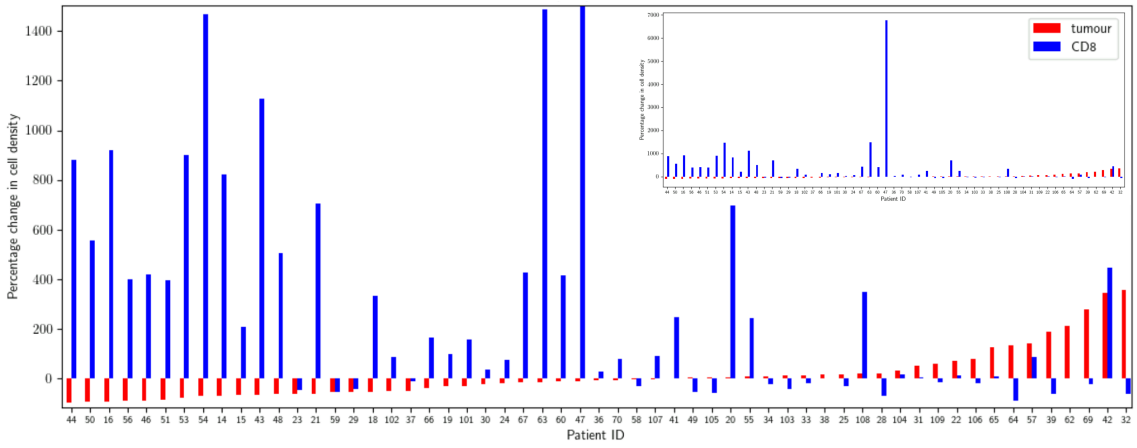


Figure 3.13: Percentage change in tumour cell density from baseline (pre-treatment) for each patient in the cohort. Inset plot shows all of the data unzoomed.

This figure highlights wide variability in the change in cell densities. We note that patients whose tumour cell density decreases, also experienced a large increase in CD8 cell density; most with at least a 300% increase. Conversely, patients whose tumour cell density increased typically experienced a decrease, or small increase, in CD8 cell density.

The change in CD8 cell density can be decomposed into proliferating and non-proliferating CD8 cells, shown in Figure 3.14. This plot shows that an increase in proliferative CD8 cells is pronounced in patients with a decrease in tumour cell density, and suggests treatment efficacy in this subset of patients, due to an expansion of cytotoxic T cells, which reduces the number of tumour cells.

However, there are a few outliers in these trends, i.e. patients who had decreases or increases in both cell types. Figure 3.15 shows that the patients may be categorised into groups defined by the percentage change in tumour and CD8 cell densities.

Most patients fall within the group of "intuitive" responders, i.e. they respond in a way that is expected given the hypothesised mechanism of action of the treatment - CD8+

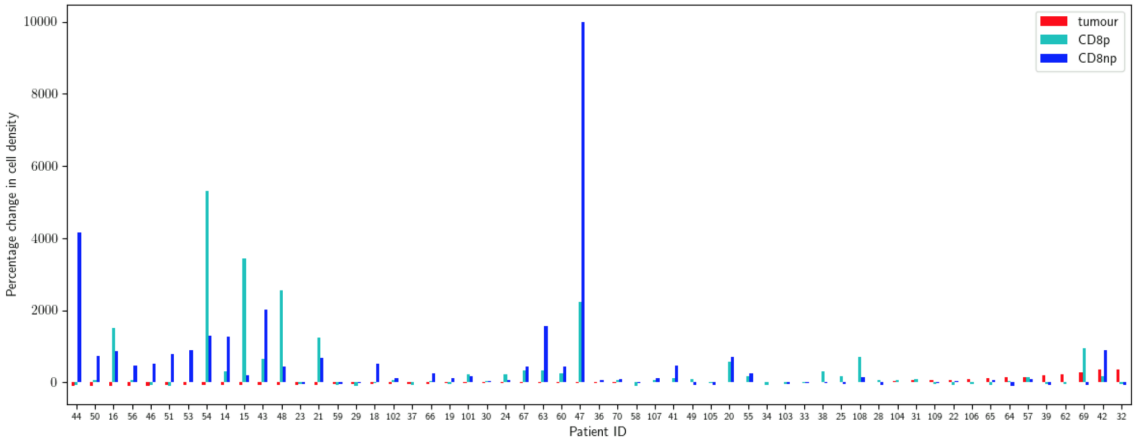


Figure 3.14: Percentage change in tumour and CD8 cell density from baseline (pre-treatment) for each patient in the cohort. CD8p = proliferating CD8 cells, CD8np = non-proliferating CD8 cells.

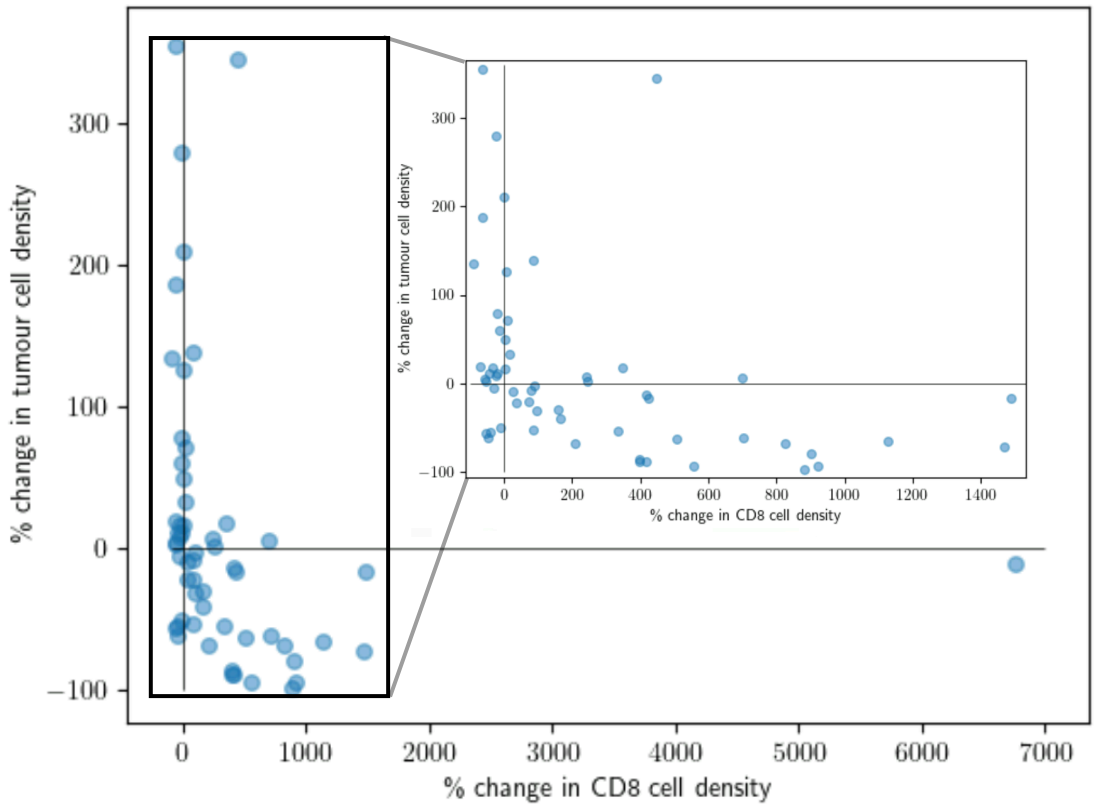


Figure 3.15: Percentage change in tumour cell density between pre- and on-treatment, plotted against percentage change in CD8 cell density. Each point represents the mean measurements for one patient, taken across the ROIs sampled from that patient's images. Inset plot shows a zoomed in view of the region bounded by the black box, for clarity.

cell density increases, and tumour cell density decreases. The second largest group contains patients with an increase in tumour cell density, and decrease in CD8+ cell density;

this group could be termed "non-responders". There are a few patients in the other two categories; in the decreased tumour density/decreased CD8+ density group, a reduction in tumour burden is seen even though the CD8 population has not expanded. Conversely, in the final group, significant CD8 expansion is observed (for most patients in this group over 100% increase) but the tumour burden also increases. In Section 3.4.7, we will use spatial statistics to investigate whether differences in the spatial locations of the different cell types may explain these findings.

A responder threshold can be defined, for example, a positive or negative response can be defined as above or below 20% change in tumour cell density, respectively. Figure 3.16 shows the percentage of patients categorised as good responders, no change, or poor responders, for different choices of the response threshold percentage. About a fifth of patients were categorised as poor responders (i.e. the proliferative tumour density increased). Most patients did not have a good response to the treatment for response thresholds of above 5%.

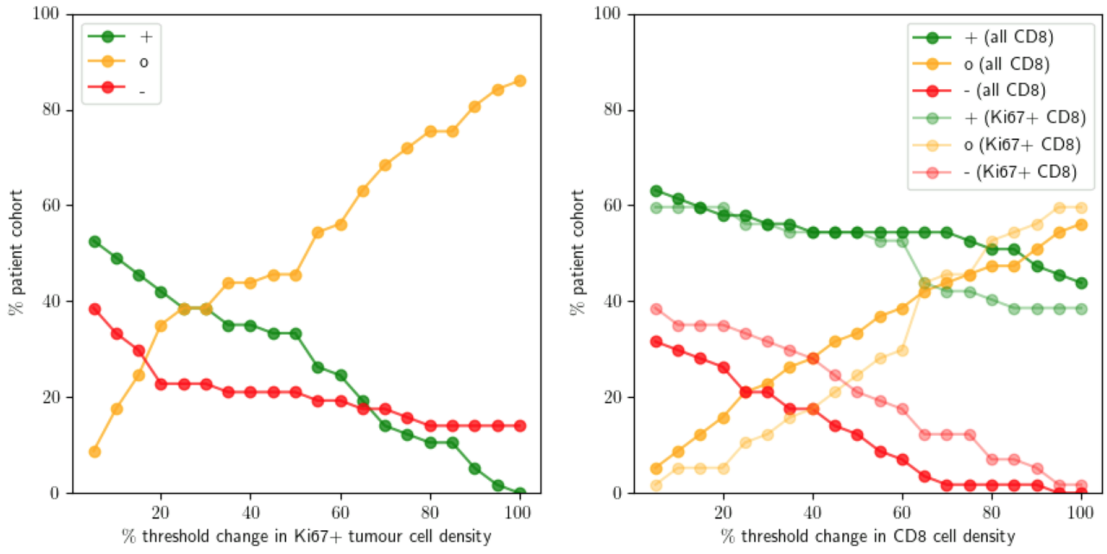


Figure 3.16: (A) Proliferating tumour cell density response - percentage of patients as a function of the choice of percentage threshold for the categorisation of patients into proliferative tumour cell density decrease (+), no significant change (o) and tumour cell density increase (-). (B) CD8 T cell density response - percentage of patients as a function of the choice of percentage threshold for the categorisation of patients into CD8 cell density increase (+), no significant change (o) and CD8 cell density decrease (-). Bold lines denote total CD8 population, and light lines denote only proliferative CD8 cell population.

3.5.3 Dependence of response on the on-treatment biopsy timepoint

When comparing patient response within the cohort, we need to take into consideration that there is a lot of variation within the dataset, in terms of primary indication, biopsy location, and the timepoint of on-treatment biopsy collection. The primary indication will be investigated in Section 3.5.6. The timepoint of the on-treatment biopsy sample (defined as days since first dose) is important as patients with earlier collection times may not have had time for the drug to take effect. Hence, direct comparison of patients' on-treatment samples taken tens of days apart may not be valid. For the patients where on-treatment sample times are available, these times range from 15 to 113 days after the first dose (see Appendix E).

In Figure 3.17 we investigate if there is a correlation between the percentage change in the tumour cell density between pre- and on-treatment samples, and the time at which the on-treatment sample was collected.

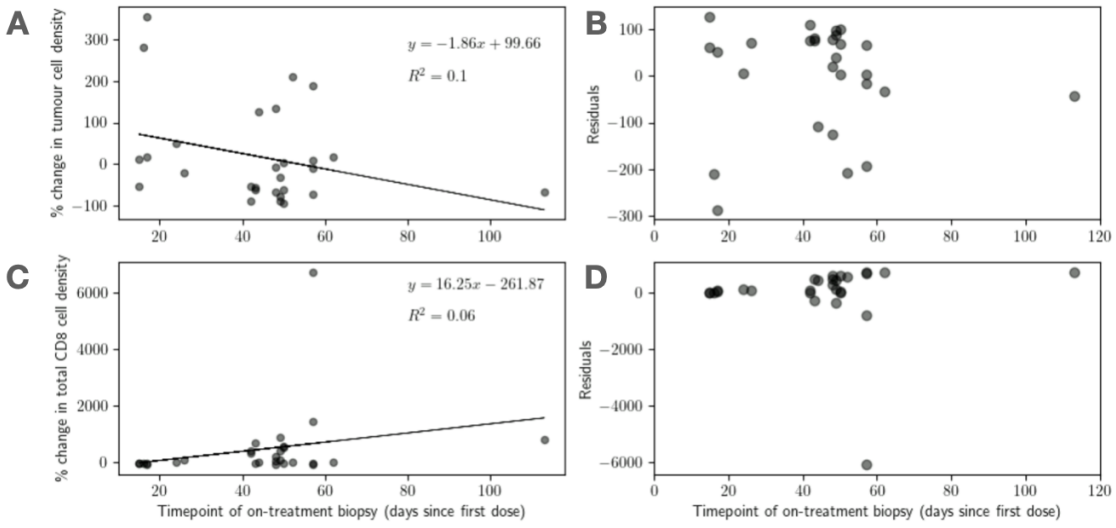


Figure 3.17: Percentage change in proliferating tumour cell density (top row) and total CD8 cell density (bottom row) as a function of timepoint of on-treatment biopsy (days since first treatment dose). Each data point represents an individual patient, and residuals of the linear fits in the left hand column plots are plotted in the right hand column.

Figure 3.17 shows a weak negative correlation between the percentage change in tumour cell density and the timepoint of on-treatment biopsy, with a low value of $R^2 = 0.1$, and linear fit equation $y = 99.66 - 1.86x$. A weak positive correlation between the percentage change of CD8 cell density and timepoint of on-treatment biopsy is found with $R^2 = 0.06$, and linear fit equation $y = -261.87 + 16.25x$. Although these correlations are weak, they suggest that for successful treatment, as treatment time increases the number of CD8+ T

cells increases and the number of tumour cells decreases. We note that the best fit lines do not pass through the origin. This is to be expected since the on-treatment biopsy timepoint is recorded as days since first dose, rather than days since the pre-treatment biopsy sample.

Figure 3.18 shows the combined change in CD8 and tumour cell density, coloured by the on-treatment biopsy timepoint. For short times between first dose and on-treatment biopsy (approx. less than 30 days), there is a decrease in CD8 density. All good responders (bottom right quadrant) had their on-treatment biopsy collected at least 30 days after the start of treatment.

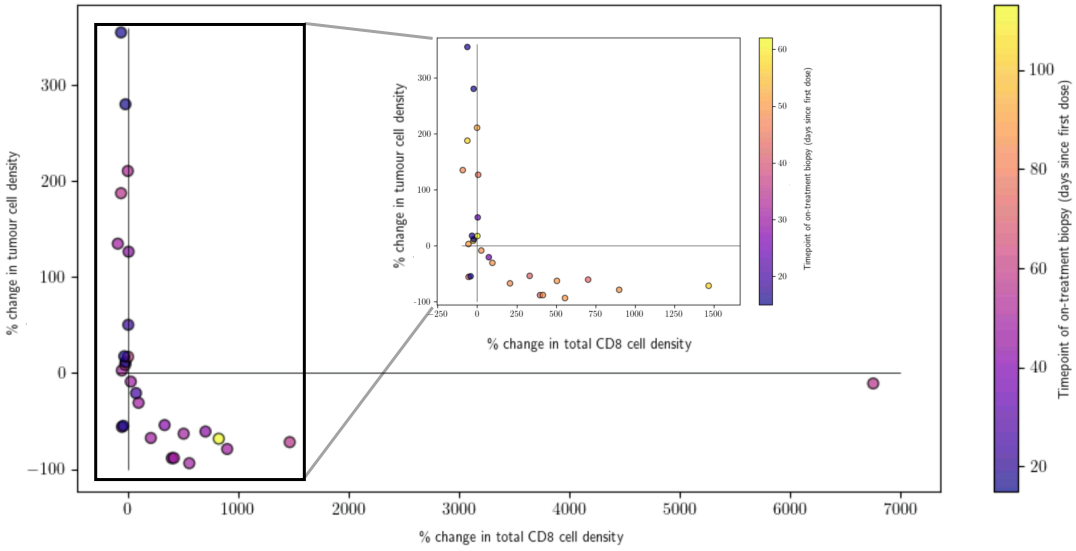


Figure 3.18: Percentage change in tumour cell density plotted against percentage change in CD8 cell density. Each point represents an individual patient, and is coloured by the timepoint of the on-treatment biopsy for that patient (days since first dose). Zoomed inset plot with adjusted colour bar is shown for clarity.

3.5.4 Clustering using cell density data

We now use the cell density data extracted from the ROIs to investigate intra-patient variability. In the previous section, mean cell densities were computed to facilitate inter-patient comparison and to establish trends across the cohort. However, in some cases there is large variability in the ROIs for a given patient. In order to determine whether the "responses" are homogeneous across a given tissue sample, or dominated by specific regions, we now analyse the ROIs from the whole patient cohort.

We use K means clustering to cluster the ROI data from the pre- and on-treatment images pooled, to see whether groups of patients cluster together, and whether these clusters

are maintained in the pre- and on-treatment data.

Figure 3.19 shows the clustering of the cell density data for pre- and on-treatment ROIs, for all patients in the cohort.

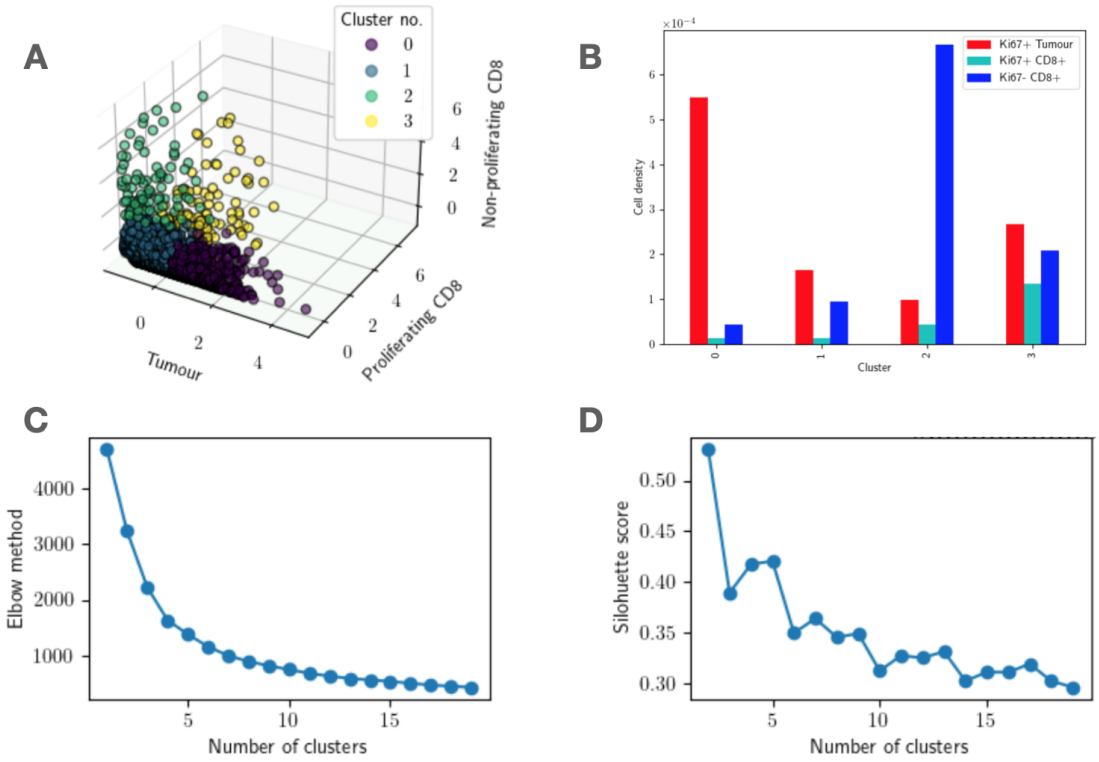


Figure 3.19: (A) ROI data clustered and labelled by cluster number. (B) Cell density composition of centroids of clusters. (C) Elbow method and (D) silhouette score used to determine optimal number of clusters (four).

ROIs cluster into 4 groups:

- Cluster 0: High density of tumour cells, not well infiltrated by CD8s
- Cluster 1 and 3: High CD8 infiltration of tumour
- Cluster 2: High density of CD8 cells, low density of tumour cells, which possibly correspond to stromal regions.

The bar charts in Figure 3.20A show the distribution of the ROIs for each patient across the clusters, for pre-treatment (labelled 1) and on-treatment (labelled 2,3,4). Figure 3.20 shows the distribution of patients by the number of clusters their ROIs are in. We see that there is a wider distribution of cluster occupation number in pre-treatment samples, however in the on-treatment samples the distribution is more uniform, such that there is more heterogeneity in the on-treatment samples compared to pre-treatment.

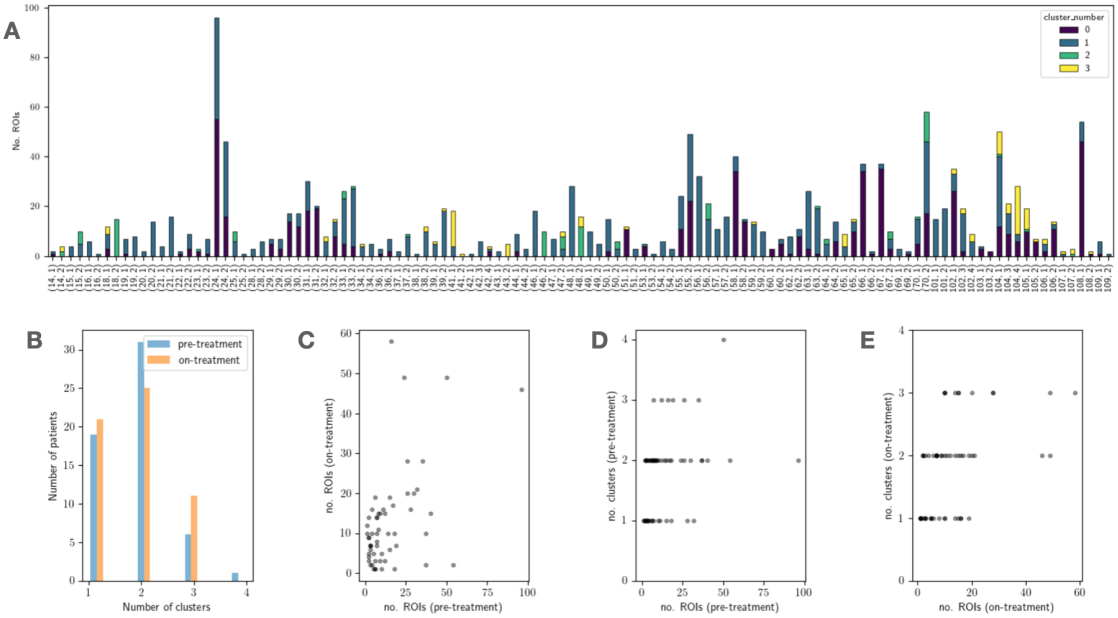


Figure 3.20: (A) Distribution of ROIs over the four clusters for each patient, pre- and on-treatment, labelled 1 and 2, respectively. (B) Distribution of number of clusters per patient, pre- and on-treatment. (C) Number of ROIs on-treatment vs. pre-treatment, (D) number of clusters a patient has ROIs in vs. number of ROIs pre-treatment and (E) number of clusters a patient has ROIs in vs. number of ROIs on-treatment.

The preliminary clustering of cell density measurements from ROIs sampled from pre-treatment and on-treatment images shows that there is a high degree of heterogeneity at the intra-patient and inter-patient levels. There is no clear mapping between pre-treatment and on-treatment clusters, indicating that with this cell density data we would not be able to predict a patient response based on the clustering of their pre-treatment ROIs. There may be several reasons why this analysis is inconclusive; firstly, analysing the clustering of cell densities may not be adequate in separating out treatment responders and non-responders; many of the ROIs are found in two distinct groups; high tumour, low CD8 and low tumour, high CD8. It is not clear whether this analysis is simply separating out tumour and stromal regions of the tissue, i.e. if an ROI mainly contained stromal tissue it may contain an increased number of CD8 cells, and be categorised into the low tumour, high CD8 cluster. Therefore we may not actually be separating the ROIs into groups which have features relevant to treatment response, and instead could simply be categorising physical regions of the tumour. This motivates the inclusion of spatial metrics in our clustering analysis later in Section 3.5.6.

3.5.4.1 Individual patient analysis

As already discussed, the patient dataset is extremely heterogeneous, with different primary indications, dosing regimes and on-treatment biopsy timepoints, which could confound the clustering analysis. For a more individualised approach we could pool the pre- and on-treatment ROIs for each patient before clustering, and then determine (a) whether the pre- and on-treatment ROIs cluster, (b) how heterogeneous the clusters are, and (c) the Euclidean distance between the cluster centroids to quantify how separated the clusters are. Figure 3.21 illustrates this approach for three different patients.

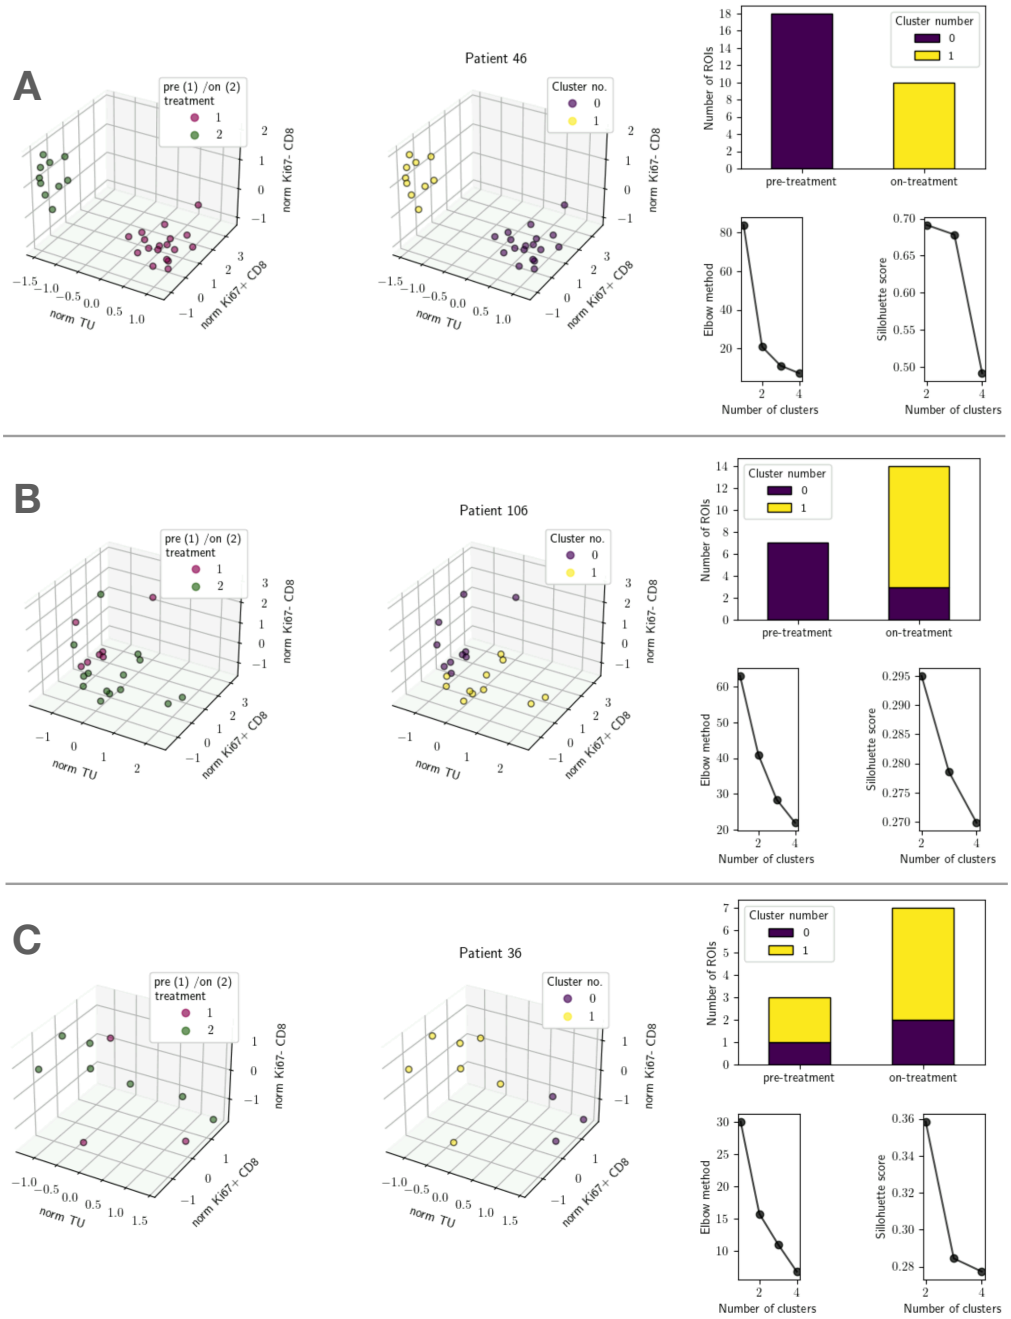


Figure 3.21: Clustering of pre- and on-treatment ROIs for three patients. Cell densities are normalised within the ROI dataset for an individual patient, clustered into n clusters informed by the inflexion point in the elbow plot and maximum values of silhouette scores. (A) Clusters separate into pre- and post-treatment data. (B) Cluster 1 is on-treatment specific, but there is some mixing of pre- and on-treatment data within clusters. (C) Clusters are mixed and the clustering is not informative, possibly due to the patient not responding to treatment, or the lack of data.

The degree of mixing of the clusters indicates which features remain unchanged between pre- and on-treatment samples. The distance between the clusters, if there is good separation of pre- and on-treatment (i.e. not a lot of mixing) indicates the strength of the response to the treatment. For example, in Figure 3.21A the pre- and on-treatment data separate into distinct clusters and there is a large relative separation between clusters. This is termed "relative separation", as the ROI data for each patient is normalised separately. As a result, distance in normalised space is not comparable between patients; comparison would necessitate transformation back to unnormalised densities. We observe a decrease in tumour cell density and increase in Ki67- CD8+ cell density between pre- and on-treatment. In Figure 3.21B, the clusters contain mixtures of pre- and on-treatment ROIs, and there is less relative separation between clusters. For this patient, we observe a decrease in CD8 cell density and an increase in tumour cell density on average between pre- and on-treatment. However, examining the composition of each cluster reveals that a small number of on-treatment ROIs are within the same cluster as the pre-treatment ROIs. This suggests that this patient is a partial responder, with some regions unchanged from the pre-treatment state. In Figure 3.21C, the mixing of pre- and on-treatment ROIs in each cluster suggests that this patient has not responded to the treatment.

3.5.5 Trends in PCF metrics

The analysis in the previous section highlighted the variability in treatment responses at the inter-patient and intra-patient levels. It showed that cell densities alone provide limited information about what is happening in the tissue. High tumour/low CD8 and low tumour/high CD8 regions logically clustered together, and ROIs for a given patient were spread over multiple clusters, which could indicate that this analysis was picking out mainly tumour and mainly stromal regions, respectively. We now investigate whether spatial analysis can help to distinguish these regions, and in particular, to understand how and why, for certain patients, both the CD8 and tumour cell densities increase between pre- and on-treatment samples.

The PCFs for CD8 and tumour cells, and the cross-PCFs between tumour and CD8 cells, were used to compile a spatial signature for each ROI (see Section 3.4.7 for details). The PCFs were smoothed using a Savitsky-Golay filter (see Section 3.4.8). Figure 3.22 shows the distribution of counts for each cell type across all ROIs in the patient cohort. Most ROIs have low cell numbers, with a positive skewed distribution. As the PCFs are sensitive to noise at low cell densities, a minimum cell threshold for PCF calculations was imposed. From examining these ROI cell count distributions, and PCFs calculated from individual ROIs, we set the minimum cell counts as 50 cells within the inner region of the

ROI (which is of size 600 by 600 distance units) - the "source" cell population (see Figure 3.8), as these are the cells from which the PCF is calculated and averaged over. There are very few ROIs with more than 50 proliferating CD8 cells, henceforth proliferating and non-proliferating CD8 cells are pooled into one cell group.

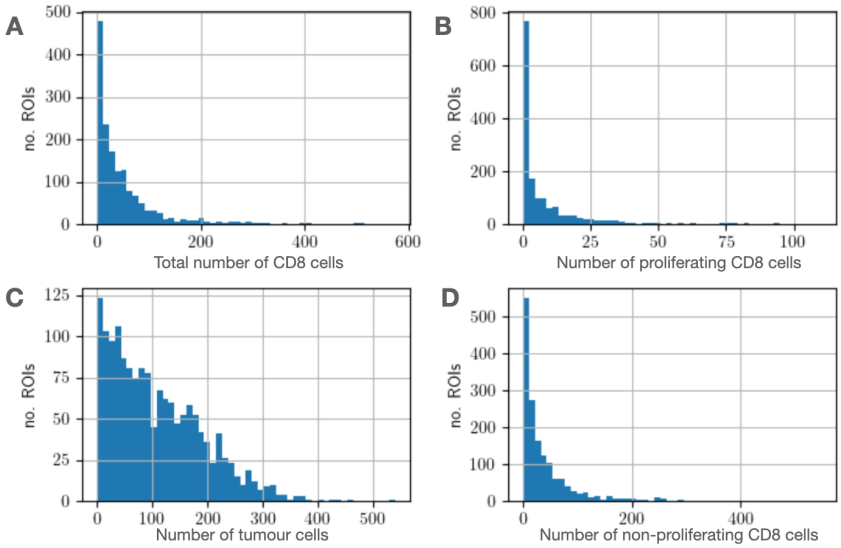


Figure 3.22: Series of histograms showing the number of ROIs containing a particular number of cells of a certain type. The number of cells in the "inner" region of the ROI is considered here - see Figure 3.8. This is because the cells in the inner region are treated as the "source" cells in the PCF calculations, and these cells are definitely included in the calculations (i.e. there may be some "target" cells within the outer border region which are not within distance r_{max} of any of the source cells, so would not be considered in the PCF calculations).

Bull et al. (2020) found that the maximum value of the PCF, and its corresponding location can be useful for characterising ROIs and summarising the PCF signature. We firstly examine some typical PCFs to ascertain if there are any intuitive metrics which can be extracted from them.

Figure 3.9 shows that for the same-cell PCFs, the maximum value of the PCF, g_{max}, and the radius at which this maximum occurs, r_{max}, may provide characteristic measurements of the PCF. For the TU-CD8 cross-PCF, the form of the PCF is not consistent in terms of the position of g_{max}, so we will instead look at the metrics for the minimum of the PCF; g_{min} and r_{min}. Our aim is to find if any of these PCF metrics change between pre- and on-treatment samples.

3.5.5.1 Tumour - CD8 cross-PCF

Figures 3.23, 3.24 and 3.25 show how the percentage change in cell densities between pre- and on-treatment samples change with g_{max} and r_{max}, respectively. g_{max} and r_{max} are measured from the average of the TU-CD8 cross PCFs calculated on the ROIs from each patient's pre- and on-treatment samples (separately). The average PCFs for each patient are calculated as described in Section 3.4.9, and statistics such as g_{max} are extracted from the mean PCF. The two average PCFs over the ROI sets for pre- and on-treatment are calculated so that the pre- and on-treatment PCF metrics can be directly compared - there are typically not the same number of ROIs in the pre- and on-treatment samples for a given patient, and there is no way of comparing a single pre-treatment ROI with an on-treatment ROI.

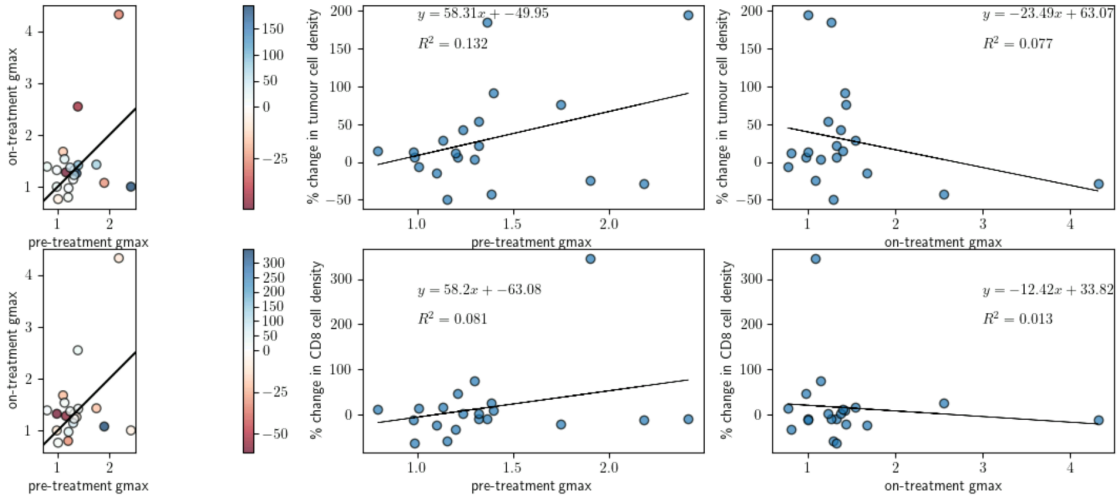


Figure 3.23: Investigating trends in g_{max} in pre- and on-treatment samples. The first row of plots explore the relationship between g_{max} and change in tumour cell density (calculated as the mean over the ROIs which were eligible for spatial analysis). The second row of plots explore the relationship between g_{max} and change in CD8 cell density (calculated as the mean over the ROIs which were eligible for spatial analysis). Each data point represents one patient.

The results presented in Figures 3.23 and 3.24 show no significant (strong) correlations between the PCF metrics and changes in cell densities (the R^2 values are small). A positive correlation ($R^2 = 0.132$) of the change in tumour cell density with g_{max}^{pre} and negative correlation ($R^2 = 0.156$) with r_{max} suggests that patients with a more uniform distribution of CD8 cells with clustering of CD8 cells at large distances from tumour cells results in an improvement in tumour burden.

Whilst the correlations are weak, they are also counter-intuitive. For example, we

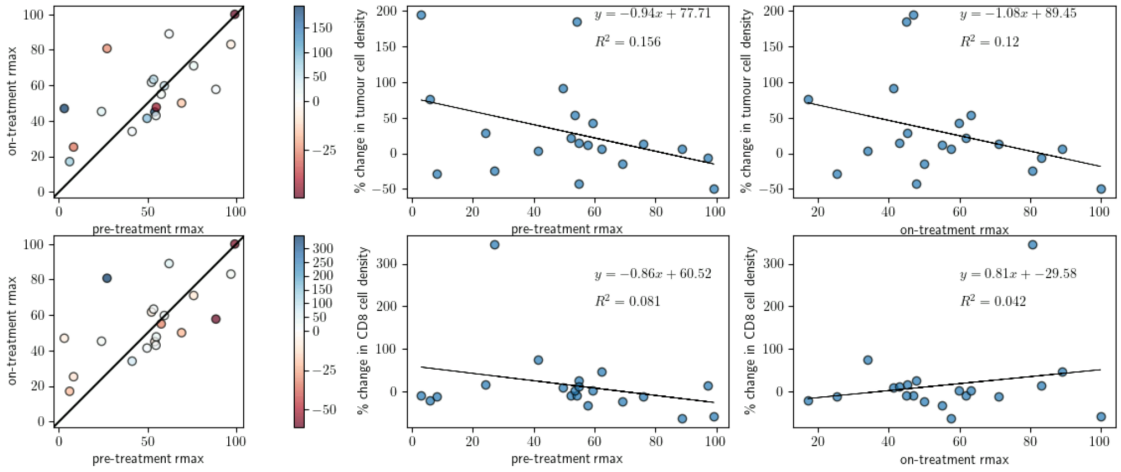


Figure 3.24: Investigating trends in r_{max} in pre- and on-treatment samples. The first row of plots explore the relationship between r_{max} and change in tumour cell density (calculated as the mean over the ROIs which were eligible for spatial analysis). The second row of plots explore the relationship between r_{max} and change in CD8 cell density (calculated as the mean over the ROIs which were eligible for spatial analysis). Each data point represents one patient.

might expect increased CD8 infiltration into the tumour (co-localisation of tumour and CD8 cells) to result in a reduction in tumour burden. However, in patients responding to the treatment, we also may expect localisation of CD8 cells within the stroma in the early stages of treatment, where there is a high density of FAP-expressing tumour associated cells, to which the molecule attaches and locally promotes CD8 cell proliferation. The CD8 cells would then need to migrate out of the stroma to the tumour cells to have a therapeutic cytotoxic effect. When we examine the shape of the TU-CD8 PCF, the g_{max} and r_{max} values vary considerably (i.e. similar values of g_{max} may be associated with low r_{max} and high r_{max} values due to the shape of the PCF). Therefore this may be confounding the results, and we now analyse the position of the minimum of this function, r_{min}, in Figure 3.25, which corresponds to the radial distance with the most dispersion of CD8 cells.

