Characterisation of Macrophage Infiltration into Solid Tumours via Image Analysis and Computational Modelling

Joshua Bull

St Hugh's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

October 2018

Abstract

Macrophages are innate immune cells that play a key role in conditions as diverse as wound healing, atherosclerosis and cancer. In cancer, macrophages may be characterised as anti-tumour, promoting a wider immune response and phagocytosing tumour cells, or as pro-tumour, suppressing the immune response and promoting blood vessel growth and metastasis. The spatial distribution of macrophages inside a tumour correlates with patient prognosis and treatment response, and is hypothesised to be predictive of macrophage phenotype. In this thesis, we use a combination of mathematical modelling and image analysis to characterise these distribution patterns and understand the mechanisms by which macrophages may infiltrate solid tumours.

We have developed an image analysis pipeline for the automated analysis of whole slide histology images, a rich source of spatial data for the investigation of immune cell distributions *in vivo*. Our pipeline automatically identifies macrophages and other cell types using a combination of superpixellation and Support Vector Machines (SVMs), and we develop summary statistics which characterise distinct macrophage distributions across different tumour indications based on detected cell locations. We validate our pipeline against real data and benchmark it against human pathologists. We then demonstrate its flexibility and robustness in a range of applications across a variety of datasets.

We investigate the observed distributions of macrophages further by using the Chaste (Cancer, Heart and Soft Tissue Environment) modelling framework to develop a new agent-based model of macrophage infiltration into solid tumours. We first use our model to simulate *in vitro* experiments in which inert microbeads infiltrate into multicellular tumour spheroids (MCTSs). This study enables us to investigate the role of passive infiltration caused by spatial variation in cell proliferation and death rates across the MCTS. The interplay between cell proliferation and death drives an advective flow from nutrient-rich, proliferative, regions towards necrotic regions. Our agent-based framework reproduces results from previous continuum models of this process and provides new insight into the observed infiltration patterns that cannot be obtained without resolving the trajectories of individual particles in an agent-based framework. We subsequently extend this model to describe macrophage infiltration into MCTSs via a combination of passive infiltration and chemotaxis in response to colony stimulating factor-1 (CSF-1). This chemoattractant, which is expressed by hypoxic tumour cells, is implicated in metastasis and is a target for immunotherapy treatments currently in Phase I trials at Roche. Finally, we discuss methods of integrating whole slide imaging with our agent-based modelling approaches and compare simulated macrophage distributions with those from *in vivo* images analysed using our imaging pipeline.

Acknowledgements

The first time I attempted to write these acknowledgements, they weren't exactly eloquent. Despite having several years to think about them, I ended up forgetting that I actually had to put them on paper until my thesis was almost finished printing. As a result, I'm sure that I missed out a lot of people who deserve my thanks. Unfortunately, I'm fairly certain that I haven't managed to solve this problem completely in this second attempt. So to those who don't get a special shout-out here: my apologies, and I can only plead that if I included everyone who should be listed here by name then this thesis would be significantly longer than it already is.

At the top of the huge list of people who deserve thanks are my supervisors, without whom this thesis would never have happened. Helen and Sarah have been amazing mentors, and I'm privileged to have benefitted from their support over the last few years. Thank you both. I have also been fortunate enough to have fantastic industrial supervisors in Tom and Franziska, and my sincere thanks go also to Vicente and Anthony, who planted the seeds of this thesis years ago. I would also like to thank Philip Macklin and Chris Pugh from the Nuffield Department of Medicine, who have embraced this work and led it in exciting new directions. I look forward to seeing where the journey takes us.

Special thanks are of course also reserved for my family, who have always been there for me. In particular I'd like to thank my Mum and Dad and sister Rachel, who have somehow managed to maintain a pretence of sanity when dealing with the rest of the world. Also to Gus the cat, who caught his final mouse in between drafts of this acknowledgements section.

I can't write this without saying thank you to Tom, who has appreciated - or, perhaps more accurately, tolerated - my dazzling sense of humour for four long years in the same office. We may have been each others' greatest source of procrastination, but I think we were far more of a help to each other than we were a hindrance. Thanks Tom - fancy a Sporcle? To the rest of Team Sparrowhawk, SABS 2014 - did anyone think we'd ever actually finish one day? Maybe we should have paid more attention and eaten less cake...! To everyone in the Wolfson Centre for Mathematical Biology: thanks for providing such a great atmosphere to do a DPhil in, and may there be many more WCMBeers in the evening.

To all the sensational musicians who've been part of Dot's Funk Odyssey during my time here: every year, you made it awesome. There are many things that are said to be "the best thing you'll do at Oxford". DFO really was. Any time you need a spare saxophonist, just let me know...

Finally, to the wonderful people of St Hugh's MCR over the past four years: there are too many of you to name, but my time here certainly wouldn't have been the same without you. Arguably I might have finished this thesis slightly faster, but let's not worry about that now!

Contents

List of abbreviations

The following list of abbreviations contains abbreviations which are commonly used throughout the text. Where an abbreviation is only used locally within the text, it is not listed here.

Abbreviation	Meaning	Context
CA	Cellular Automata	A type of agent-based modelling framework
CAIX / CA9	Carbonic anhydrase 9	Transmembrane protein used as a biomarker of hypoxia
CP	Cellular Potts	A type of agent-based modelling framework
CCL2	Chemokine (C-C motif) ligand 2 (also MCP-1)	A macrophage chemoattractant
CCL5	Chemokine (C-C motif) ligand 5	A macrophage chemoattractant
CD8	Cluster of Differentiation 8	A marker used in pathology to identify T cells
CD68	Cluster of Differentiation 68	A marker used in pathology to identify macrophages
CSF-1	Colony stimulating factor-1	A macrophage chemoattractant
CSF-1R	Colony stimulating factor-1 receptor	Receptor found on monocytes and macrophages that binds with CSF-1
EGF	Epidermal growth factor	Factor produced by macrophages that promotes cancer cell invasion
FoxP3	Forkhead Box P3	A marker used in pathology to identify T regulatory cells
H&E	Haematoxylin and Eosin	A commonly used staining combination in histopathological images
LPS	Lipopolysaccharides	Large molecules found in the outer membranes of many common bacteria that signal the presence of foreign cells to macrophages
ICAM-1	Intercellular adhesion molecule-1	Transmembrane protein which plays a role in intercellular adhesion
IHC	Immunohistochemistry	Imaging technique where antibodies are used to stain cells expressing specific proteins