Of note in these results is in the top left plot in Figure 3.25, which shows that for most of the patients who have a reduction in tumour cell density, the r_{min} value decreases between pre- and on-treatment samples (green box drawn on plot to refer to these points). This means that an increase in co-localisation of CD8 cells nearer tumour cells between pre-treatment and on-treatment may be indicative of a good tumour response. Note that in this analysis, the minimum value for r_{min} is 3, as this is the δr value used in the PCF calculations. There are no significant trends in the r_{min} data in Figure 3.25; the largest R^2 value (0.409) corresponds to the apparent positive correlation between change in tumour

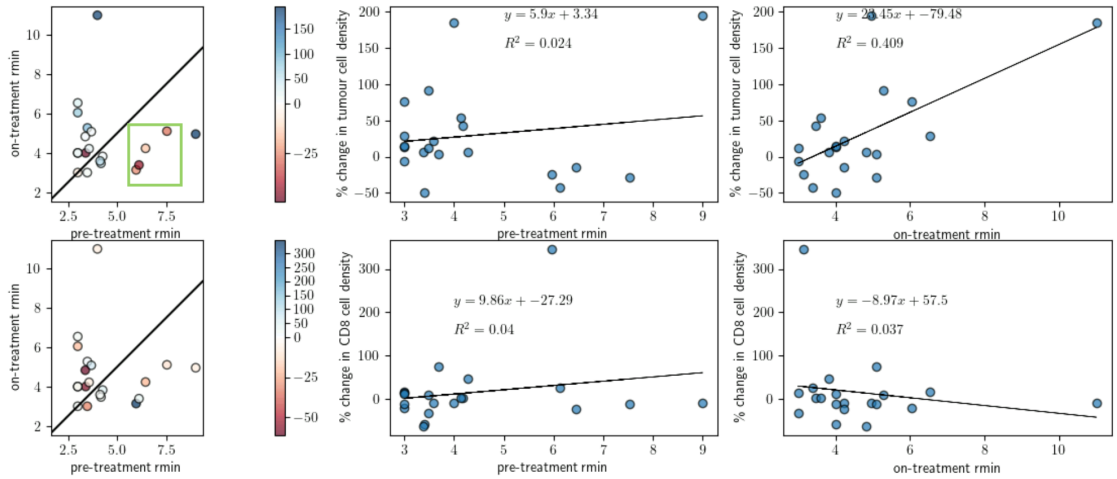


Figure 3.25: Top row: Change in tumour cell density. Bottom row: Change in CD8 cell density. Investigating trends in r_{min} in pre- and on-treatment samples for the TU-CD8 cross PCF. The first row of plots explore the relationship between r_{min} and change in tumour cell density (calculated as the mean over the ROIs which were eligible for spatial analysis). The second row of plots explore the relationship between r_{min} and change in CD8 cell density (calculated as the mean over the ROIs which were eligible for spatial analysis). Each data point represents one patient.

cell density and on treatment r_{min}.

There is a lot of variability in these data and it is clear that no single PCF metric can identify any trends. This may be partly due to the shape of the PCF but also, as discussed previously, that the data are very disparate. These findings lead us to the next section of analysis, where we consider the whole PCF signature using K means clustering, and separate patients into disease types (for those patients for whom we have such data) to attempt to group the widely varying patient cohort and analyse the response of each primary disease type to the treatment. We also note that this section has analysed average ROI statistics; in the next section we will consider intra-patient heterogeneity by considering the PCF signatures of each individual ROI.

3.5.6 Analysis considering primary tumour indication

In the previous section, we investigated whether using metrics extracted from the PCFs could identify any trends in the spatial distribution of cells pre- and on-treatment. No significant trends were identified. In this section we will separate the data by primary indication (given in Appendix E), and use the entire PCF signatures of all ROIs as input for K means clustering. Through this analysis we aim to investigate whether treatment responses differ between primary indications. We consider the four largest groups of primary

indication:

- RCC - Renal cell carcinoma
- MEL - Melanoma
- BAC - Breast adenocarcinoma
- SCC - Squamous cell carcinoma

Figure 3.26 shows the percentage change in tumour and CD8 cell densities averaged over all ROIs from a given patient, coloured by each patient's primary indication. We find that melanoma (MEL) is the only primary indication for which there is a constant reduction in tumour cell density, across all melanoma patients. Only one of the five melanoma patients experiences a decrease in CD8 cell density; all other melanoma patients exhibit an increase in CD8 cells.

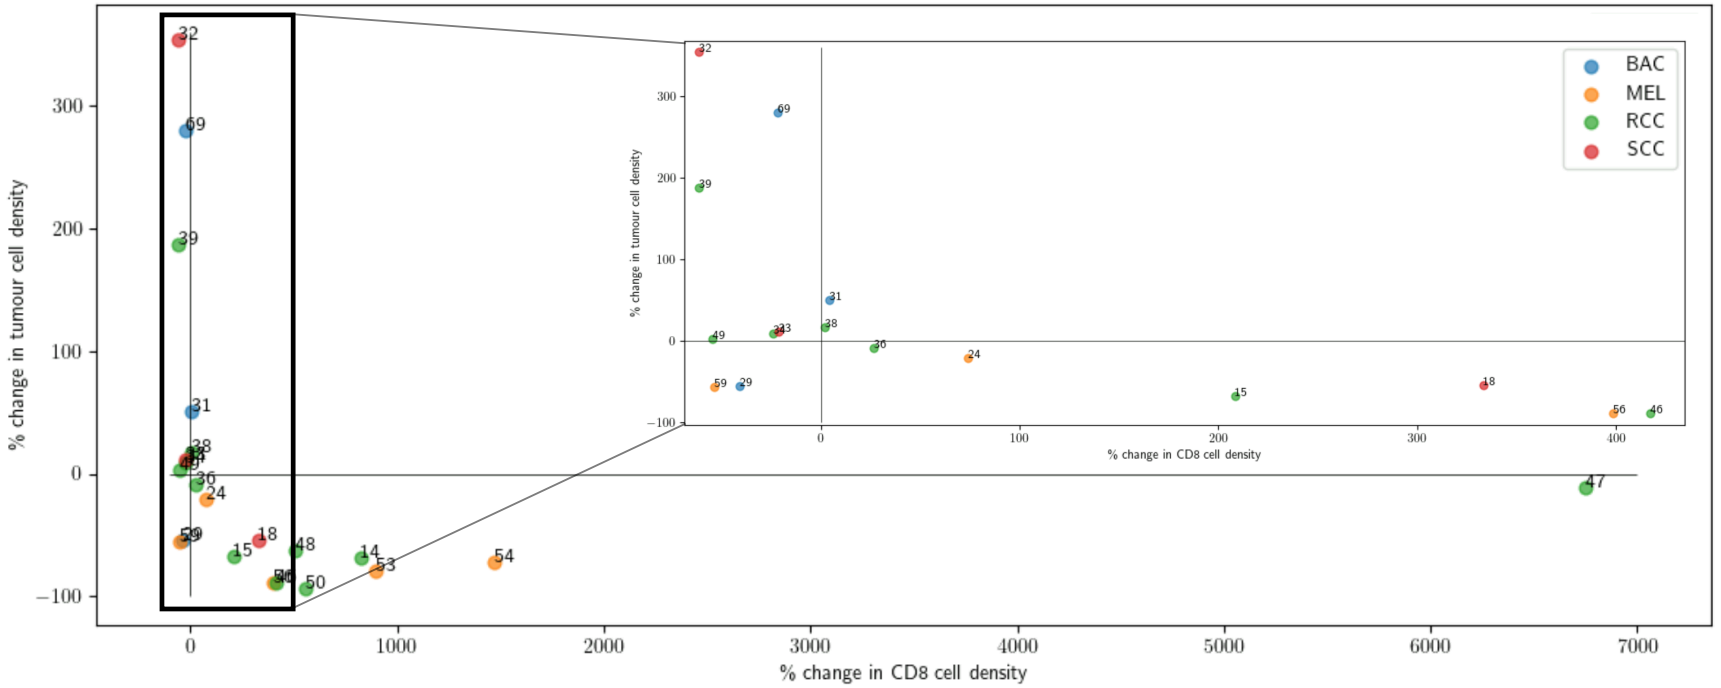


Figure 3.26: Percentage change in tumour cell density from pre- to on-treatment plotted against percentage change in CD8 cell density. Each data point corresponds to the mean of ROIs for an individual patient. Each data point is coloured by the primary indication of the patient's tumour, and labelled by patient ID. Primary indication key: BAC (Breast adenocarinoma), MEL (melanoma), RCC (renal cell carcinoma), SCC (squamous cell carcinoma).

Renal cell carcinoma (RCC) is the largest primary indication group within this dataset, with 11 patients. We see a wide range of responses, with a poor response (decrease in CD8 and increase in tumour cells) in 3 out of 11 patients, and a good response (increase in CD8 and decrease in tumour cells) in 7 out of 11 patients. One patient (patient 38) showed an increase in both tumour and CD8 cell densities on-treatment compared to pre-treatment.

Recall that all melanoma patients showed a decrease in tumour cell density. By contrast, the tumour responses of the RCC patients were varied, with some tumour cell densities decreasing in the on-treatment sample, and others increasing by more than 20%. Conversely, the CD8+ T cell densities in two of the melanoma patients did not show significant change on-treatment, whereas the majority of the RCC patients who responded in a decrease in tumour cell density, exhibited increases in CD8+ T cell density of 4-5 times their pre-treatment levels.

We clustered these ROI cell density measurements to investigate whether ROIs from patients with the same primary indication cluster together, and whether the tumour composition varies between primary indications, pre- and on-treatment. Four clusters were identified for the pre- and on-treatment samples (Figures 3.27 and 3.28, respectively).

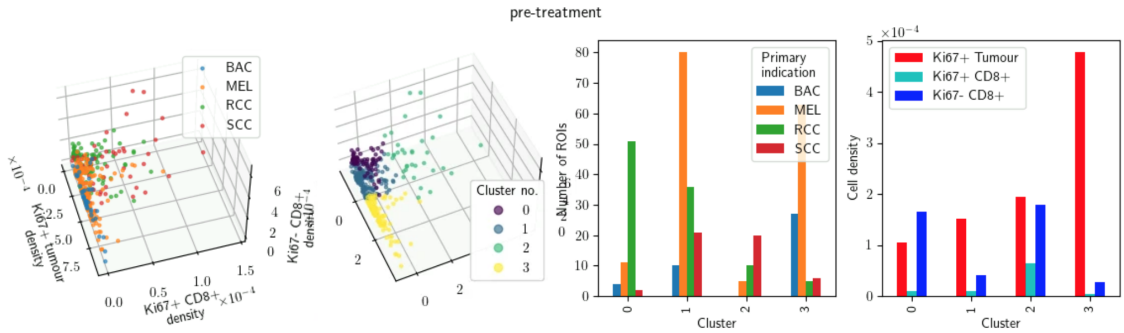


Figure 3.27: Clustering of cell density measurements for the four largest groups of primary indication, for pre-treatment samples. Each data point corresponds to a single ROI. The distribution of ROIs across the clusters, with the cell density measurements corresponding to the cluster centres, is plotted.

Pre-treatment, each cluster contains a mixture of primary indications, with ROIs from all indications present in all clusters, except cluster 2. This cluster corresponds to ROIs with high CD8 proliferation and does not contain any ROIs from Breast Adenocarcinoma (BAC) patients. This finding is also reflected in the absence of BAC ROIs in on-treatment clusters 1 and 3, both of which contain high densities of CD8 cells. These results suggest that either this disease type does not respond to the treatment, or that the treatment has not had enough time to take effect; the three BAC patients' on-treatment samples were taken at 15, 16 and 24 days after the first dose, compared to a mean time of 44 days over

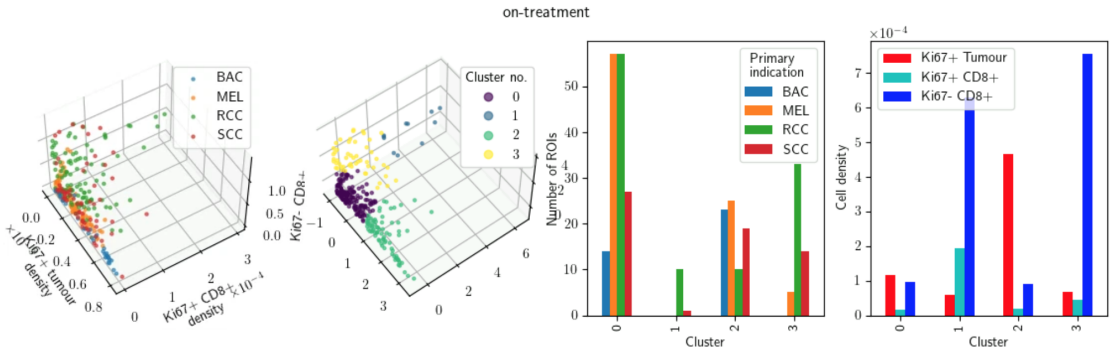


Figure 3.28: Clustering of cell density measurements for the four largest groups of primary indication, for on-treatment samples. Each data point corresponds to a single ROI. The distribution of ROIs across the clusters, with the cell density measurements corresponding to the cluster centres, is plotted.

the cohort (see Figure 3.29). However, with the small sample sizes from these groups of primary indications there is risk of sampling error.

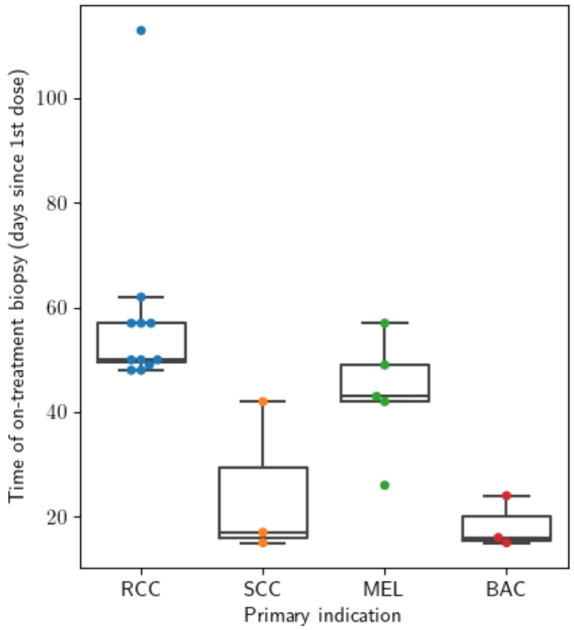


Figure 3.29: Time of on-treatment biopsy (days since first dose), by primary indication. Each data point represents an individual patient.

We will now analyse these ROIs spatially using the PCFs and cross-PCF, to investigate whether CD8+ T cells localise closer to tumour cells in melanoma patients than in RCC patients. We calculate the PCFs and cross-PCFs for each ROI that contains at least 25 cells in the "target" and "source" cell population. If there are fewer than 25 cells then the PCF is not calculated and is set to the mean of the other ROIs, so that it does not affect the

PCA. The three PCF and cross-PCF vectors are concatenated with the cell density data into a single "feature" vector for each ROI. The feature vectors from the melanoma and RCC patients are partitioned into two clusters, and projected onto the first three principal components in Figure 3.30.

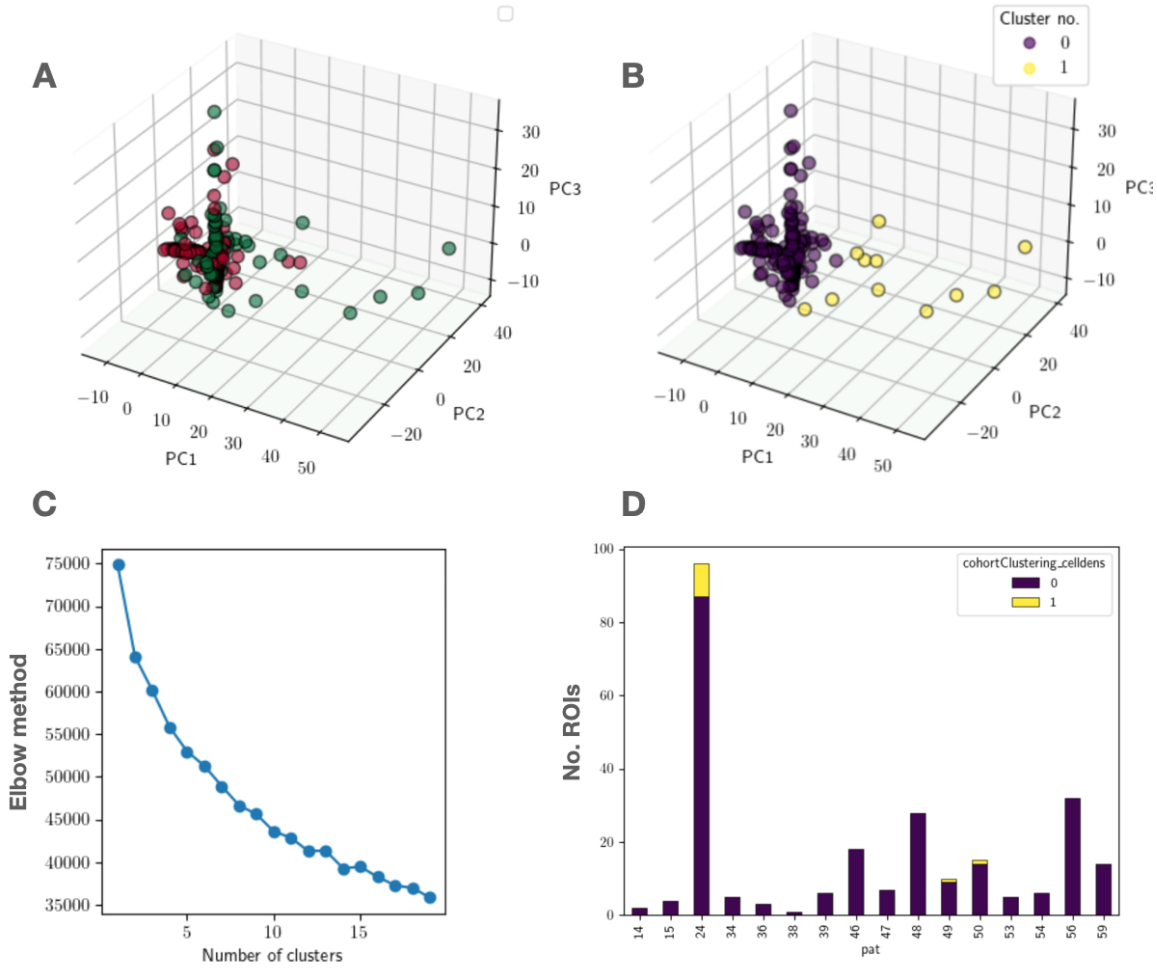


Figure 3.30: (A) Clustered ROI data for melanoma and RCC patients projected onto the first three principal components. Coloured by primary indication: Melanoma (green), RCC (red). (B) ROI data coloured by cluster membership. (C) Elbow plot does not clearly show an optimal number of clusters, so two clusters were chosen, to reflect the two groups of primary indications. (D) Distribution of ROIs in each cluster for each patient.

Figure 3.30(A) shows the ROIs coloured by primary indication. Some melanoma ROIs are at distance from the main cluster which contains a mixture of RCC and melanoma ROIs. Figure 3.31 shows the PCFs and cell densities of the cluster centroids, and we see that the melanoma ROIs at a distance from the main cluster are characterised by strong colocalisations of CD8 cells (shown by a larger peak in CD8-CD8 PCF) and exclusion of CD8 cells at short distances from tumour cells (shown by the TU-CD8 PCF having values

less than one). Perhaps this spatial feature means that the melanoma patients' CD8 cells are located away from the tumour cells (perhaps in the stromal regions) pre-treatment, so therefore are optimally located to be stimulated by the treatment as it localises within the stroma.

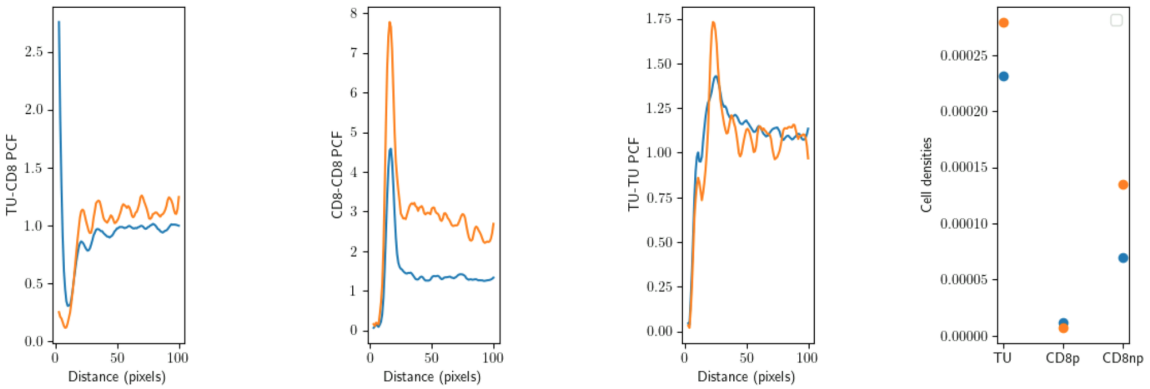


Figure 3.31: PCFs and cell densities of the centroids of the clusters. Blue PCF=purple cluster, orange PCF = yellow cluster.

The variation of the data explained by the first five principal components (PCs) in order are {20%, 15%, 7.8%, 5.9% and 3.6%}, and we note that these PCs only represent 50% of the variance in the data. Interestingly, when we analyse the distribution of the loadings of the first five principal components (PCs) over the feature vector of concatenated cell densities and PCFs in Figure 3.32, we see that the first three PCs are each correlated with one of the PCFs. For example, PC1, the PC which explains the most variation in the data, is dominated by the CD8-CD8 PCF. This suggests that the distribution of CD8 cells relative to each other is the most important feature in separating the pre-treatment data for RCC and melanoma patients. This could potentially be detecting stromal regions in the ROIs, with clustered regions of CD8 cells, as in the cell density analysis in Section 3.5.4.

Figures 3.33, 3.34 and 3.35 correspond to the clustering of the on-treatment ROIs. In the on-treatment clustering we find the opposite result to that of the pre-treatment clustering - most ROIs have CD8 cells located further away from tumour cells (cluster 0 and 2) - this could be due to the localisation of the drug in the stromal regions causing CD8 proliferation at these larger distances from tumour cells. Some patients have co-localised CD8 and TU cells (cluster 1), which may be those patients for which the CD8 cells have left the stroma and are well-infiltrated in the tumour. Some of the patients with ROIs in cluster 1 correspond to those who experienced a decrease in tumour cell density without large expansion of the CD8 cell population, for example patient 24.

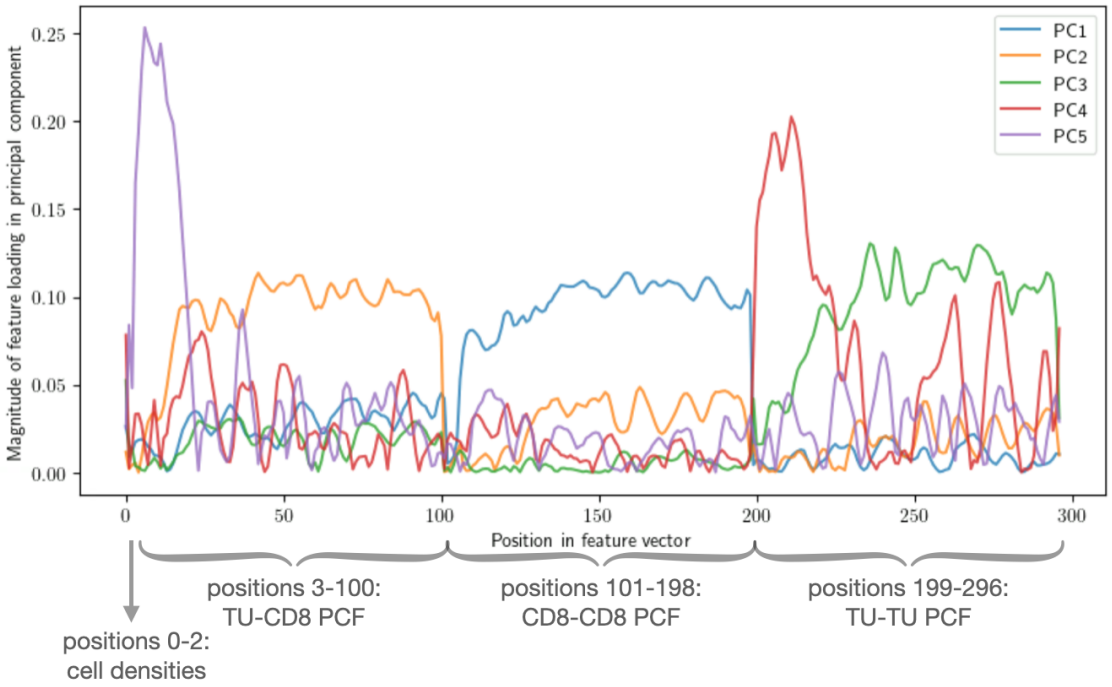


Figure 3.32: Magnitude of the PC loadings with respect to each feature in the feature vector (concatenated cell densities and PCFs) for the first 5 PCs of the dataset containing RCC and melanoma samples, pre-treatment.

3.6 Robustness of findings to changes in methodology

Many decisions regarding the methodology used in this analysis were based on limitations of the data. For example, square ROIs of dimension 800 distance units were chosen from inspection of the ROI sampling for different ROI sizes. This size of ROI enabled more patients to be included in the analysis, as the size of the ROIs were often limited by the shortest axis length of the tissue pieces, in order to exclude tissue edges in the sampling. Smaller ROIs enabled more of the tissue area to be sampled, but these ROIs contained fewer cells, which affected the statistical significance of the results. Conversely, larger ROIs containing more cells per ROI produced more robust statistics, but in many cases less of the tissue area could be sampled due to the irregular shapes of the tissue samples.

As there is no clear "optimal" ROI size, it is not clear whether the ROI dimension we have used provides a representative sampling of the original tissue sample. To investigate how representative the ROIs are, we now calculate the average statistics for each image and compare them to the statistics computed for the set of sampled ROIs for that image. There are several ways to calculate image-averaged statistics. For the cell density measurements, it is straightforward to use the tissue annotations to find the total area of tissue in the

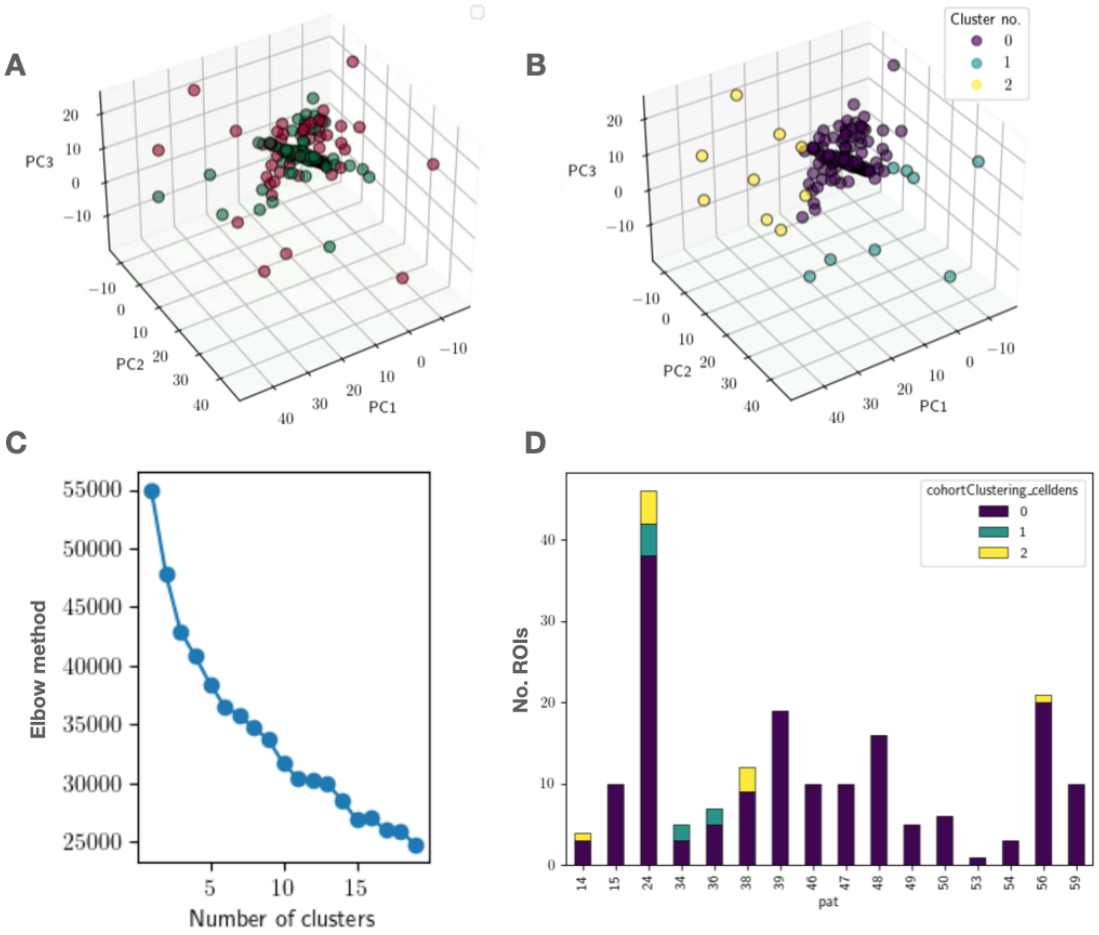


Figure 3.33: (A) Clustered ROI data for melanoma and RCC patients projected onto the first three principal components. Coloured by primary indication: Melanoma (green), RCC (red). (B) ROI data coloured by cluster membership. (C) Elbow plot shows a "bend" at 3 clusters. (D) Distribution of ROIs in each cluster for each patient.

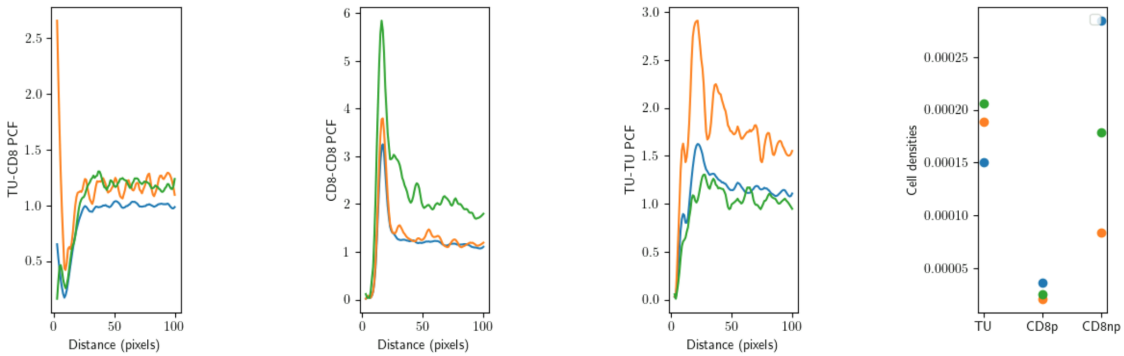


Figure 3.34: PCFs and cell densities of the centroids of the clusters. Blue=cluster 0, orange = cluster 1, green = cluster 2

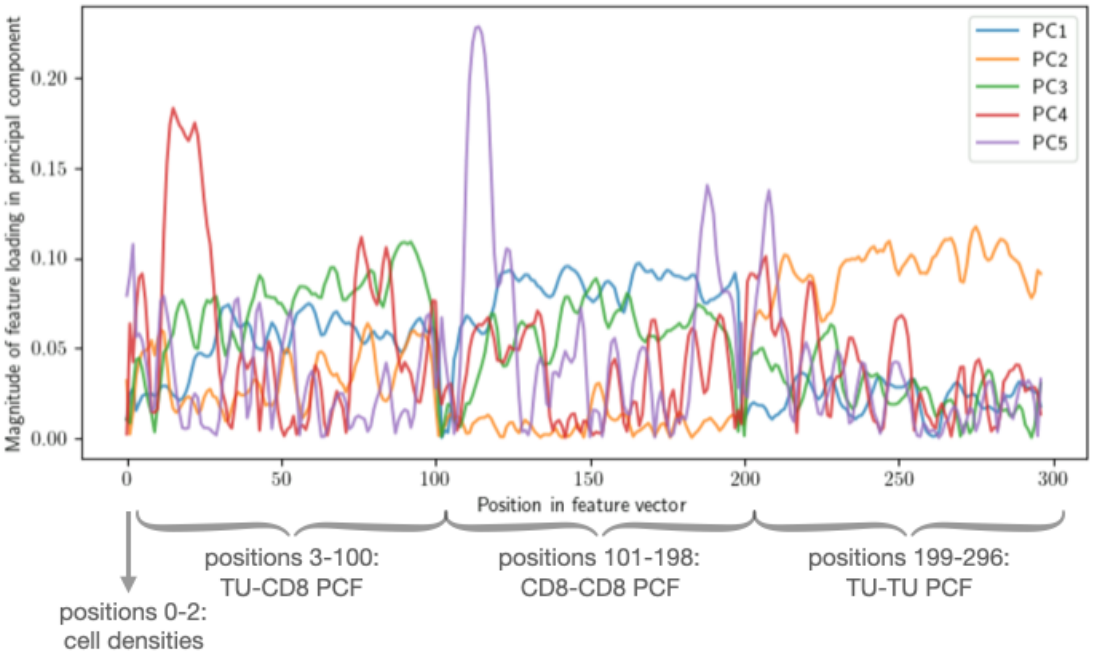


Figure 3.35: Magnitude of the PC loadings with respect to each feature in the feature vector (concatenated cell densities and PCFs) for the first 5 PCs of the dataset containing RCC and melanoma samples, on-treatment.

image and total number of cells of each type. However, edge effects and artefacts near the tissue edges can confound the results.

3.6.1 Image-averaged statistics

The image-averaged cell density, ρ_T, for cell type T, is calculated using Equation 3.4, for an image containing N_p tissue pieces, with each tissue piece having area A_i and containing N_i cells of type T.

$$\rho_T = \frac{\sum_{i=1}^{i=N_p} N_i}{\sum_{i=1}^{i=N_p} A_i}. \tag{3.4}$$

PCF calculations for an image are calculated using Equation 3.5:

$$g_{ST}(r) = \sum_{n=1}^{N_P} \frac{N_S}{N_{S_{total}}} \frac{1}{N_S} \frac{1}{N_T/A} \sum_{i=1}^{N_S} \sum_{j=1}^{N_T} \frac{I(|\mathbf{x}_i - \mathbf{x}_j|)}{A_r(\mathbf{x}_i)}, \tag{3.5}$$

where $A_r(\mathbf{x}_i)$ is defined as the intersection of the annulus of radius r with the tissue-tumour intersection shape, i.e. the area of the annulus that lies within the tissue sample - see Figure 3.7. This modification accounts for tissue edge effects.

3.7 ROI representativeness of whole image

As many non-overlapping and tissue-edge-excluding ROIs as possible were sampled from each image. Figure 3.36 shows the distribution of the number of ROIs sampled over the cohort. Each data point represents the number of ROIs sampled from the pre-treatment and on-treatment images for an individual patient, showing that there is a large variation in the number of ROIs which were accepted for samples across the dataset.

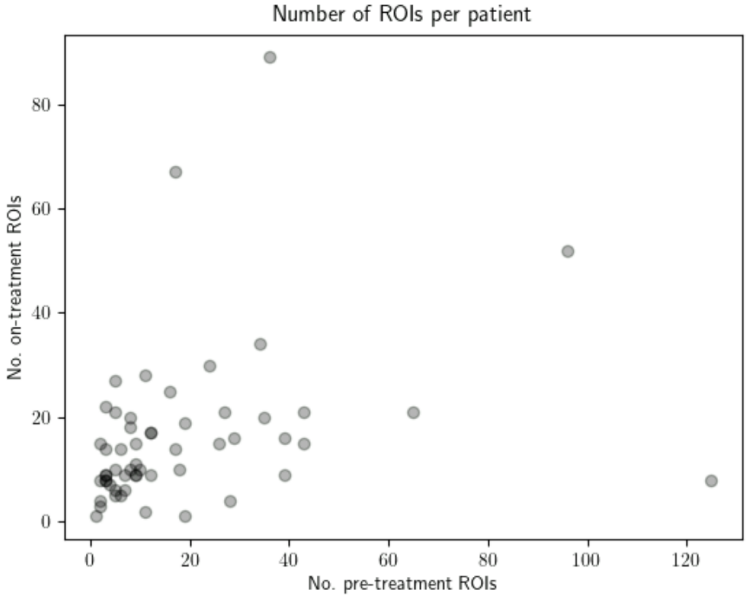


Figure 3.36: Distribution of number of ROIs sampled from each patient, for pre- and on-treatment; each point represents an individual patient. There is large variability in the number of ROIs that are able to be sampled from a given image due to the disparity in the size of the samples.

In Figures 3.37 and 3.38 we compare image-averaged change in cell densities with mean change in cell densities over the sampled ROIs, between pre- and on-treatment. The data suggests that for the majority of the patients, ROIs are a good representation of the whole sample, although there are a few outliers. The outliers may correspond to tissue samples where there were too few ROIs to capture the intra-patient heterogeneity, or samples which contain necrotic regions (with low cell densities) which were not marked by the annotations.

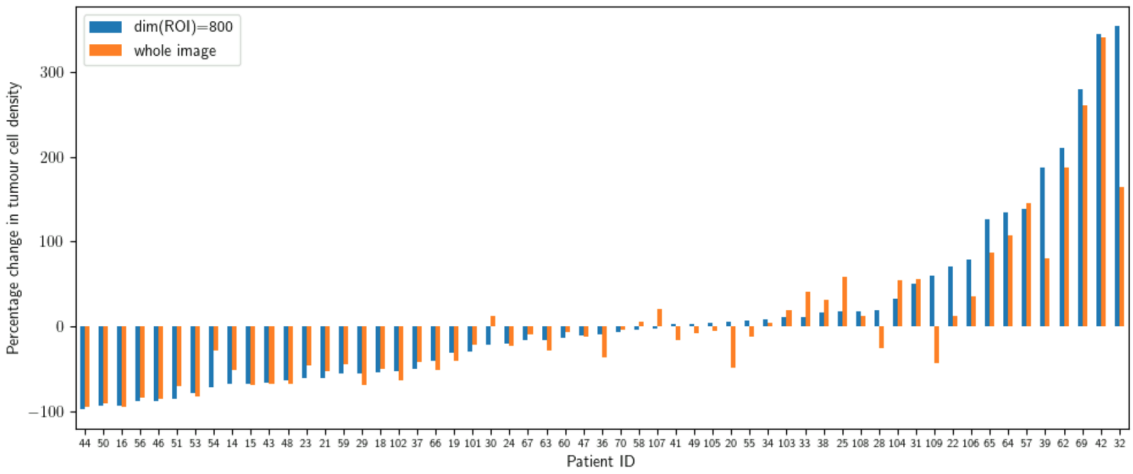


Figure 3.37: Image averaged change in density of proliferating tumour cells (which takes into account all of the data in the image) and the mean value of proliferating tumour cell density change of the ROIs sampled from that image.

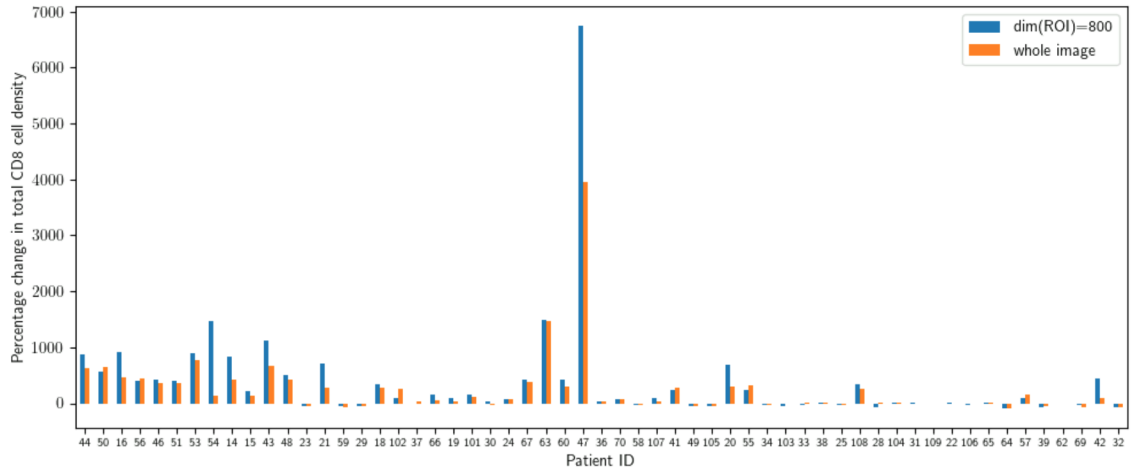


Figure 3.38: Image averaged change in density of total CD8+ T cells (which takes into account all of the data in the image) and the mean value of total CD8+ T cell density change of the ROIs sampled from that image.

3.8 Addressing some of the issues with the data suitability for these analyses

The analysis in this chapter has highlighted specific points relating to the suitability of core-needle biopsy samples to these analysis methods. Firstly, many tissue samples have a long, thin morphology, making it difficult to extract square ROIs which are small enough to fit into the tissue without edge inclusion, and large enough to contain an adequate number of cells on which to calculate the PCFs.

Other studies of core needle biopsy image data have used machine learning methods

to identify spatial correlations and patterns within the data, and usually divide the image up into a grid. In these studies, the datasets usually contain more samples and so tissue edge effects in the ROI sampling may be less significant, or the samples were larger, so ROIs containing mostly empty space could be discarded whilst maintaining a large enough number of ROIs to analyse. The small sample sizes in our dataset mean that our ROI sampling method must capture as much of the tissue area as possible. However, at the same time this needs to be balanced with capturing the heterogeneity of the sample by sampling multiple ROIs. Nir et al. (2018) discuss this in their analysis of prostate cancer biopsy images, and adopt the method of sampling ROIs at different scales and analysing these ROIs in their machine learning pipeline. However, due to the small dataset in our analysis, machine learning methods may not be suitable due to the need to have a sufficient training dataset, and many of the samples could not accommodate ROIs greater than 800 x 800 pixels, which is our default ROI size in this work.

This issue calls for an alternative ROI sampling method which maximises the information contained in biopsy samples, whilst minimising the inclusion of tissue edges. We have developed a method to generate user-defined ROIs and calculated statistics on these ROIs. Figure 3.39 shows an example of a user-generated ROI.

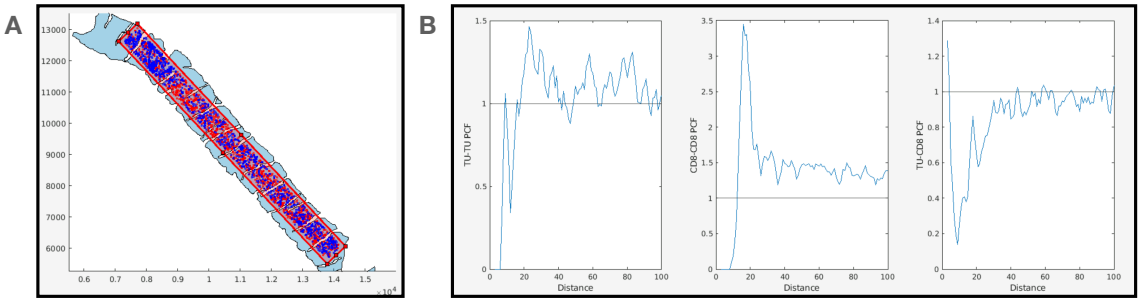


Figure 3.39: (A) Example of a user-defined ROI specific to the morphology of the tissue sample. (B) PCFs corresponding to the user-defined ROI.

The ROI is interactively drawn and re-sized within the image, and rotated to align with the orientation of the tissue sample to maximise the amount of tissue sampled, then double-clicked to generate the PCFs for the cells contained within the ROI. The method in Section 3.4.7 of sampling an ROI (inner region) with an outer margin was adapted to generalise to any rectangular shape. The code for the ROI analysis was adapted to account for any tissue edge inclusions within the ROI, such that if there was empty space in the ROI, then the PCF was weighted for the cells at the tissue edges, by calculating the overlap between the annuli and the tissue boundary. The method checks for multiple tissue fragments within the ROI, and will not perform the spatial analysis if this is the case.

The image analysed in Figure 3.4 illustrates some of the limitations in our automated ROI sampling algorithm: for small, sparse samples, the algorithm may not sample the ROIs which would optimise the amount of data extracted. The user-defined ROI method presented here would mitigate this by allowing the user to have control over which specific regions they wish to be analysed.

Another limitation of our automated ROI sampling algorithm is that it is not optimised for typical samples of this dataset with tears and long, thin morphology - the square ROI algorithm is adapted to allow for edge tolerance and overlap, but cannot be optimised for individual samples. Using our method, a clinician would be able to define which areas are of particular interest to further explore. This larger ROI could be divided up into a grid if investigating heterogeneity within the tissue. As this method involves manual input, it would only be suitable for use on a case-by-case basis, or in the case of small clinical trial cohorts, such as those considered in this work. This custom ROI method for optimising the area of tissue sampled for this particular image illustrates how the sampling method may affect the spatial statistics calculated. Figure 3.40 shows the PCFs calculated on the single ROI which was accepted for the sample, following the square ROI tiling protocol.

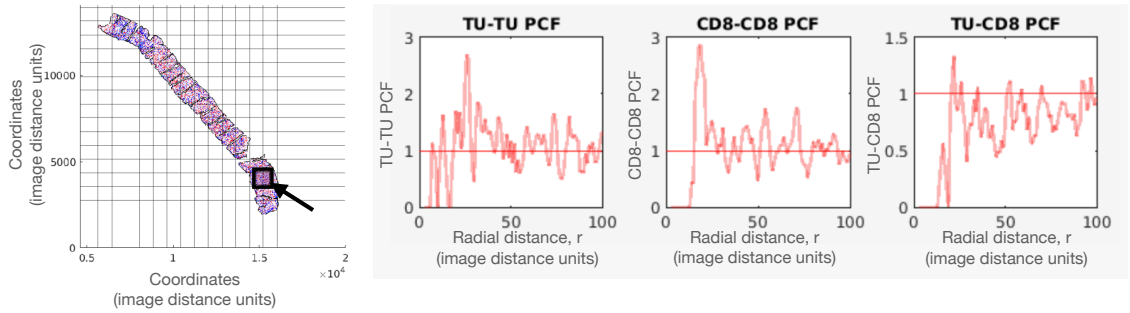


Figure 3.40: PCFs for the single square ROI (indicated by the arrow in first plot) which was accepted for the sample which is also analysed in Figure 3.39.

Comparison with Figure 3.39 shows a difference in the qualitative behaviour of the TU-CD8 PCF, with short-range co-localisation found in the custom ROI analysis, but not in the automated square ROI which samples only a small fraction of the total tissue area. This highlights that the choice of ROI sampling method impacts the spatial analysis of the ROIs. However, the custom ROI contains some empty space and tissue edges, and further work is needed to quantify their effect on the PCFs. Although the custom ROI method can include tissue edges within the sampled region, the user can control which edges are included.

Future work on this custom ROI method could involve developing an interactive ROI, such that the user could scan the ROI across a sample, with dynamically updating PCFs

plotted. This could complement other metrics, such as cell densities, and other spatial statistics displayed alongside the image.

3.9 Discussion

In this chapter we analysed image data from core-needle biopsies sampled from patients on Phase I clinical trials of Simlukafusp, an immunotherapy designed to promote proliferation of CD8+ T cells within the tumour microenvironment. We developed a pre-processing pipeline to analyse the image data, incorporating pathologist annotation data to identify tissue regions suitable for analysis. This pipeline also extracts from the samples non-overlapping ROIs which do not include tissue edges or holes. Our pipeline can analyse images which contain multiple tissue fragments, as is often the case with such samples. Our algorithm generates ROIs that are better suited for subsequent analysis than ROIs obtained by simply partitioning the tissue sample using a grid format. Our pipeline calculates cell densities and PCFs for ROIs and generates averaged statistics for each image.

Paired pre-treatment and on-treatment samples from each patient were initially analysed using cell densities and categorised into three groups; good responders, no change and poor responders, according to the change in proliferating tumour cell density between pre- and on-treatment samples. We found that, for a threshold response of more than 10% decrease in proliferating tumour cell density, fewer than half of the patients exhibited a good response, indicating that the drug was not effective in most of the patients.

Analysis of the data was challenging for several reasons. Firstly, we had no information regarding the locations and abundance of non-proliferating tumour cells. Secondly, the dataset contained many different primary indications and on-treatment sampling time-points. Furthermore, due to the biopsy collection method, the samples were small, and artefacts such as tears and folds in the tissue can arise as a result of the collection, processing and imaging processes of the sample, leading to exclude annotations which cannot be reliably analysed. We found that some images could not be analysed due to the lack of tissue, or because of artefacts.

As many of the tissue samples are typically long and thin in shape, increasing the ROI size would necessitate the inclusion of tissue edges, which may introduce artefacts to the statistics. Further work is needed to quantify the effects of including tissue edges in the analysis of the ROIs. Further work is also needed to investigate how changes in the methodology affect the conclusions of the analysis. For example, the minimum cell threshold for ROI analysis and PCF calculations may be optimised such that more data can be included in the analysis, whilst minimising the inclusion of noise. This threshold could be determined by generating synthetic point clouds with different spatial patterns (i.e.

different PCFs), and subsequently removing points at random, decreasing the cell number with which to calculate the PCF and observe to what extent this affects the PCF.

If further data with larger samples were available then other spatial statistics, graph-based methods, or topological data analysis (TDA) could be used to enrich the spatial characterisation of the tissue samples.

The work in this chapter has highlighted many questions regarding whether this type of data (core-needle biopsies) are suitable for this analysis, and the extent to which it is possible to reliably classify treatment responders. This exploratory analysis has highlighted additional types of data which could improve the analysis. As the drug is designed to bind to FAP-expressing cells, most of which are found in the stroma, it would be useful to have a stromal marker that distinguishes tumour nests and stromal regions. These areas could then be sampled and analysed separately, in order to determine if there is an increase in CD8 cell proliferation within the stromal regions. It would also help to characterise the heterogeneity in the spatial distribution of tumour and stromal regions for each patient, and see if there is a correlation between this distribution and treatment efficacy.

This work has raised questions about how large a tissue sample is needed for ROI sampling and spatial analysis, and whether one on-treatment biopsy sample is sufficient to assess treatment efficacy. Synthetic datasets generated from mathematical models are a valuable tool for studying such questions, particularly when there is limited clinical data available. In the next chapter, we will use an ABM to generate synthetic data with the aim of addressing the above questions.

Chapter 4

Exploring the robustness of spatial analysis methodology using synthetic and real data

4.1 Overview

In this thesis we have explored the themes of modelling and data analysis separately. In this chapter, we consider how to integrate these components such that biopsy data may inform mathematical modelling, and vice versa. In Chapter 2 we investigated how modelling assumptions can affect the emergent dynamics in an on-lattice ABM of tumour-T cell interactions. We used PCFs to quantify the spatial distribution of the cells at 50% tumour burden and showed that spatial analysis of ABMs can explain behaviour that cannot be determined from cell counts alone. We showed that spatial analysis allows for a qualitative comparison of the behaviours exhibited by the model as parameter values vary. This analysis raised several issues regarding the model's suitability for simulating tumour-immune interactions, and we suggested possible improvements to the model. In Chapter 3 we conducted an exploratory analysis on a dataset of paired biopsy images from a clinical trial of Simlukafusp, an immunotherapeutic which is intended to localise in the tumour microenvironment and stimulate CD8+ T cell proliferation. We highlighted the challenges of analysing disparate data sets which are often generated in clinical trials, and the importance of the on-treatment biopsy timepoint. In this chapter, we extend the work from Chapters 2 and 3 to investigate how spatial analysis can be used to predict ABM dynamics while also exploring the potential for integrating biopsy data with mathematical models for clinical applications.

As discussed in Chapter 3, the number of core needle biopsies that can be sampled from a patient during a clinical trial are limited due to their invasive nature. Usually, patients are limited to one on-treatment biopsy sample, therefore this should be collected at a time that

best demonstrates whether the treatment is having a therapeutic effect. If the on-treatment biopsy is taken too early then a patient may not have started to respond to the treatment. Conversely, if the on-treatment biopsy is taken too late, a non-responding patient may lose vital time before the unsuccessful treatment is halted and they are offered an alternative treatment. Clearly there is a need for a more comprehensive and systematic approach to biopsy scheduling in this setting.

Hutchinson and Grimm (2022) proposed a pipeline which integrates sampled ROIs from pre- and on-treatment images, and uses them to optimise the parameters of an ABM. The ABM was then simulated to produce a synthetic timecourse of the spatial dynamics of tumour cells and CD8+ T cells to inform the scheduling of on-treatment biopsies. A subset of the paired patient data was used to optimise selected parameters of the ABM by comparing the CD8+ T cell PCF of the on-treatment and simulated ROIs at the on-treatment timepoint. The optimised parameters were those pertaining to CD8 proliferation, CD8 (natural) death, bias in CD8 random walk, and CD8 killing of tumour cells. All other parameters were fixed at default values. These optimised "population parameters" were then combined with other patients' pre-treatment samples to simulate their tumours' growth dynamics. The predicted cell distributions at the on-treatment timepoint were compared to the on-treatment data for these patients to assess the accuracy of their method for reproducing the on-treatment spatial distribution of CD8 cells, and they found agreement in 70% of cases.

The work presented in this chapter approaches the biopsy scheduling problem from a different angle, which does not require optimisation of ABM parameters. In the pipeline presented by Hutchinson and Grimm (2022), model parameters were optimised using a training set of paired pre- and on-treatment images, so optimisation of the model to the on-treatment biopsies is carried out retrospectively for some patients. In a clinical trial you would not necessarily initially have the data to train the model before needing a quantitative decision for the on-treatment biopsy. Our method aims to inform on-treatment biopsy scheduling in cases where there is no training set available, for example in early-stage clinical trials. In later trials, clinicians could then choose to supplement our method with optimisation of ABM parameters (such as in Hutchinson and Grimm, 2022) using data from previous clinical trials.

We focus on the following questions:

- Is the PCF robust across multiple stochastic realisations of an ABM for a given parameter set at individual timepoints?

- How does the PCF change as ABM parameters vary? Can the PCF distinguish stochastic realisations from different parameter sets? Are there any particular PCF metrics which are indicative of simulated tumour outcome?
- How does the PCF evolve over time? At what timepoint are the synthetic tumour outcomes distinguishable, and to what accuracy? How does the separation timepoint depend on patient-specific initial conditions sampled from real data?

This chapter is structured as follows. In Section 4.2 we recap the ABM structure used in Chapter 2, describe the model adaptations we implement, and outline the methodologies we use to analyse the ABM data. In Section 4.3, we conduct a coarse parameter sweep of the model to assess the robustness of the PCF analysis across a range of parameter regimes which capture the different tumour outcomes that the ABM exhibits. In Section 4.4 we investigate the time evolution of the spatial analysis of the ABM data, and ascertain a minimum timepoint at which tumour outcomes separate with a specified accuracy, based on classifiers trained with both spatial and cell count data. In Section 4.5 we then take multiple ROIs sampled from a single patient's pre-treatment biopsy image and analyse the variation in predicted tumour outcome in the ABM for a single parameter set. In Section 4.6 we discuss how this methodology may be integrated with real data, such as the Simlukafusp dataset, and demonstrate that the accuracy of the ABM predictions is dependent on the initial condition. Finally, we summarise our findings and discuss future work in Section 4.8.

4.2 Methods

4.2.1 Protocol

In this chapter we propose a new protocol for determining when to schedule the on-treatment sampling of core needle biopsies in oncology clinical trials, based on an existing method for identifying the 3 E's of immunoediting in the output of an ABM (Bull and Byrne, 2023). This protocol is outlined in Figure 4.1. We aim to tackle the problem of predicting the optimal timepoint for on-treatment biopsy scheduling by investigating the limitations of using a given model in predicting tumour outcomes. Instead of fitting parameters using limited timepoint data, we simulate the ABM, sampling from the possible parameter space, and use spatial statistics to train a classifier at different timepoints. In so doing, we aim to determine whether the metrics extracted from the ABM data can distinguish between tumour outcomes, and if so, at what timepoint(s) this separation occurs.

The ABM is not parameterised using patient data. Instead, ABM simulations are generated across a wide range of realistic parameter values. However, this approach maintains patient personalisation as the ABM is initialised using ROIs sampled from each pre-treatment

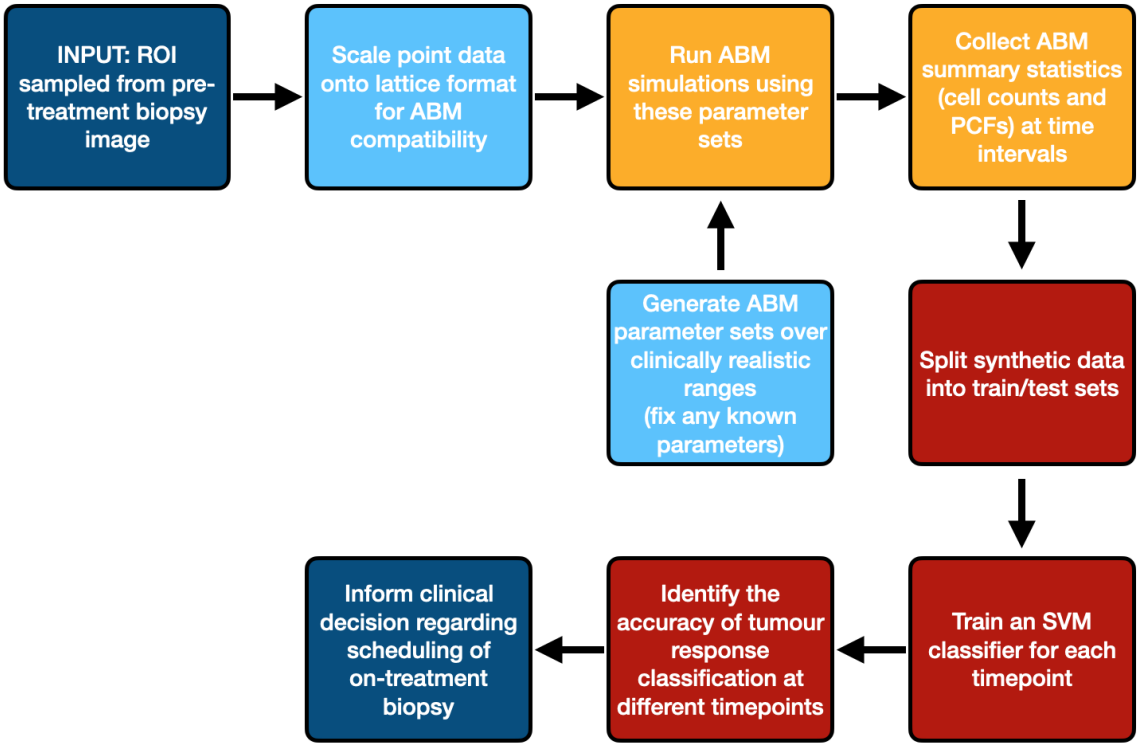


Figure 4.1: Schematic of proposed method for simulating ABM timecourses initialised by biopsy ROIs, to ascertain the minimum timepoint at which tumour outcomes separate up to a specified accuracy.

biopsy image. Sampling across parameter space can be achieved using techniques such as Latin Hypercube sampling. This method optimises the number of samples needed whilst ensuring the parameter space is adequately sampled. The model considered in this work includes nine parameters, resulting in a high dimensional parameter space to sample. However, our method should not be more computationally expensive than optimising ABM parameters as both methods require sampling from parameter space. Known parameters (for example, parameter values ascertained from *in vitro* studies) may be fixed to reduce the dimensions of the parameter space being sampled.

4.2.2 The model

We want to use a simple model to perform this analysis. Recalling that the model analysed in Chapter 2 was developed to specifically compare to the results of Hutchinson and Grimm (2022), we now make some changes to the model in Chapter 2. In summary, we change the timestep to be equal to one CD8+ cell migration step, and we use a Gaussian filter over the tumour cell positions to build a map to inform the CD8+ cell biased migration. These changes are detailed as follows.

Occupancy rules: Single occupancy rules allowing only one cell per lattice site are retained to maintain model simplicity. A 100 by 100 cell diameter lattice is used.

Initial condition: The same initial condition with 8000 tumour cells and 500 CD8+ T cells randomly distributed on the lattice is used to initialise the model.

Biased motion of CD8 cells towards tumour cells: We decided to use a Gaussian filter from the image processing toolbox in Python, and applied it to the lattice sites occupied by tumour cells, yielding an interaction kernel to inform CD8 movement. We note that other approaches, such as modelling a chemoattractant using a PDE, with tumour cells as sources and CD8+ cells as sinks could also have been used, however we justify this approach to make the interaction kernel more generalisable and simple to implement. The map value computed by the kernel at each lattice site is between 0 and 1. We set the standard deviation of the Gaussian kernel to 1, and truncated the kernel such that CD8+ cells can only detect values above 0.00001, to incorporate a threshold of sensitivity of the CD8+ cells to the signal. See Appendix D for more details. Before a CD8+ cell migrates, if it is marked to move in a biased way towards tumour cells, it will select an empty neighbouring lattice space with the largest value in the Gaussian kernel. Otherwise, the CD8+ cell randomly selects an empty neighbouring lattice space to move into.

Timestep: We employ a timestep on the scale of the fastest process in the model, which is CD8+ T cell migration. Each migration step of one cell diameter takes approximately 0.12hr (Boldajipour et al., 2016). We refer back to Chapter 2 for the algorithm used to simulate the model.

Simulation runtime: We run simulations for a timeframe of 150 days or, in the case of tumour elimination, until the simulation is terminated. The simulation end time of 150 days was chosen to limit the computational time and make the results relevant to the on-treatment biopsies; these were all sampled less than 150 days after the first treatment dose in the Simlukafusp dataset.

Model parameters: Parameters are taken from the 0.12 hr time step case in Chapter 2, and are provided in Table 4.1. All probabilities were scaled down for the 0.12hr timestep as described in Chapter 2, except the CTL proliferation probabilities, which were approximated by linear scalings for clarity purposes when plotting; the order of magnitude of these probabilities is still relevant to using other parameter scalings.

4.2.3 Labelling synthetic tumour outcomes

We define the tumour outcomes for the synthetic data, within the clinically relevant timeframe of 150 simulated days, as follows:

Parameter	Description	Values considered in parameter sweep
r_T	Tumour cell proliferation probability	$(1+0.5)^{1/100}-1$
m_T	Tumour cell migration probability	0.0035
d_T	Spontaneous tumour cell death probability	$1-(1-0.12)^{1/100}$
r_E	CTL proliferation probability	{0.0002, 0.00025, 0.0003, 0.00035, 0.0004}
m_E	CTL migration probability	0.8
d_E	Spontaneous CTL death probability	$1-(1-0.015)^{1/100}$
κ	Probability of CTL killing of a tumour cell	0.009
β	Probability of bias in CTL random walk towards closest tumour cell	{0.25, 0.5, 0.75, 1}
η	Speed of CTL movement, corresponding to the number of lattice spaces moved across on a given time step	1

Table 4.1: Baseline parameter values and the range of values considered in the two parameter sweep.

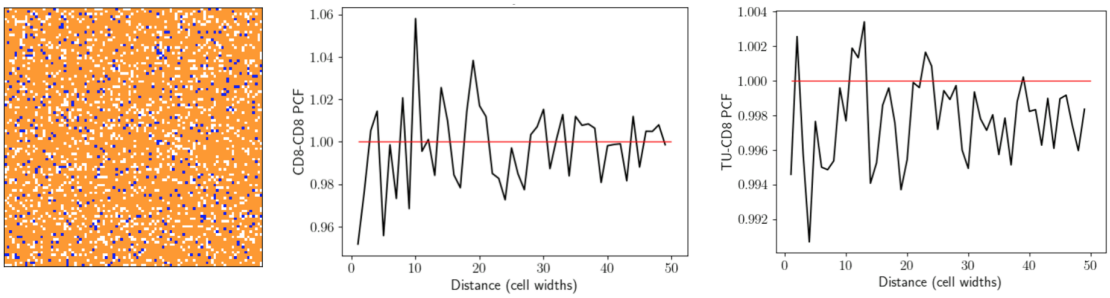


Figure 4.2: Initial condition used to initialise the ABM, which is the same condition used in the work of Chapter 2. The corresponding CD8-CD8 PCF and TU-CD8 PCF are plotted to illustrate the random distribution of each cell type.

- Responder: the tumour cell burden is reduced by at least 50% of its initial value. In these simulations, all synthetic tumours start with 8000 cells at time $t = 0$ days, so we set the responder threshold to be less than 4000 tumour cells at time $t = 150$ days.
- Non-responder: the tumour cell burden is greater than 50% of the initial value (4000

or more tumour cells) at time $t = 150$ days.

Figure 4.3 shows the tumour cell population trajectories for 100 stochastic realisations of the ABM using one parameter set. Responders and non-responders at day 150 are indicated.

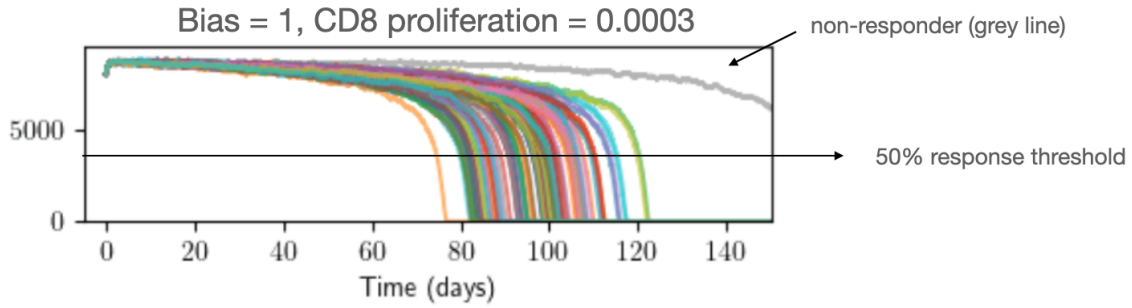


Figure 4.3: Example of tumour cell population number over time for 100 realisations of the ABM for one parameter set ($\beta = 1, r_E = 0.0003$). For this parameter set, all but one realisation is eliminated within 150 days. Each line represents the tumour cell timecourse for one stochastic realisation of the model.

We note that there are other ways to denote tumour response; we use a 50% reduction in tumour burden here as a simple starting point.

4.3 2-parameter sweep, varying CD8+ T cell proliferation rate and bias in CD8+ T cell motion

As discussed in Chapter 3, Bull and Byrne (2023) developed an off-lattice overlapping spheres ABM of tumour-macrophage interactions which simulated the 3Es of immunoediting, and analysed the spatial distribution of cells using the PCF and weighted PCF to distinguish these behaviours. We adopt a similar approach here, with the aim of investigating the robustness of PCFs for cell distributions generated from repeat simulations using the same parameter set. Our parameters of interest are the CD8+ T cell proliferation probability and the bias in CD8+ T cell motion towards tumour cells.

We simulate the ABM for 150 days over 100 stochastic repeats, for each of the 20 parameter sets (the parameter values are stated in Table 4.1). We use the same random initialisation of cells at the beginning of each simulation, instead of randomising the distribution for each repetition. This constrains the degrees of freedom of the system and enables us to focus on variation in the metrics associated with stochastic effects, rather than variation induced by different initial conditions. Figures 4.4 and 4.5 show the time evolution of the tumour and CD8+ cell populations respectively, over the parameter sweep.

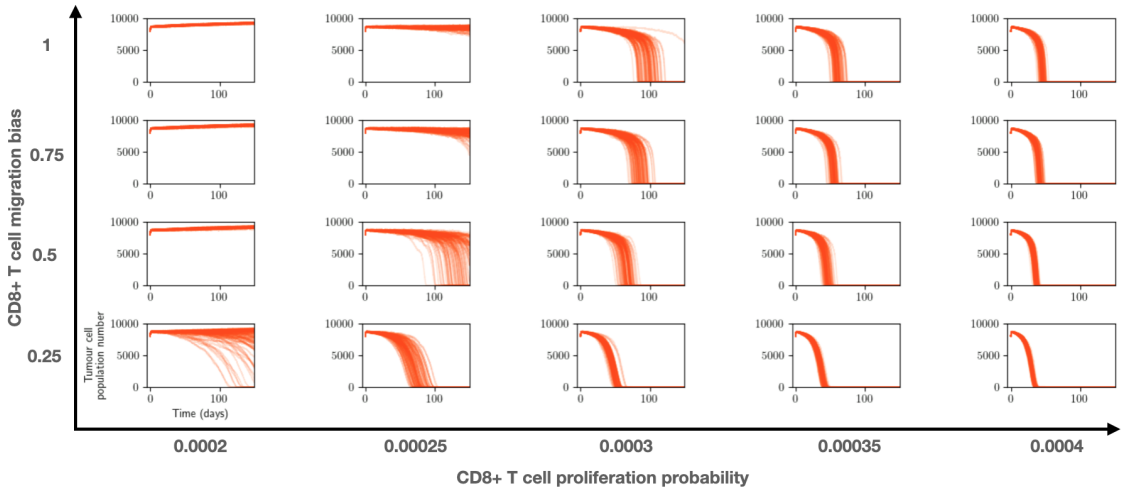


Figure 4.4: Timecourses of tumour cell population for 20 parameter sets, varying T cell proliferation probability and bias in T cell movement. Each line corresponds to the tumour cell population timecourse for a single model realisation.

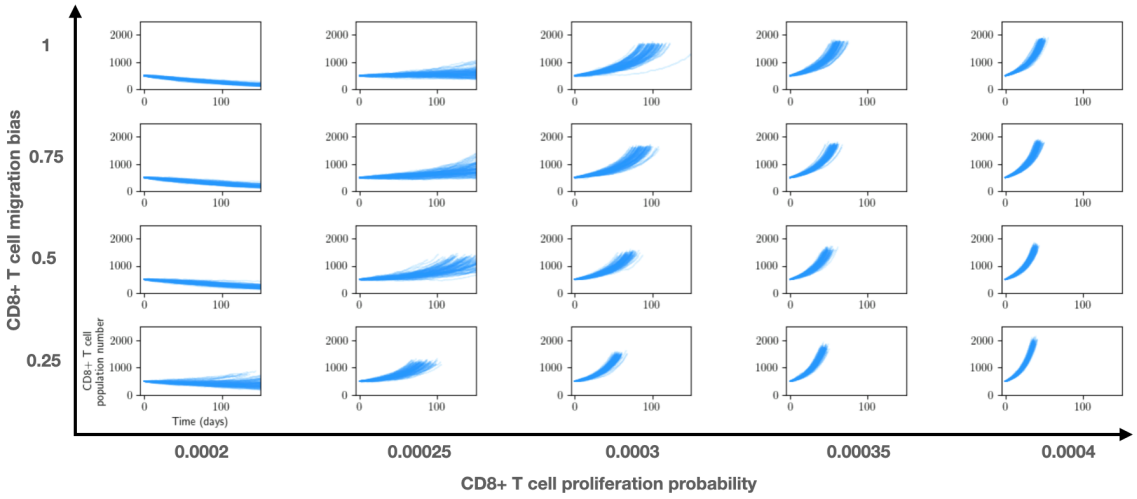


Figure 4.5: Timecourses of CD8+ T cell population for 20 parameter sets, varying T cell proliferation probability and bias in T cell movement. Each line corresponds to the CD8+ T cell population timecourse for a single model realisation.

The timecourses of the tumour cell and CD8+ T cell populations in Figures 4.4 and 4.5 show a wide range of qualitative behaviours when the parameters governing CD8+ T cell proliferation and biased motion are varied (see Figure 4.6).

For parameter sets with low CD8+ proliferation and high motion bias values, we consistently see tumour escape, with increasing tumour cell numbers and a gradual elimination of the CD8+ cell population. Tumour elimination is found in all realisations with high CD8+ proliferation, with movement bias having little effect on the tumour elimination time, as

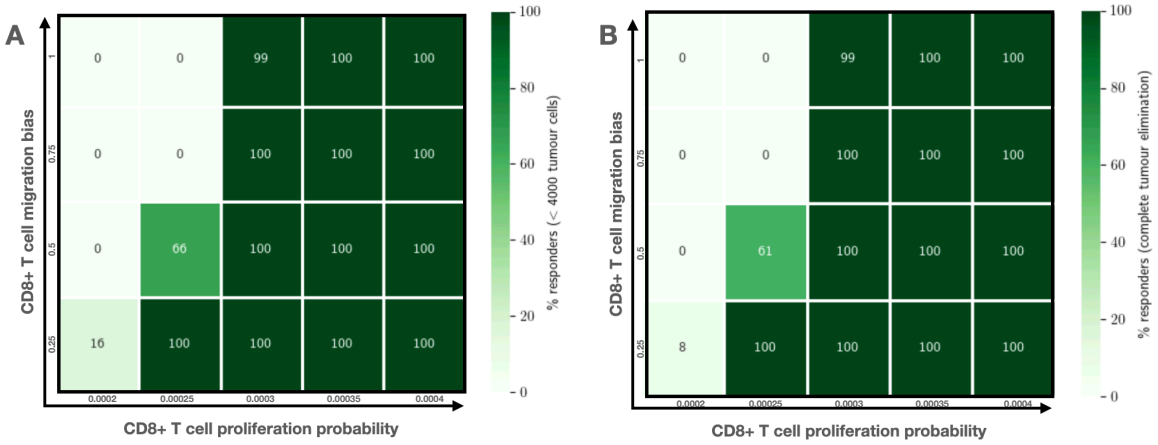


Figure 4.6: % responders for each parameter set

seen in the results of Chapter 2. For intermediate values of CD8+ proliferation, we see bistability in the tumour response across the multiple realisations of the same parameter set. These behaviours are reflected in the CD8+ population timecourses, where we see population growth and equilibrium in the CD8+ T cell population when we see bistability in the tumour response.

Figure 4.7 shows examples of snapshots taken either at day 145, or just before tumour elimination if this occurred before day 145, showing complex spatial cell distributions which will be analysed using the PCF and cross-PCF.

Each panel in Figure 4.8 shows how the CD8-CD8 PCF changes over time for a single realisation of the ABM and as the T cell proliferation and movement bias are varied.

As the movement bias increases, co-localisation of CD8 cells decreases as CD8 cells migrate more readily to regions in which tumour cells are dense. For parameter sets with less movement bias, we observe that the clustering of CD8 cells is at short distances, due to proliferation. The clustering becomes more pronounced as CD8 proliferation is increased. For low movement bias, the time evolution of the CD8-CD8 PCFs reveals increased short range clustering of CD8 cells over time. Some PCF signatures show dispersion at distances just greater than the local clustering, indicating clustering at multiple length scales. Interestingly, for high bias values, the CD8-CD8 PCFs do not show significant variation as the CD8 proliferation probability is increased, even though the tumour outcomes are very different (ranging from tumour escape at low proliferation probability to tumour elimination at high proliferation probability). This is reflected in the TU-CD8 PCFs in Figure 4.9, which indicate that the CD8 cells are uniformly distributed around the tumour cells when the movement bias is high, regardless of the CD8 proliferation probability. We conclude that for high bias values, the spatial distribution of the cells may not be an important factor

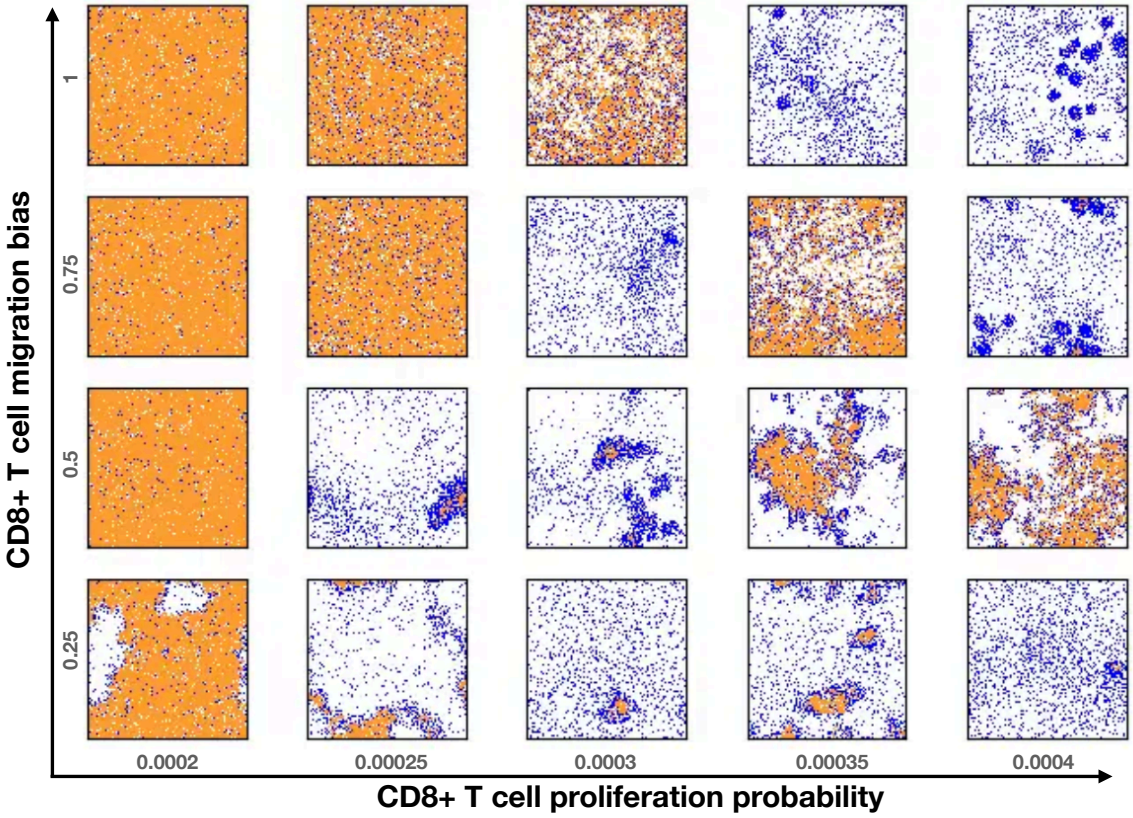


Figure 4.7: Snapshots of ABM at either $t = 145$ days, or the timepoint up to 5 days before tumour elimination occurs (as simulation snapshots are saved every 5 days).

in tumour outcome.

In Figure 4.10 we show how the CD8-CD8 PCF changes over time and how it varies as the CD8 movement bias and proliferation probability is changed. In Figure 4.11 we present similar plots for the TU-CD8 PCFs. In both cases, at early timepoints the value of the PCFs at small r are most varied across the parameter sweep. Over time, the range of r values for which there is variation in the PCFs increases, which suggests that at different timepoints, spatial features at different lengthscales are important.

4.3.1 Variation in PCF metrics over simulation timecourse

In Figure 4.12, we show how metrics extracted from the CD8-CD8 PCF change over time as the CD8 proliferation probability varies. At early timepoints (before day 20) the distribution of $g_{max} = \max_{\{r\}} g(r)$ for each parameter set is indistinguishable. However, at longer times for each parameter set, g_{max} attains a maximum value and then decreases. The time at which g_{max} starts to decrease (attains its maximum value) corresponds to the time at which the tumour cell population starts to decrease, therefore indicating treatment response.

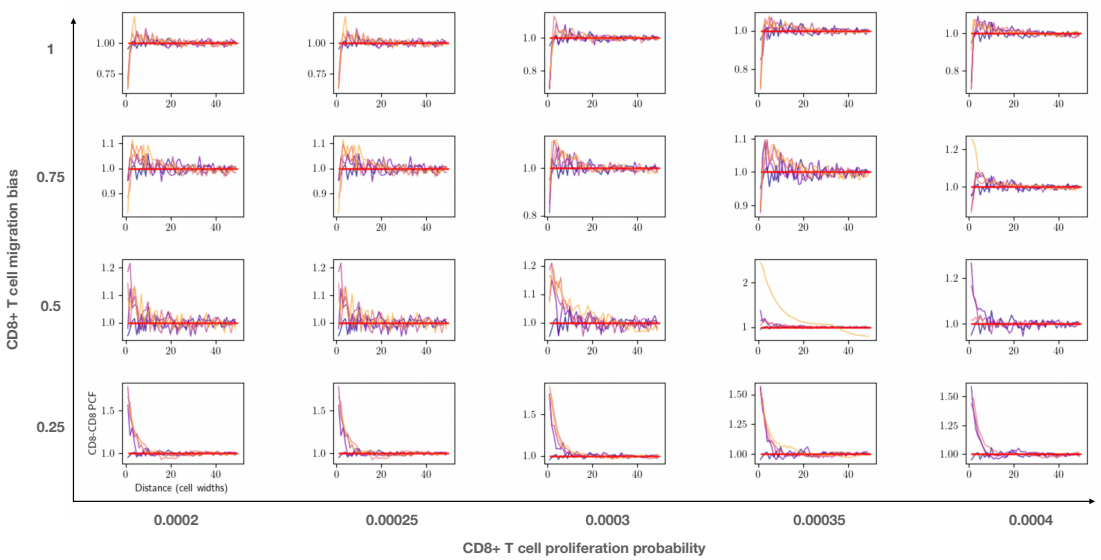


Figure 4.8: CD8-CD8 PCFs for one realisation per parameter set. Timecourse of the CD8-CD8 PCFs are shown by shaded lines, varying from purple to yellow, corresponding to different timepoints. Note that the y axes are on different scales for visualisation purposes in each plot.