Abbreviation	Meaning	Context
IL-4	Interleukin-4	Cytokine which can promote a pro-tumour phenotype in macrophages
IL-6	Interleukin-6	Cytokine which can upregulate tumour cell proliferation and survival
IL-10	Interleukin-10	Cytokine which can promote a pro-tumour phenotype in macrophages
MCP-1	Monocyte chemoattractant protein-1 (also CCL2)	A macrophage chemoattractant
M-CSF	Macrophage colony stimulating factor	An alternate name for CSF-1
MCTS	Multicellular tumour spheroid	An *in vitro* experimental model of oxygen-limited tumour growth
OS	Overlapping Spheres	A type of agent-based modelling framework
PCF	Pair Correlation Function	A form of second-order statistic used for analysis of point clouds
RITA	Roche Immunohistochemical Tissue Atlas	A dataset containing whole slide histological images
ROC	Receiver Operating Characteristic	A curve for assessing the performance of a binary classifier
SLIC	Simple Linear Iterative Clustering	An algorithm for efficiently generating superpixels
SVM	Support Vector Machine	A machine learning classification algorithm
TAM	Tumour associated macrophage	Macrophage which is inside a tumour
TNFα	Tumour necrosis factor α	Inflammatory cytokine produced by macrophages which stimulates tumour growth and angiogenesis
TGF-β	Transforming growth factor β	Cytokine which can promote a pro-tumour phenotype in macrophages
VMR	Variance/mean ratio	Metric for comparing the mean and variance of a distribution
VEGF-A	Vascular Endothelial Growth Factor-A	Growth factor promoting the growth of new blood vessels

Abbreviation	Meaning	Context
VB	Vertex Based model	A type of agent-based modelling framework
VT	Voronoi Tesselation	A type of agent-based modelling framework
WSI	Whole Slide Image	Digital image containing a complete histology slide

Chapter 1

Introduction

1.1 Thesis aims and structure

Macrophages are a key cell in the innate immune system, performing a wide range of tasks in a diverse array of biological contexts. In non-pathological tissue, these roles include the disposal of damaged cells, destroying foreign objects such as bacteria, regulating the adaptive immune system and playing a key part of the wound healing process. As a result of the diverse processes they play a role in, macrophages can express a wide range of phenotypes which depend on the cues they receive from their microenvironment.

In cancer, cells which have accumulated a series of mutations proliferate uncontrollably. Macrophages are usually able to recognise foreign or damaged cells and attack them, destroying many mutant cells. However, the ability to evade this type of immune surveillance is one of the characteristics of cancer cells, and has been recognised as one of the modern hallmarks of cancer (Hanahan and Weinberg, 2011). One of the ways through which cancer cells evade macrophages is by producing signalling molecules called cytokines (Dranoff, 2004). These cytokines influence the macrophage phenotype, so that cancer cells cause tumour associated macrophages (TAMs) to express pro-tumour phenotypes more usually associated with the wound healing process. Some of these cytokines also recruit macrophages to the tumour via chemotaxis, defined as the movement of cells up a chemical concentration gradient - cytokines which act as a chemoattractant are also known as chemokines (Mantovani et al., 2004). We elaborate on these processes in Section 1.2. Many chemokines which recruit macrophages to tumours are also responsible for altering macrophage phenotypes. As a result, it is hypothesised that the spatial distribution of macrophages within a tumour is related to the phenotype of those macrophages. Support for this comes from the experimental observation that the spatial distribution of macrophages is correlated with patient prognosis and treatment response across a wide range of different cancer types, suggesting that tumour infiltrating macrophages have a pro-tumour phenotype (Pyonteck et al., 2012; Sica et al., 2015; Chanmee et al., 2014; Scholl et al., 1994).

Immunotherapies, treatments which attempt to manipulate the immune system to target diseases, have become a key current focus of the pharmaceutical industry. By preventing tumour cells from evading the immune system, some treatments produced extraordinary results. Our industrial collaborators at Roche Diagnostics GmbH (henceforth abbreviated to Roche) are developing treatments which aim to manipulate macrophage phenotype to achieve novel tumour killing behaviour, which we describe in Section 1.2.3. These treatments, in combination with the observation that macrophage distribution correlates with patient prognosis, require that we understand the biological and physical factors that determine the spatial localisation of macrophages in a tumour. This requires tools which can assign quantitative metrics to the spatial distribution of macrophages in a tumour, and an understanding of the forces which drive cell movement within tumours.

This thesis aims to provide tools for quantifying spatial distributions of immune cells within cancer based on histological images of tumours. In Chapter 2 we develop an automated image analysis pipeline for the analysis of whole slide histology images. Histology samples show cell locations and morphology for a slice through a tumour, containing detail from the sub-cellular scale to the tissue scale. They are a common source of information and are routinely used in universities and in the clinic for research, diagnosis and treatment planning. Historically these images have been interpreted by skilled pathologists using microscopes, but recent trends in the field of digital pathology have introduced new opportunities. Slide scanners provide an affordable means for pathologists to digitize entire tissue slides to form whole slide images (WSIs). WSIs provide a rich dataset for the spatial analysis of immune cells, but the analysis of these images typically relies on metrics developed before digitization allowed detailed interrogation of the images. As such, automated methods to extract detailed quantitative information from WSIs have the potential to make a large impact on current clinical and research protocols by enabling more complex analyses of WSIs. The pipeline we present in Chapter 2 is designed to provide a means of flexibly interrogating WSIs to extract spatial information about cell locations.

In Chapter 3, we apply our pipeline to histological images to obtain cell locations for several immune cell subtypes within a tumour. We also show how our pipeline can be used to align histological images which come from consecutive slices of the tumour to combine information from multiple slides. This results in maps of macrophage locations within different tumours, combined with information about the oxygen environment of the tumour. We apply spatial statistics to the macrophage distributions obtained from our pipeline to identify patterns of macrophage localisation. Using statistics which measure clustering and dispersion, we show that combinations of quantitative spatial statistics are able to group tumour images into qualitatively similar groups. We also use our pipeline to determine spatial statistics which describe the immune cell distributions in a tumour, and investigate how accurately summary statistics obtained from a small sample of a tumour, such as a biopsy, can approximate the true value of those statistics across the whole tumour.

Chapters 2 and 3 provide tools for quantitatively describing macrophage distributions from histological data. In Chapter 4, we develop a mathematical model to understand the underlying mechanisms resulting in the differing distributions. Our model is motivated by experimental data described by Dorie et al. (1982) in which infiltration experiments were conducted on tumour spheroids. Multicellular tumour spheroids (MCTSs) are a model system of avascular tumour growth, in which a cluster of tumour cells is sustained by a nutrient which diffuses into the spheroid. As the spheroid grows a nutrient gradient is established which deprives cells in the core of nutrient and oxygen. This causes spheroids to grow to a maximum steady state size at which proliferation of cells at the spheroid edge is balanced by the death of cells in the spheroid core. In (Dorie et al., 1982), inert plastic microbeads are added to the boundary of MCTSs. Over time, these beads migrate towards the spheroid core, demonstrating the presence of an advective flow from the proliferative rim to the centre. Previous continuum models have reproduced this flux and attempted to explain the dynamics of the microbead distribution (McElwain and Pettet, 1993; Thompson and Byrne, 1999; Pettet et al., 2001). We use an alternative agent-based model, implemented in the Chaste (Cancer, heart and soft tissue environment) framework (Mirams et al., 2013), to simulate oxygen limited spheroid growth and the infiltration of microbeads into the spheroid. Our approach demonstrates that the dynamics observed experimentally can only be understood by examining the trajectories of individual microbeads, and clarifies the impact that the composition of a spheroid has on the speed of the advective flow.

In Chapter 5, we extend the model described in Chapter 4 to include macrophage infiltration into MCTSs. We contrast the “passive” migration which causes inert microbeads to infiltrate spheroids with “active” migration of macrophages. This active migration is the result of macrophage chemotaxis towards CSF-1, a chemokine produced by hypoxic tumour cells in the spheroid core. Results from our model demonstrate that the speed at which an individual macrophage moves towards the spheroid centre is affected more by the spheroid composition than by macrophage sensitivity to chemotaxis. However, increased infiltration is observed when macrophages are sensitive to chemotaxis because it causes them to spend less time clustering at the spheroid boundary. In Chapter 5 we also present preliminary results linking our agent-based model with *in vivo* tumour geometries obtained from the image analysis pipeline described in Chapters 2 and 3. These results indicate that chemotaxis plays a similar role in *in vivo* tumours as it does in tumour spheroids, with macrophages spending more time in the stroma in the absence of chemokines. Chemokines such as CSF-1 also bias macrophages towards a pro-tumour phenotype, so our model results are consistent with the hypothesis that patients whose tumours have high densities of infiltrating macrophages have a bad prognosis because these macrophages are more likely to have a pro-tumour phenotype.

We conclude in Chapter 6 by summarising our findings and discussing possible directions for further research.

The remaining Sections of this Chapter are dedicated to a review of existing techniques in relevant areas. We begin with a review of the role of macrophages and the wider immune system in non-cancerous tissue before discussing their behaviour in the context of cancer and their potential as targets for immunotherapy. We discuss the process of making histological images, and review existing software designed to analyse them. This is followed by a review of existing software designed for digital pathology, as well as image analysis techniques that can be used to identify individual cells from histological images. We conduct a review of mathematical models of macrophages in cancer, and discuss different agent-based modelling methodologies and software packages designed to implement them.

1.2 Review of biological background

1.2.1 Macrophages and the immune system

The human immune system consists of two interacting subsystems: the adaptive immune system and the innate immune system (de Visser et al., 2006; Sompayrac, 1999). The innate immune system includes macrophages, neutrophils and natural killer cells; these are the first cells to respond to dangers such as bacteria which manage to pass into the body. The adaptive immune system includes B-lymphocytes and T-lymphocytes (de Visser et al., 2006), and is responsible for developing a tailored response to threats over a longer timescale of several days. It is this system which produces antibodies against specific receptors on hostile cells and "remembers" previous threats to ensure a faster response when the host is challenged with a similar threat.

One of the key cell types in the innate immune system is the macrophage (Wynn and Vannella, 2016; Lewis and Pollard, 2006; Wynn et al., 2013). Macrophages are resident in healthy tissue and ordinarily fulfil a "garbage collection" role, destroying debris left behind after cell death through a process of phagocytosis, in which the macrophage engulfs the fragments of dead cells before lysing them internally (Murray and Wynn, 2011; Sompayrac, 1999).

When a tissue is initially damaged, for example by a wound breaking the skin, tissue-resident macrophages are present in the exposed area. Macrophages express "toll-like receptors" (TLR) which bind to "pathogen-associated molecular patterns" (PAMPs), structures commonly conserved between microbes such as the lipopolyscharides (LPS) commonly found in bacterial cell walls, or bacterial DNA (Harwood et al., 2012; Wynn and Vannella, 2016). On encountering signs of infection, macrophages are "activated" to become more phagocytic and, by immediately attacking the bacteria, are able to prevent bacterial infections from becoming uncontrollable during the period before the cells of the adaptive immune system are aware of the threat (Sica et al., 2015; Wynn et al., 2013). Macrophages use a similar process to phagocytose damaged cells.