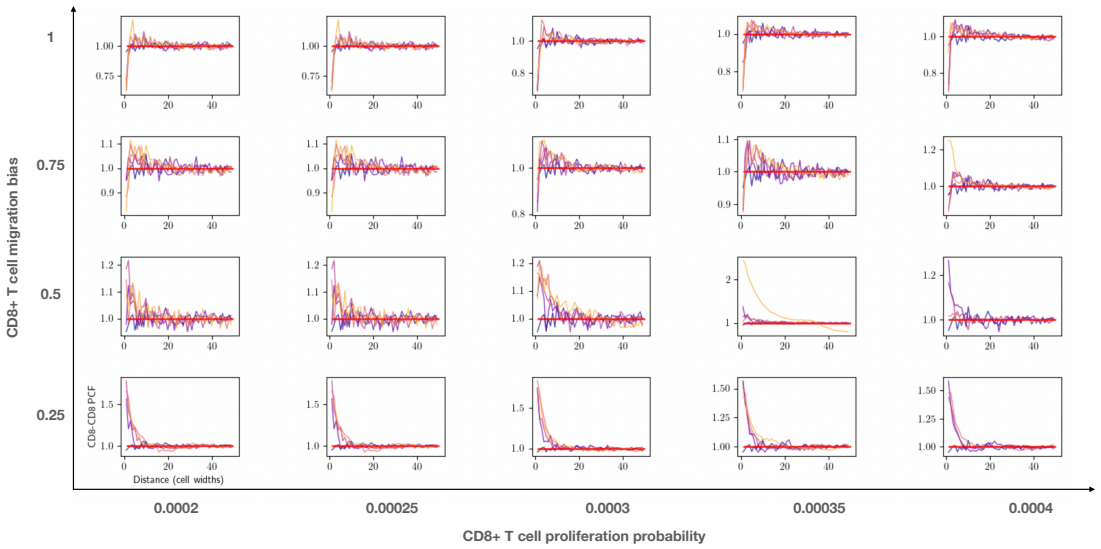


Figure 4.9: TU-CD8 PCFs for one realisation per parameter set. Timecourse of the TU-CD8 PCFs are shown by shaded lines, varying from purple to yellow, corresponding to different timepoints. Note that the y axes are on different scales for visualisation purposes in each plot.

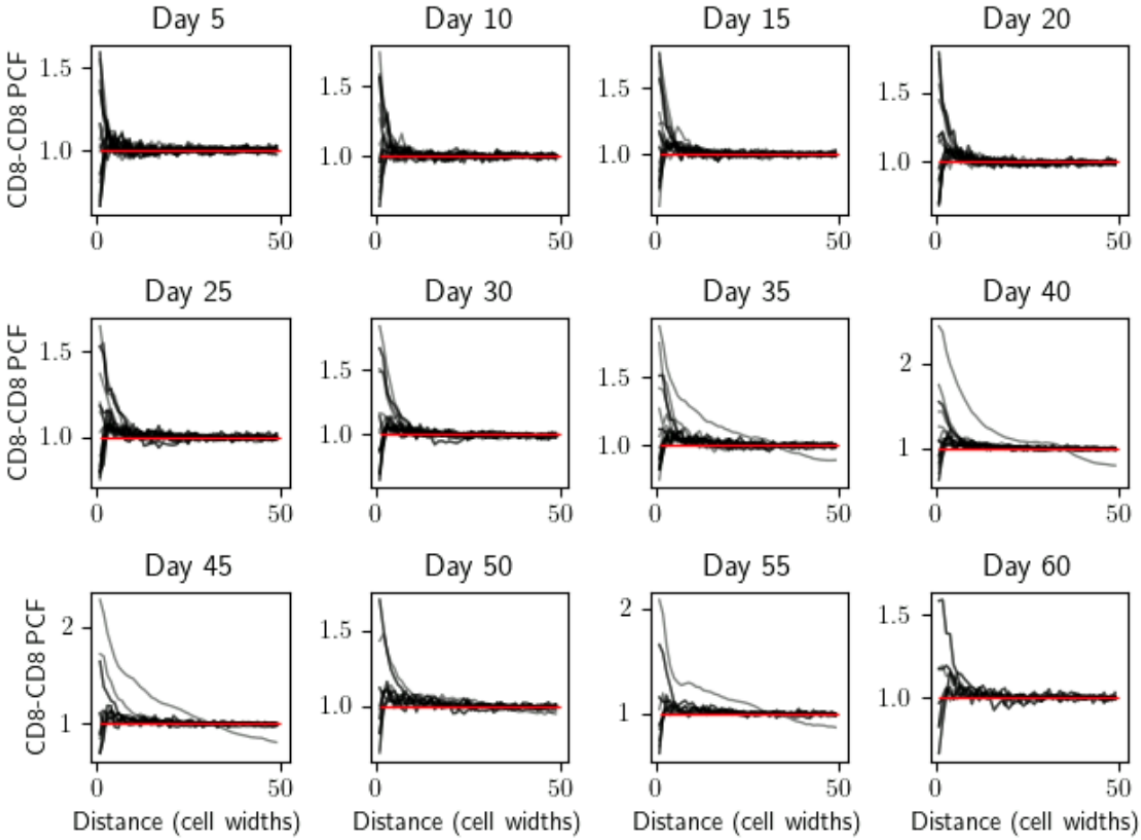


Figure 4.10: CD8-CD8 PCFs plotted at timepoints in intervals of 5 days. Each plot shows the CD8-CD8 PCF from one stochastic realisation for each of the 20 parameter sets considered in the parameter sweep.

The value of r_{max} remains around 1 cell width over time, due to CD8 proliferation causing localised clustering of CD8 cells. r_{max} increases slightly just as the tumour is about to be eliminated. The other PCF metrics, g_{min}, r_{min} and $r(PCF = 1)$ show large variation within parameter sets, and are mostly indistinguishable between parameter sets over the timecourse.

Figure 4.13 is complementary to the previous plot, instead varying the migration bias of the CD8 cells. In these plots, we see more distinct values of the PCF metrics between parameter sets, but less dependence of the PCF metrics on time, suggesting that CD8 cell proliferation is important in determining the spatial characteristics of the simulated tumours.

Figures 4.14 and 4.15 show the time evolution of PCF metrics extracted from the TU-CD8 PCF as CD8 proliferation and CD8 migration bias are varied, respectively. We see that the g_{max} value remains constant up until the time of tumour elimination. The radius at which the PCF first crosses 1 increases over time to a maximum value due to rapid tumour

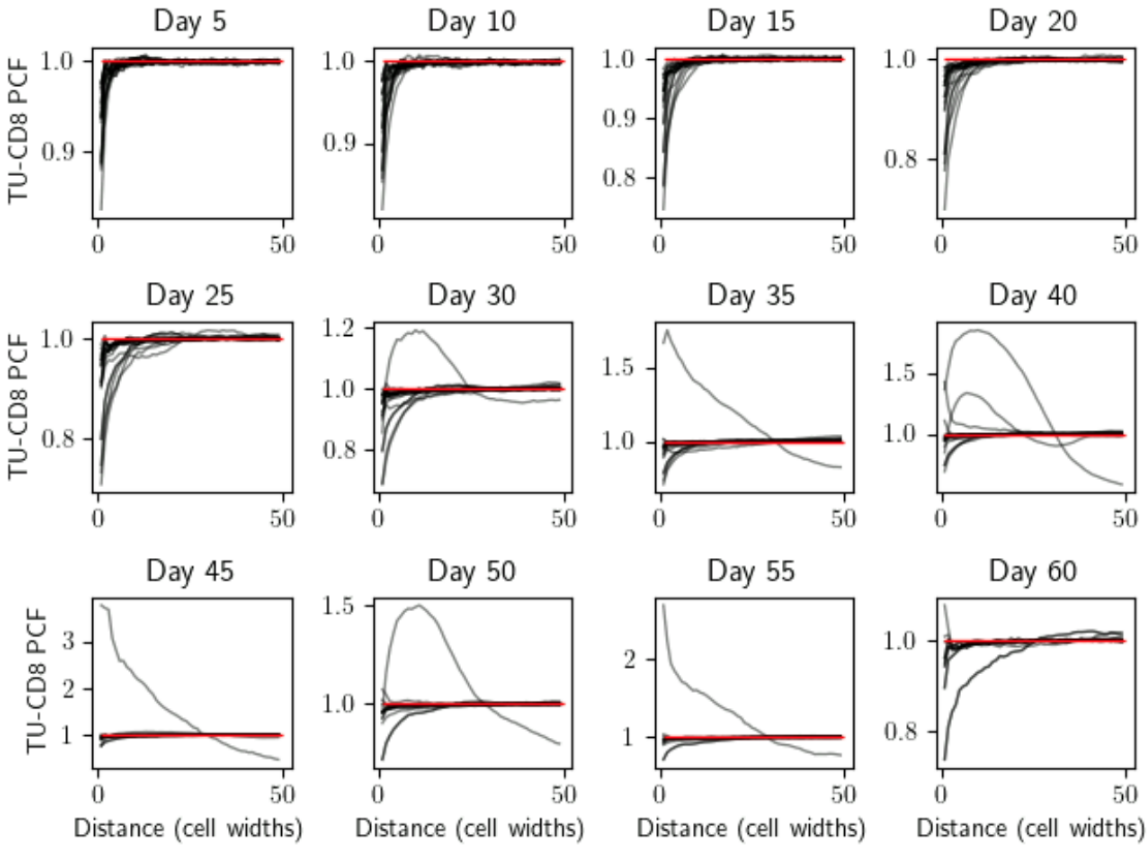


Figure 4.11: TU-CD8 PCFs plotted at timepoints in intervals of 5 days. Each plot shows the TU-CD8 PCF from one stochastic realisation for each of the 20 parameter sets considered in the parameter sweep.

proliferation causing tumour cells and CD8 cells to separate. This radius value remains at a plateau if the tumour is not eliminated, or rapidly decreases for parameter sets where the tumour is eliminated. The variation in both of these metrics is high, indicated by the large overlapping error bars.

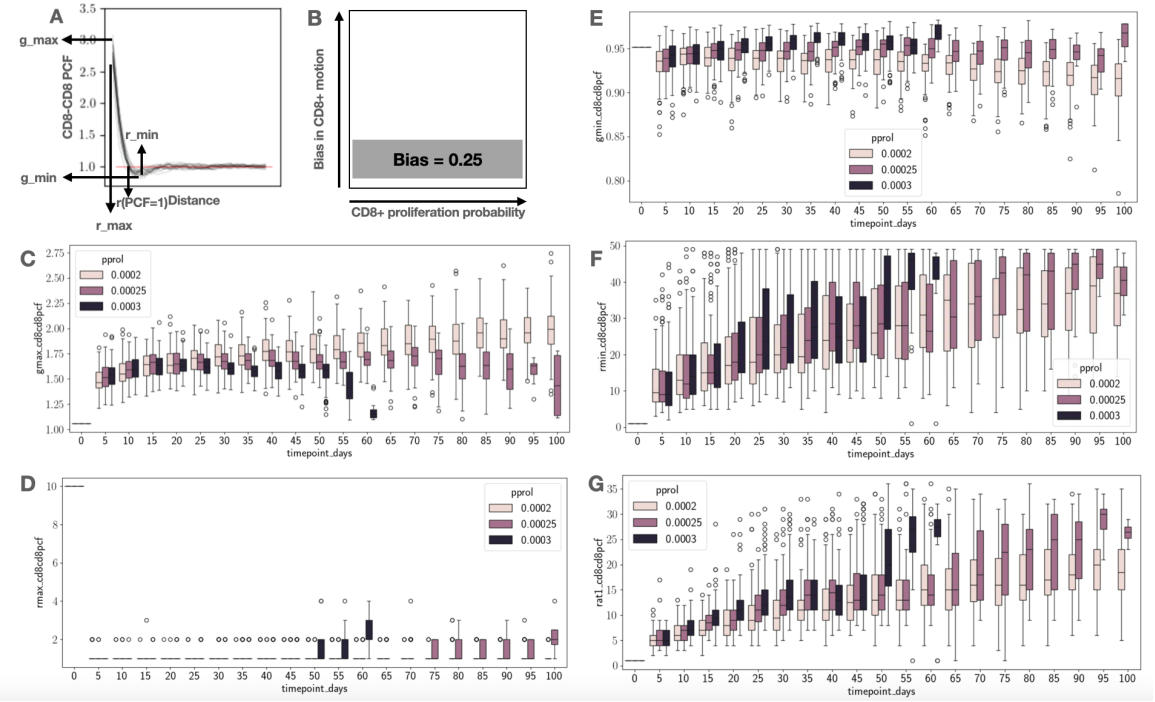


Figure 4.12: Time evolution of metrics extracted from the CD8-CD8 PCF, as CD8+ proliferation probability varies. (A) Schematic illustrating PCF metrics. (B) CD8+ cell migration bias is fixed at 0.25, and CD8+ proliferation probability is varied. Metrics considered are (C) g_{max}, (D) r_{max}, (E) g_{min}, (F) r_{min}, (G) Radius at first crossing of PCF through 1.

4.4 Predicting tumour outcome from a single timepoint

By generating synthetic data from the ABM, we have a rich timecourse for each stochastic realisation. However, in clinical trials, such as those discussed in Chapter 3, there may only be data from one on-treatment timepoint. In this section, we investigate our ability to predict qualitative tumour outcome at day 150 based on data from a single timepoint. We seek to determine when to take an on-treatment sample: is it better to wait as long as possible, or whether the indicator(s) for tumour outcome may be present at earlier timepoints, in order to determine the best time to take the on-treatment biopsy for a given initial condition from a pre-treatment sample. To do this, we take synthetic data from specific timepoints, and use them to train a support vector machine (SVM), with the synthetic tumours labelled as responder or non-responder, based on the tumour cell counts at day 150. We split the data into a training set and testing set, in a 80:20 ratio, so the SVM is trained on 1600 synthetic tumours, and tested on 400.

The classifier takes an input feature vector, one for each tumour, which comprise summary statistics at a specific timepoint. We consider two alternative feature vectors; one

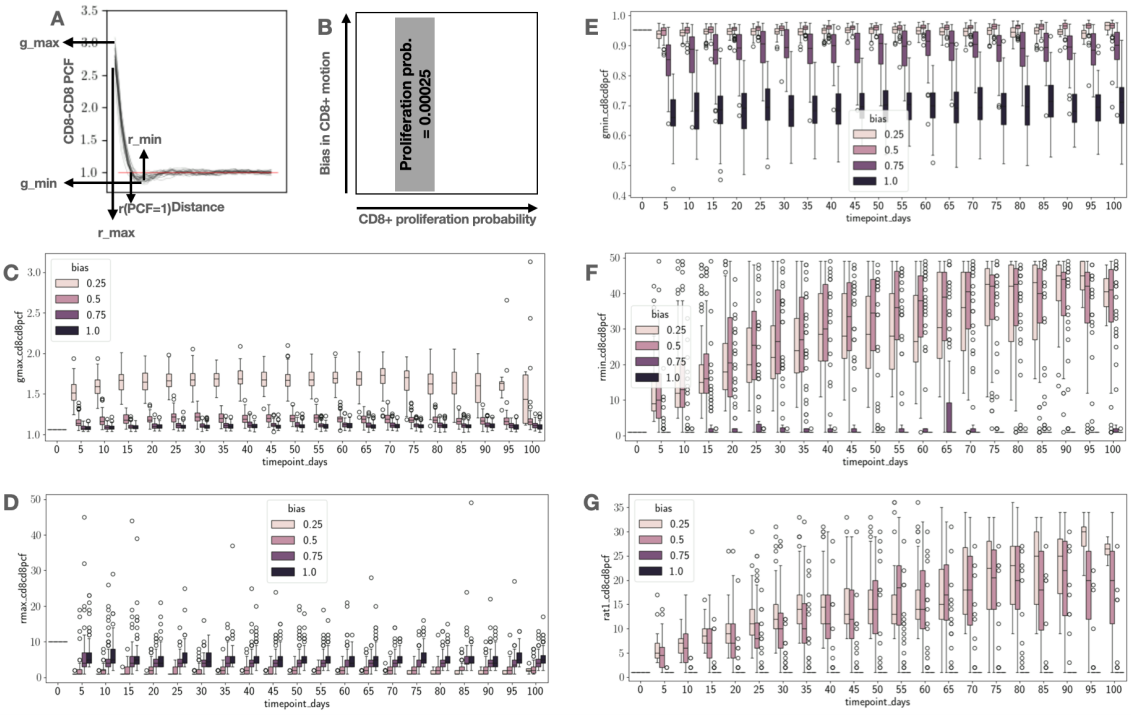


Figure 4.13: Time evolution of metrics extracted from the CD8-CD8 PCF, as CD8+ migration bias varies. (A) Schematic illustrating PCF metrics. (B) CD8+ proliferation probability is fixed at 0.00025, and CD8+ migration bias is varied. Metrics considered are (C) g_{max}, (D) r_{max}, (E) g_{min}, (F) r_{min}, (G) Radius at first crossing of PCF through 1.

using cell counts only ($\mathbf{a}_{scalar}$), and the other supplementing cell count data with PCFs ($\mathbf{a}_{spatial}$). We aim to ascertain whether supplementing the cell count data with the PCFs improves the accuracy of the classifier. The feature vectors have length 2 and 100, respectively, and are of the form:

$$\mathbf{a}_{scalar} = [\rho_T, \rho_E]$$

$$\mathbf{a}_{spatial} = [\rho_T, \rho_E, g_{TE}(\mathbf{r}), g_{EE}(\mathbf{r})]$$

where ρ_T and ρ_E are the cell counts of tumour cells and CD8+ T cells, respectively. $g_{AB}(\mathbf{r}) = [g(r_0), g(r_1), ..., g(r_{max})]$ represents the PCF in the case of $A \equiv B$ or cross-PCF otherwise, calculated on source cell type A and target cell type B.

Two separate classifiers are trained for each of the specific timepoints at 5 day intervals. One classifier is trained on the cell count data only; the second is trained on cell count and PCF data. If any of the tumours have already been eliminated by the timepoint under consideration, the data for this tumour is removed from the training/test set and is not considered. We define the following measures of accuracy:

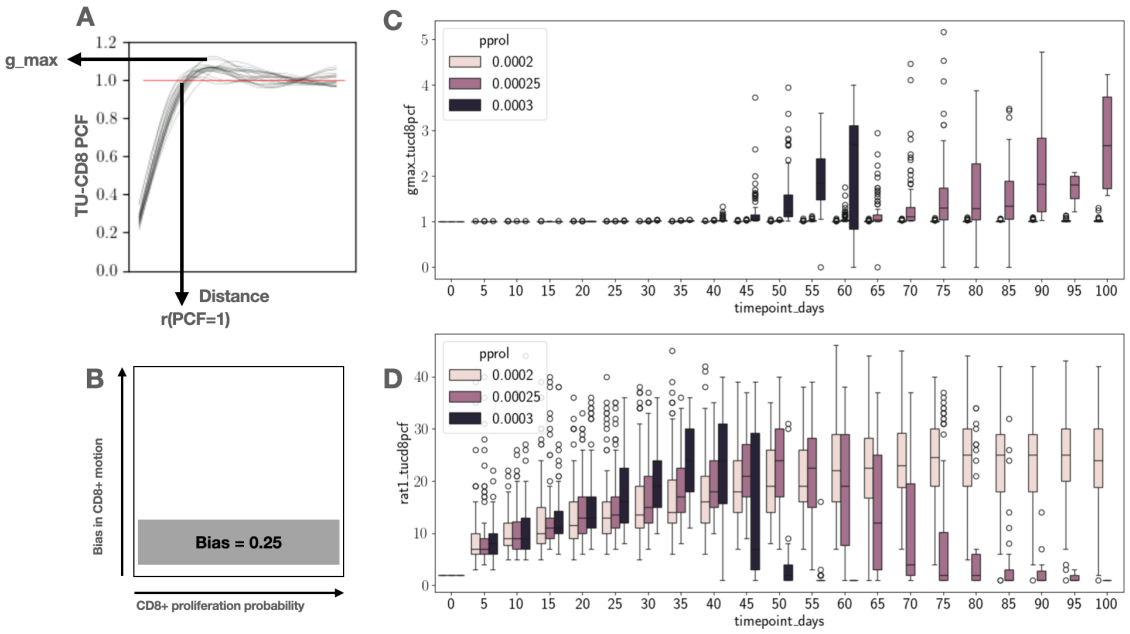


Figure 4.14: Time evolution of metrics extracted from the TU-CD8 PCF, as CD8+ proliferation probability varies. (A) Schematic illustrating PCF metrics. (B) CD8+ cell migration bias is fixed at 0.25, and CD8+ proliferation probability is varied. Metrics considered are (C) g_{max} and (D) Radius at first crossing of PCF through 1.

- Overall % accuracy = $\frac{\text{\# correct classifications}}{\text{\# samples in test dataset}}$ x 100 %
- True responder % = $\frac{\text{\# samples classified responder} - \text{\# samples incorrectly classified responder}}{\text{\# samples classified responder}}$ x 100 %
- True non-responder % = $\frac{\text{\# samples classified non-responder} - \text{\# samples incorrectly classified non-responder}}{\text{\# samples classified non-responder}}$ x 100 %
- Responder % accuracy = $\frac{\text{\# correct responder classifications}}{\text{\# responders in test dataset}}$ x 100 %
- Non-responder % accuracy = $\frac{\text{\# correct non-responder classifications}}{\text{\# non-responders in test dataset}}$ x 100 %

Figure 4.16 shows the accuracy of the classifier at predicting the tumour outcomes of the test data set, for two different tumour responder thresholds: reduction of tumour burden by 50%, and 100% (complete elimination).

The results presented in Figure 4.16 suggest that the prediction accuracy does not significantly change as the responder threshold is varied. Therefore, henceforth we focus our analysis on a 50% tumour burden threshold responder. We find that the cell counts and spatial classifiers have similar overall accuracy for early timepoints ($t \leq 20$), with the accuracy of both classifiers increasing with time, and the cell counts slightly outperforming

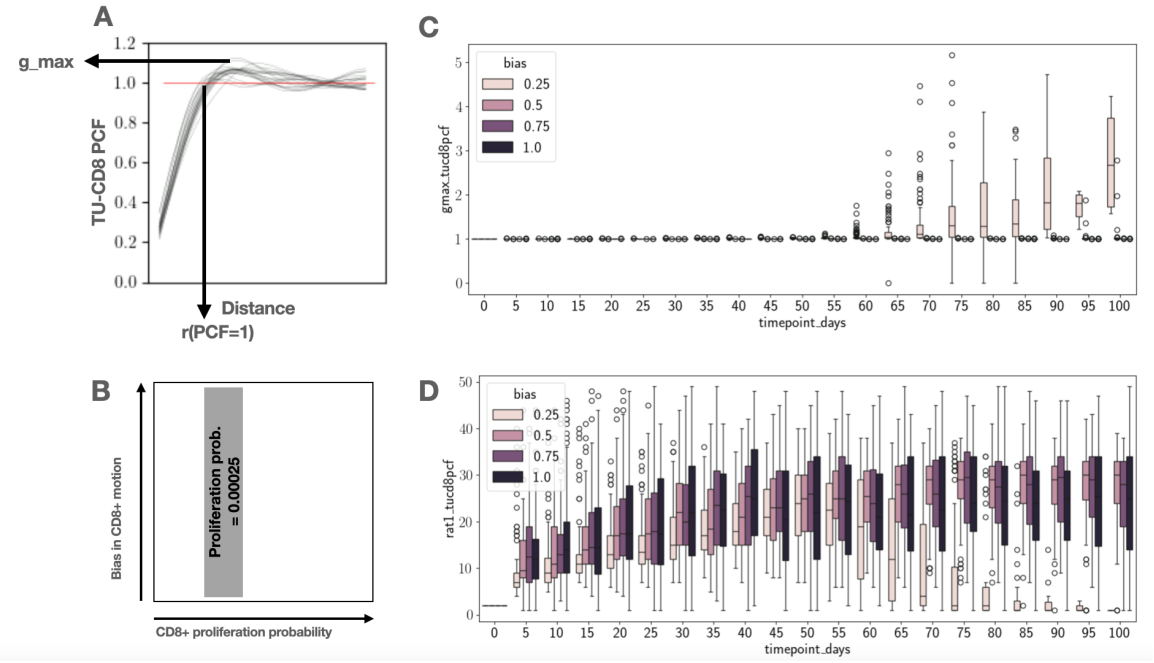


Figure 4.15: Time evolution of metrics extracted from the TU-CD8 PCF, as CD8+ migration bias varies. (A) Schematic illustrating PCF metrics. (B) CD8+ proliferation probability is fixed at 0.00025, and CD8+migration bias is varied. Metrics considered are (C) g_{max} and (D) Radius at first crossing of PCF through 1.

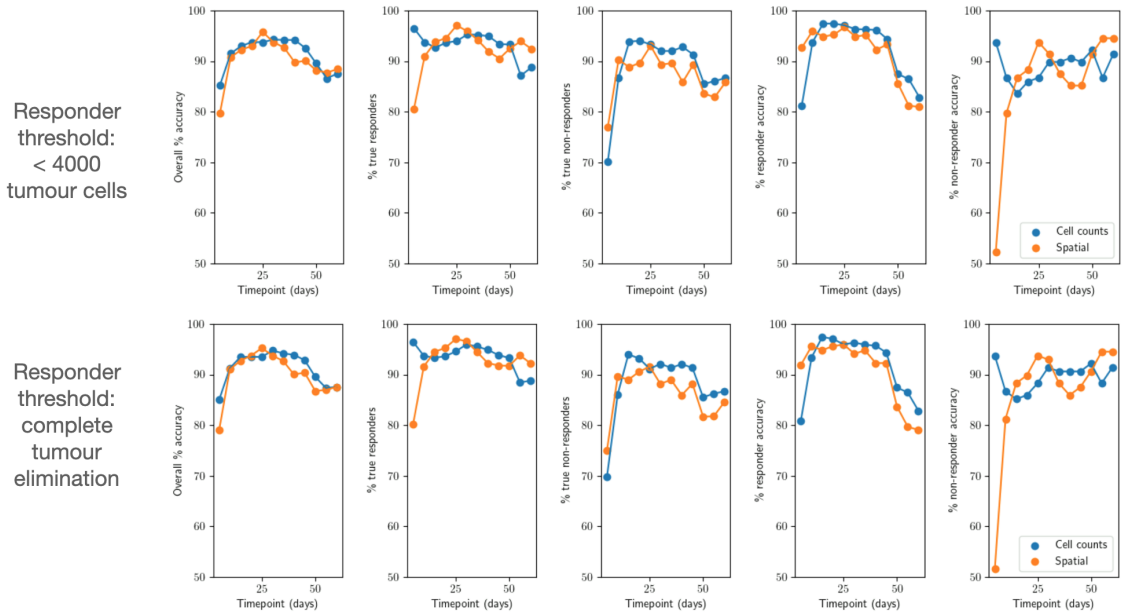


Figure 4.16: Accuracy in tumour outcome predictions for SVM classifiers trained at each timepoint, for two different responder thresholds.

the spatial classifier. At day 25, the spatial classifier achieves a maximum accuracy of over 94%, reaching this maximum level of accuracy at an earlier timepoint than the cell count

classifier. If we refer to Figure 4.12, then we see that between day 20 and day 25 the spatial metrics start to diverge between parameter sets, which is why at this timepoint the spatial information is important in improving the tumour outcome predictions.

4.5 Multiple on-treatment biopsy scheduling

Although most patients received only one on-treatment biopsy, several patients in the FAP-IL2v dataset had two on-treatment biopsies. An example of how the cell count data from one such patient cluster by the timepoint of the three biopsy samples is shown in Figure 4.17.

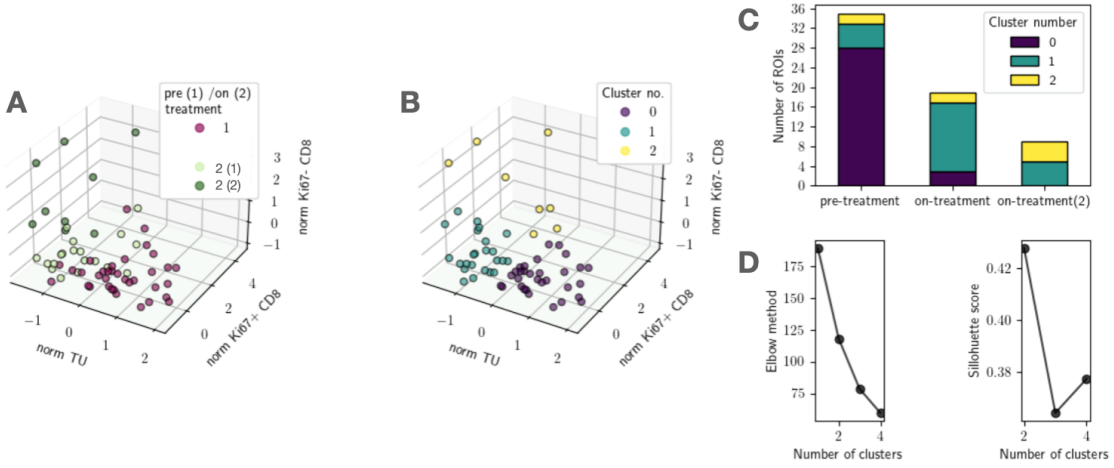


Figure 4.17: Example of ROIs from one patient for which there were two on-treatment biopsies taken. The plots show how the cell count data from this patient cluster by timepoint (i.e. by pre/on-treatment). The cell count data are normalised before the clustering (axes correspond to normalised (norm) cell counts).

Our method could be used to identify the best times to take the first and second on-treatment biopsies. Additionally, data from the first on-treatment biopsy could be used to inform the time at which the second one is taken. In the first case, the two timepoints are fixed in advance, and in the second case, you adapt the second timepoint based on the information from the first timepoint. In the following analysis we consider the first case, in which both timepoints are fixed in advance.

This method could be adapted to incorporate clinical constraints. For example if there was a limit on the time at which the second sample could be taken, several timepoints for the first on-treatment timepoint could be simulated to find which combination of timepoints could maximise the prediction accuracy of tumour outcome.

4.5.1 Predicting tumour outcome using data from two timepoints

We have found that the timepoint at which an on-treatment biopsy is taken, and hence the spatial data provided by the PCFs, affects the accuracy with which we can predict tumour outcome. In particular, for our ABM with all parameters fixed except CD8 proliferation and biased motion, supplementing cell count data with the PCF data at day 25 yielded more rapid diagnosis with a higher degree of accuracy than that obtained from cell count data at a later timepoint (day 30) - see Figure 4.16. We now investigate whether a second on-treatment biopsy would further improve the prediction accuracy, and, if so, how the two on-treatment biopsies should be coordinated to maximise the prediction accuracy. To this end, we consider combinations of pairs of timepoints, and concatenate the data from the two timepoints, to form a feature vector of length 4 (in the case of just cell counts) or length 200 (in the case of cell counts and PCF data). For each combination of timepoints, we train an SVM classifier to predict tumour outcome, using the data at these two timepoints.

Figure 4.18 shows the results of the predictions using data from combinations of two timepoints, for cell counts only (top row) and cell counts with spatial data (middle row), and the increased percentage error for each timepoint combination when using the spatial prediction, compared to the prediction based only on cell count data (bottom row).

We find that, for all timepoint combinations, except for days 10 and 25, the cell count classifier outperforms the spatial classifier. This is perhaps counter-intuitive, as you would assume that having information from another timepoint would improve the classification.

Including spatial metrics in feature vectors when they are not important has a negative impact on the classifier predictions - these parameter sets may show similar spatial distributions to other parameter sets where the spatial features are actually important. For example, in the CD8-CD8 PCFs, most of the variation between parameter sets is within small values of r.

So, we retrained the classifiers for each timepoint combination, considering only regions of parameter space in which the spatial effects were found to be most significant in our earlier analysis of the PCFs. We consider data generated from parameter values with CD8 proliferation probability, $r_E = \{0.0002, 0.000225, 0.00025\}$, and a finer sampling of CD8 motion bias values, $\beta = \{0.125, 0.25, 0.325, 0.5, 0.625, 0.75, 0.875, 1\}$. See Appendix G for supplementary results, including cell count timecourse data. Figure 4.19 shows a higher error in the classification accuracy in both cell count and spatial classifiers compared to analysing the whole parameter sweep data. This is due to bistability in the tumour outcomes when CD8 proliferation probability is small. However, within these parameter sets, using the PCF data improves the classification accuracy obtained using cell counts only for many timepoint combinations.

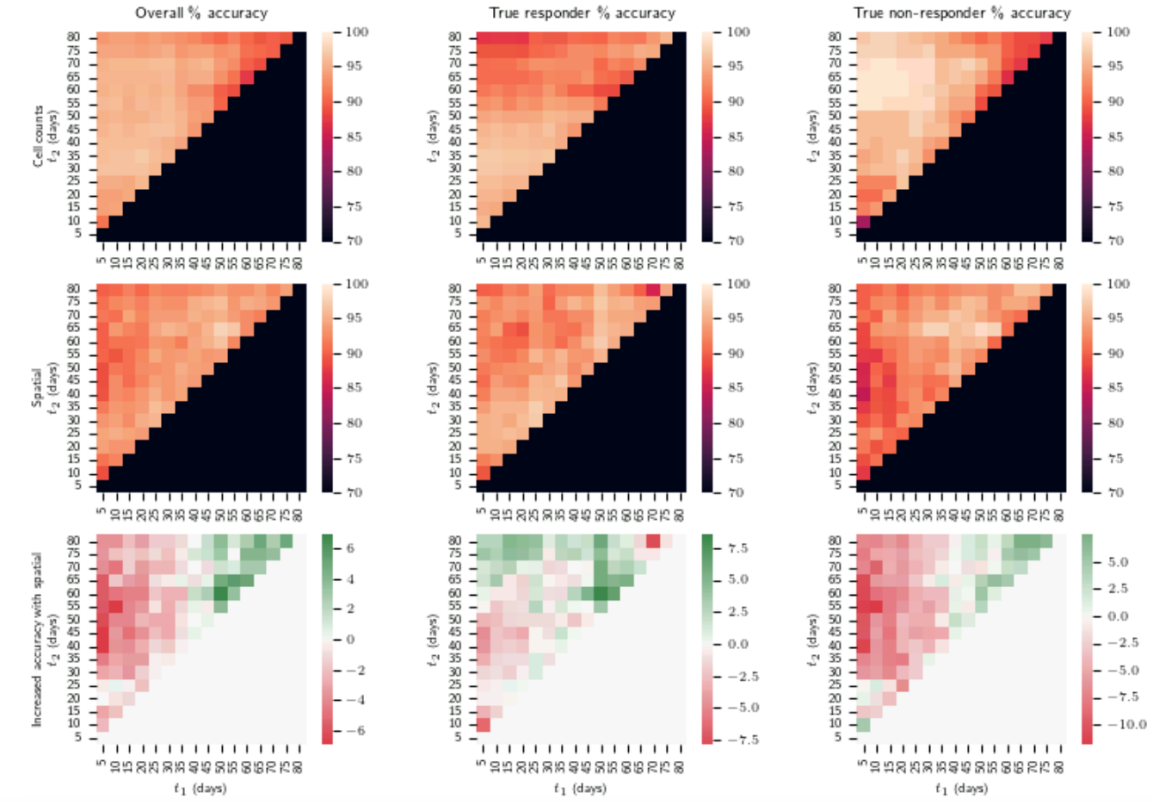


Figure 4.18: Overall percentage error in responder classification for combinations of on-treatment biopsy timepoints, for cell count data (top row) and cell count and PCF data (middle row). (Bottom row) Increased percentage accuracy for spatial classifier compared to cell count classifier.

4.5.1.1 Predictions using PCF metrics

We found in Section 4.3.1 that the maximum value of the CD8-CD8 PCF, g_{max}, showed a time dependence (increased to a maximum, before decreasing just before tumour elimination), and may be an indicator of tumour outcome. We focus on g_{max} as other PCF metrics were found to either be uninformative in separating parameter sets, or had a large variance within parameter sets. When training the SVMs in the previous analysis, we input the PCFs for each synthetic tumour data point. In practice however, there will be some uncertainty in the position and value of g_{max} in this feature vector between stochastic realisations. Therefore, in this section we train the spatial SVMs using feature vectors containing the cell counts of each population, and the g_{max} value of the CD8-CD8 PCF, rather than the whole PCF signature.

Figure 4.20 shows the prediction accuracy of the classifiers for each combination of timepoints.

Comparing Figure 4.20(A) with Figure 4.18(A) shows that the classification error is

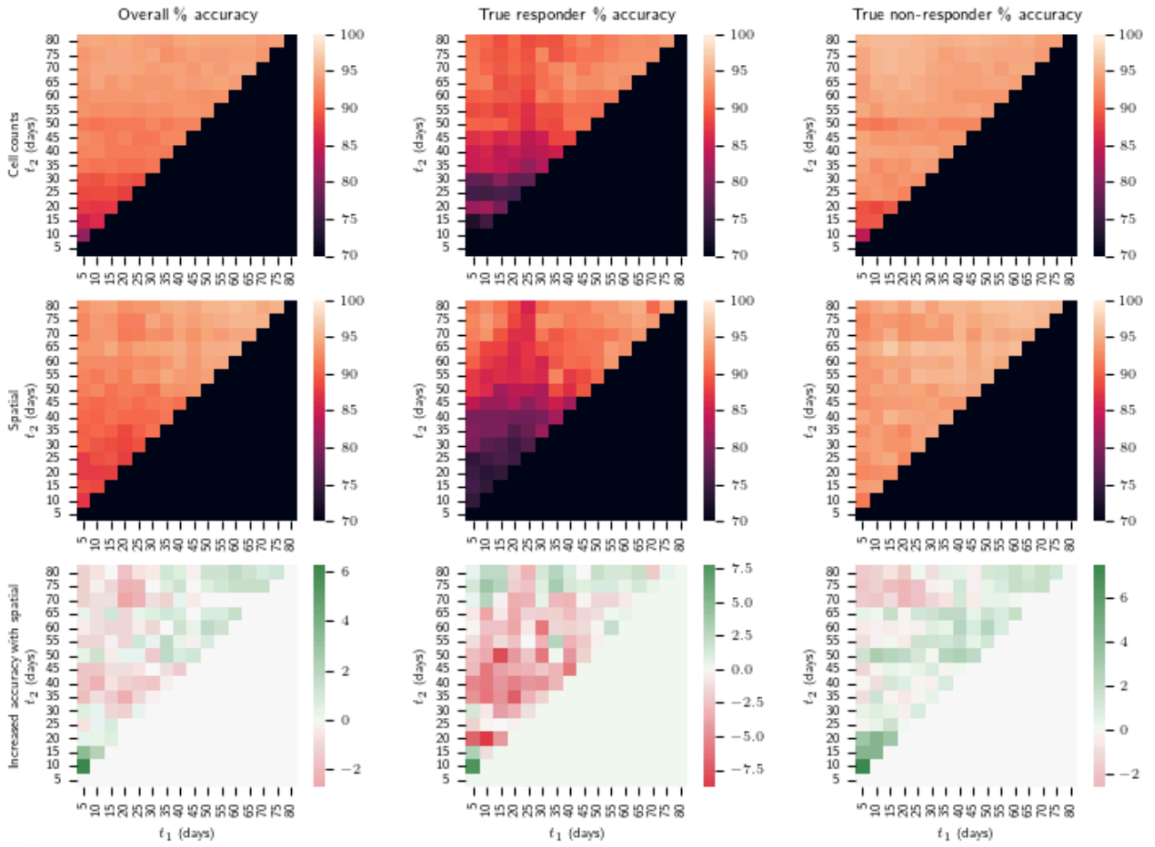


Figure 4.19: Overall percentage error in responder classification for combinations of on-treatment biopsy timepoints, for cell count data (top row) and cell count and PCF data (middle row). (Bottom row) Increased percentage accuracy for spatial classifier compared to cell count classifier. Classifiers trained on subset of data, where CD8 proliferation probabilities = 0.0002, 0.000225, 0.00025. Green indicates combinations for which spatial data improves predictions.

reduced when using g_{max} rather than the whole PCF to characterise the spatial distribution. We also find that the predictions show the greatest improvement in accuracy (compared to the cell count data) when the second on-treatment biopsy is taken at 30 days. From Figure 4.12, we note that the g_{max} values for the tumour elimination parameter sets reach their maximum values for $t < t = 30$ days. We hypothesise that two on-treatment biopsies taken before and after the time at which g_{max} attains its maximum value would be optimal for predicting tumour elimination in this model. This would allow detection of an initial increase, and then subsequent decrease, in g_{max}, something which one on-treatment biopsy would not be able to capture.

Figure 4.21 shows results in which we once again use g_{max} in the classifier training, but now only for the parameter sets with $r_E = 0.0002, 0.00025$. We see an overall increase in prediction error because the bistable parameter sets comprise a greater proportion of the

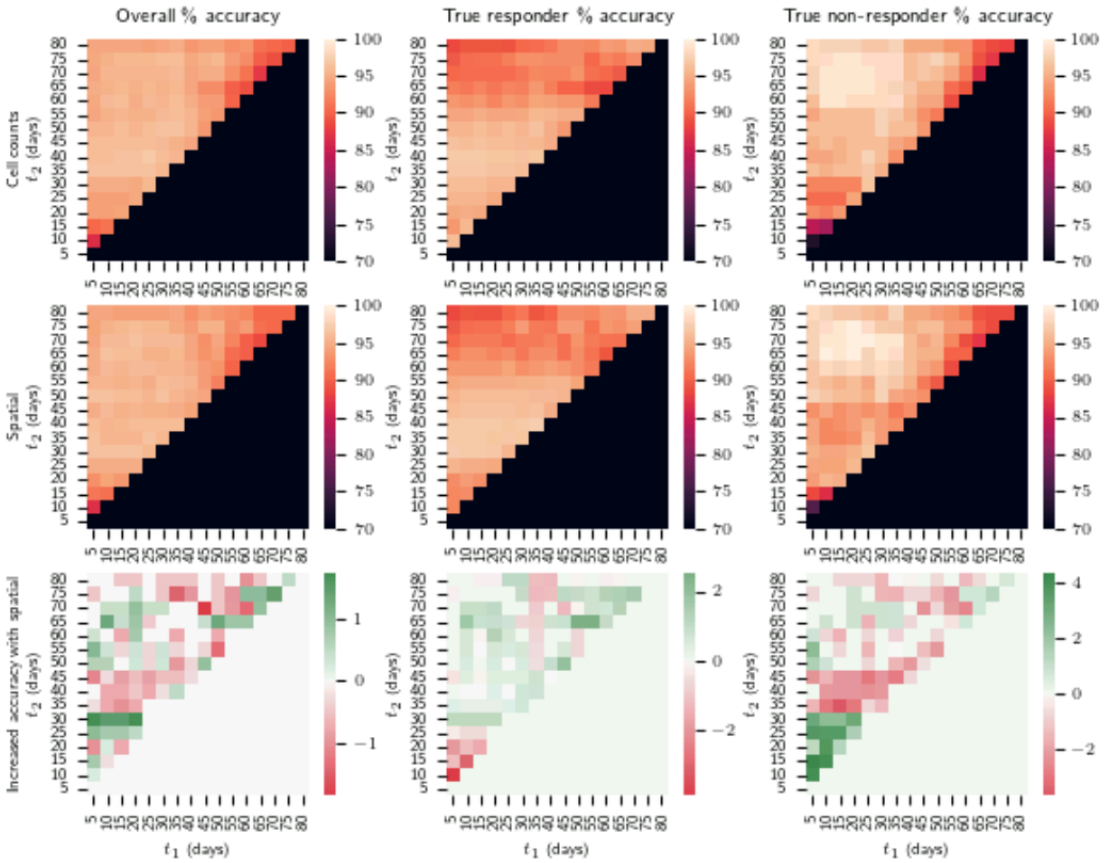


Figure 4.20: Overall percentage error in responder classification for combinations of on-treatment biopsy timepoints, for cell count data (top row) and cell count and g_{max} data (middle row). (Bottom row) Increased percentage accuracy for spatial classifier compared to cell count classifier.

overall training data set. However there is a significant improvement in predictions across almost all timepoint combinations when the cell count data are supplemented with g_{max} values. This is because at low CD8 proliferation, spatial effects are more important for determining tumour elimination. As the CD8 proliferation probability increases, spatial effects become less important as the CD8 cell population reaches a threshold number at which it is able to take control of the tumour.