Macrophages can express a range of phenotypic traits, enabling them to contribute to a variety of processes (Lewis and Pollard, 2006; Wynn et al., 2013). There is a spectrum of macrophage phenotypes which are determined by the immediate tissue microenvironment and a host of signalling molecules (cytokines) influence their phenotype, including interleukins IL-4, IL-13, and IL-10, colony stimulating factor-1 (CSF-1), transforming growth factor-β (TGF-β) and interferon γ (IFNγ) (Beck et al., 2009; Lewis and Pollard, 2006; Mosser and Edwards, 2008; Sica et al., 2015; Wynn and Vannella, 2016). The phenotypic spectrum is often simplified as two distinct and extreme phenotypes denoted M_1 and M_2, although in reality macrophages exhibit a spectrum of phenotypes (Sica et al., 2015; Mosser and Edwards, 2008). Table 1.1 summarises the key differences between M_1 and M_2 macrophages and lists key growth factors known to promote expression of an M_1 or M_2 phenotype.

The primary role of M_1 macrophages is to destroy hostile cells (Wynn et al., 2013). They do this in two ways. The first mechanism is phagocytosis: here an M_1 macrophage engulfs a cell. The macrophage also presents fragments of destroyed cells as antigens to prime other immune cells (Lewis and Pollard, 2006). M_1 macrophages also release inflammatory chemicals such as Nitric Oxide (NO) or reactive oxygen intermediates (ROIs) (de Visser et al., 2006; Sica et al., 2015; Wynn and Vannella, 2016) which damage both invading cells and healthy tissue, causing

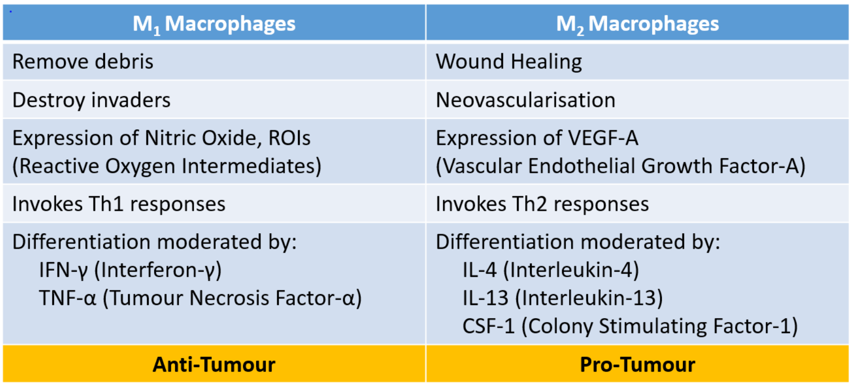

M_1 Macrophages	M_2 Macrophages
Remove debris	Wound Healing
Destroy invaders	Neovascularisation
Expression of Nitric Oxide, ROIs (Reactive Oxygen Intermediates)	Expression of VEGF-A (Vascular Endothelial Growth Factor-A)
Invokes Th1 responses	Invokes Th2 responses
Differentiation moderated by: IFN-γ (Interferon-γ) TNF-α (Tumour Necrosis Factor-α)	Differentiation moderated by: IL-4 (Interleukin-4) IL-13 (Interleukin-13) CSF-1 (Colony Stimulating Factor-1)
Anti-Tumour	**Pro-Tumour**

Table 1.1: Table comparing properties of M_1 and M_2 macrophage phenotypes. Some cytokines which are known to promote the expression of a particular phenotype are also shown.

leakage of cellular fluid into the tissue and leading to inflammation.

CSF-1 is a powerful chemoattractant for macrophages which attracts them from the surrounding tissue (Beck et al., 2009) and prompts monocytes, the precursor cells of macrophages, to extravasate from the bloodstream and differentiate into macrophages (Ginhoux and Guilliams, 2016). CSF-1 also causes M_1 macrophages to express an M_2 phenotype (Murdoch et al., 2004; Wynn et al., 2013), characterised by reduced phagocytic abilities and expression of growth factors such as vascular endothelial growth factor-A (VEGF-A), which promote angiogenesis from existing vasculature (Sidibe et al., 2018; Murdoch et al., 2004; Wynn et al., 2013; Wynn and Vannella, 2016). While M_1 macrophages promote a pro-inflammatory immune response (Th1 response), M_2 macrophages promote an anti-inflammatory (Th2) response (Sica et al., 2015; Mills et al., 2000).

In practice, macrophages cannot be decomposed into distinct subpopulations. Instead, it is more accurate to view a macrophage as "polarized" or "activated" to exhibit behaviours characteristic of those described above (Murray et al., 2014; Murray and Wynn, 2011). The "functional phenotype" of a macrophage depends on its microenvironment (Stout and Suttles, 2004; Stout et al., 2005) and may change through a process termed "functional plasticity" (Belgiovine et al., 2016). It is possible to target the specific pathways involved in this process and thereby inhibit phenotypic switching (Kloepper et al., 2016; Peterson et al., 2016; Mylonas et al., 2009; Hagemann et al., 2008).

1.2.2 The role of macrophages in cancer biology

Macrophages play an important role during tumour growth, and may constitute up to 50% of the mass of some tumours (Murdoch et al., 2004; Franklin et al., 2014). The two extreme macrophage phenotypes have opposing effects on tumour growth: M_1 macrophages have an anti-tumour effect, attacking tumour cells via phagocytosis and the expression of lysing factors, as well as promoting an inflammatory immune response (Th1 response) by recruiting other immune system cells hostile to the tumour to the site. M_2 cells have a pro-tumour effect: their expression of vascular growth factors stimulates the growth of the disordered vascularisation associated with tumours (Sica et al., 2015; Sidibe et al., 2018; Lewis et al., 1999; Lamagna et al., 2006). New vessels bring oxygen and nutrients to hypoxic tumour regions and provide resources that enable the tumour to grow. M_2 macrophages also exert an immunosuppressive anti-inflammatory effect and each macrophage expressing an M_2 phenotype has a greatly reduced ability to attack the tumour (Munn and Bronte, 2016).