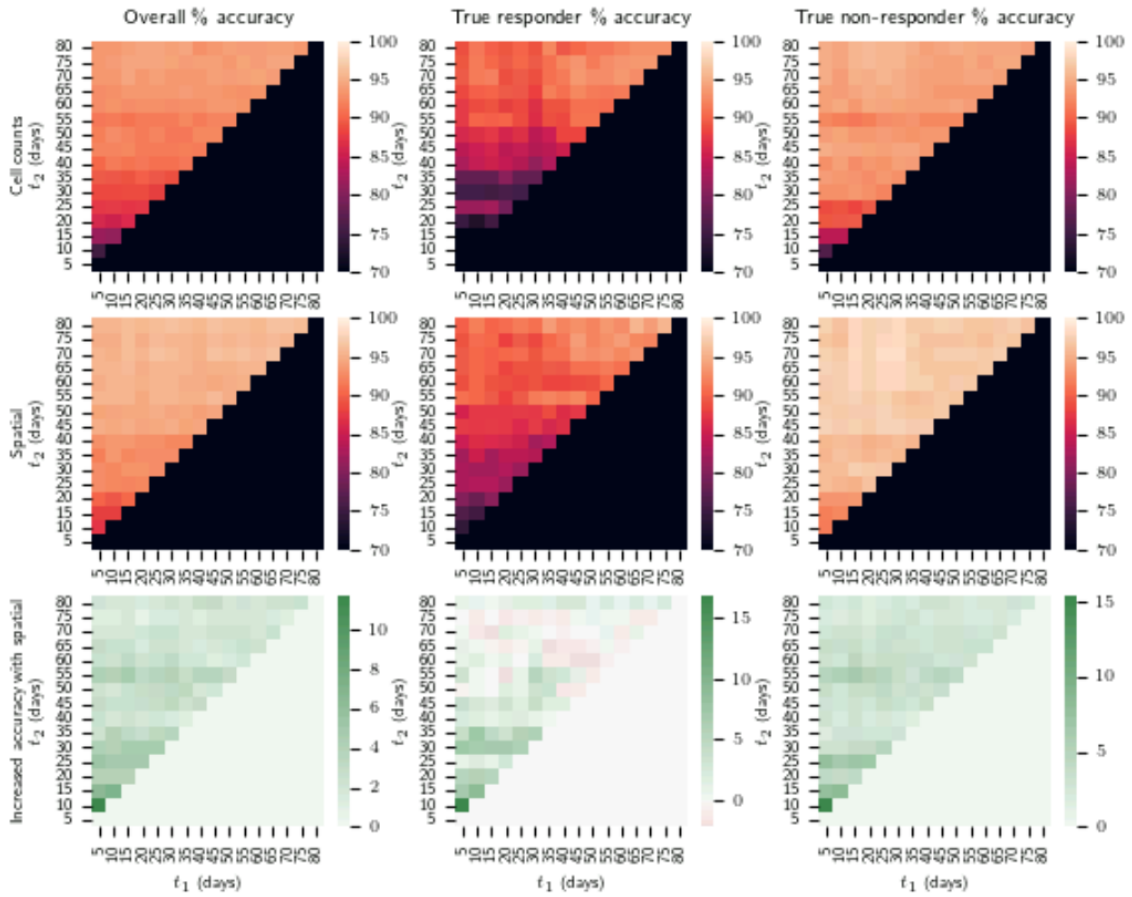


Figure 4.21: Overall percentage error in responder classification for combinations of on-treatment biopsy timepoints, for cell count data (top row) and cell count and g_{max} data (middle row). (Bottom row) Increased percentage accuracy for spatial classifier compared to cell count classifier. Classifiers trained on subset of data, where CD8 proliferation probabilities = 0.0002, 0.000225, 0.00025.

4.6 Varying initial condition

4.6.1 Converting biopsy data into lattice format

After considering one initial condition with randomly distributed cells to generate preliminary results, we now initialise the model with ROIs from a patient pre-treatment biopsy image and investigate how the model predictions depend on the initial spatial distribution of the cells. To convert the ROIs sampled from the image into a compatible on-lattice format for initialising the ABM, cell coordinates were scaled by a factor of 6. This rescaling was universally found across the images to be the minimum distance between tumour cell centres. We estimated the diameter of a single tumour cell to be approximately 6 image distance units (pixels) and each cell lattice site represents the size of a single tumour cell.

If there were tumour cell overlaps (i.e. more than one cell was scaled down to the same lattice site) then one tumour cell was moved to a neighbouring site. For all other cell-cell combinations, the CD8+ T cell was moved to a neighbour site. The case of overlapping cells (cells co-located within 6 pixels) was explored in the previous chapter, and without the raw image data it is not possible to ascertain whether these are in fact co-localised cells, or an artefact of the tissue data. For simplicity, in this work we assume that these features are not artefacts, and do not discard any data from the images.

4.6.2 Inferring non-proliferating tumour cell population

As previously stated, we only have data for the Ki67+ stained proliferating tumour cells. We do not know the positions of the non-proliferating tumour cells. The fraction of Ki67+ tumour cells in a sample can vary depending on the tumour type. However, the lack of data on the locations of non-proliferating tumour cells poses a challenge regarding how these data should be integrated with a mathematical model. The locations of the non-proliferating tumour cells could be inferred in different ways, for example:

- Assuming a uniform distribution of tumour cells within the intersection of tumour and tissue regions.
- Estimate the fraction of proliferating tumour cells based on the tumour type, and placing non-proliferating tumour cells within a prescribed radius around the locations of proliferating tumour cells.
- Position Gaussians around the location of proliferating tumour cells and place non-proliferating tumour cells within this Gaussian field.
- Apply image analysis convolutions to create a mask for the tumour region using the locations of the proliferating tumour cells, and then fill this area with a prescribed density of uniformly distributed tumour cells.

Any of the above methods of inferring the locations of the Ki67- tumour cells will add some uncertainty or artefacts into the data. In this work, we adopt a simple approach, "filling in" the non-proliferating tumour cells uniformly (and randomly) throughout the ROI, up to a specified density, which is arbitrarily chosen. However, this assumption neglects the presence of stroma within the tissue, and also results in loss of information in the spatial variation of tumour density. The other methods listed above would preserve the variation in local tumour density to a degree, however for this proof of concept study we adopted the simplest approach. With the availability of the raw image data, all tumour cells (not

limited to just the proliferating tumour cells) could be identified through cell segmentation methods.

Other considerations to make are that the data set does not include information regarding the distribution of oxygen or other diffusive factors within the tissue, such that necrotic regions may not be identifiable. There is also no information about the distribution of stromal tissue within the sample, which would affect how the tumour tissue grows and how the T cells migrate within the tissue. We note that for some samples, certain tissue regions are characterised by a high density of proliferative tumour cells, while other regions have no tumour cells but a high density of T cells, see Figure 4.22.

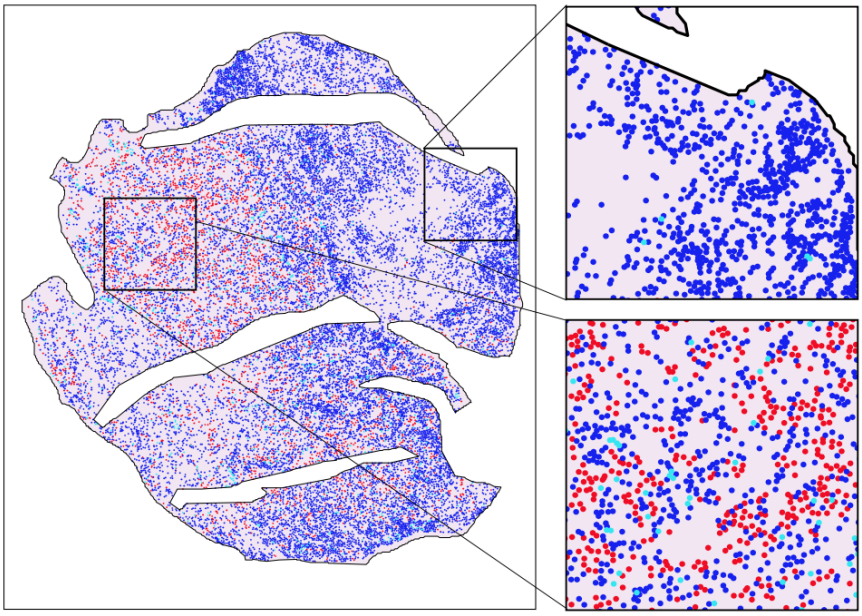


Figure 4.22: An example of a biopsy sample where there are regions of high tumour density, and other regions of high T cell density, with no/few tumour cells. Proliferating tumour cells shown as red points, proliferating CD8+ T cells shown as cyan points, non-proliferating CD8+ T cells shown as dark blue points.

The latter regions may be viewed as stromal regions, as it is known that T cells migrate and reside within the stroma. However, the presence of these T cell-rich, tumour-poor regions could be due to the cell cycles of the tumour cells syncing locally within the tissue, thus expressing Ki67 in dense regions Mattfeldt et al. (2009), i.e. tumour cells are present in other regions but are not detected.

4.6.3 Comparison of spatial distributions of real and synthetic data

In Figure 4.23 we compare the functional form of individual PCFs between the real and synthetic data.

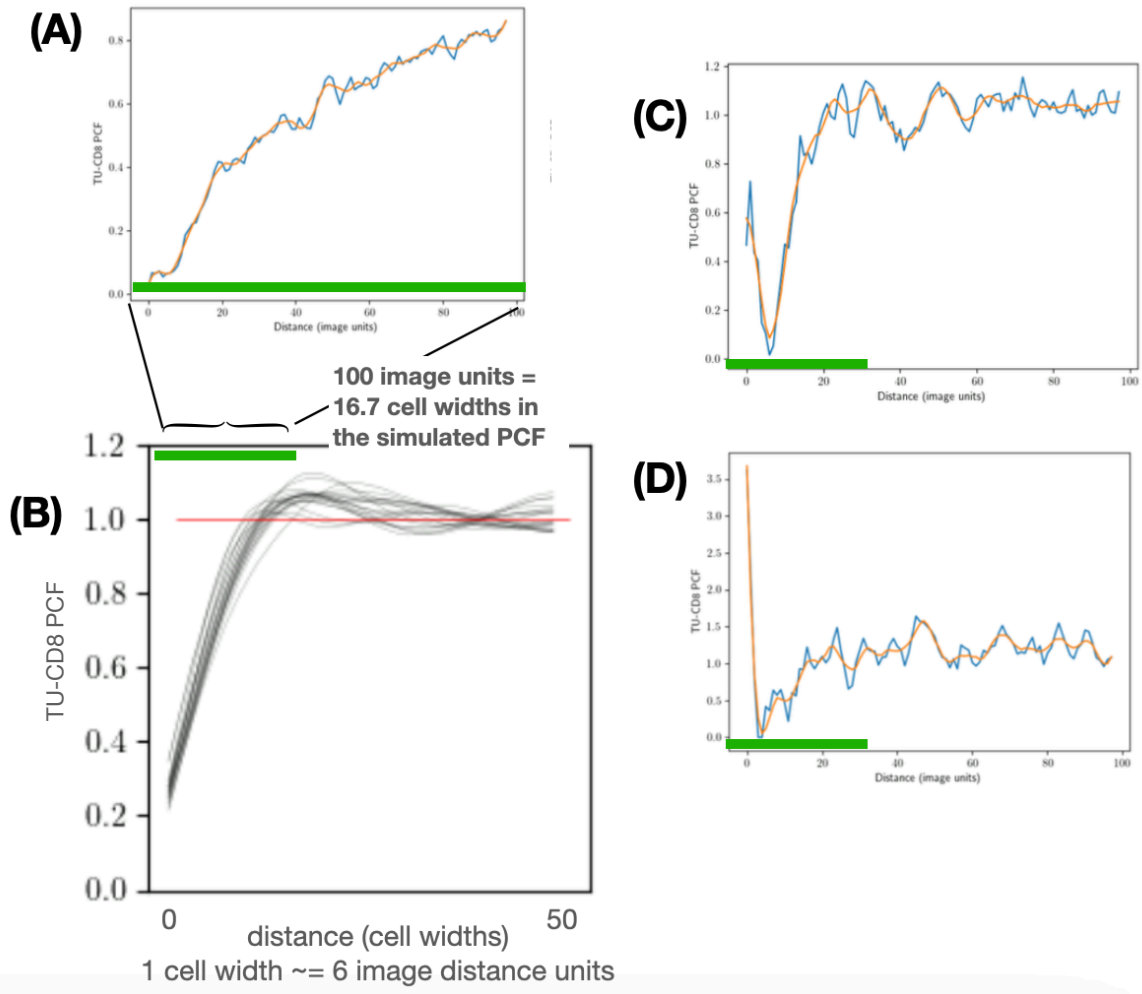


Figure 4.23: Qualitative comparison of functional form of tumour-CD8 cell PCF for real data (A,C,D) and synthetic data for 100 stochastic realisations of the ABM (B). The raw PCFs for the real data plots are shown in blue and the smoothed PCFs in orange. The green lines indicate equivalent distances. The real data PCFs are calculated on the raw point data, prior to scaling onto the lattice.

We calculate the PCF over a shorter distance scale and finer resolution for the real data compared to the ABM, due to constraints in the size of ROI that can be sampled from the images without including tissue edges. The green lines in Figure 4.23 represent equal distances. The ABM PCF qualitatively agrees with the real PCF example in Figure 4.23(A) on the short lengthscale, however does not show the same functional form as other example PCFs (Figure 4.23(C, D)) from the real data because the tumour proliferation rate in the model is very high and there are few free spaces in the lattice, therefore the CD8 cells cannot migrate through the tumour mass because of the single occupancy rule.

The CD8-CD8 PCFs from the ABM have two qualitative functional forms, as depicted in Figure 4.24(B) and (C).

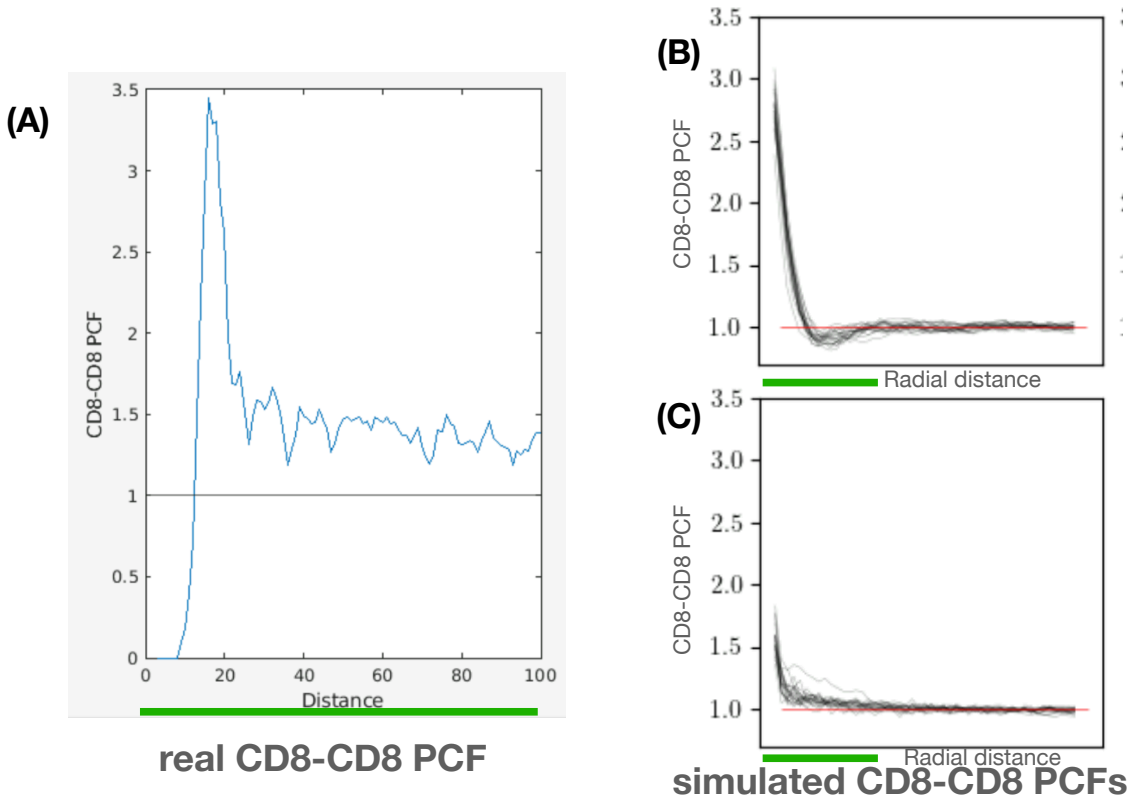


Figure 4.24: Qualitative comparison of functional form of CD8-CD8 cell PCF for real data (A) and synthetic data (B,C). The green lines indicate equivalent distances. The real data PCFs are calculated on the raw point data, prior to scaling onto lattice.

Figure 4.24(B) shows clustering of CD8 cells at short distances, with subsequent oscillations of the PCF around 1 (corresponding to complete spatial randomness, CSR), before tending to 1 at long distances. Figure 4.24(C) shows clustering at all radii (PCF ≥ 1), until the PCF tends to 1 at long distances. The green lines indicate the comparable distance scales (1 ABM cell width approximately equals 6 real image distance units). The typical real data CD8-CD8 PCF function better resembles Figure 4.24(C). Figure 4.24(C) is from a simulation where the CD8 cells are more dispersed in the tumour compared to Figure 4.24(B) - again this points to ABM grid occupancy effects in the PCF.

4.6.4 The effect of ROI size on the spatial evolution of the ABM

In our analysis we have considered ROIs of size 100 x 100 cell widths. In this section, we will explore whether the ABM output is robust to the ROI dimension chosen. As discussed in Chapter 3, core needle biopsies typically yield small tissue samples, such that in many of the images in our dataset the largest ROI dimension which could fit within the tissue boundaries was of dimension 100. So, for this analysis we select a patient with a sufficiently

large tissue sample to enable sampling ROIs larger than 100 x 100 cell widths. We sample an ROI of 400 x 400 cell widths, and consider corresponding ROI dimensions of 200 and 100 cell widths by splitting the large ROI up as shown in Figure 4.25.

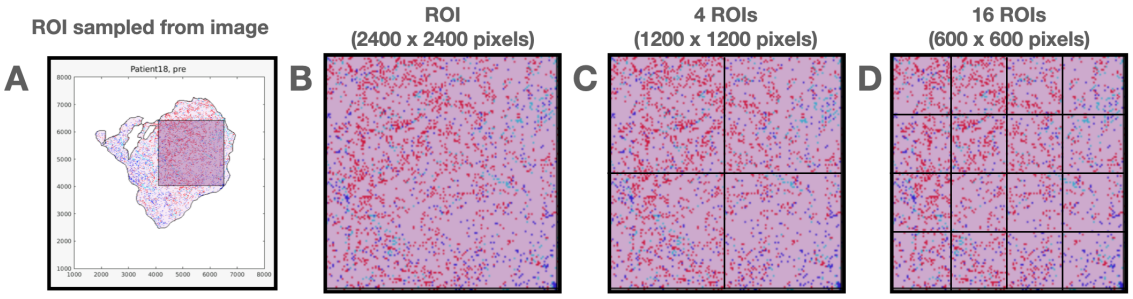


Figure 4.25: (A) Sampling of 2400 x 2400 pixels ROI (B), which is divided up into (C) 4 ROIs of dimension 1200 x 1200 pixels and (D) 16 ROIs of dimension 600 x 600 pixels.

Figure 4.26 shows the corresponding PCFs for the initial conditions of different ROI sizes, illustrating the inter-ROI variation of smaller ROIs within the 400 x 400 ROI.

We initialised the ABM using these ROIs of different dimensions as initial conditions to see if the ABM output is robust to decreasing the ROI size and sampling the heterogeneity of the tissue by taking many, smaller ROIs. We collected summary statistics (cell densities, PCFs and cross-PCF) every 5 simulated days. Each ROI was simulated over 10 stochastic realisations, for a single parameter set (CD8 proliferation = 0.000325, CD8 kill probability = 0.005). We set the CD8+ motion bias to zero, such that the CD8 cells moved randomly, to reduce spatial effects and examine whether even in this parameter set there are differences between ROI sizes. To account for the unknown positions of non-proliferating tumour cells, we filled 70% of the empty spaces in the domain (i.e. those not already occupied by a tumour cell or CD8 cell) with tumour cells. In order to directly compare data from simulations of different ROI dimensions, we calculated cell densities, and calculated the PCFs up to a radius of 49 cell widths for all ROIs.

Figure 4.27 shows timecourses of tumour and CD8+ cell populations for an example initial condition of size 400, 200, and 100 cell widths. We see a sharp initial expansion of the tumour cell population during the burn-in period of the ABM. For this parameter set, we do not see a large variation in tumour outcome, with most tumours in equilibrium with the immune response. A few of the tumours are starting to be eliminated, shown by a sharp reduction in tumour cell population number, which is coupled with an expansion of CD8+ T cells. There is a range of responses in the CD8+ cell population between stochastic realisations. We see that decreasing the ROI size increases the variation between stochastic realisations.

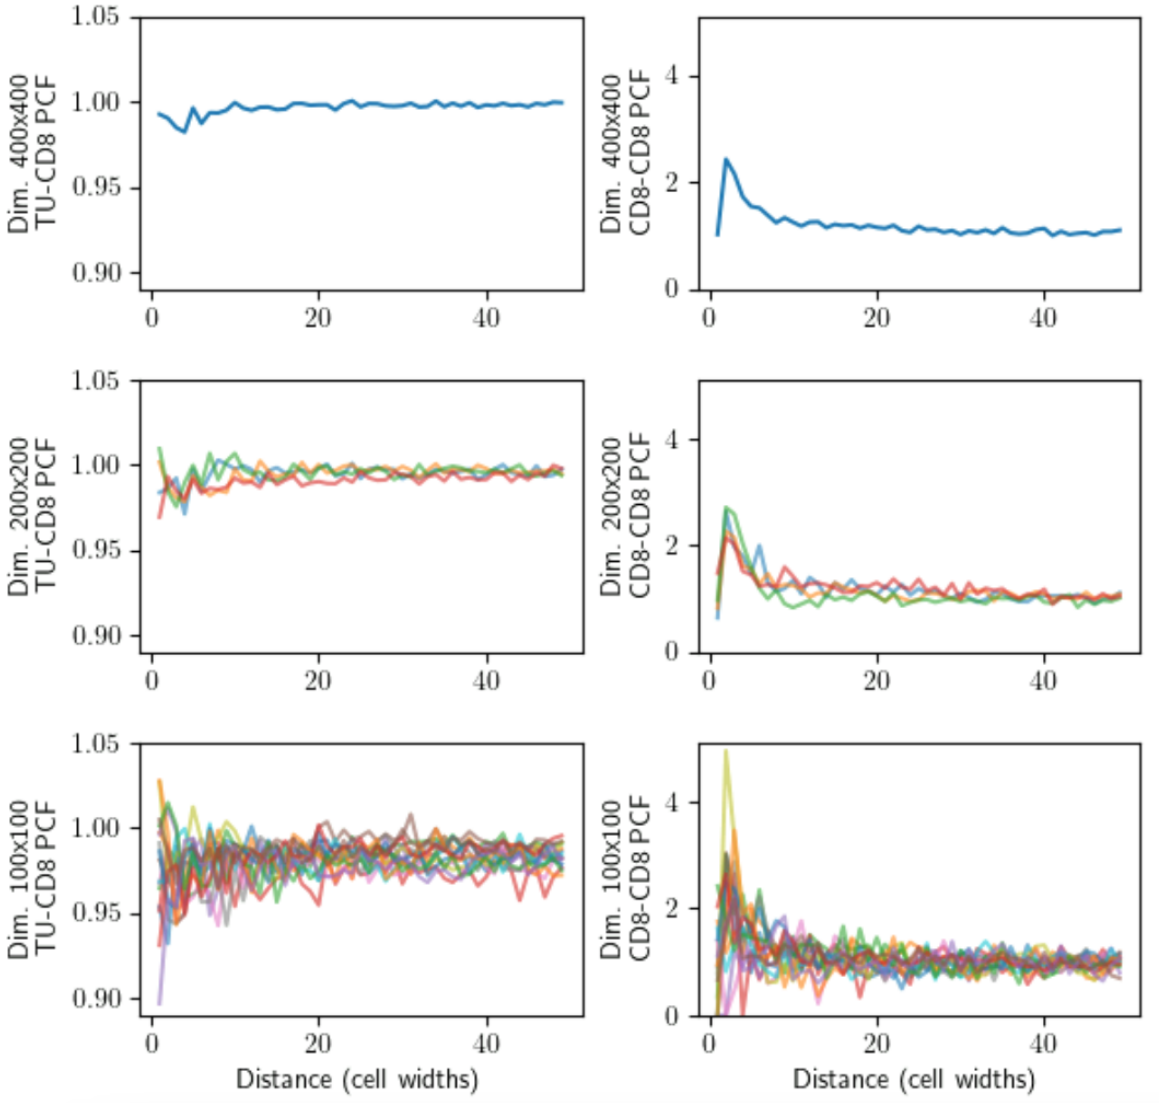


Figure 4.26: On-lattice PCFs of different dimension ROIs sampled from patient 18 pre-treatment.

We concatenate all data at a specific timepoint for each stochastic realisation (tumour and CD8 cell densities, TU-CD8 PCF and CD8-CD8 PCF), and apply principal component analysis to the data set. Figure 4.28 shows the data projected onto the first three principal components for timepoints 50, 100 and 145 days. Each data point represents a stochastic realisation of the model, and each point is coloured by the dimension of the domain. Note that the principal components are not the same between timepoints, and are timepoint-specific.

We find that data generated from different sized ROIs all cluster together in the same region of PCA space, i.e. ROI sizes are not distinguishable from each other, however there is more variation in the data generated from smaller ROIs. This suggests that, on average,

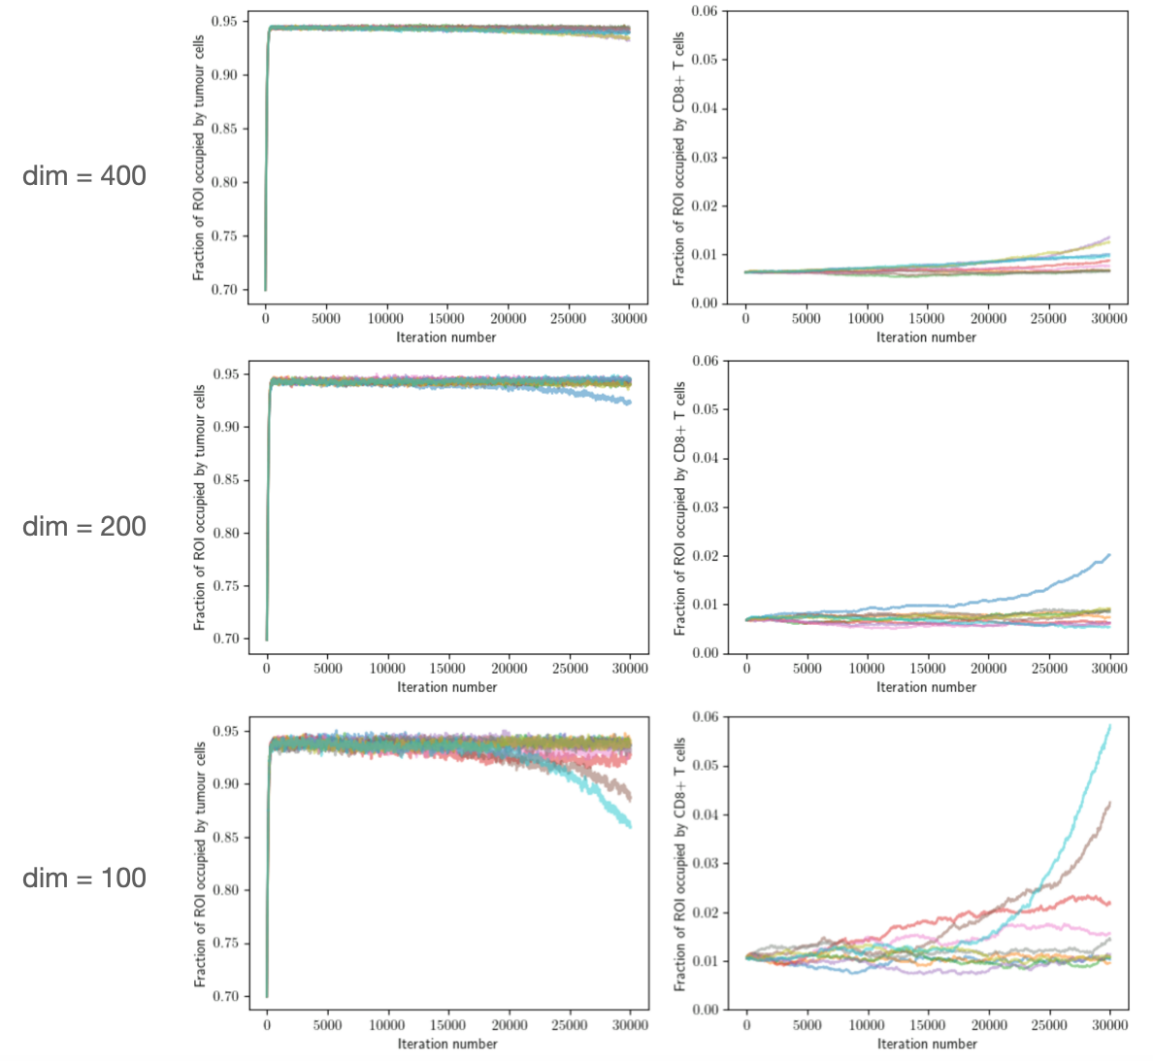


Figure 4.27: Example of timecourses of tumour cell and CD8+ T cell populations for three different ROI sizes. Each row shows results for one ROI of the specified size, and lines represent the timecourses from each of the 10 stochastic realisations. Note that the ROIs with the most variation in model output have been selected to demonstrate the maximum variation of results between different ROI sizes.

the model predictions are robust to ROI size for this particular parameter set, although we must note that as found earlier, there are a few outliers. Of course, we have only considered one parameter set in this case, which may not show significant spatial dynamics due to the random motion of the CD8+ cells. We have shown that even in this case, there are some ROIs of size 100 which may not be representative of the whole sample (size 400), however there are some size 100 ROIs which lie close to the size 400 ROIs in PCA space, and may produce representative model predictions. This would be vital to reduce computational costs whilst considering the representativeness of the ROI used to initialise the model. We

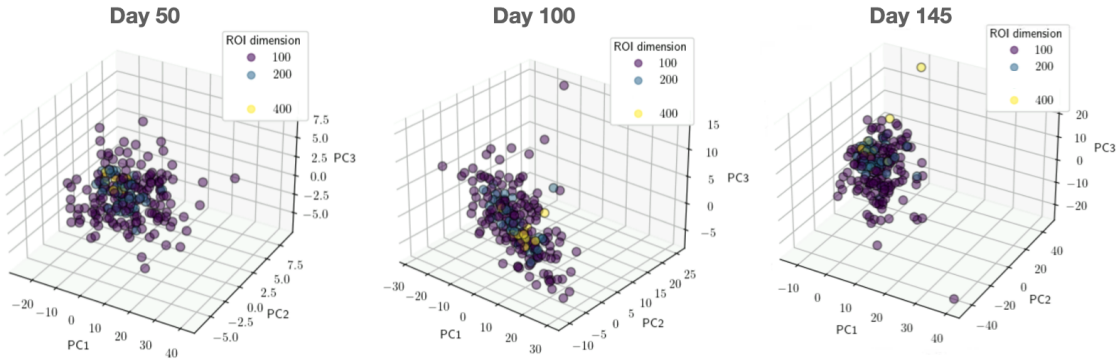


Figure 4.28: PCA timecourse for three different timepoints, for simulated data with initial conditions of different sizes.

propose that in future work, given a tissue sample, this procedure of sampling different sized ROIs could be done for random parameter sets sampled from parameter space, as it is computationally expensive to simulate large ROIs. From these simulations, you may be able to identify one or several smaller ROIs which yield the same predictions as the large ROI. These selected smaller ROIs could then be used to initialise the model for a more comprehensive parameter sweep to make predictions of tumour outcome.

4.6.5 Dependence of classifier prediction accuracy on initial condition

We now take three of the ROIs of size 100 cell widths, initialise the ABM with each ROI, and conduct a parameter sweep over the CD8+ cell proliferation probability ($r_E = \{0.0002, 0.00025, 0.0003, 0.00035, 0.0004, 0.00045, 0.0005\}$) and CD8+ migration bias ($\beta = \{0.25, 0.5, 0.75, 1\}$). We simulate the model for 10 realisations per parameter set. We then follow the same procedure as earlier, and train a classifier at each timepoint, for the synthetic tumours generated from each initial condition, thereby producing a set of classifiers for each initial condition. Figure 4.29 shows the accuracy of these classifiers over time, for each of the three initial conditions.

From analysing the overall classification accuracy plots in Figure 4.29, we find that for a given timepoint, depending on the initial condition used to generate the synthetic tumours, there is a discrepancy as to whether cell counts alone, or cell counts combined with PCF data are more informative.

4.7 Clinical relevance

The primary aim of the work in this chapter was to propose a method to find an upper bound on the accuracy of predictions made from a single timepoint of data, in a system which is

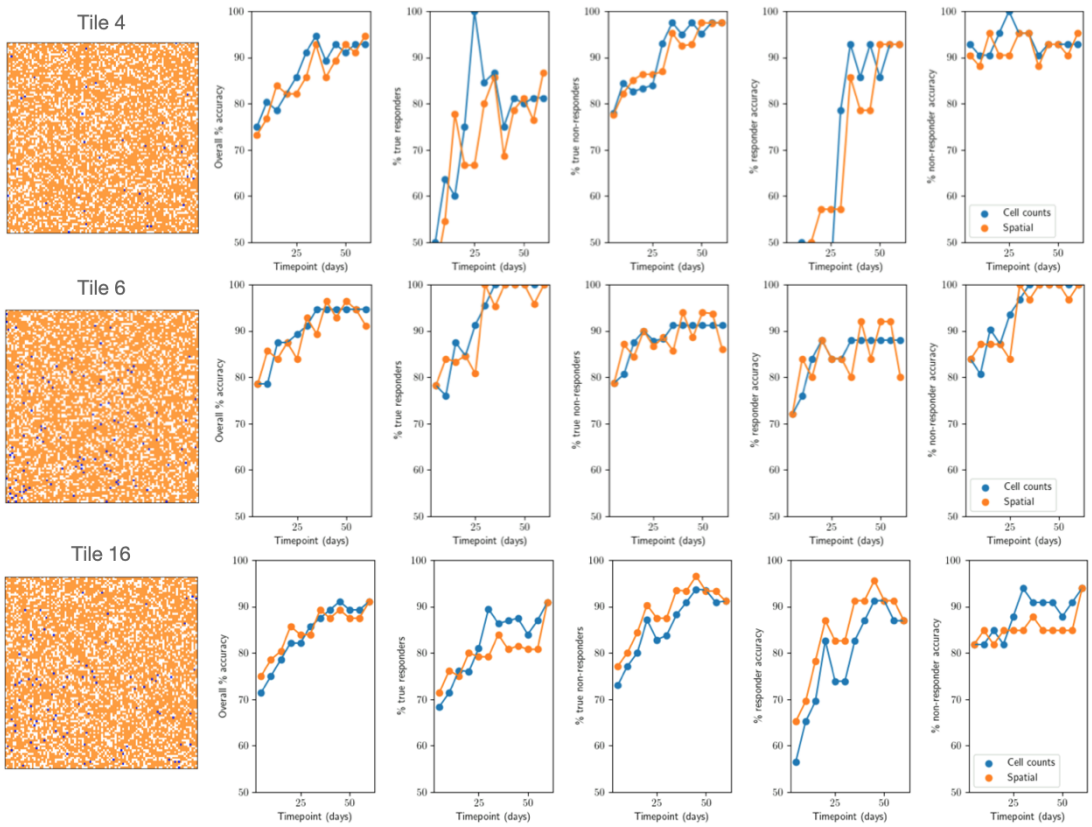


Figure 4.29: SVMs trained at single timepoints for three different initial conditions.

fixed, by varying two model parameters to produce different tumour outcomes. In this method we can sample the time course data at small time intervals, in addition to knowing the outcome at the end timepoint. Therefore, if there is uncertainty in the predicted tumour outcomes in this highly constrained system, where the classifiers were trained using the same model that produced the synthetic data, then this would suggest that data from a single timepoint is insufficient to make predictions in a clinical setting. While our method worked when most model parameters were held fixed, in order to fully assess its suitability of our method for clinical applications, a more comprehensive sampling of the global parameter space, and inspection of the resulting classifier accuracies, is required.

Depending on the clinical tolerance of uncertainty in the model predictions, our protocol may need to be adapted for use in a clinical setting. A key limitation of the method is uncertainty in the patient-specific model parameters. Although we have mitigated this by proposing to sample within the physiologically relevant parameter space, this may be too large a space to sample, and data from in vivo and in vitro experiments (such as known doubling rates for the tumour type) may be needed to reduce the parameter space.

In our preliminary analysis, supplementing cell counts with spatial information from the

PCF and cross-PCF worked well to distinguish tumour response, particularly in bistable regions of parameter space. In these regions, where multiple different outcomes result from the same parameter set, the spatial configuration of the CD8+ T cells is important in determining the outcome, demonstrated by the increased accuracy of the spatial classifier compared to the cell counts classifier. In regions of parameter space which do not exhibit bistability (in our example, high proliferation regimes, and high immune migration bias regimes), the timepoint at which the tumour outcome is obvious from the cell population data is easily identified in the cell count time course figures, and there is low variation in the cell population time courses between stochastic realisations.

As it is the bistable regions of parameter space in which we are most interested, due to the spatial configuration of cells driving the tumour outcome, these are the regions of parameter space with which we would want to train the classifiers in a clinical setting. Then, the minimum timepoint for a suitable level of accuracy in tumour classification would be determined from the bistable regions classifier, and the minimum outcome-distinguishing timepoint extracted from the regions of parameter space which are primarily driven by cell numbers. Otherwise, depending on the choice of the sampling intervals for each parameter, we may not pick up on these bistable regimes, and the classifier would be biased towards identifying the "clear" outcomes, i.e. the parameter sets which always result in tumour elimination or escape. This is demonstrated in our results in this chapter, where a finer sampling of the CD8+ T cell proliferation probability parameter was required to produce synthetic tumour data in the bistable regions of the parameter space.

In summary, this method may not be clinically applicable in its current form, but it provides a method of assessing if, given full knowledge of the biological mechanisms and dynamics, it is possible to make an accurate predictions using data from a single timepoint, for a system simulated by an ABM.

4.8 Discussion

In this chapter, we have investigated how biopsy data such as the Simlukafusp data set may be integrated with an ABM for clinical applications. We aimed to investigate whether we can use image analysis and mathematical modelling to aid decisions about the timing of on-treatment biopsies within a clinical trial. We used the ABM to ascertain whether there is a timepoint(s) at which the simulated tumour outcome prediction accuracy is maximised, and whether analysis at this timepoint is enhanced by incorporating spatial data extracted from the model.

To investigate whether spatial information extracted from the ABM by calculating the PCF can shed light on the tumour population dynamics, we started by performing a coarse sweep over two ABM parameters, keeping all other parameters fixed. We varied CD8 proliferation and CD8 biased motion, as they are relevant to Simlukafusp, an immunotherapy which is designed to stimulate CTL proliferation.

We found that the g_{max} value for the CD8-CD8 PCF initially increased, reached a maximum and then rapidly decreased, for parameter sets which resulted in tumour elimination. We reasoned that this was because increased CD8 proliferation probabilities in these parameter sets increases the size of the CD8 population and the co-localisation of CD8 cells. These CD8 cells kill the tumour cells and drive expansion of the CD8 clusters. The clusters eventually become so large that the CD8 cells migrate randomly into any empty spaces, causing g_{max} to decrease at later times. These temporal changes in the PCF metrics motivated our focus on PCFs in the remainder of the chapter.

Simulations for each parameter set were manually classified as a responder or non-responder, based on whether the tumour burden reduced by 50% at day 150. A rich timecourse of cell count and PCF data was recorded for each simulation, and concatenated into a vector at distinct timepoints. To mimic the process of sampling the on-treatment biopsies, we extracted data from a single timepoint, and trained an SVM classifier to predict tumour outcomes at day 150 based on the data at that timepoint. The SVM classifiers at each timepoint were tested. Including the PCF data was found to provide better classifier accuracy than using cell counts alone when selecting a timepoint of 25 days. This timepoint corresponds to the time at which the spatial dynamics start to diverge, illustrated by the timecourse of the CD8-CD8 g_{max} values.

We then investigated how multiple on-treatment biopsies may be predictive of tumour response. Several patients in the Simlukafusp clinical trials received two on-treatment biopsies, although we do not know if these were taken at the same time, or at different times. Assuming these biopsies were taken at different times, it raised the question about the extent

to which taking multiple on-treatment biopsies improves the accuracy of tumour outcome predictions. Using SVM classifiers trained using combinations of timepoint data, we identified combinations of on-treatment timepoints that maximise the classification accuracy. We also showed that the accuracy of the spatial classifier could be improved by including the g_{max} value in the data rather than the whole PCF signature. An interesting question for future work would be to consider whether multiple biopsies at the same timepoint, or single biopsies at different timepoints may be more informative. In our approach, we assumed that the two on-treatment timepoints are fixed prior to any on-treatment biopsy being sampled. We could extend our method to use the first on-treatment sample to inform and adapt the scheduling of the second on-treatment sample (Browning et al., 2022).

We then integrated biological data from the Simlukafusp data set into our ABM, to investigate whether the size of the ROI had an effect on the spatial dynamics. We found more variation in cell population timecourses in smaller ROIs, indicating that there is a sensitivity to the size of the sample which is simulated. We hypothesise this could be due to the boundary conditions, which become more significant for smaller ROIs.

Although we used biological data to initialise our model, we postpone comparison of on-treatment ROIs with simulated ROIs for future work; the point of our study is not to build a model that accurately reproduces the spatial features of the data set, but to test our methodology of tumour outcome prediction. By developing an ABM, we were able to generate timecourses of spatial and cell count data points. Predicting tumour outcomes using such synthetic data we are classifying provides the "perfect" scenario, in that we know that the underlying rules of the model are the same as those underpinning the synthetic data. Even with this advantage, at some timepoints the classifier does not accurately predict the tumour outcomes well with a single timepoint of data. In reality, we would not be so confident that the model captures every biological process driving the real tumour dynamics, and be able to justify fixing all parameters but two, as we have done in this study. Therefore, more of the parameter space would need to be sampled to capture all clinically possible behaviours. Future work would entail using latin hypercube sampling to conduct a more comprehensive sampling of biologically-feasible parameter space (Hamis et al., 2021), and to investigate how the tumour outcome predictions change with increasing uncertainty of ABM parameter values. Through this future work we would also seek to identify any parameters for which knowledge of their value results in the largest reduction in the prediction uncertainty.

Chapter 5

Discussion

The power of using quantitative approaches to personalise medicine is becoming increasingly realised, with the growing use of mathematical modelling and the development of sophisticated imaging techniques, particularly in the study of cancer. These highly resolved tissue data, coupled with advances in computing, are increasingly enabling the integration of data and models (Stéphanou et al., 2020), particularly in the case of spatial models in which complex dynamics are simulated, requiring rich data sets for model validation and parametrisation. The aims of this thesis were to develop mathematical modelling of tumour-T cell interactions, and histology image analysis methods, and demonstrate how these concepts can be integrated to provide insight in the context of biopsy scheduling in immunotherapy clinical trials. We developed an ABM of tumour-T cell interactions in Chapter 2 to identify key drivers of emergent model behaviours, and conducted a thorough investigation of how underlying modelling assumptions impact tumour outcome. In Chapter 3, we used spatial analysis methods, notably the PCF and cross-PCF, to describe and quantify the spatial dynamics underpinning mixed cell density responses to the immunotherapy drug Simlukafusp, a molecule which binds to FAP-expressing tumour-associated cells, and stimulates CD8+ cells to proliferate. We identified key considerations that must be taken into account in the analyses of cell-coordinate data from core needle biopsies, such as how to sample the images and deal with edge effects, and developed a workflow to analyse these data. We then combined these modelling and image analysis methods in Chapter 4 to propose a methodology to investigate whether these data can be integrated with an ABM to aid on-treatment biopsy scheduling decisions, which was the starting motivation behind the work of this thesis.

In this chapter, we summarise the research presented in this thesis, highlighting key findings and contextualising our results within the wider literature. We then discuss future directions of this work and identify questions that merit further investigation.

5.1 Research summary

In Chapter 2, we took an existing model of tumour-T cell interactions which exhibits bistability and temporal oscillations in the tumour cell population, and simplified it to identify a minimal model that exhibits the same qualitative dynamics.

We found oscillatory behaviour of the tumour cell population in our simple ABM. We used the PCF to analyse the spatial structure of the model simulations, and, in so doing, identified the emergent behaviour in the T cell population that drives these oscillations in the tumour population. Through sufficient proliferation and tumour killing, we found that T cells were able to form clusters within the tumour tissue which gradually expanded to erode the tumour. For some parameter regimes, the T cells formed rings which expanded through tumour killing, creating voids in the domain. If a break in the T cell ring occurred due to T cell death or migration, voids were subsequently invaded by migrating and proliferating tumour cells, resulting in oscillations in the tumour cell counts. We conducted a parameter sweep over four parameters controlling the T cell response, namely proliferation, tumour killing, migration and biased migration towards tumour cells. We found that the spatial distribution of T cells within the tumour varies as the parameters governing their behaviour are varied, and in particular we used the PCF to quantify the different spatial distributions T cells form to achieve 50% reduction in tumour burden in different regions of parameter space. We found that tumour elimination time plateaus as T cell proliferation is increased, due to crowding effects which mean that T cells marked to proliferate may not be permitted to, as we assumed single cell occupancy of the lattice sites, and also restricting T cell infiltration into high density regions of tumour cells, meaning that tumour killing is limited to the tumour edges which are easily accessed by the T cells. We also found an interesting non-monotonic relationship between the tumour elimination time and the T cell migration bias, in which there is an optimal bias value between 0 and 1 (random and fully biased motion, respectively), for minimising time to tumour elimination. This behaviour is particularly pronounced at low T cell proliferation probabilities, as fully biased motion causes T cells to localise in dense regions of tumour cells, meaning that proliferation space is restricted under the single cell occupancy assumption. Our results suggest that in the case of our on-lattice model, the best strategy for T cells to adopt is a combination of random and directed movement towards tumour cells. As this was an unexpected result, we explored how robust this result was to changes in the modelling assumptions underpinning our ABM. We found a significant dependence on the choice of timestep, and the underlying kernel governing the biased movement of the T cells. Our findings point to the need for a systematic testing of how sensitive the results of an ABM are to the underlying assumptions.

In Chapter 3, we used the PCF and cross-PCF to analyse the spatial structure of paired pre- and on-treatment biopsy image data from Phase I clinical trials of Simlukafusp, provided by our collaborators at Roche. First, we developed a pipeline to process the cell coordinate and pathologist annotation data, and to sample ROIs from the images, after assessing several possible methods of sampling the data. Then, we computed the changes in total cell densities and spatial variation in cell densities between pre- and on-treatment samples and found large variability in patient responses to the treatment. We did not find any significant trends in individual PCF metrics such as g_{max}, so considered the whole PCF signature (the PCF vector calculated from $r = 0$ to $r = r_{max}$) in a clustering analysis to utilise the spatial data for a larger distance range. We found that ROIs separated into two groups; one with co-localisation of tumour and CD8 cells, and the other with dispersion, with clustering of CD8 cells at large distances from tumour cells. Recalling the proposed mechanism of action of Simlukafusp (i.e. it localises in the stroma and promotes localised CD8+ T cell proliferation), we hypothesised that differences in the spatial distribution of the stroma relative to the tumour microenvironment may explain why there is such a wide variety in the spatial localisation of tumour and CD8+ T cells and, hence, the patient responses to Simlukafusp. We also developed a custom ROI tool which allows the user to define a specific ROI to analyse within the image. With further work, this would be a useful tool to analyse an image with a dynamically-updating dashboard of spatial statistics.

In Chapter 4, we considered the question which motivated this project, namely, whether mathematical modelling could be used to optimise the decision as to when the on-treatment biopsy should be taken within a clinical trial. Most patients in the Simlukafusp data set had only one on-treatment biopsy, because the procedure is highly invasive. Scheduling this biopsy at a timepoint at which a response is most likely to be visible is critical such that if a patient is not responding then the clinician can re-assess the treatment plan, and if the patient is a responder, that the treatment has had sufficient time for its effect to be visible. We approached this problem of optimising on-treatment biopsy scheduling from the angle of a "perfect" system, in which we used synthetic data from an ABM of tumour-T cell interactions to train and test the accuracy of a classifier (motivated by the work of Bull and Byrne (2023)) that predicts tumour outcomes from ABM data sampled at a single timepoint. We fixed all parameters except CD8+ proliferation probability and CD8+ migration bias. By varying these parameters across a range of values, we generated different qualitative tumour outcomes at day 150. We trained and tested SVM classifiers using data from a single timepoint. We found that using only cell count data yielded more accurate predictions at most timepoints than supplementing with PCF data. However at a specific timepoint (day 25), the accuracy of the spatial classifier (cell counts and PCF data)

attained its maximum value (over 94%) at an earlier timepoint than the cell counts classifier. By training classifiers using data from two timepoints (in the range of 5 to 80 days), we identified biopsy scheduling times for which cell counts alone, or cell counts combined with spatial data, were most accurate for predicting tumour outcomes at $t = 150$ days. Our novel adaptation of an existing method of spatial analysis of an ABM allows us to make conclusions as to the accuracy of a tumour outcome classifier which is trained on limited data, i.e. from one or two timepoints. Our method does not rely on the ability to fit the ABM parameters, but instead simulates all realistic biological outcomes for a given initial condition and specified parameter space. We have identified regions of parameter space where cell counts alone outperform a spatial classifier, and other regions where spatial effects are important for determining tumour outcomes. We have shown that even when using synthetic data generated from the ABM which is used to train the classifier, and fixing all but two parameters, bistable behaviour of the model can result in uncertainty in tumour outcome predictions. We expect that when our method is used in a more realistic setting, for example when most of the parameters are not exactly specified, one or two timepoints of data may not be sufficient to have confidence in the outcome predictions generated by the model.

5.2 Future directions

Having summarised the findings of this thesis, we now highlight new questions stimulated by our work, and discuss possible directions for future work. We consider separately extensions to the two main themes of our work, and then suggest how their integration could be further developed, given suitable data.

5.2.1 Extension of modelling approach

In this thesis, we focussed on developing a simple ABM of tumour-T cell interactions which minimised the number of parameters required, to identify key mechanisms driving the qualitative behaviours found such as oscillations in the density of the tumour cell population. Furthermore, we sought to minimise the model complexity such that it could be applicable to integration with histology data, which was limited. There are several avenues for future model development.

First, further investigation into the timestepping algorithm used to implement the simulation of the ABM, incorporating engagement time for cell actions (for example, cell death), and cells occupying the grid space for longer while these processes are happening would be an interesting extension to the current model, and would mitigate the current "instantaneous" nature of birth and death processes. Another area for development would be to

account for differences in the sizes of tumour cells and T cells. This could be achieved by relaxing the single cell occupancy rule, and allowing multiple occupancy of lattice sites. This would reduce crowding effects in the lattice by allowing more CD8+ cells to occupy lattice sites. Some behaviours seen in the single occupancy model, such as the T cell rings, may be less robust, as these structures arise as a result of T cells being blocked from infiltrating the tumour by the single cell occupancy rule.

Even if the occupancy rules are relaxed, on-lattice models are constrained in terms of the biological realism they can simulate. An off-lattice model would allow more complex migration behaviour of the T cells. However, these added complexities increase the computational costs. A possible avenue of work could be to compare on-lattice models which have different occupancy rules with off-lattice models to investigate which on-lattice model best replicates the spatiotemporal dynamics of the off-lattice models, and to identify conditions (or parameter regimes) for which more detailed models are required. By comparing off-lattice and on-lattice models to identify simple and computationally efficient models that emulate complex biological processes, methods such as those proposed in Chapter 4 for predicting on-treatment timepoints may become more computationally feasible.

5.2.2 Development of image analysis

In this thesis we have focussed on using the PCF to characterise the spatial distribution of cells, however the methods we have used are easily extended (and already have been in the literature) to incorporate other spatial statistics, or techniques such as network-based methods or TDA. It would be interesting to investigate whether the tumour outcome predictions in Chapter 4 could be improved by using an optimum combination of spatial statistics.

A significant avenue for future work would be to stain for stromal markers in the biopsy samples, as this would provide vital information as to whether the Simlukafusp treatment is working in some patients who demonstrated an apparent response in cell density measurements, and would allow us to analyse whether CD8+ T cell proliferation is increased in these stromal regions. We would also be able to analyse the spatial distribution of stromal regions within the tumour, and ascertain whether this distribution is important in treatment response.

5.2.3 Integration of mathematical modelling and image data

In Chapter 4, we demonstrated how image data may be integrated into a modelling framework to make predictions of tumour outcome. This work can be extended in several ways.

First, a wider sampling of parameter space must be conducted and the accuracy of predictions assessed, as in our proof of concept study in Chapter 4, we fixed all but two parameter values; in reality there would be an associated biologically feasible range of values for each parameter which should be sampled from. A systematic study of ABM predictions using initial conditions of different ROI sizes may provide insight into how to best sample an image for integration with our method, for example, whether a large tissue sample can be represented by a smaller ROI which is computationally more efficient to simulate. Finally, other responder definitions could be investigated, for example, expansion of the CD8 cell population, to ascertain whether the accuracy of using cell counts alone or in combination with spatial data is dependent on the responder definition, i.e. what characteristics the clinician is specifically looking for.

5.3 Concluding remarks

In this thesis, we have built upon existing mathematical modelling and spatial analysis methods, in the context of informing biopsy scheduling in oncology clinical trials. Our work has uncovered complex behaviours in our on-lattice model of tumour-T cell interactions, and has highlighted the need for consideration of the effect of modelling assumptions on the emergent dynamics of such models, to establish whether the emergent behaviour is reflective of the biology or the modelling framework used. Our work has pointed to developments to tailor our model more closely to the biological mechanisms of tumour-T cell interactions, for example, by considering differences in cell size and T cell migration mechanisms. Through our analysis of the Simlukafusp dataset, we have demonstrated a set of methods by which core-needle biopsy data may be analysed, and discussed how to account for disparities in the data, for example, differences in on-treatment biopsy timepoints. We have demonstrated how spatial analysis of these data may help to explain dynamic trends in the cell density data. We have presented a methodology to ascertain the extent to which an ABM can be used to optimise on-treatment biopsy scheduling with limited data. Our work reflects the increasing integration of mathematical modelling and imaging data using spatial analysis, and highlights key areas of consideration when using these methods.

Appendix A

Initial condition

In this Appendix we detail the rationale behind the choice of initial condition for our model, and results of preliminary simulations used in this decision. To ascertain a benchmark for the behaviour of our model, we chose to initialise our model using a random distribution of cells within the domain. We want the initial condition and subsequent simulation results to be representative of a random initial condition, however we also want the initial condition spatial distribution to be preserved, because if we run 100 repeats using a different random configuration each time, as it is a finite grid, then the spatial distribution of the initial condition will be slightly different each time. However, we also want to make sure that the random configuration we choose as the initial condition is indeed representative of a randomly distributed population of cells. For example, there is a non-zero probability that random allocation of cell positions could result in clumps of CTLs, even if the probability of this happening is very small.

A.1 Initial cell numbers

In this model, we assume that our initial condition is a region of interest within the tumour tissue, i.e. we assume that the tumour is already established and therefore that the ROI does not contain the tumour margin. Hence, it makes sense to assume that the ROI has a high cell density. Hutchinson and Grimm (2022) set the initial tumour cell density to 70% of the domain carrying capacity, and from examination of their sensitivity analyses, many simulations reached a steady state value of 8000 tumour cells (80% domain carrying capacity). Therefore we decided on an initial population of 8000 tumour cells. Regarding the CTL population, Hutchinson and Grimm (2022) considered initial conditions with a wide range of initial CTL population sizes, ranging from 7 to 768 cells in their sensitivity analysis. We hence chose an initial CTL population of 500 cells from preliminary simulations

of our model, in order to have enough cells to analyse their behaviour, but note that future work should consider varying the initial number of cells in the domain.

A.2 Procedure of selecting initial condition

After selecting our initial distribution and number of cells, we wanted to examine how using the same initial condition (Figure 2.1) for all repeat realisations for a given parameter set may differ from using a different random configuration of cells for each single realisation. Parameter values used for the simulations contained in this Appendix are given in Table A.1

Parameter	**Value**
r_T	0.5
m_T	0.35
d_T	0.12
r_E	0.025
m_E	0.8
d_E	0.015
κ	1
β	1
η	100

Table A.1: Parameter values used in the studies contained in this Appendix.

Figure A.1 shows the tumour cell and CTL populations over time for 100 realisations of the model using one initial condition (shown in black) and 100 realisations each using a different initial condition (shown in red and blue for the tumour and CTL populations respectively).

The mean and median cell numbers are consistent between the two cases, and the standard deviations (shaded regions) overlap. We note that there is an outlier in the single initial condition data set (shown in black). We also considered the spatial distribution of the cells, comparing the fixed initial condition case and the case where the initial condition is randomly generated for each realisation of the model. Figure A.2 shows the CTL-CTL PCF and tumour-CTL PCF for the fixed initial condition case (shown in black) and multiple initial conditions case (shown in blue). As the cells are randomly initialised within the domain, we expect that the value of the PCFs will be close to 1 at all radial distances,

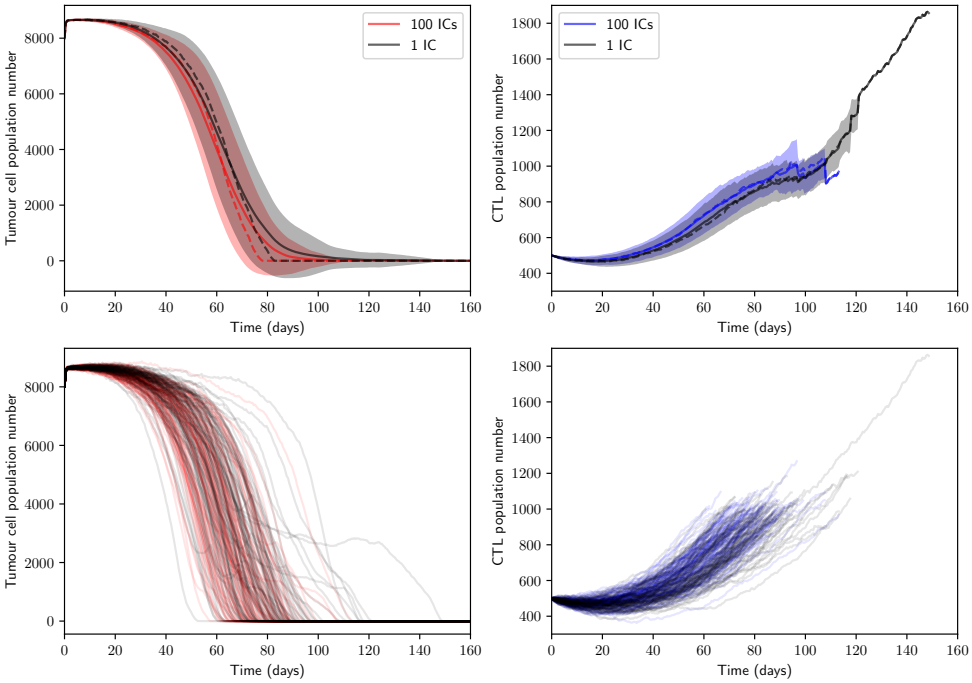


Figure A.1: Tumour cell population and CTL population numbers over time for the single initial condition case (black lines) and multiple initial condition case (red and blue lines for tumour and CTL populations respectively). The top row plots show the mean (solid line), median (dashed line) and standard deviation about the mean (shaded regions). The bottom row plots show results of individual realisations of the model.

which is what we find. Figure A.2 shows that the single initial condition mostly falls within the standard deviation of the PCF values over multiple initial conditions, and within 5% of a value of 1 for the CTL-CTL PCF, and within 1% of a value of 1 for the TU-CTL PCF.

A.2.1 Burn-in period

We note that there is a burn-in period shown in the tumour cell population in Figure A.1, which is a common feature of ABMs. However, this burn-in period is transient (is significant within the first few days) and has a shorter timescale than the timing between biopsies or tumour elimination times, which are the timescales for which this work is relevant. Furthermore, there is not an obvious procedure to completely eliminate this short-time behaviour. One method could involve growing only tumour cells within the domain, allowing the population to reach a steady state, and then using the average number of cells in this steady state as the tumour cell number in our initial condition. However, when the CTL population are introduced, this would transiently affect this steady state value, depending

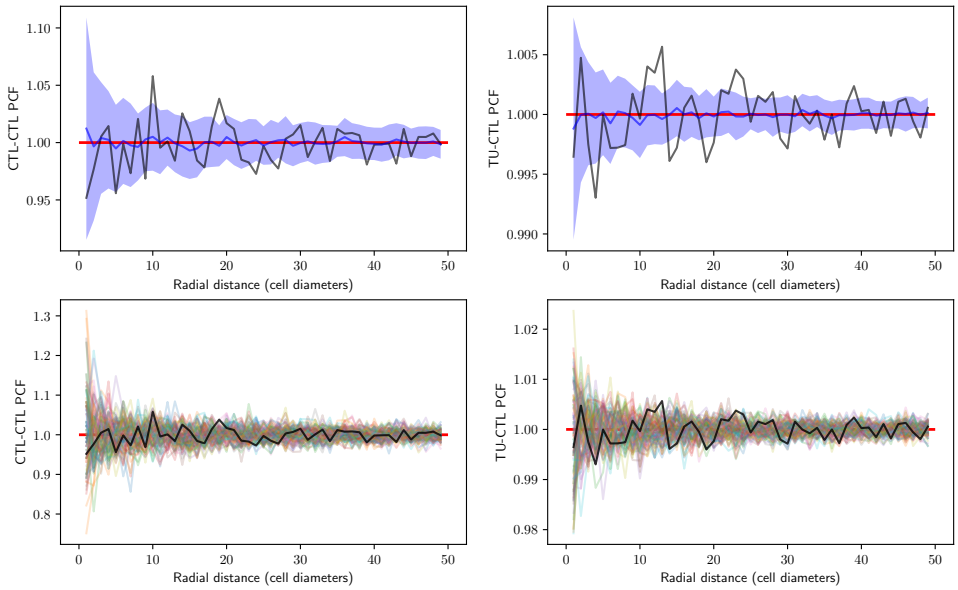


Figure A.2: Spatial distribution of initial conditions for the single initial condition case (black line), and 100 initial conditions case (blue line with shaded region corresponding to standard deviation of the 100 PCFs).

on the CTL parameter values and the number of CTLs. Methods for minimising the effect of this burn-in period would be scope for future work.

Appendix B

Distance transform map for CTL sensing of tumour cells

In this Appendix, we detail how the distance transform map is calculated such that it is used in the biased random motion of the CTLs towards the tumour cells, and investigate how the time taken to reach a single tumour cell in a simple setup is sensitive to changing the assumptions made in constructing this map. The choice to use a distance map to implement this biological mechanism was maintained from the Hutchinson and Grimm (2022) model for consistency. The map does not model the diffusion of a chemokine explicitly, but rather is used to identify neighbouring cells for CTLs to migrate to that are closest to the nearest tumour cell in order for co-localisation and interaction of the cells.

B.1 Distance transform map

The distance_transform_cdt function from the scipy ndimage package is used to calculate the distance transform map using the Chebychev distance, as shown in Figure B.1. Periodic boundary conditions are implemented by replicating the grid in a 3 x 3 configuration and calculating the distance transform, and then recovering the distance transform of the central grid.

B.2 CTL biased random walk

In our model, CTLs migrate via a biased random walk towards tumour cells. At each migration step, a random number is generated which is compared to the CTL migration bias parameter β to determine whether the CTL migrates to a randomly chosen empty neighbouring space, or whether it migrates to the empty neighbour that is closest to the nearest tumour cell, by consulting the distance transform map.

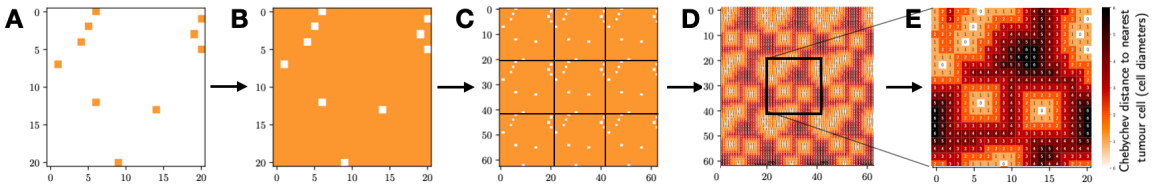


Figure B.1: Process of generating the distance transform map for the spatial distribution of tumour cells in the domain. The grid in (A) containing the tumour cells (orange squares) is inverted in (B) such that grid spaces containing a tumour cell have a value of 0 and spaces containing CTLs or that are empty have a value of 1. The domain is then replicated in a 3x3 configuration in (C). The distance transform function is applied to (C) such that the periodic boundaries can be taken into account. The resulting distance transform in (D) is trimmed such that the central panel is recovered, which is the distance transform of the tumour cells in (A) with periodic boundaries applied.

We test the effect of this chemotaxis map on the motion of CTLs using the simple setup shown in Figure B.2, with a single CTL sensing one tumour cell. For the baseline set of parameters given in Table B.1, the bias parameter β is increased from 0 (completely random CTL motion) to 1 (fully directed CTL motion towards tumour cell). Figure B.4 shows a selection of 10 trajectories of the CTL for different realisations of the simulation, for increasing values of β. These experiments were repeated for four different distance transform cases given below, to compare the effect of the domain boundary conditions and distance metric on a single CTL migration trajectory.

(i) Non-periodic boundary conditions and Chebychev distance (NPBC Chebychev)

(ii) Non-periodic boundary conditions and Euclidean distance (NPBC Euclidean)

(iii) Periodic boundary conditions and Chebychev distance (PBC Chebychev)

(iv) Periodic boundary conditions and Euclidean distance (PBC Euclidean)

Figure B.3 shows the time for the CTL to reach the tumour cell t_{final} as a function of the bias in CTL motion, for 100 realisations of the model for each case and parameter set.

The general trend shows a monotonic decrease in time to reach the tumour as the CTL motion becomes more directed, which is as expected. For each value of β, the error bars for each case overlap, except for the case of NPBC Chebychev and $\beta = 0$, in which the CTL did not reach the tumour cell in 500 days for any of the model realisations. The difference in the mean behaviour between the four cases decreases as β increases, suggesting that the choice of boundary conditions and distance metric is more important for low values of β, i.e. $\beta < 0.2$. For $\beta \geq 0.7$, the difference in t_{final} between all four cases is indistinguishable.

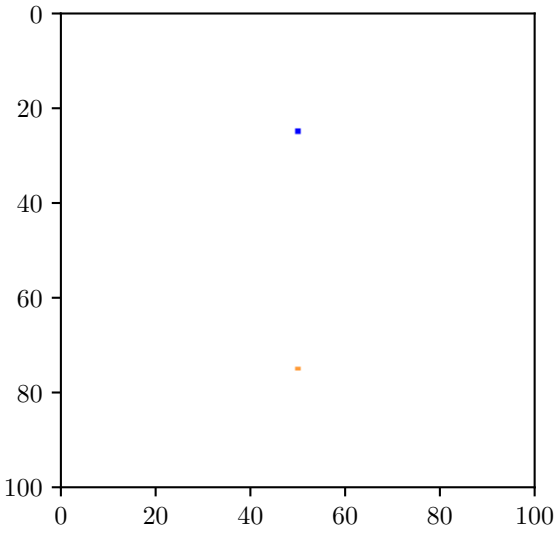


Figure B.2: Initial condition for the bias testing. One CTL is shown by the blue square and a single tumour cell by the orange square. The relevant boundary conditions are imposed at the edges of the domain. The distance between the CTL and tumour cell through the domain is equal to the distance between them by traversing over the periodic boundary.

Parameter	**Value**
r_T	0
m_T	0
d_T	0
r_E	0
m_E	1
d_E	0
κ	1
β	[0,1] in 0.1 intervals
η	1

Table B.1: Parameter values used for the studies presented in this Appendix.

Comparing the NPBC and PBC cases, using PBCs decreases t_{final} for a given distance metric compared to the NPBC case for $\beta < 0.2$.

Comparing distance metrics, using the Euclidean metric for a given boundary condition and β results in smaller t_{final}. This can be explained by the CTL trajectories shown in Figure B.3, in which there is more variance about the shortest-distance line between the

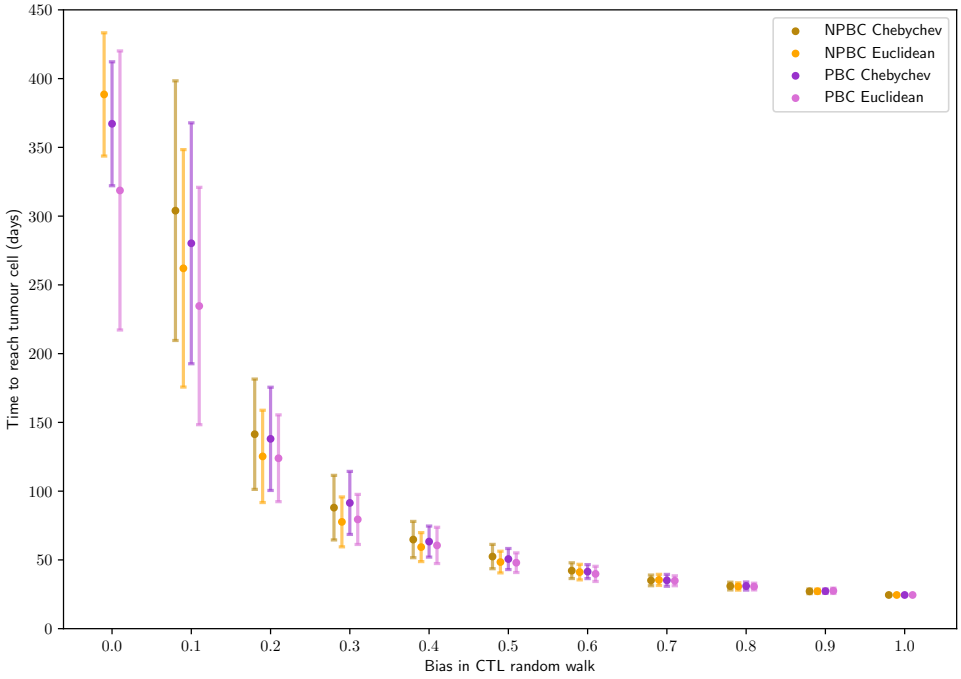


Figure B.3: Time (in days) for the CTL to reach the tumour cell (reaching tumour cell is defined as occupying one of the neighbouring sites of the tumour cell, such that the tumour cell may be killed by the CTL). PBC and NPBC denote periodic and non-periodic boundary conditions respectively. Data points show the mean time and error bars denote the standard deviation for 100 repeat simulations. The simulation was run for 500 days. In the case of the non-periodic Chebychev distance map and $\beta = 0$, the CTL did not reach the tumour cell within the 500 days in any of the 100 simulation runs.

initial CTL position and the tumour cell. This is because with the Chebychev distance, there is less precision in the way the CTLs sense if a space is closer to the tumour cell, i.e. as the distance increases in concentric squares from the tumour cell, a movement along the horizontal axis does not change the Chebychev distance from the tumour cell but does change the Euclidean distance.

In summary, this very simple study shows that the effect of the choice of boundary conditions and distance metric becomes less significant as the bias in the CTL random biased walk, β, increases. For small β, periodic boundary conditions decrease t_{final} for a given distance metric compared to employing non-periodic boundary conditions, and using the Euclidean distance decreases t_{final} compared to using the Chebychev distance.

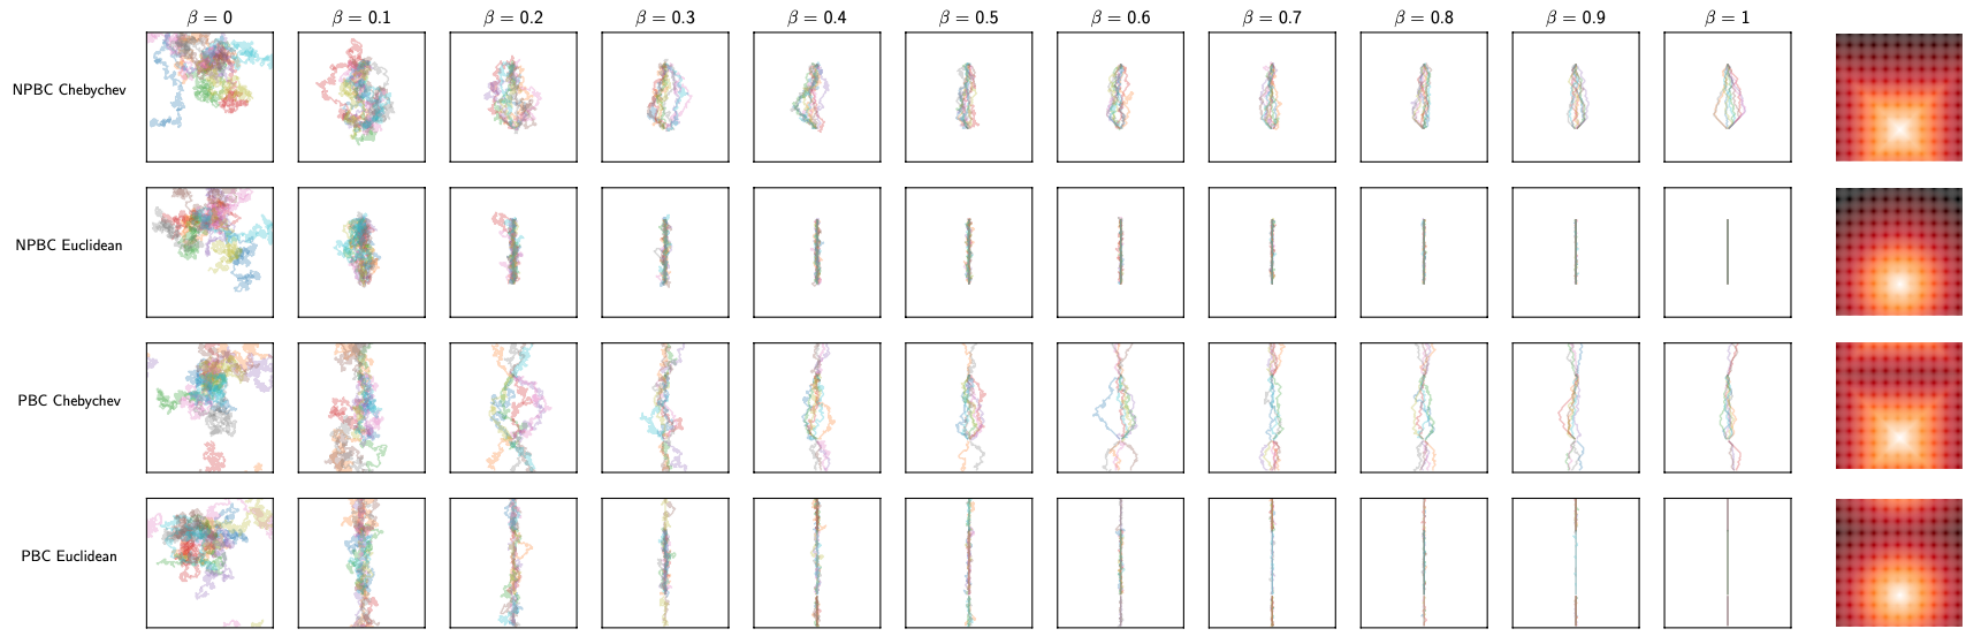


Figure B.4: Trajectories of CTLs for a sample of 10 repeat simulations. The trajectory of each CTL is shown in a different colour. Panels left to right show CTL trajectories for increasing values of the bias parameter β. Panels top to bottom show four different methods of generating the distance map. The furthermost right panels show the distance transform map used to direct the CTL towards the tumour cell.

Appendix C

Parameter scaling

In Chapter 2 we investigate the effect of the choice of timestep on the predictions of the model. In order to scale the parameter values appropriately, a series of supporting simulations were conducted to verify the scaling factors used, which are presented in this Appendix. We consider a simple non-spatial model of cell proliferation and death to provide the reasoning behind our scaling choices.

C.1 Notation

Symbol	Definition
r_T	Cell proliferation probability
d_T	Cell death probability
ξ	Random number $\in [0, 1]$
$\mathcal{P}_{\tau_i}(A)$	Probability of event A occurring during timestep of length $\tau_i = \Delta t$
$\mathcal{P}_{\tau_i}(A')$	Probability of event A not occurring during timestep of length $\tau_i = \Delta t$

Table C.1: Parameter scaling notation

C.2 Non-spatial stochastic model

We firstly consider a non-spatial stochastic model to determine the scaling factor appropriate when reducing the timestep, with the aim of preserving the resultant cell population dynamics. We consider a single cell population with prescribed proliferation and death probabilities, $r_T = 0.5$ and $d_T = 0.12$, respectively, which are the same as the tumour

proliferation and death probabilities in the simulations considered in Chapter 2. An initial population of cells n_0 is considered, and at each time step, a random number ξ is generated for each cell which determines its action. Similar to our CA model in Chapter 2, if $\xi \leq r_T$ then a cell is added to the population. If $r_T < \xi \leq r_T + d_T$, then the cell is removed from the population.

We consider three methods of rescaling the parameters with decreasing timestep. Method 1 ("linear scaling") involves simply linearly scaling the probability parameter value by the factor of timestep reduction, i.e. if the timestep is reduced by a factor of 100, the probability parameter value is also scaled down by a factor of 100.

$$\mathcal{P}_{\tau_1}(A) = \frac{\mathcal{P}_{\tau_2}(A)}{\tau_r} \tag{C.1}$$

where $\tau_r = \frac{\tau_1}{\tau_2}$, τ_1 and τ_2 are timesteps of different sizes, and $\tau_1 > \tau_2$.

Method 2 ("conditional scaling") considers the probability that an event A does not occur over the number of smaller timesteps, according to the following:

$$\mathcal{P}_{\tau_1}(A) = 1 - \mathcal{P}_{\tau_1}(A') \tag{C.2}$$

$$\mathcal{P}_{\tau_1}(A) = 1 - (1 - \mathcal{P}_{\tau_2}(A))^{\tau_r}. \tag{C.3}$$

After rearranging terms,

$$\mathcal{P}_{\tau_2}(A) = 1 - (1 - \mathcal{P}_{\tau_1}(A))^{\frac{1}{\tau_r}}. \tag{C.4}$$

Method 3 ("interest rate scaling") equates the expected cell population size after one large timestep with τ_r small timesteps.

$$n_0(1 + \mathcal{P}_{\tau_2}(A))^{\tau_r} = n_0(1 + \mathcal{P}_{\tau_1}(A)). \tag{C.5}$$

Cancelling each side by a factor of n_0 and rearranging gives,

$$\mathcal{P}_{\tau_2}(A) = \sqrt[\tau_r]{1 + \mathcal{P}_{\tau_1}(A)} - 1. \tag{C.6}$$

C.2.1 Proliferation probabilities

Firstly, we considered a proliferative cell population of size $n_0 = 100$, without cell death. Cells proliferate and are added to the population at each timestep as described in Section C.2. Relative timesteps $\tau_r = \{0.001, 0.005, 0.01, 0.05, 0.1, 0.25, 0.5, 1\}$ were considered. The model was simulated up to time $t = 10$, and the cell population size at this time was recorded. The simulations were repeated 100 times, and the mean and standard deviation of the final population size at time $t = 10$ are shown in Figure C.1.

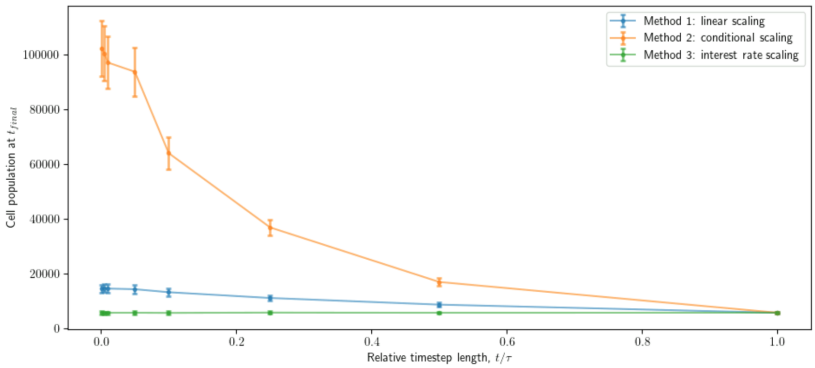


Figure C.1: Cell population at t_{final} as a function of relative timestep length for the three scaling methods, in a population with only proliferation, and no death events. Data points show mean over 100 repeat simulations, and error bars indicate standard deviation.