Tumours can subvert anti-tumour M_1 macrophages and cause them to switch to a pro-tumour, M_2 phenotype: hypoxic tumour cells secrete cytokines including CSF-1 which stimulate macrophages to stop attacking the tumour, and to instead produce VEGF-A to promote vascular growth (Sidibe et al., 2018; Stockmann et al., 2014; Sica et al., 2008). Treatments for a range of cancer types have been proposed that target pathways associated with switching to an M_1 phenotype (Pradel et al., 2018; Edwards V et al., 2018; Fujimura et al., 2018; Mantovani et al., 2017; Peterson et al., 2016; Kloepper et al., 2016; Cassier et al., 2015; Noy and Pollard, 2014; Cannarile et al., 2014).

Macrophages are also implicated in tumour metastasis (Qian et al., 2011; Vasiliadou and Holen, 2013; Keklikoglou and De Palma, 2009; Yuan et al., 2014; Qian and Pollard, 2010; Kadioglu and de Palma, 2015; Qian et al., 2009). A prominent candidate for the mechanism through which macrophages promote metastasis is the "tumour microenvironment of metastasis" (TMEM) (Laoui et al., 2014; Harney et al., 2015; Hughes et al., 2015; Kadioglu and de Palma, 2015). A TMEM is formed when an endothelial cell, macrophage and cancerous cell come into close proximity, with the macrophage causing transient permeability of the vessel wall and permitting the cancerous cell to enter the blood vessel (Harney et al., 2015).

The spatial distribution of tumour associated macrophages (TAMs) has been shown to be an important indicator of treatment viability. For instance, in solid tumours the extent of macrophage infiltration has been identified as a poor prognostic factor, correlating with negative treatment outcome (Pyonteck et al., 2012; Ries et al., 2014; Sica et al., 2015). TAM infiltration positively correlates with tumour cell proliferation (Lewis and Pollard, 2006). In non-small cell lung cancer, the extent of M_1 macrophage infiltration into the tumour islets has been claimed to have a positive effect on survival time (Ma et al., 2010). The spatial distribution of macrophages varies widely between both tumour type and prognosis, and understanding and manipulating this distribution forms a key part of current research.

1.2.3 Macrophages as targets for cancer immunotherapy

Immunotherapy is a rapidly expanding field of cancer treatment which aims to encourage the immune system to attack cancerous cells. It is “experiencing a renaissance” (Hoos, 2016) with therapies that involve almost all cells of the adaptive and innate immune systems currently under development (Khalil et al., 2016; Hoos, 2016; Cassier et al., 2015). Understanding the ways in which the cells of the immune system interact with cancer and the feasibility of manipulating these cells to induce tumour death are crucial to the development of effective immunotherapeutic treatments.

Preventing tumour manipulation of macrophage phenotype is an area of active research in immunotherapy. The CSF-1 receptor (CSF-1R) has been proposed as a target for manipulating the functional phenotype of macrophages (Pradel et al., 2018; Rovida and Sbarba, 2015; Van Overmeire et al., 2015). For example, antibodies which can block the CSF-1 receptor to prevent CSF-1 interacting with macrophages may bias the population to adopt an M_1 phenotype in patients with diffuse-type giant cell tumours (Ries et al., 2014).

1.2.4 Imaging tumours using histology slides

Histology is the study of anatomy on a microscopic scale, and is one of the oldest and most common methods used for medical imaging (Slaoui and Fiette, 2011). When studying diseased tissue, the field is also known as histopathology. At its most general level, histology involves using a light microscope to make detailed observations of a tissue sample at the scale of individual cells, typically in the form of a 2D slice through the tissue. The process of preparing a tissue sample for histological imaging generally proceeds according to the following steps (Slaoui and Fiette, 2011):

1. Fixation
2. Embedding
3. Sectioning
4. Staining
5. Slide examination

We describe these stages in the following paragraphs in the context of tumour tissue, although the same stages are required to process any other tissue sample.

Fixation The aim of tissue fixation is to prevent the degradation of proteins within the tissue. Once the sample has been surgically removed, cells within it immediately begin degrading. Fixation halts this process, ensuring that the state of the histological image is as close as possible to that of the *in vivo* sample. In situations when the sample must be analysed quickly, such as a biopsy taken during surgery which is to be examined to ensure that sufficient tissue has been removed that the invasive margin of the tumour has been extracted, fixation can be achieved by flash freezing, which can damage the tissue (Buchwalow and Böcker, 2010). For the images used in this thesis, the more common method of fixing using formalin was used (Slaoui and Fiette, 2011). After formalin fixing the tissue must be dried, which can cause morphological deformations such as the shrinking of necrotic tissue that cannot be easily prevented.

Embedding Once fixed, the tumour sample is embedded within a paraffin block. The sample must be soaked in paraffin to allow it to infiltrate the tissue, and it is then placed in a mould filled with liquid paraffin. This block is then cooled until it becomes solid. Embedding the tissue in paraffin preserves the structure and allows the block to be sliced thinly without deforming the tissue.

Sectioning Sectioning is the process of cutting a tissue sample into $4\,\mu$m - $5\,\mu$m thick slices, or sections. This is achieved using a microtome, a precision knife that can replicably produce extremely thin sections (Slaoui and Fiette, 2011). Each slice is floated onto a water bath and a glass microscope slide placed underneath. The section is then removed directly onto the slide and allowed to dry. The sectioning process can damage the section in a variety of ways which must be accounted for when analysing the image, such as tearing the slice or trapping bubbles between the tissue and the slide.