Method 2 is not suitable for the proliferation probability scaling, as it assumes that if a cell proliferates, then it will not be selected to proliferate again within a 12hr timestep, however we do not have memory in our model, and a cell could be selected to proliferate more than once within a 12hr period in the case of a smaller timestep $<$ 12hr. Method 3 must be used to scale proliferation in the case of this model to produce the same expected cell population size within the same time period, regardless of the timestep size.

C.2.2 Death probabilities

We next considered a population of initial size $n_0 = 10000$, which only underwent death events (i.e. there is no proliferation in the population). Cells die and are removed from the population at each timestep as described in Section C.2. Relative timesteps $\tau_r = \{0.001, 0.005, 0.01, 0.05, 0.1, 0.25, 0.5, 1\}$ were considered. The model was simulated up to time $t = 10$, and the cell population size at this time was recorded. The simulations were repeated 100 times, and the mean and standard deviation of the final population size at time $t = 10$ are shown in Figure C.2. In this case, conditional scaling (Method 2) of the death rates results in convergence in final population size as the timestep is reduced.

C.2.3 Killing probabilities

To determine an appropriate scaling for the killing probability, we consider the case of a non-proliferating tumour cell population of size $n_0 = 8000$, without natural tumour cell death. We allow the CTL population of initial size $n_0 = 500$ to kill the tumour cells, with a probability of 0.5 in the 12 hour timestep case. Figure C.3 shows that linear scaling is most appropriate for the killing probability, such that the timestep choice does not affect the final tumour population size.

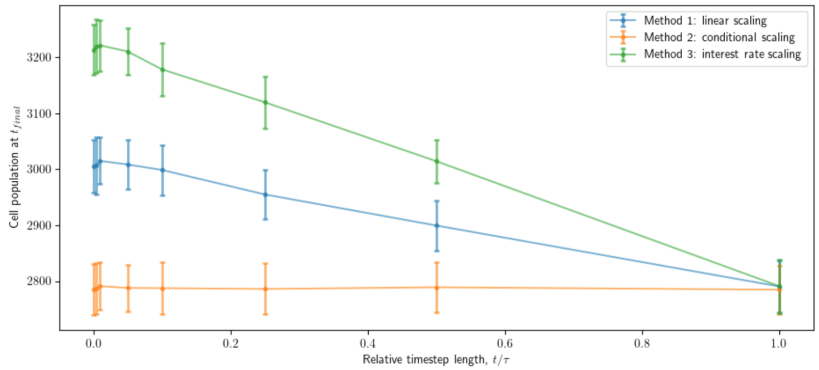


Figure C.2: Cell population at t_{final} as a function of relative timestep length for the three scaling methods, in a population with only death, and no proliferation events. Data points show the mean over 100 repeat simulations, and error bars indicate standard deviation.

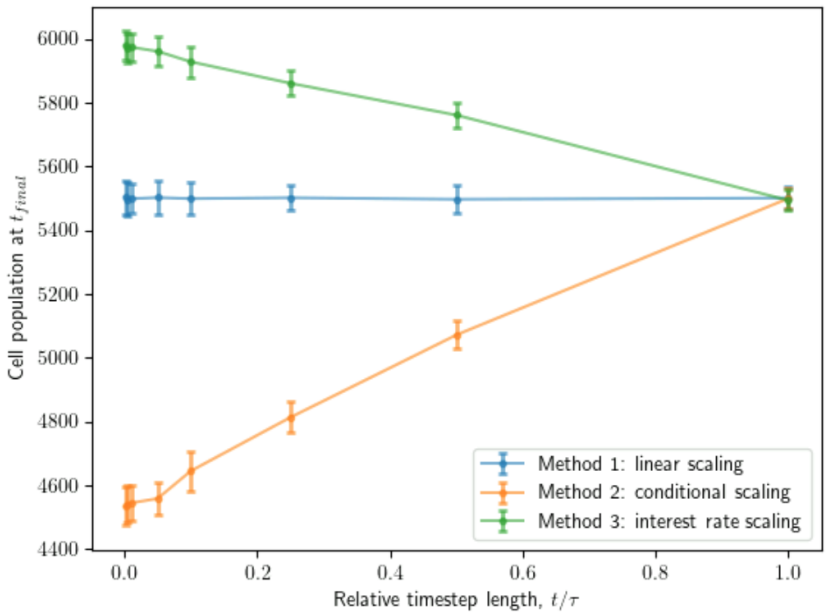


Figure C.3: Cell population at t_{final} as a function of relative timestep length for the three scaling methods, in a population with only death, and no proliferation events. Data points show the mean over 100 repeat simulations, and error bars indicate standard deviation.

In Figure C.3, we consider the case of a proliferating and dying CTL population. As the rates of these quantities are low, the linear scaling of the killing probability still holds.

C.2.4 Migration probabilities

There are two migration probability scalings to consider in our model; the probability a cell migrates in a timestep, and the number of lattice spaces a cell can migrate within that timestep. For the tumour cell population, the latter is fixed at one step, and for

the CTL population, cells may move up to 100 lattice spaces within 12 hours so as to model the relative speeds of tumour cells and CTLs. We scale down the number of CTL migration steps permitted within a timestep linearly, i.e. in 0.12 hours a CTL moves one lattice space. As we were scaling the migration steps, we keep the migration probability parameter fixed for all timestep choices. As discussed in the main text, this does not model the same dynamics in terms of the number of CTLs which migrate some distance within 12 hours, however there is not an obvious way of scaling the migration so as to exactly replicate the dynamics in the 12 hour timestep case, without introducing a "memory" to the cells, and constraining each cell to perform one action per 12 hours. We reason that our approach conserves the total number of migration steps available to the CTL population as a whole, per 12 hours. For the tumour cell population, we linearly scale the migration probability with the timestep. Similarly to the CTL population, this scaling will not result in exactly the same dynamics, however in both populations our scaling approach should yield migration probabilities with an appropriate order of magnitude to the timestep. As we are less interested in quantitatively comparing the dynamics between timestep choices, and more interested in the qualitative behaviour of whether an optimal bias value exists for smaller timesteps, we justify our simple scaling approach here.

Appendix D

CD8+ biased motion implementation

In our model we adopt a similar procedure of implementing CD8 biased motion to that used in Kather et al. (2017) and Hutchinson and Grimm (2022) to maintain consistency with these models. As detailed in Appendix B, a distance transform is used to find the Chebychev distance at each grid site to the closest tumour cell. The CTLs use this map when undergoing their biased random walk towards tumour cells. However, a limitation in this method is that the concentration of tumour cells is not considered in the construction of the map. For example, if a CTL is situated at equidistance between a single tumour cell and a group of tumour cells, then the CTL will "see" these two tumour masses as having an equivalent chemotactic "pull", as only the distance to the closest tumour cell is considered. In practice, it could be possible that the chemokine concentration secreted by tumour cells would be summed over tumour cells in a compact mass.

A preliminary study to explore the effect of the biased motion map assumptions on the model dynamics is presented in this Appendix. As a next step in accounting for the concentration of tumour cells in the domain, we use a Gaussian kernel over the tumour cell population, and convert this into a map. The Gaussian can be adjusted through the number of standard deviations (parameter σ) to adjust how diffuse the signal is from each tumour cell. Figure D.1 shows the distribution of tumour cells in the initial condition used in the main text of the chapter, and the distance transform maps and Gaussian filter maps for different values of σ which are used as the chemotaxis maps for the following simulations in this Appendix. Figure D.2 shows the profile of the map for different values of σ for the Gaussian filter, for a cross-section through the centre of the domain. As the value of σ is increased, the profile becomes more smooth. From Figure D.1 it can be seen that a Gaussian filter with a low value of σ, for example $\sigma = 0.1$, gives a similar map to using the distance transforms. This demonstrates that the assumption that Kather et al. (2017)

and Hutchinson and Grimm (2022) made in which CTLs move towards the closest tumour cell is approximately considering a strong short range attraction of the tumour cells, with a long range tail which drops off quickly. This may not be realistic to model the diffusion of a chemoattractant from the tumour cells, but it could model the way CTLs can extend finger-like protrusions and sense their surroundings.

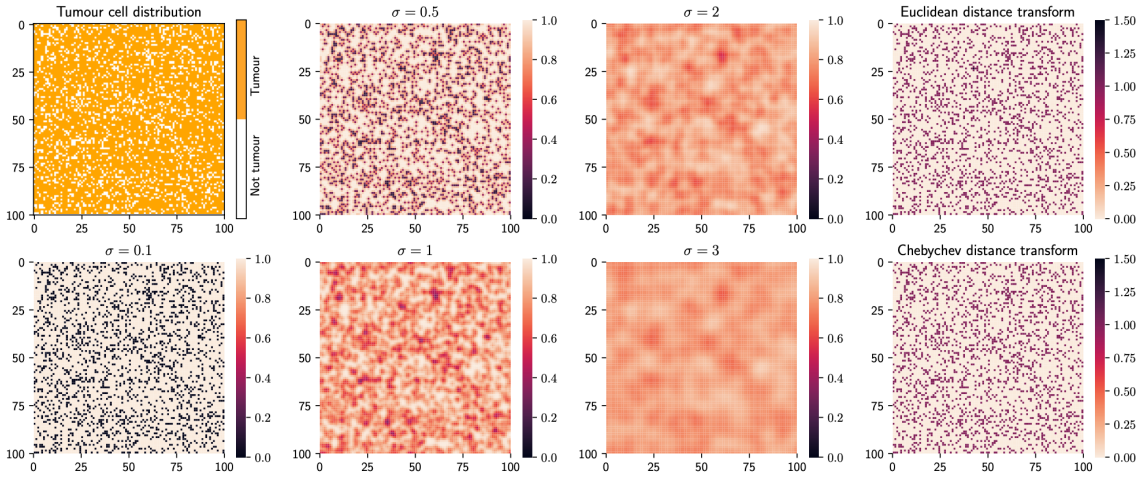


Figure D.1: Top left: Distribution of tumour cells (shown in orange) in the initial condition used to initialise the model in the main chapter analysis. The remaining plots show the maps generated using a Gaussian filter on the tumour cell distribution (all plots are indicated with the corresponding σ value of the Gaussian). The far right plots show two distance transforms; Euclidean (top right) and Chebychev (bottom right).

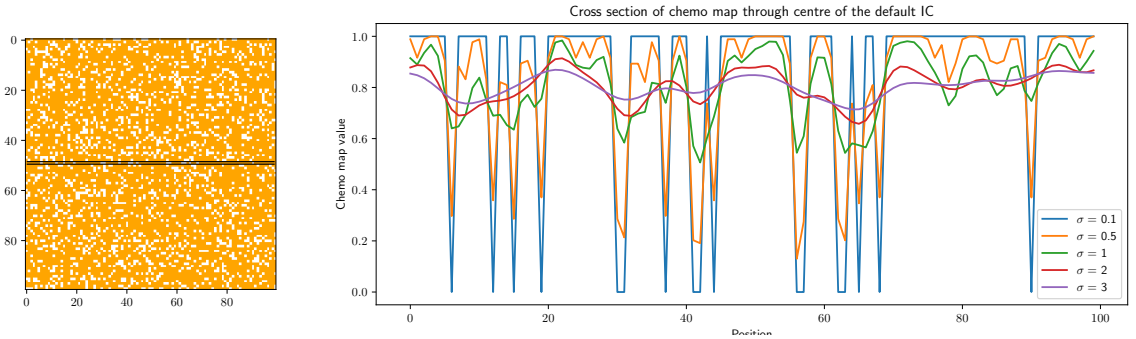


Figure D.2: Gaussian filter to generate chemotaxis map. Left: distribution of tumour cells within the initial condition. Right: Distribution of the map values for a cross section through the initial condition, indicated in the left hand plot, for Gaussian filters generated using different values of σ.

Simulations were run for the parameter values given in Table D.1, sweeping over values of the CTL migration bias parameter β for a low and high value of CTL proliferation probability. We used a timestep of 12 hours.

Parameter	Value
r_T	0.5
m_T	0.35
d_T	0.12
r_E	0.025, 0.04
m_E	0.8
d_E	0.015
κ	0.9
β	[0,1] in 0.1 intervals
η	100

Table D.1: Parameter values used in the studies contained in this appendix.

For each parameter set, 100 stochastic realisations of the model were simulated, and the mean and standard deviation of those simulations are presented in Figures D.3 and D.4.

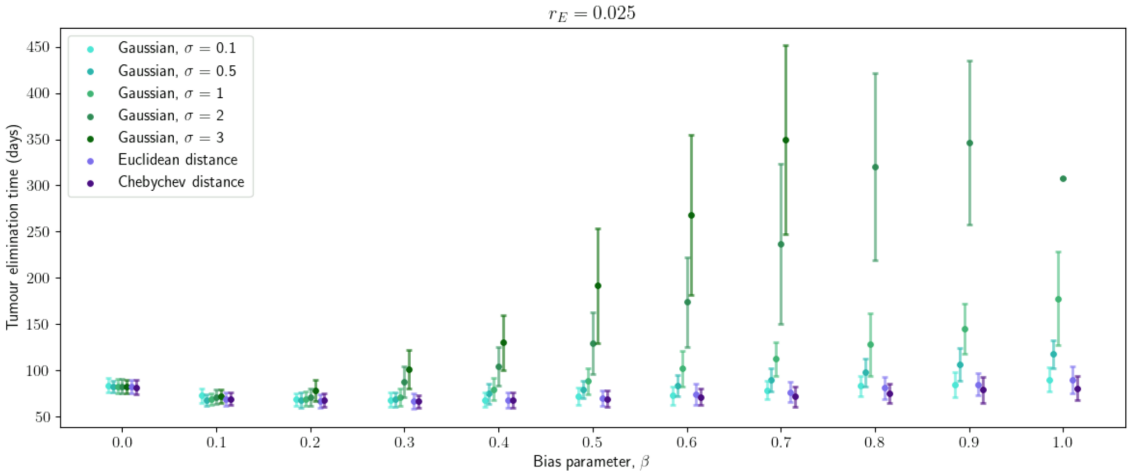


Figure D.3: Low CTL proliferation probability case: Tumour elimination time (days) as a function of bias in CTL biased random walk β, for biased motion modelled using different distance transform maps (shown in purple and indigo) and Gaussian filters using different σ values (shown in green). Data points show mean behaviour over 100 realisations and error bars indicate the standard deviation. Note that for the case of $\beta = 1$ and Gaussian filter with $\sigma = 2$, only one realisation resulted in tumour elimination, such that there are no error bars around this data point.

The "control" case of $\beta = 0$ (no bias, fully random CTL motion) shows consistent tumour elimination times across all map implementations, which is expected considering

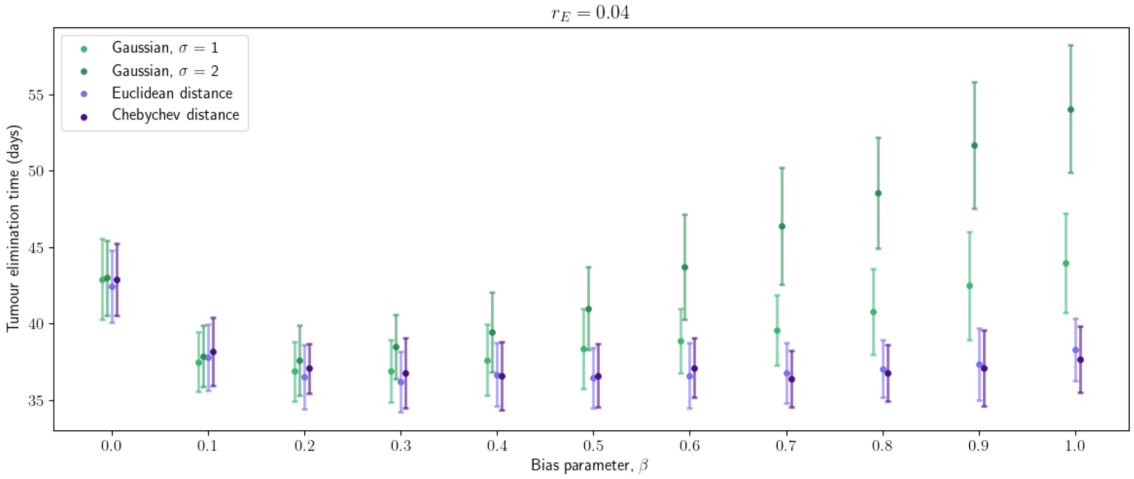


Figure D.4: High CTL proliferation probability case: Tumour elimination time (days) as a function of bias in CTL biased random walk β, for biased motion modelled using different distance transform maps (shown in purple and indigo) and Gaussian filters using different σ values (shown in green). Data points show mean behaviour over 100 realisations and error bars indicate the standard deviation.

there is no biased motion in this parameter set. As β is increased, t_{elim} for the Gaussian filter maps increases, with a greater increase for larger σ values. This could be explained by considering that at high σ, the distribution of the map becomes increasingly homogeneous, such that there is increasingly less specificity in the CTL directed motion. What is unclear however is why $t_{elim}(\beta = 1)$ does not tend towards the value of $t_{elim}(\beta = 0)$ (the case of random motion), and in future work the implementation of a PDE to model the chemotaxis could be compared to this method of Gaussian distributions over the tumour cells.

This study shows that there is not a significant difference between the different distance map metrics as the error bars overlap, and this is particularly relevant with these results as this is run with a high value of CTL migration probability, in which this region of parameter space would be most sensitive to the implementation of the biased motion. This shows that the choice of distance metric used in our model is not significant in the outcome of the results. These results also show that the existence of an optimal value of β is not a result of the assumption made in how the biased motion is implemented.

Appendix E

Patient data set

E.1 Primary indications and biopsy location

Table E.1: Patient cohort, with primary indication, baseline and on-treatment biopsy location and time of on-treatment biopsy (days). Key: RCC - renal cell carcinoma, LN - lymph node. N/A refers to data which were not available.

Patient ID	**Primary indication**	**Baseline biopsy location**	**On-treatment biopsy location**	**Time of on-treatment biopsy (days since first dose)**
14	RCC	LN	LN	113
15	RCC	Stomach	Stomach	48
18	Squamous cell carcinoma of head and neck	Skin	N/A	42
19	Adenocarcinoma of prostate	Prostate gland	N/A	49
21	Adenocarcinoma	Liver	N/A	43
24	Melanoma skin	Skin	Skin	26
25	Invasive ductal breast cancer	Lung	Lung	17
29	Breast adenocarcinoma	LN	N/A	15
31	Breast adenocarcinoma	N/A	LN	24
32	Squamous cell carcinoma	LN	N/A	17
33	Squamous cell carcinoma	Head	N/A	15
34	RCC	Lung	Lung	57
36	RCC	Liver	Liver	48

Table E.1 – continued from previous page

Patient ID	Primary indication	Baseline biopsy location	On-treatment biopsy location	Time of on-treatment biopsy (days since first dose)
38	RCC	Lung	Lung	62
39	RCC	LN	LN	57
46	RCC	Soft tissue	Soft tissue	49
47	RCC	Abdominal cavity	Abdominal cavity	57
48	RCC	Liver	Liver	50
49	RCC	LN	LN	50
50	RCC	Kidney	Kidney	50
53	Melanoma	Liver	N/A	49
54	Melanoma	Liver	N/A	57
56	Melanoma	N/A	N/A	42
59	Melanoma	Liver	N/A	43
62	Adenocarcinoma of esophagus	Liver	N/A	52
64	Adenocarcinoma of the cervix	Lung	N/A	48
65	Adenocarcinoma lung	Liver	N/A	44
69	Breast adenocarcinoma	Lung	Lung	16

Ethics statement

All participants provided written informed consent to take part in the study and the study was approved by the ethics committee of Hoffman-La Roche.

Data usage statement

The data were extracted from the Roche in-house database of biopsy samples and are used with permission from project teams.

Data availability statement

For up to date details on Roche's Global Policy on the Sharing of Clinical Information and how to request access to related clinical study documents, see here:

https://go.roche.com/data_sharing

Appendix F

Co-localisation of cells at sub-cell diameter distances

F.1 Overview

Previous analysis of the minimum distance between tumour cells suggested that the average diameter of a tumour cell is 6 pixels in these images, however we do not have a lengthscale or the raw image data set, so cannot confirm this. Pairs of tumour cells and CD8+ T cells were found at distances less than one cell diameter (6 pixels). This analysis attempts to ascertain if this is an artefact in the data and should not be considered in, for example, calculations of the PCFs and cross-PCFs.

The Roche data set provides coordinates $\{x, y\}$ for $x, y \in \mathbb{Z}$ for the following cell types:

- Ki67+ CD8- (proliferating tumour cells)
- Ki67+ CD8+ (proliferating CD8+ T cells)
- Ki67- CD8+ (non-proliferating CD8+ T cells)

F.2 Distributions of nearest neighbour distances

Figure F.1 shows the distribution of nearest neighbour distances for different cell types, for patient 18. Tumour cells are located at least 6 pixels apart. The distribution of CD8+ T cell nearest neighbours is smoother at low distances compared to the tumour-tumour distribution, but they are not found immediately next to each other. In contrast, there are pairs of tumour cells and CD8+ T cells with distances less than 6 pixels, and a few pairs which completely overlap (i.e. have the same coordinates).

Figure F.2 shows the same histograms plotted for patient 59. This patient's images had the highest frequency of perfect overlaps (3 and 9, respectively). On examination of the bottom row of plots (the nearest neighbour distances between tumour and CD8 cells), it

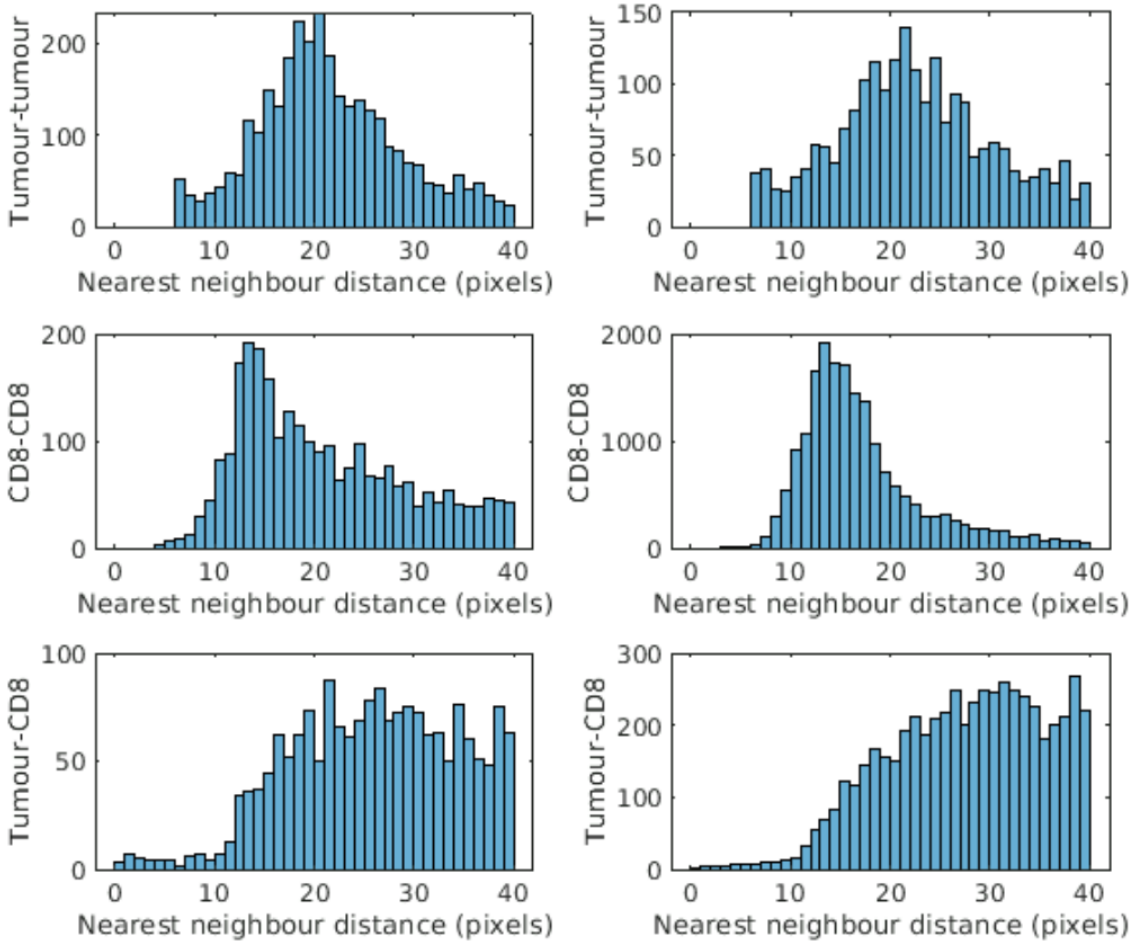


Figure F.1: Patient 18 - Histograms of nearest neighbour distances between cells. Columns correspond to pre-treatment and post-treatment images respectively, and rows correspond to distances between tumour cells (top row), between CD8+ T cells (middle row), and between tumour and CD8+ T cells (bottom row).

can be seen that compared to patient 18, most of the tumour and CD8 cells are located very close to each other, with a significant number within 6 pixels of each other. As this signature is in both samples for this patient, it would suggest that this co-localisation is a biological feature rather than an artefact.

F.3 Hypotheses for existence of cell-cell pairs separated by distance less than a cell diameter

We generate hypotheses as to how pairs of cells could be located at less than a cell diameter from each other, in order to find out if the data needs to be adjusted before analysis. We define *perfectly overlapping* cells as those which have the same coordinates, and we also consider pairs of cells located within x pixels of each other, i.e. they are not completely

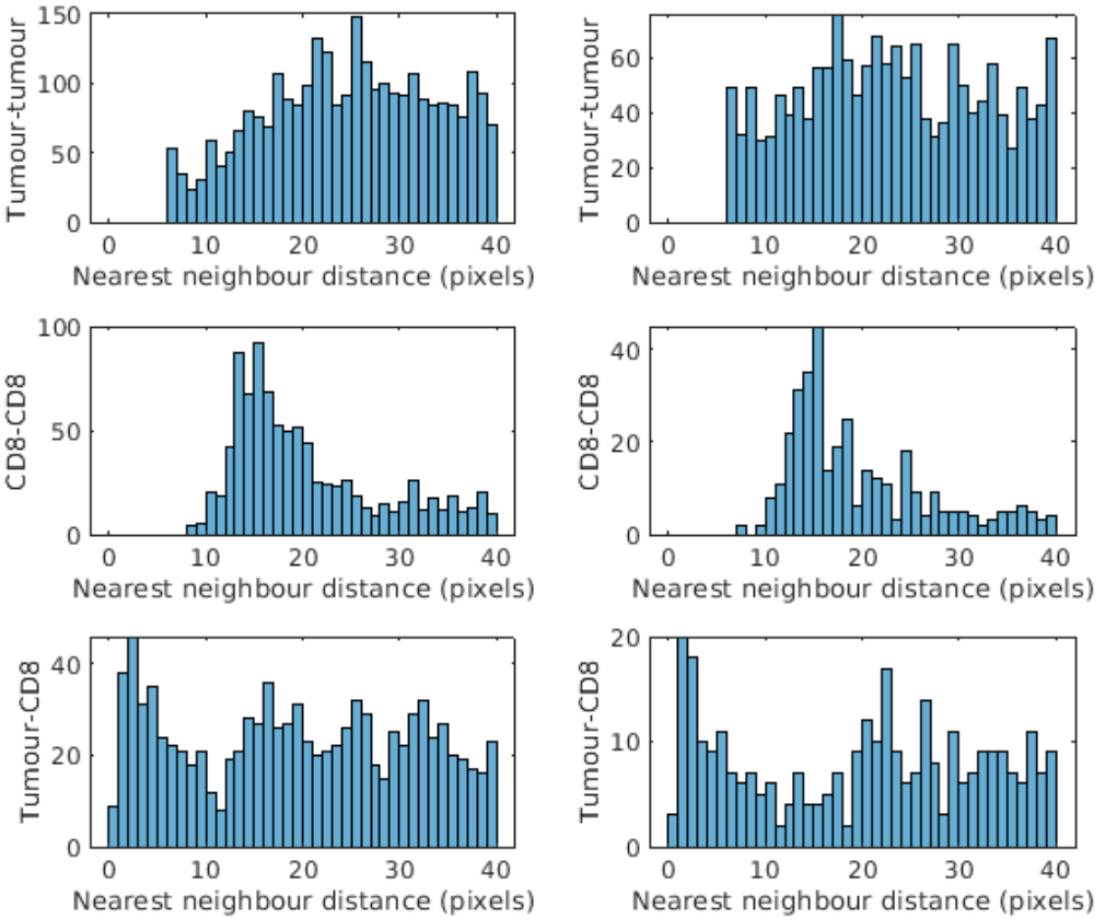


Figure F.2: Patient 59 - Histograms of nearest neighbour distances between cells. Columns correspond to pre-treatment and post-treatment images respectively, and rows correspond to distances between tumour cells (top row), between CD8+ T cells (middle row), and between tumour and CD8+ T cells (bottom row).

overlapping, but the two cells are closer than a cell radius apart.

- **H1:** Cell segmentation algorithms are not 100% accurate, so these overlaps could be due to one cell being "detected" as two separate cells - this could happen if the algorithm is based on threshold of marker presence. **Test:** *More likely that a proliferating T cell may be confused as a proliferating tumour cell - tumour cells do not express CD8 but do express Ki67. If most of the cell-cell pairs are one Ki67+CD8+ proliferating T cell and one Ki67+CD8- proliferating tumour cell, then this would support this theory.*

- **H2:** The cells have formed a complex, and one is being killed by the other. **Test:** *Pairs of cells would be composed of one tumour cell and one T cell. Would perhaps*

see more of these "complexes" in on-treatment samples, although this relies on the assumption that the treatment is working.

- **H3:** There may be more than a single layer of cells in areas of the tissue, possibly due to distortion/folds. **Test:** *If very close cell pairs are found near the tissue edges or exclude regions (where folds are), or separate cell pairs are situated close together, this may suggest the presence of folds in the tissue. Also, if the tissue is formed of many fragments or has a very irregular perimeter, i.e. has a large perimeter to area ratio, then you would expect to see more pairs of closely located cells.*

- **H4:** Cells are not solid spheres and may change morphology. **Test:** *There is no clear way to determine this from the images, although you would not expect to see completely overlapping cells. If the coordinates are the cell nuclei or the cell centroids then this makes a difference to what extent you can justify this assumption.*

- **H5:** Our assumption that tumour cell diameter is approximately 6 pixels may be incorrect - perhaps the closest tumour cells are ones which are actively dividing. **Test:** *There is not a clear way to determine this, just a point to keep in consideration.*

F.4 Testing H1: Mislabelling of segmentation algorithm

Hypothesis 1 (H1) poses that one cell has been mislabelled as two separate cells, causing "perfect" overlaps, where two cells have the same coordinates. Most of the images (94/131, 72%) did not contain any pairs of perfectly overlapping cells. A total of 58 pairs of perfectly overlapping tumour and CD8+ cells were found across the image dataset. Of these, 32 (55%) pairs contained one tumour cell and one proliferating CD8+ cell, and 26 (45%) pairs contained one tumour cell and one non-proliferating CD8+ cell. As the frequency of these overlaps is so small compared to the total number of cells in these samples, the difference in these percentages is not statistically significant, and we can conclude that it is approximately equally likely for a proliferating or non-proliferating T cell to perfectly overlap with a tumour cell in the case of a perfectly overlapping pair of cells.

F.4.1 Relationship between number of overlaps and cell density

If the overlapping cell locations is a result of a systematic labelling error, we might expect the frequency of these cell pairs to increase linearly with the total number of cells, i.e. the proportion of mislabelled cells would be constant across images.

Figures F.3 and F.4 show the number of perfectly overlapping pairs of cells (distance between cells is zero) and pairs of cells within 3 pixels distance vs. the total number of

CD8+ T cells in the sample, respectively. Each data point is one tissue piece, as some images contain multiple pieces of tissue with a large degree of heterogeneity between the pieces. We plot the total number of CD8+ T cells, as in this hypothesis, the T cells are being mislabelled as two cells - a T cell and a tumour cell.

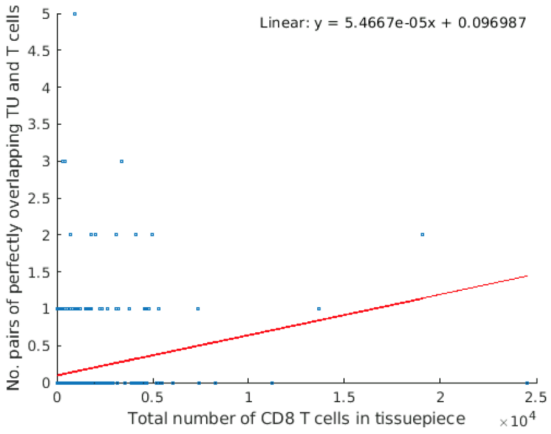


Figure F.3: Frequency of perfectly overlapping cell pairs, against total number of CD8 T cells in the sample. Each data point represents a single piece of tissue (some images contain many tissue pieces).

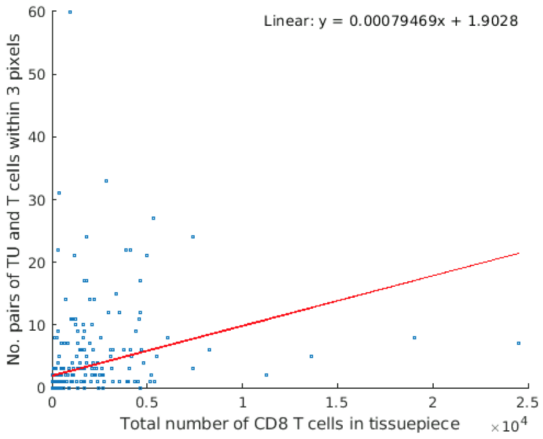


Figure F.4: Frequency of cell pairs located within 3 pixels of each other, against total number of CD8 T cells in the sample. Each data point represents a single piece of tissue (some images contain many tissue pieces).

In both cases, a linear fit to the data shows a positive relationship between the frequency of overlapping cell pairs and the total number of T cells, however the data have a large variance. A positive relationship between the frequency of these co-localisation (overlap) events and the total number of T cells could support the hypothesis of a systematic error, which causes an increase in the frequency of these co-localisation pairs as the number of T cells being identified and labelled increases. However, conversely it could be argued that in

samples with increased numbers of T cells, the immune system may be more efficacious in removing the tumour cells, and thus there may be more localisation of T cells and tumour cells, simply through killing events and increased T cell infiltration. Therefore these results cannot categorically confirm hypothesis 1.

F.5 Testing H2: Cells in co-localised pairs have formed a complex

We do not observe any perfectly overlapping same-type cell pairs, suggesting that either the overlapping tumour and CD8 cells have formed a complex, or that a proliferating CD8 cell has been mislabelled as a proliferating TU cell and a CD8 cell, as there is not a marker used solely to distinguish tumour cells in this dataset. There is not a significant change in the fraction of overlapping cell pairs in pre- vs on-treatment samples, however if the treatment was not working then we may not see a universal killing response. 6 out of the 64 patients had overlapping cells in their pre- and on-treatment samples, which could suggest it was a common biological feature, however as the number of overlaps in each sample is so low, i.e. some samples contain just one overlap, it is not something we can have confidence in concluding.

F.6 Testing H3: Spatial distribution of cell-cell overlaps

F.6.1 Plotting location of overlaps

Figure F.5 shows the pre- and on-treatment samples for three example patients. We define a metric for tissue morphology as the perimeter to area ratio of a tissue piece. The higher the value of this ratio, the more irregular the polygon defining the tissue shape is; long thin pieces of tissue have high P/A values. The idea is that the more irregular the tissue sample is, the more likely there may be folds and distortions of the tissue.

Many of the perfectly overlapping (plotted by bullseye symbol) and pairs of cells less than 3 pixels apart (plotted in red) are located near tissue edges. Also of note is that closely located pairs of cells are found in exclude regions - these are shown as red points outside of the tissue area. This may support the hypothesis that these co-localised cells are due to folds in the tissue, as these exclude regions denote where the folds are.

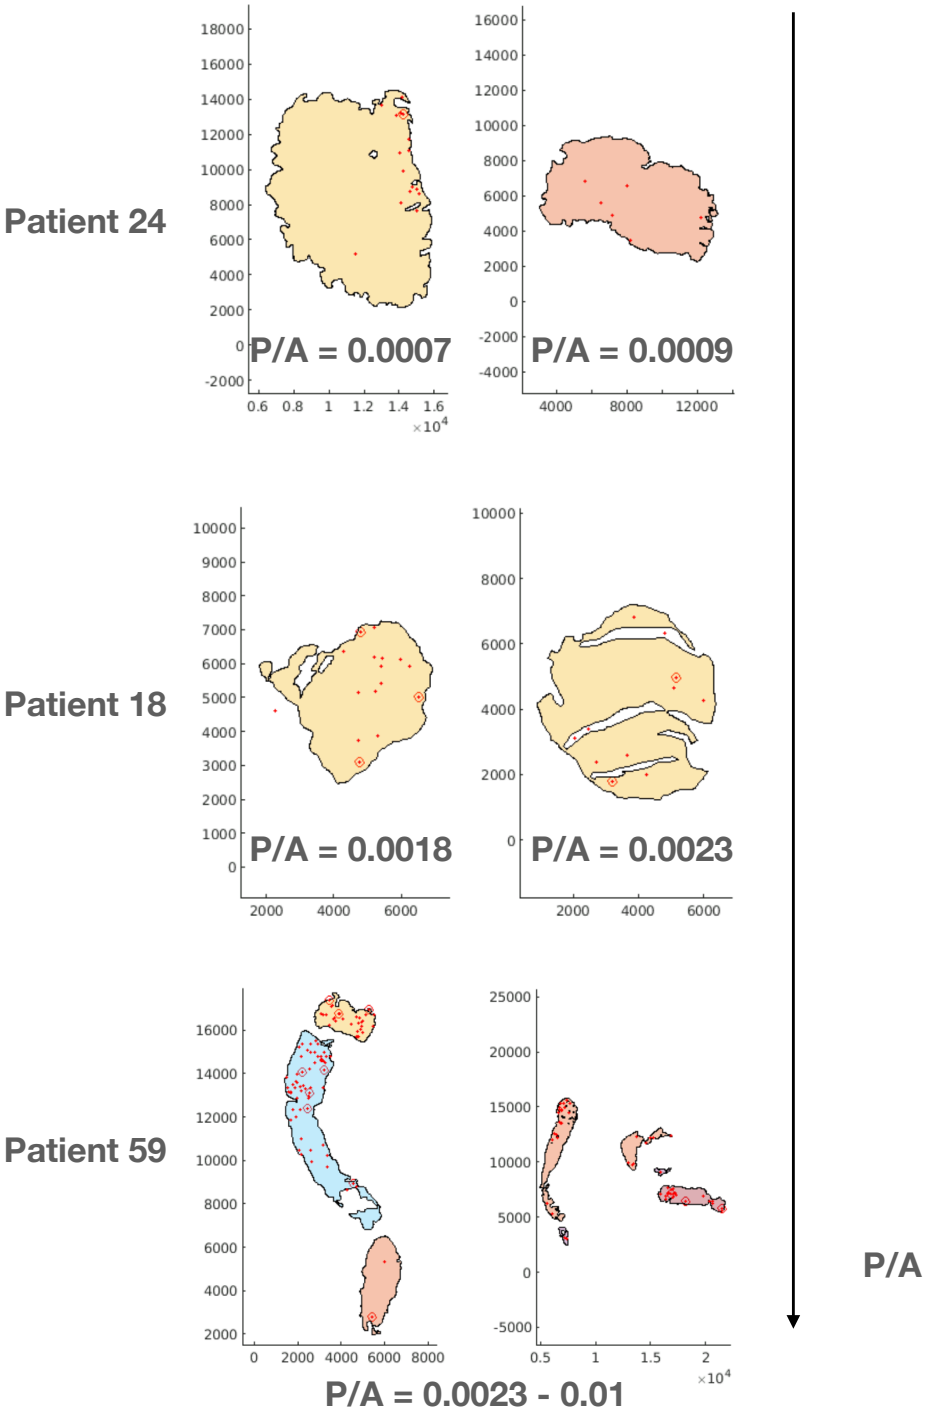


Figure F.5: Perfectly overlapping cell pairs (bullseye points) and cell pairs located within 3 pixels of each other (red points) for 6 images (paired pre-treatment (column 1) and on-treatment (column 2) for 3 example patients). P/A = tissue perimeter/tissue area.

F.6.2 Relationship between cell-cell localisation and tissue morphology

H3 assumes that folds and artefacts at the edges of the tissue are causing the overlapping cell coordinates. To test this, we plot the fraction of overlapping cell pairs over the total number of cells against the P/A ratio for each tissue piece. Here we look at the total number of cells (tumour and T cells), instead of just T cells, as we are considering that the tumour and T cells are both present in the tissue, they are just located on top of each other due to tissue folding. Tissue pieces with fewer than 500 cells were not considered in the analysis in this section due to the infrequency of these overlap artefacts. Tissue pieces which contained no tumour cells or T cells were also removed from the analysis.

Figures F.6 and F.7 plot the number of overlapping pairs (perfectly and within 3 pixels respectively) divided by the total number of cells against P/A for each tissue piece.

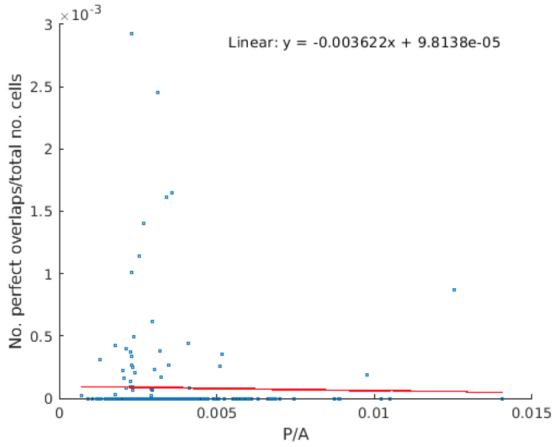


Figure F.6: Fraction of number of perfect overlaps to total number of cells plotted against tissue piece morphology measure, P/A. Each data point represents one tissue piece. P/A = tissue perimeter/tissue area.

Although there appears to be a negative correlation between the fraction of cells which are in these co-localised pairs and the P/A value, there is again a lot of variation in the data and it is not a strong correlation. There is also the confounding factor that the number of cells in a piece of tissue generally decreases as the P/A value increases - this is because highly irregular tissue samples tend to be small and so do not contain many cells - see Figure F.8.

F.6.3 Location of overlaps relative to distance to tissue edge

A circle is drawn around each of the overlap points and expanded. The circle radius is recorded when it intersects with the tissue edge which gives an estimate for the distance of the point from the edge. Figure F.9 shows the distribution of perfectly overlapping pairs

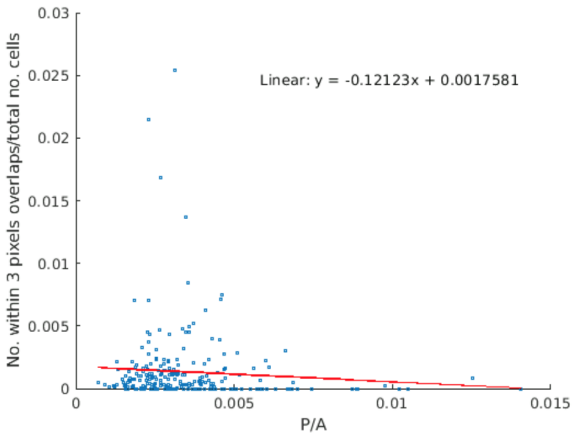


Figure F.7: Fraction of number of pairs of cells that are separated from each other by less than 3 pixels to total number of cells plotted against tissue piece morphology measure, P/A. Each data point represents one tissue piece. P/A = tissue perimeter/tissue area.

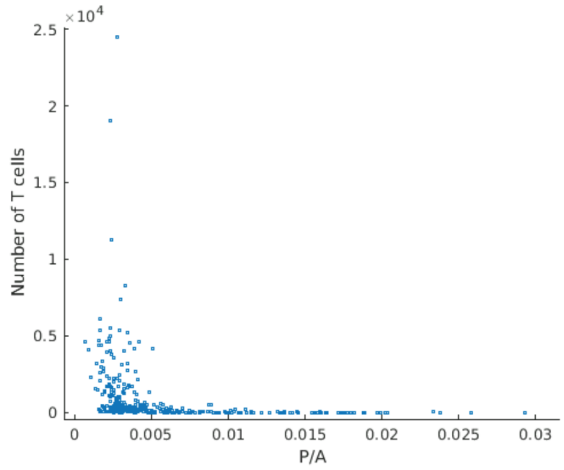


Figure F.8: Number of T cells against tissue piece morphology measure, P/A. Each data point represents one tissue piece. P/A = tissue perimeter/tissue area.

of cells with respect to the distance from the tissue edge. All of the overlapping pairs were located less than 1000 pixels from the tissue edge, with most located within 400 pixels of the edge. This corresponds to approximately 67 tumour cell diameters. There is a peak within 50 pixels (approx. 7 tumour cell diameters), so it is possible that these overlaps are being caused by edge effects.

Next, pairs of cells located within 3 pixels of each other were measured with respect to their location from the tissue edge to investigate whether these cells located very close together may be due to edge effects and tissue distortions. Figure F.10 shows the distribution of these pairs of cells as a function of distance from tissue edge. This plot shows that there is a higher frequency of these pairs located near the tissue edge, however, caution needs to be taken when interpreting these results, as the samples are all different sizes, for example

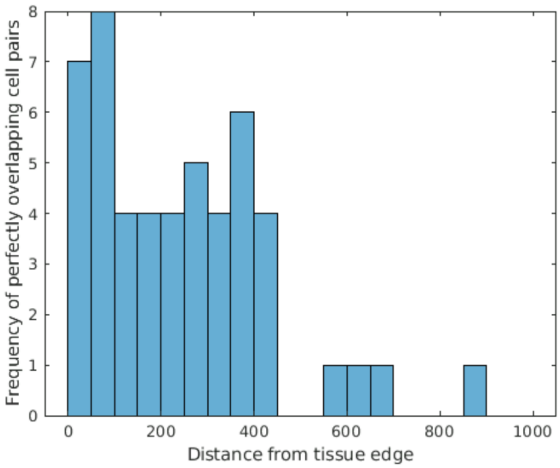


Figure F.9: Frequency of perfectly overlapping cell pairs with increasing distance from the tissue edge.

some samples may have long thin morphology and have a width of 500 pixels; not all of the tissue pieces are greater than 2000 pixels in width, so we would expect to see this skew in the distribution at short distances from the edge. However, this does not negate the fact that these pairs are still very close to the edges and from visual inspection of some of the images (e.g. patient 24 in Figure F.5) it is evident that edge effects may be playing a role in the apparent co-localisation of cells.

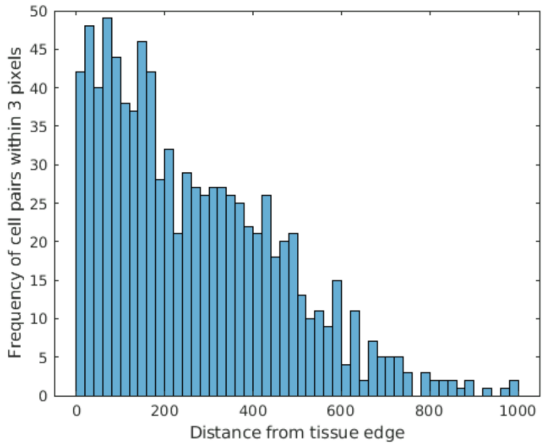


Figure F.10: Frequency of overlapping cell pairs (within 3 pixels of each other) with increasing distance from the tissue edge.

Of note is that we do not observe overlaps between cells of the same type - if the co-localisation was due to tissue folding we might expect to see overlaps regardless of cell type, however perhaps the overlap of tumour cells and T cells is more easily picked up as they express different markers.

F.7 Comments on H4 and H5

We do not see same-cell types co-localised so closely, which suggests that the cells do not get that close and it is a systematic labelling classification error, or this could equally support the hypothesis that cells are only this close because they are a tumour cell - T cell complex.

F.8 Nearest neighbour distance threshold

Figures F.11 and F.12 show the cross-PCFs of tumour and CD8+ T cells for the pre-treatment and post-treatment images respectively, of one patient. This patient had the highest frequency of overlapping tumour and CD8+ T cells (9 pairs in the pre-treatment image and 3 pairs in the post-treatment image). There was also significant localisation of tumour cells and CD8+ T cells at distances less than 6 pixels (see corresponding nearest neighbour histogram, Figure F.2).

The cross-PCF of tumour and CD8+ T cells was plotted for three different threshold values, corresponding to the radius at which you begin counting cells within the annuli. For example, a threshold value of zero corresponds to not altering the data; a threshold value of 6 means any cell-cell distances below 6 pixels are not considered in the PCF calculation.

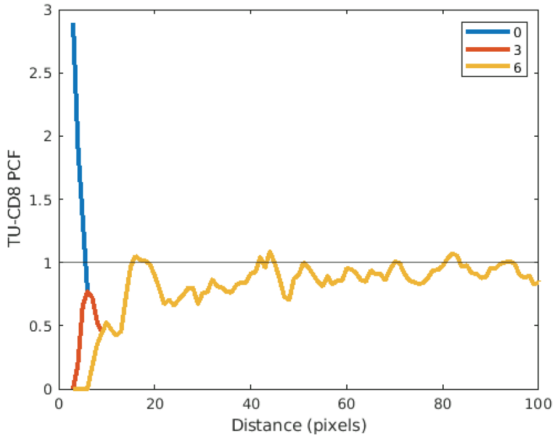


Figure F.11: Cross-PCF of tumour cells and CD8+ T cells for the pre-treatment image for patient 59. PCFs are calculated for different nearest neighbour distance thresholds.

Figures F.11 and F.12 show that, as expected, the threshold value greatly affects the PCF at small distances. This is important as whatever decision we make regarding the threshold we adopt produces a markedly different PCF signature at small distances, and may introduce artefacts, such as shifting the value of the distance at which the PCF is a maximum.

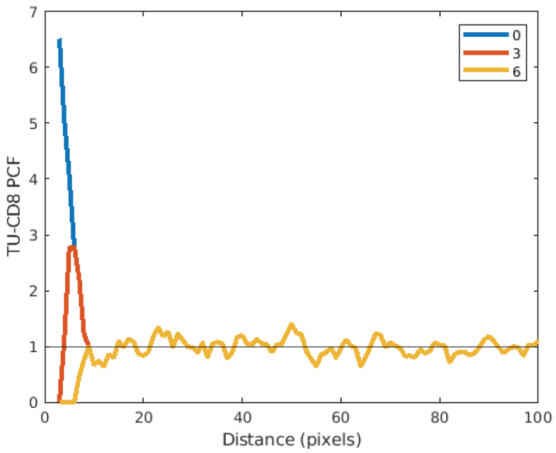


Figure F.12: Cross-PCF of tumour cells and CD8+ T cells for the post-treatment image for patient 59. PCFs are calculated for different nearest neighbour distance thresholds.

Appendix G

Supplementary results for fine-grain parameter sweep over CD8+ T cell biased motion and proliferation probabilities

In Section 4.5, we identified a region of parameter space in which spatial features enable increased tumour outcome prediction accuracy. This region of the parameter space encompassed low values of the CD8+ T cell proliferation probability, r_E, for a range of biased motion values, β. In this Appendix, we present the timecourses of the tumour and CD8+ T cell populations over this fine-grain parameter sweep, and show examples of snapshots at tumour elimination or $t = 145$ days. Parameter values considered are $\beta = \{0.125, 0.25, 0.375, 0.5, 0.625, 0.75, 0.825, 1\}$ and $r_E = \{0.0002, 0.000225, 0.00025\}$.

Figures G.1 and G.2 show the tumour and CD8+ T cell population numbers over time, for each parameter set across the parameter sweep. Figure G.3 shows the corresponding snapshots from one realisation of the ABM, taken either just before tumour elimination, or at day 145.

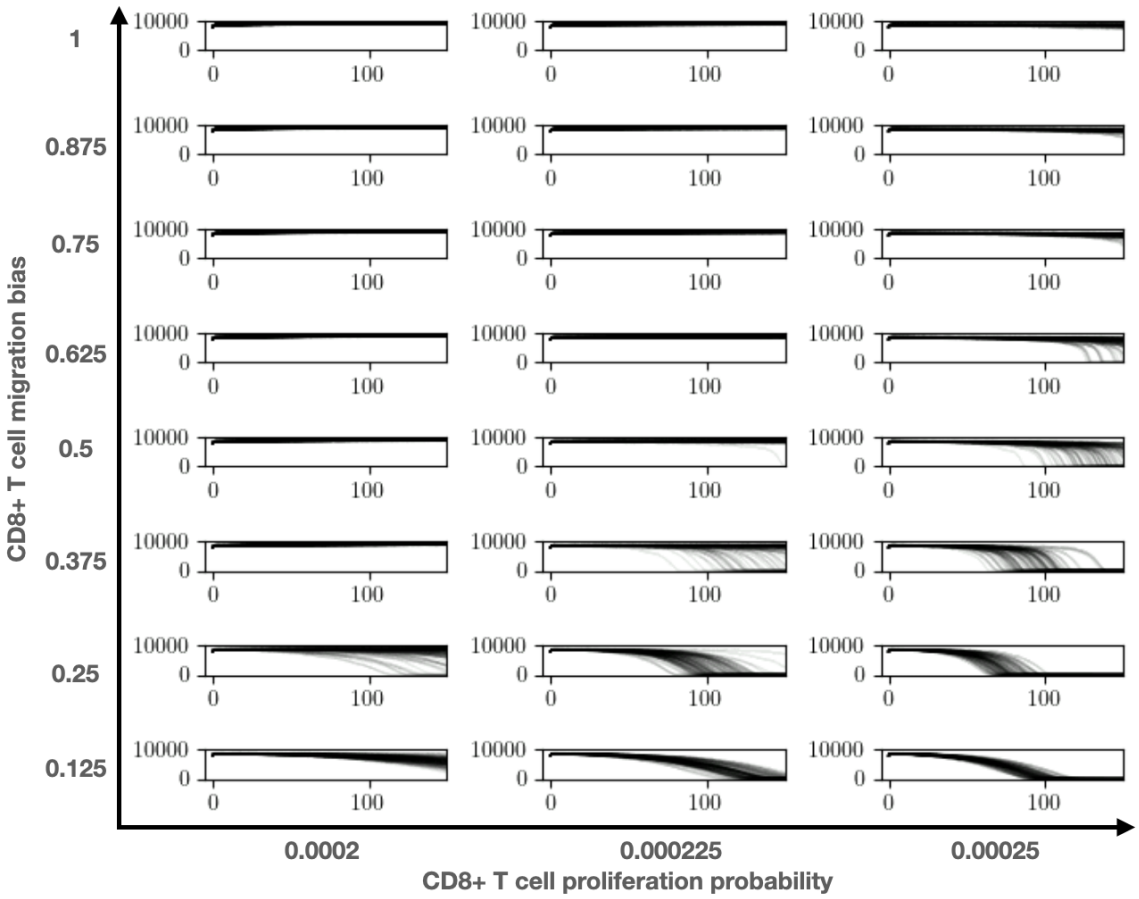


Figure G.1: Timecourses of the tumour cell population for 24 parameter sets, varying CD8+ T cell proliferation probability and bias in T cell movement. Each line corresponds to the tumour cell population timecourse for a single model realisation.

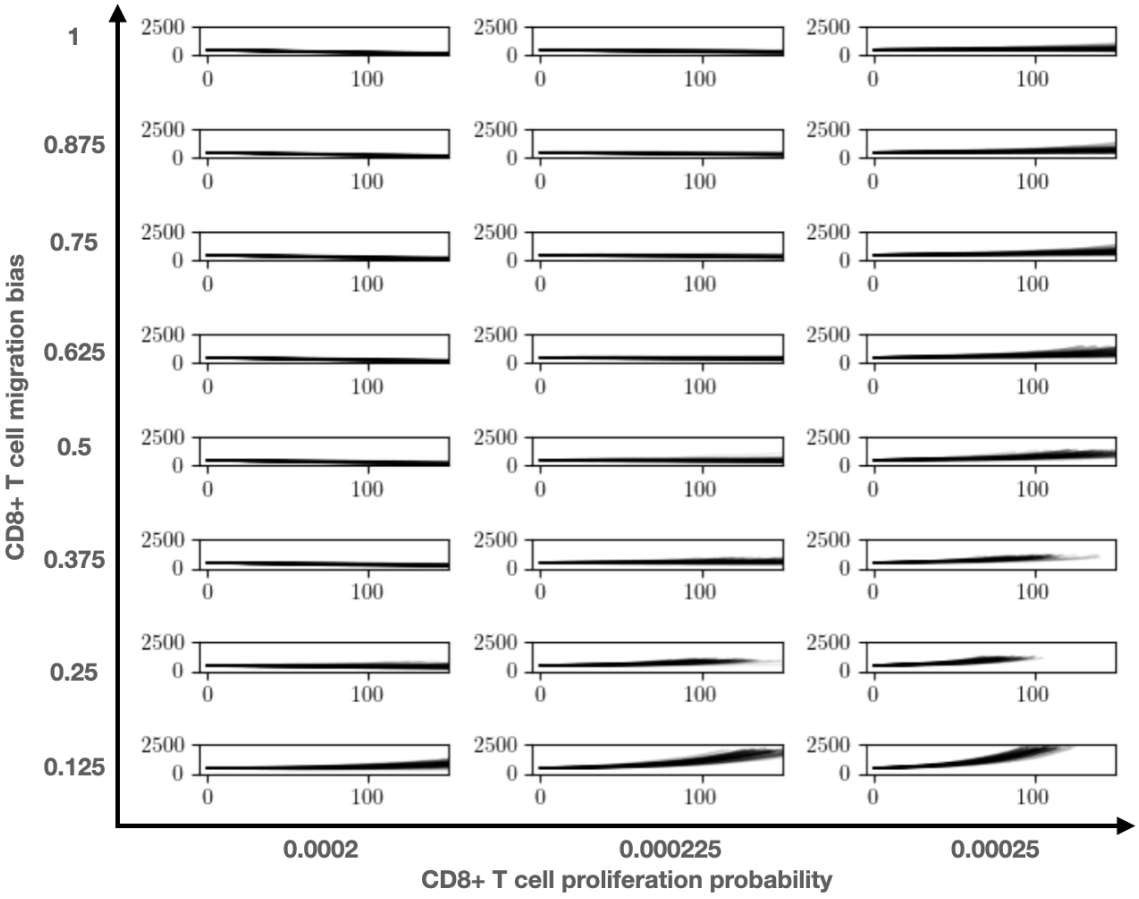


Figure G.2: Timecourses of CD8+ T cell population for 24 parameter sets, varying T cell proliferation probability and bias in T cell movement. Each line corresponds to the CD8+ T cell population timecourse for a single model realisation.

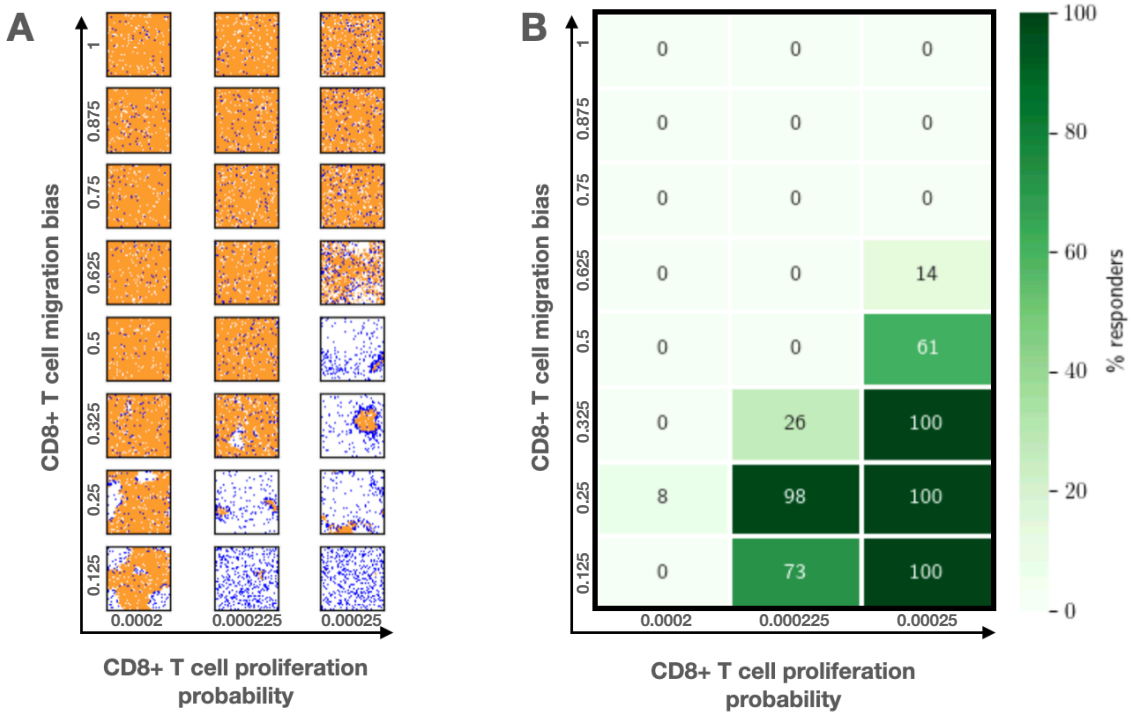


Figure G.3: (A) Snapshots over parameter sweep for a single ABM realisation, taken either just before tumour elimination, or at day 145. (B) Percentage of responders (tumour burden reduction by over 50%, i.e. less than 4000 tumour cells remaining) over the 100 ABM realisations for each parameter set.

References

Agnew, D., Green, J., Brown, T., Simpson, M. and Binder, B. (2014), 'Distinguishing between mechanisms of cell aggregation using pair-correlation functions', *Journal of Theoretical Biology* **352**, 16–23.

Alarcón, T., Byrne, H. M. and Maini, P. K. (2003), 'A cellular automaton model for tumour growth in inhomogeneous environment', *Journal of Theoretical Biology* **225**(2), 257–274.

Bacac, M., Fauti, T., Sam, J., Colombetti, S., Weinzierl, T., Ouaret, D., Bodmer, W., Lehmann, S., Hofer, T., Hosse, R. et al. (2016), 'A novel carcinoembryonic antigen T-cell bispecific antibody (CEA TCB) for the treatment of solid tumors', *Clinical Cancer Research* **22**(13), 3286–3297.

Binder, B. J. and Simpson, M. J. (2015), 'Spectral analysis of pair-correlation bandwidth: application to cell biology images', *Royal Society Open Science* **2**(2), 140494.

Binder, B. and Simpson, M. (2013), 'Quantifying spatial structure in experimental observations and agent-based simulations using pair-correlation functions', *Physical Review E* **88**(2), 022705.

Boissonnas, A., Fetler, L., Zeelenberg, I. S., Hugues, S. and Amigorena, S. (2007), 'In vivo imaging of cytotoxic T cell infiltration and elimination of a solid tumor', *The Journal of Experimental Medicine* **204**(2), 345–356.

Boldajipour, B., Nelson, A. and Krummel, M. (2016), 'Tumor-infiltrating lymphocytes are dynamically desensitized to antigen but are maintained by homeostatic cytokine', *JCI Insight* **1**(20).

Bravo, R. R., Baratchart, E., West, J., Schenck, R. O., Miller, A. K., Gallaher, J., Gatenbee, C. D., Basanta, D., Robertson-Tessi, M. and Anderson, A. R. (2020), 'Hybrid Automata Library: A flexible platform for hybrid modeling with real-time visualization', *PLoS Computational Biology* **16**(3), e1007635.

Breart, B., Lemaître, F., Celli, S., Bousso, P. et al. (2008), 'Two-photon imaging of intratumoral cd8+ t cell cytotoxic activity during adoptive t cell therapy in mice', *The Journal of clinical investigation* **118**(4), 1390–1397.

Browning, A. P., Lewin, T. D., Baker, R. E., Maini, P. K., Moros, E. G., Caudell, J., Byrne, H. M. and Enderling, H. (2022), 'Predicting radiotherapy patient outcomes with real-time clinical data using mathematical modelling', *arXiv e-prints* pp. arXiv–2201.

Bughda, R., Dimou, P., D'Souza, R. R. and Klampatsa, A. (2021), 'Fibroblast activation protein (FAP)-targeted CAR-T cells: Launching an attack on tumor stroma', *Immuno-Targets and Therapy* **10**, 313.

Bull, J. (2018), Characterisation of macrophage infiltration into solid tumours via image analysis and computational modelling, DPhil thesis, University of Oxford.

Bull, J. A. and Byrne, H. M. (2023), 'Quantification of spatial and phenotypic heterogeneity in an agent-based model of tumour-macrophage interactions', *PLOS Computational Biology* **19**(3), e1010994.

Bull, J. A., Macklin, P. S., Quaiser, T., Braun, F., Waters, S. L., Pugh, C. W. and Byrne, H. M. (2020), 'Combining multiple spatial statistics enhances the description of immune cell localisation within tumours', *Scientific Reports* **10**(1), 18624.

Byrne, H., Cave, G. and McElwain, D. (1998), 'The effect of chemotaxis and chemokinesis on leukocyte locomotion: A new interpretation of experimental results', *Mathematical Medicine and Biology: A Journal of the IMA* **15**(3), 235–256.

Byrne, H. M. and Owen, M. R. (2004), 'A new interpretation of the Keller-Segel model based on multiphase modelling', *Journal of Mathematical Biology* **49**, 604–626.

Carstens, J. L., Correa de Sampaio, P., Yang, D., Barua, S., Wang, H., Rao, A., Allison, J. P., LeBleu, V. S. and Kalluri, R. (2017), 'Spatial computation of intratumoral T cells correlates with survival of patients with pancreatic cancer', *Nature Communications* **8**(1), 1–13.

Chen, D. and Mellman, I. (2013), 'Oncology meets immunology: the cancer-immunity cycle', *Immunity* **39**(1), 1–10.

Chiu, S. N., Stoyan, D., Kendall, W. S. and Mecke, J. (2013), *Stochastic geometry and its applications*, John Wiley & Sons.

Christophe, C., Müller, S., Rodrigues, M., Petit, A., Cattiaux, P., Dupré, L., Gadat, S. and Valitutti, S. (2015), 'A biased competition theory of cytotoxic T lymphocyte interaction with tumor nodules', *PLOS One* **10**(3), e0120053.

Corthay, A. (2009), 'How do regulatory T cells work?', *Scandinavian Journal of Immunology* **70**(4), 326–336.

Cosgrove, J., Butler, J., Alden, K., Read, M., Kumar, V., Cucurull-Sanchez, L., Timmis, J. and Coles, M. (2015), 'Agent-based modeling in systems pharmacology', *CPT: Pharmacometrics & Systems Pharmacology* **4**(11), 615–629.

de Pillis, L., Radunskaya, A. and Wiseman, C. (2005), 'A validated mathematical model of cell-mediated immune response to tumor growth', *Cancer Research* **65**(17), 7950–7958.

Dini, S., Binder, B., Fischer, S., Mattheyer, C., Schmitz, A., Stelzer, E., Bean, N. and Green, J. (2016), 'Identifying the necrotic zone boundary in tumour spheroids with pair-correlation functions', *Journal of The Royal Society Interface* **13**(123), 20160649.

Dini, S., Binder, B. and Green, J. (2018), 'Understanding interactions between populations: Individual based modelling and quantification using pair correlation functions', *Journal of Theoretical Biology* **439**, 50–64.

Dunn, G. P., Old, L. J. and Schreiber, R. D. (2004), 'The three Es of cancer immunoediting', *Annu. Rev. Immunol.* **22**, 329–360.

Eisenhauer, E. A., Therasse, P., Bogaerts, J., Schwartz, L. H., Sargent, D., Ford, R., Dancey, J., Arbuck, S., Gwyther, S., Mooney, M. et al. (2009), 'New response evaluation criteria in solid tumours: revised recist guideline (version 1.1)', *European journal of cancer* **45**(2), 228–247.

Ferreira Jr, S., Martins, M. and Vilela, M. (2002), 'Reaction-diffusion model for the growth of avascular tumor', *Physical Review E* **65**(2), 021907.

Fozard, J., Kirkham, G., Buttery, L., King, J., Jensen, O. and Byrne, H. (2011), 'Techniques for analysing pattern formation in populations of stem cells and their progeny', *BMC Bioinformatics* **12**(1), 396.

Gavagnin, E., Owen, J. and Yates, C. (2018), 'Pair correlation functions for identifying spatial correlation in discrete domains', *Physical Review E* **97**(6), 062104.

Ghaffarizadeh, A., Heiland, R., Friedman, S., Mumenthaler, S. and Macklin, P. (2018), 'PhysiCell: An open source physics-based cell simulator for 3-D multicellular systems', *PLoS Computational Biology* **14**(2), e1005991.

Gong, C., Anders, R. A., Zhu, Q., Taube, J. M., Green, B., Cheng, W., Bartelink, I. H., Vicini, P., Wang, B. and Popel, A. S. (2019), 'Quantitative characterization of CD8+ T cell clustering and spatial heterogeneity in solid tumors', *Frontiers in Oncology* **8**, 649.

Gong, C., Milberg, O., Wang, B., Vicini, P., Narwal, R., Roskos, L. and Popel, A. (2017), 'A computational multiscale agent-based model for simulating spatio-temporal tumour immune response to PD1 and PDL1 inhibition', *Journal of The Royal Society Interface* **14**(134), 20170320.

Hamis, S., Stratiev, S. and Powathil, G. G. (2021), 'Uncertainty and sensitivity analyses methods for agent-based mathematical models: An introductory review', *The Physics of Cancer: Research Advances* pp. 1–37.

Harris, T. H., Banigan, E. J., Christian, D. A., Konradt, C., Tait Wojno, E. D., Norose, K., Wilson, E. H., John, B., Weninger, W., Luster, A. D. et al. (2012), 'Generalized Lévy walks and the role of chemokines in migration of effector CD8+ T cells', *Nature* **486**(7404), 545–548.

Hutchinson, L. and Grimm, O. (2022), 'Integrating digital pathology and mathematical modelling to predict spatial biomarker dynamics in cancer immunotherapy', *npj Digital Medicine* **5**(1), 1–13.

Iarosz, K., Martins, C., Batista, A., Viana, R., Lopes, S., Caldas, I. and Penna, T. (2011), On a cellular automaton with time delay for modelling cancer tumors, *in* 'Journal of Physics: Conference Series', Vol. 285, IOP Publishing, p. 012015.

Johnston, S., Simpson, M., McElwain, D. S., Binder, B. and Ross, J. (2014), 'Interpreting scratch assays using pair density dynamics and approximate Bayesian computation', *Open Biology* **4**(9), 140097.

Jolliffe, I. T. and Cadima, J. (2016), 'Principal component analysis: a review and recent developments', *Philosophical transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **374**(2065), 20150202.

Kather, J., Charoentong, P., Suarez-Carmona, M., Herpel, E., Klupp, F., Ulrich, A., Schneider, M., Zoernig, I., Luedde, T., Jaeger, D. et al. (2018), 'High-throughput screening of

combinatorial immunotherapies with patient-specific in silico models of metastatic colorectal cancer', *Cancer Research* **78**(17), 5155–5163.

Kather, J., Poleszczuk, J., Suarez-Carmona, M., Krisam, J., Charoentong, P., Valous, N., Weis, C., Tavernar, L., Leiss, F., Herpel, E. et al. (2017), 'In silico modeling of immunotherapy and stroma-targeting therapies in human colorectal cancer', *Cancer Research* **77**(22), 6442–6452.

Kirschner, D. and Panetta, J. (1998), 'Modeling immunotherapy of the tumor-immune interaction', *Journal of Mathematical Biology* **37**(3), 235–252.

Kuznetsov, V., Makalkin, I., Taylor, M. and Perelson, A. (1994), 'Nonlinear dynamics of immunogenic tumors: parameter estimation and global bifurcation analysis', *Bulletin of Mathematical Biology* **56**(2), 295–321.

Lai, X., Geier, O., Fleischer, T., Garred, Ø., Borgen, E., Funke, S., Kumar, S., Rognes, M., Seierstad, T., Børressen-Dale, A.-L. et al. (2019), 'Towards personalized computer simulation of breast cancer treatment: a multi-scale pharmacokinetic and pharmacodynamic model informed by multi-type patient data', *Cancer Research* pp. 4293–4304.

Lambert, B., MacLean, A., Fletcher, A., Combes, A., Little, M. and Byrne, H. (2018), 'Bayesian inference of agent-based models: a tool for studying kidney branching morphogenesis', *Journal of Mathematical Biology* **76**(7), 1673–1697.

Li, L., Jiang, G., Chen, Q. and Zheng, J. (2015), 'Ki67 is a promising molecular target in the diagnosis of cancer', *Molecular Medicine Reports* **11**(3), 1566–1572.

Li, X., Gruosso, T., Zuo, D., Omeroglu, A., Meterissian, S., Guiot, M.-C., Salazar, A., Park, M. and Levine, H. (2019), 'Infiltration of CD8+ T cells into tumor cell clusters in triple-negative breast cancer', *Proceedings of the National Academy of Sciences* **116**(9), 3678–3687.

Liu, F., Qi, L., Liu, B., Liu, J., Zhang, H., Che, D., Cao, J., Shen, J., Geng, J., Bi, Y. et al. (2015), 'Fibroblast activation protein overexpression and clinical implications in solid tumors: a meta-analysis', *PloS one* **10**(3), e0116683.

Liu, T., Han, C., Wang, S., Fang, P., Ma, Z., Xu, L. and Yin, R. (2019), 'Cancer-associated fibroblasts: an emerging target of anti-cancer immunotherapy', *Journal of Hematology & Oncology* **12**(1), 1–15.

López, Á. G., Seoane, J. M. and Sanjuán, M. A. (2017*a*), 'Destruction of solid tumors by immune cells', *Communications in Nonlinear Science and Numerical Simulation* **44**, 390–403.

López, Á. G., Seoane, J. M. and Sanjuán, M. A. (2017*b*), 'Dynamics of the cell-mediated immune response to tumour growth', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **375**(2096), 20160291.

Macfarlane, F., Lorenzi, T. and Chaplain, M. (2018), 'Modelling the immune response to cancer: an individual-based approach accounting for the difference in movement between inactive and activated T cells', *Bulletin of Mathematical Biology* **80**(6), 1539–1562.

Macfarlane, F. R., Chaplain, M. A. and Lorenzi, T. (2019), 'A stochastic individual-based model to explore the role of spatial interactions and antigen recognition in the immune response against solid tumours', *Journal of Theoretical Biology* **480**, 43–55.

Mallet, D. and De Pillis, L. (2006), 'A cellular automata model of tumor-immune system interactions', *Journal of Theoretical Biology* **239**(3), 334–350.

Marino, S., Hogue, I. B., Ray, C. J. and Kirschner, D. E. (2008), 'A methodology for performing global uncertainty and sensitivity analysis in systems biology', *Journal of Theoretical Biology* **254**(1), 178–196.

Mattfeldt, T., Eckel, S., Fleischer, F. and Schmidt, V. (2009), 'Statistical analysis of labelling patterns of mammary carcinoma cell nuclei on histological sections', *Journal of Microscopy* **235**(1), 106–118.

Matzavinos, A., Chaplain, M. and Kuznetsov, V. (2004), 'Mathematical modelling of the spatio-temporal response of cytotoxic T-lymphocytes to a solid tumour', *Mathematical Medicine and Biology* **21**(1), 1–34.

Mirams, G., Arthurs, C., Bernabeu, M., Bordas, R., Cooper, J., Corrias, A., Davit, Y., Dunn, S.-J., Fletcher, A., Harvey, D. et al. (2013), 'Chaste: an open source C++ library for computational physiology and biology', *PLoS Computational Biology* **9**(3), e1002970.

Muñoz-Aguirre, M., Ntasis, V. F., Rojas, S. and Guigó, R. (2020), 'PyHIST: a histological image segmentation tool', *PLoS Computational Biology* **16**(10), e1008349.

Nir, G., Hor, S., Karimi, D., Fazli, L., Skinnider, B. F., Tavassoli, P., Turbin, D., Villamil, C. F., Wang, G., Wilson, R. S. et al. (2018), 'Automatic grading of prostate cancer in digitized histopathology images: Learning from multiple experts', *Medical Image Analysis* **50**, 167–180.

Norton, K.-A., Gong, C., Jamalian, S. and Popel, A. (2019), 'Multiscale agent-based and hybrid modeling of the tumor immune microenvironment', *Processes* **7**(1), 37.

Owen, J. P., Kelsh, R. N. and Yates, C. A. (2020), 'A quantitative modelling approach to zebrafish pigment pattern formation', *Elife* **9**, e52998.

Pitt-Francis, J., Pathmanathan, P., Bernabeu, M., Bordas, R., Cooper, J., Fletcher, A., Mirams, G., Murray, P., Osborne, J., Walter, A. et al. (2009), 'Chaste: a test-driven approach to software development for biological modelling', *Computer Physics Communications* **180**(12), 2452–2471.

Rejniak, K. and Anderson, A. (2011), 'Hybrid models of tumor growth', *Wiley Interdisciplinary Reviews: Systems Biology and Medicine* **3**(1), 115–125.

Robertson-Tessi, M., El-Kareh, A. and Goriely, A. (2012), 'A mathematical model of tumor-immune interactions', *Journal of Theoretical Biology* **294**, 56–73.

Rousseeuw, P. J. (1987), 'Silhouettes: a graphical aid to the interpretation and validation of cluster analysis', *Journal of Computational and Applied Mathematics* **20**, 53–65.

Sauls, R. S., McCausland, C. and Taylor, B. N. (2018), 'Histology, T-Cell Lymphocyte'.

Savitzky, A. and Golay, M. J. (1964), 'Smoothing and differentiation of data by simplified least squares procedures.', *Analytical Chemistry* **36**(8), 1627–1639.

Scholzen, T. and Gerdes, J. (2000), 'The Ki-67 protein: from the known and the unknown', *Journal of Cellular Physiology* **182**(3), 311–322.

Siebert, N., Leopold, J., Zumpe, M., Troschke-Meurer, S., Biskupski, S., Zikoridse, A. and Lode, H. N. (2022), 'The immunocytokine FAP-IL-2v enhances anti-neuroblastoma efficacy of the anti-Gd2 antibody dinutuximab beta', *Cancers* **14**(19), 4842.

Sottoriva, A., Vermeulen, L. and Tavaré, S. (2011), 'Modeling evolutionary dynamics of epigenetic mutations in hierarchically organized tumors', *PLoS Computational Biology* **7**(5), e1001132.

Stéphanou, A., Ballet, P. and Powathil, G. (2020), 'Hybrid data-based modelling in oncology: successes, challenges and hopes', *Mathematical Modelling of Natural Phenomena* **15**, 21.

Strobl, M. A., Gallaher, J., West, J., Robertson-Tessi, M., Maini, P. K. and Anderson, A. R. (2022), 'Spatial structure impacts adaptive therapy by shaping intra-tumoral competition', *Communications Medicine* **2**(1), 46.

Sunnåker, M., Busetto, A., Numminen, E., Corander, J., Foll, M. and Dessimoz, C. (2013), 'Approximate Bayesian Computation', *PLoS Computational Biology* **9**(1), e1002803.

Swat, M., Thomas, G., Belmonte, J., Shirinifard, A., Hmeljak, D. and Glazier, J. (2012), Multi-scale modeling of tissues using CompuCell3D, *in* 'Methods in Cell Biology', Vol. 110, Elsevier, pp. 325–366.

Szabó, A. and Merks, R. (2013), 'Cellular Potts modeling of tumor growth, tumor invasion, and tumor evolution', *Frontiers in Oncology* **3**, 87.

Tiwari, A., Oravecz, T., Dillon, L. A., Italiano, A., Audoly, L., Fridman, W. H. and Clifton, G. T. (2023), 'Towards a consensus definition of immune exclusion in cancer', *Frontiers in Immunology* **14**, 1084887.

Tzedakis, G., Tzamali, E., Marias, K. and Sakkalis, V. (2015), 'The importance of neighborhood scheme selection in agent-based tumor growth modeling', *Cancer Informatics* **14**, CIN–S19343.

Vasconcelos, Z., Müller, S., Guipouy, D., Yu, W., Christophe, C., Gadat, S., Valitutti, S. and Dupré, L. (2015), 'Individual human cytotoxic T lymphocytes exhibit intraclonal heterogeneity during sustained killing', *Cell Reports* **11**(9), 1474–1485.

Vignali, D. A., Collison, L. W. and Workman, C. J. (2008), 'How regulatory T cells work', *Nature Reviews Immunology* **8**(7), 523–532.

Waldhauer, I., Gonzalez-Nicolini, V., Freimoser-Grundschober, A., Nayak, T. K., Fahrni, L., Hosse, R. J., Gerrits, D., Geven, E. J., Sam, J., Lang, S. et al. (2021), Simlukafusp alfa (FAP-IL2v) immunocytokine is a versatile combination partner for cancer immunotherapy, *in* 'MAbs', Vol. 13, Taylor & Francis, p. 1913791.

Weigelin, B., den Boer, A. T., Wagena, E., Broen, K., Dolstra, H., de Boer, R. J., Figdor, C. G., Textor, J. and Friedl, P. (2021), 'Cytotoxic T cells are able to efficiently eliminate cancer cells by additive cytotoxicity', *Nature Communications* **12**(1), 1–12.

Weiner, L., Dhodapkar, M. and Ferrone, S. (2009), 'Monoclonal antibodies for cancer immunotherapy', *The Lancet* **373**(9668), 1033–1040.

Wilkie, K. and Hahnfeldt, P. (2013), 'Mathematical models of immune-induced cancer dormancy and the emergence of immune evasion', *Interface Focus* **3**(4), 20130010.

Wilson, S. and Levy, D. (2012), 'A mathematical model of the enhancement of tumor vaccine efficacy by immunotherapy', *Bulletin of Mathematical Biology* **74**(7), 1485–1500.

Zhang, Z., Tao, J., Qiu, J., Cao, Z., Huang, H., Xiao, J. and Zhang, T. (2023), 'From basic research to clinical application: targeting fibroblast activation protein for cancer diagnosis and treatment', *Cellular Oncology* pp. 1–21.

Zhu, J. and Paul, W. E. (2010), 'Heterogeneity and plasticity of T helper cells', *Cell Research* **20**(1), 4–12.