Staining Paraffin embedded slides typically have very low contrast when viewed through a microscope. To increase contrast and to highlight different features of cells, the field of immunohistochemistry (IHC) has been developed in order to selectively stain cells which express specific proteins (Buchwalow and Böcker, 2010). IHC involves the creation of antibodies which selectively bind with a target protein, or marker, in the tissue, and which also contain a dye molecule. The process of staining for a particular receptor involves creating antibodies which will bind with this receptor. The histology slide is then washed in a buffer containing these antibodies. Once the antibodies have been allowed to bind with the target receptors in the tissue, excess buffer is washed away. Typically a counterstain is added which identifies cell nuclei, such as DAPI or hematoxylin, to ensure that the tissue surrounding stained cells can be seen (Buchwalow and Böcker, 2010). Finally, stained slides are covered with another slip of glass to protect the tissue sample.

Although most markers are endogenous - they are proteins which are naturally expressed by the cell - some markers are not naturally present. One such marker present in the datasets examined in this thesis is pimonidazole. Pimonidazole is not itself expressed by cells, but is a compound which binds selectively with thiol-containing proteins in hypoxic cells (Varia et al., 1998). As such, pimonidazole must be injected in to the patient before a tissue sample is taken so that it can bind with these proteins. Antibodies selected to bind with pimonidazole can then be applied to histology slides following the same process as for an endogenous marker.

Slide examination Once the slice is mounted on a microscope slide, it can be examined using light microscopy. The most common way of examining histology slides is by looking through such a microscope, but numerous slide scanners now exist for taking high-resolution digital images of these slides (García Rojo et al., 2003; Abels and Pantanowitz, 2017). Typically, these digitised slides are assessed by a pathologist with no further computational assistance - digitising the image is generally seen as a way of making slide examination easier, rather than as a prerequisite for more advanced computational analysis of the slide (Williams et al., 2018b).

1.3 Image analysis methods for histology

1.3.1 Existing software for computational histology

A recent report to the UK government from the life sciences sector highlights pathology as an area which is “ripe for innovation”, and in which algorithms based on artificial intelligence could provide “prognostic insights that are not currently available through conventional methodology” (Bell, 2017). Several commercial platforms have been developed to address this need (Webster and Dunstan, 2014), but their use is not widespread (Williams et al., 2018b). In a survey of academic and clinical histopathology departments, 96% of respondents identified the cost of these systems as a major factor restricting the uptake of digital pathology systems (Williams et al., 2018b). Among the commercial systems in common use are Visiopharm (Holten-Rossing et al., 2017), Definiens (Baatz et al., 2009), Aperio and Halo (Bueno et al., 2016; Webster and Dunstan, 2014; Stack et al., 2014).

Academic development of open source software for whole slide image analysis is growing in response to demand for non-commercial systems. ImageJ is a commonly used tool for medical image analysis research (Schindelin et al., 2015), and plugins designed to manipulate WSIs are available (Deroulers et al., 2013; Blom et al., 2017). An advantage of ImageJ is the familiarity that many medical researchers have with the interface and system, and its modular nature means that researchers who use it for other image analysis features can integrate whole slide analysis into their workflows. However, as a generic image manipulation tool ImageJ is not specifically designed for the analysis of whole slide images, which means that additional plugins must be installed to perform simple operations (Deroulers et al., 2013). Recently, QuPath has been designed to help meet the need for open source digital pathology software (Bankhead et al., 2017; Hamilton et al., 2014; Kather et al., 2018). In this thesis, we aim to provide a tool which directly integrates mathematical and statistical analysis of histology images into the image analysis process.

1.3.2 Image analysis of histology data

Superpixellation

One way to reduce image size while retaining object boundaries is through the use of *superpixels* (Ren and Malik, 2003). Superpixels are clusters of pixels formed of adjacent pixels with similar colour intensities. They are designed to reduce the size of an image while minimising the amount of information lost. By grouping neighbouring pixels which have similar colour profiles, superpixel segmentation dramatically reduces the effective number of pixels in an image which must be iterated over during classification. Although there is no accepted formal definition of a superpixel, they are “perceptually meaningful” regions which “should contain pixels that are similar in color, texture, etc., and therefore are likely to belong to the same physical world object” (Veksler et al., 2010). Stutz et al. (2018) identify six properties that most authors require of superpixel algorithms:

- Partition (every pixel should be in exactly one superpixel),
- Connectivity (superpixels should represent connected sets of pixels),
- Boundary adherence (superpixels should conserve the boundaries of objects within an image),
- Compactness (in the absence of boundaries during oversegmentation, superpixels should be smooth, compact and regularly sized),
- Efficiency (algorithms should not be overly computationally expensive),
- Controllable number of superpixels.

There are many algorithms for superpixel generation, each producing superpixels with different characteristics. One class of algorithms uses graph-based methods: each pixel is regarded as a node of a graph and two nodes are connected if the corresponding pixels in the image are adjacent. The edges are then weighted according to the similarity in intensity or colour between nodes. Superpixels are then generated by minimising a cost function which is defined over this graph. Examples of graph-based methods include algorithms by Veksler et al. (2010), Felzenszwalb and Huttenlocher (2004) and the Normalised Cuts algorithm of Shi and Malik (2000). In their comparison of superpixel algorithms, Achanta et al. (2012) note that graph-based methods become increasingly computationally expensive as image sizes grow, since more nodes must be added to the graph. This limits their utility to analysis of histological images which contain approximately 10^9 pixels.

A second class of superpixel algorithms, based on gradient ascent, rely on an initial rough clustering of superpixels. Clusters are refined iteratively until a convergence condition is met. Gradient-ascent methods include the TurboPixels of Levinshtein et al. (2009), Quick Shift (Vedaldi and Soatto, 2008) and Mean Shift (Comaniciu and Meer, 2002). Simple Linear Iterative Clustering (SLIC) (Achanta et al., 2012; Lucchi et al., 2010) is another gradient ascent based method. On its first iteration, SLIC assigns a "seed" pixel to act as the centre for each superpixel, arrayed in an evenly spaced grid across the image. It then iterates over each pixel in the image and assigns it to a superpixel seed based on its proximity to surrounding seeds in a five dimensional space defined by three colour channels and the cartesian distance from the seed. New superpixel seeds are then calculated as the mean of the pixels within the superpixel, and the process is repeated until no cluster centre varies by more than a given threshold between iterations. Each pixel in SLIC is only tested against "local" pixels: if the expected extent of a superpixel is $S \times S$ pixels then each pixel will only be assigned to a cluster if the centre falls within a search region of size $2S \times 2S$. This significantly reduces the number of distance calculations which must be made and gives SLIC a speed increase over other superpixel algorithms, with a complexity that is linear in the number of pixels in the image N and independent of the number of superpixels ($O(N)$) (Achanta et al., 2012). In Figure 1.1 we compare a representative result of the output of the SLIC algorithm, in which each superpixel has been coloured according to the mean colour value of the pixels within itself, with the original image. The image used was chosen to demonstrate how SLIC preserves boundaries between distinctly coloured regions while reducing unnecessary detail within them.

Classification of cell types

Many commonly used techniques for the classification of cell types are classed as machine learning, and can broadly be split into two categories: supervised learning and unsupervised learning (Litjens et al., 2017). In both techniques, an algorithm called a "classifier" evaluates characteristics of an object and uses a series of rules to

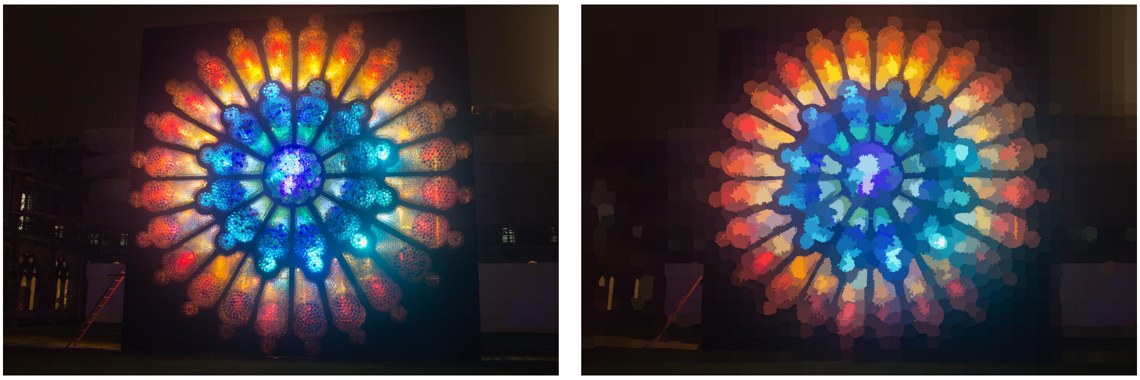

(a) Original Image (b) SLIC superpixel segmentation

Figure 1.1: The SLIC superpixel algorithm applied to an art installation. Superpixels have been coloured according to the colour of the superpixel seed pixel. Each superpixel in Figure 1.1b is approximately 2048 times larger than a pixel in Figure 1.1a.

decide on which of a series of "classes" the object belongs. This object could be an individual pixel or superpixel, or an entire unprocessed image. Supervised learning algorithms must be given a set of input/output pairs from which the algorithm will infer a set of classification rules which best describe this training data. When faced with new data, the algorithm will use these rules to decide on the new class. In order to infer the rules, supervised learning algorithms must be provided with a set of "features" upon which to base its decision. Defining these features is a crucial part of training a strong classifier. Supervised learning techniques include support vector machines (SVMs) (Cortes and Vapnik, 1995; Mohri et al., 2012; Hsu and Lin, 2002), decision trees and random forests (Tin Kam Ho, 1995), and linear regression. Learning algorithms which do not require a preset list of features are perhaps more popular for medical image analysis (Litjens et al., 2017). In particular, convolutional neural networks and deep learning perform well at cell and region identification tasks (Sirinukunwattana et al., 2016; Fukushima, 1980; Krizhevsky et al., 2012; Litjens et al., 2017).

In this Section we focus on defining SVMs, which we use as the basis of the classifiers in this thesis. The pipeline which we present in Chapter 2 is designed such that the user can provide or train a different type of classifier if desired.

Support vector machines

The method we use for superpixel classification relies on the idea of a support vector machine (SVM) (Cortes and Vapnik, 1995; Mohri et al., 2012; Hsu and Lin, 2002). SVMs have previously shown promise in segmenting images from biological imaging data including electron microscope images (Lucchi et al., 2010, 2015), mammogram images (El-Naqa et al., 2002) and Magnetic Resonance (MR) images (Chaplot et al., 2006). They have also been used successfully for cell detection in histology images (Arteta et al., 2012).

To define a classifier, we first define a "feature space" for the set of superpixels. This space is typically high-dimensional, with some machine learning strategies using hundreds or even thousands of features for each point to be classified, although in practice a much lower number of features can produce similar results with a much lower computational cost. The size of the feature space is defined by the number of features defined on each superpixel, so a classification which looked at the mean and variance of each of the three colour channels over all pixels within a superpixel would be defined in a 6-dimensional space. The motivation behind a classifier is that since each superpixel defines a point in this feature space, selecting a feature space in which superpixels with a similar known labelling cluster together can permit the classification of a new superpixel based on its proximity to known clusters. A classification algorithm is then used to determine which regions of the feature space correspond to superpixels with a given classification. Since each feature defines a new dimension in feature space, the selection of appropriate features to describe each superpixel for classification is an important decision. The optimal combination of features will be one in which the data is clearly separated into classes while the complexity of the feature space is kept as low as possible.

A SVM consists of a set of vectors in feature space which define a hypersurface

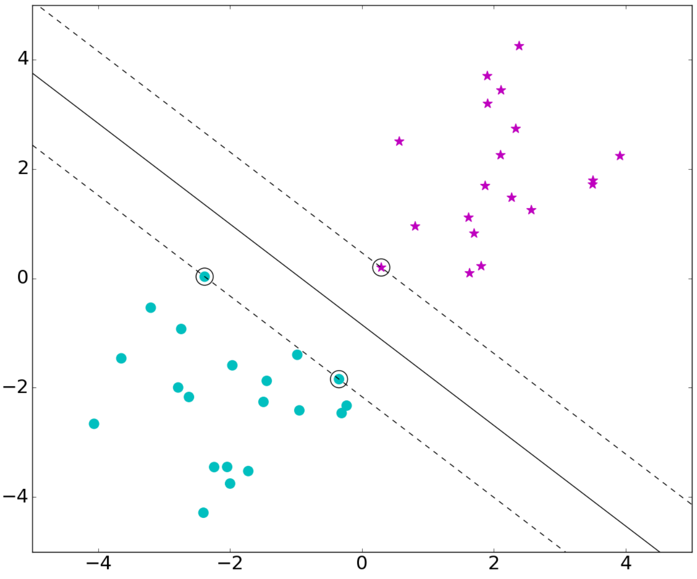

Figure 1.2: An example of a support vector machine maximum-margin separating hyperplane. The cyan circles and magenta stars represent two separable classes in a feature space defined by two features, the x and y coordinates of the points. The circled vectors are the support vectors and define the location of the hyperplane (solid line). The dashed lines show the extent of the margin - note that by adjusting the angle of the hyperplane, the margin between one or more of the support vectors and the hyperplane will be reduced.

dividing two labellings of the data. This is demonstrated in Figure 1.2, which shows a data set of feature vectors labelled with two different colours and defined in a two-dimensional feature space. The solid line dividing the clusters is a hyperplane defined in a two-dimensional feature space.

In this example, the two labellings are linearly separable, and so a hyperplane can be found which classifies this training data with complete accuracy. Since there are infinitely many hyperplanes which separate the data correctly, the optimum solution is taken to be the "maximum-margin separating hyperplane". The margin of a SVM is the distance between the hyperplane and the nearest feature vector, and those vectors which lie on this margin are defined to be "support vectors". In Figure 1.2, the support vectors are surrounded by a solid ring, and the dashed lines are separated from the hyperplane by the margin. This hyperplane is defined by the training data shown in the image. When a new data point with unknown labelling is to be classified, it is simply assigned a label based on which side of this hyperplane its feature vector resides. Maximising the margin therefore minimises the chance that a new observation will be misclassified by falling on the wrong side of the hyperplane.

A hyperplane is the simplest surface to define since it is linear and can therefore be defined fully for an n-dimensional space using only n points, thus carrying a relatively low computational cost to train. However, it is not always possible to cleanly separate training data using a hyperplane. In this case, a penalty will be assigned for each misclassified point and the SVM is considered to be sufficiently well trained when the overall penalty is below some given tolerance. This can often be useful to minimise overfitting, as extensions to SVM which use nonlinear hypersurfaces can become overly complicated if the tolerance is too low, resulting in poor predictive power.

A nonlinear SVM (El-Naqa et al., 2002; Lucchi et al., 2015) operates on the same principles as the linear SVM described above, but instead of seeking an optimal hyperplane the separating surface is described as a more general nonlinear surface. Rather than explicitly defining a surface and classifying vectors by calculating their location relative to it, it is common to define a kernel, or kernel function, which compares the relative similarity of two input vectors without explicitly calculating their coordinates in the feature space. The kernel is weighted to provide structure on the feature space, so the choice of kernel can be thought of as defining the shape of the hypersurfaces at the boundary of each classification.

While SVMs with a linear kernel are very quick to train due to the comparatively low number of parameters needed to describe a hyperplane, in theory a nonlinear kernel could use arbitrarily many parameters. Consider the case of a polynomial kernel trained to separate a collection of vectors in a two-dimensional feature space. If a polynomial of sufficiently high degree d is chosen, a highly convoluted separating surface will be obtained with d parameters which can be fine-tuned to separate the training data. However, just as when fitting a polynomial of high degree to a collection of points, this is very likely to lead to overfitting. An overcomplicated kernel has a higher likelihood of being overfitted to the training data and therefore is unlikely to classify new points with acceptable accuracy. To avoid this danger, higher degree polynomials are generally penalised when assessing how good a classifier is and a kernel of lower degree which fits the training data to a given tolerance would be preferable.

A popular alternative to a polynomial kernel is the Radial Basis Function (RBF) kernel, or Gaussian kernel. In this case, the kernel is defined using a combination of gaussian distributions. This can lead to a decision boundary which is a closed curve rather than an infinite surface, which is necessary when the choice of feature space leads to one of the classes being surrounded by another.

1.4 Mathematical modelling

Mathematical modelling of tumour-immune interactions is an active area of research, with techniques from across the field of modelling being applied to many problems involving interactions between components of the tumour-immune system (Eftimie et al., 2016).

Many mathematical models examine the role of the adaptive immune system in tumour growth and focus on the behaviour of tumour-infiltrating lymphocytes and T cells (de Pillis et al., 2009; Robertson-Tessi et al., 2012; Wilkie and Hahnfeldt, 2013; López et al., 2014; Friedman and Liao, 2015). Cappuccio et al. (2006) and de Pillis et al. (2009) use ordinary differential equations (ODEs) to describe immunotherapy strategies based on interactions between T cells and natural killer (NK) cells. Other ODE models describe tumour cell interactions with dendritic cells (Robertson-Tessi et al., 2012), or include a generic immune component (López et al., 2014; Wilkie and

Hahnfeldt, 2013). In this Section, we describe several key models which describe interactions between macrophages and tumours.

1.4.1 Mathematical models of macrophages in cancer

Owen and Sherratt (1998) developed an ODE model for the invasion of small, avascular tumours by macrophages. Their model focusses on macrophage lysis of tumour cells, but includes macrophage chemotaxis in response to a chemical which also regulates macrophage activation and proliferation. While the model is not spatially resolved, macrophage chemotaxis towards a generic chemical regulator is modelled as an additional influx of macrophages due to extravasation. Macrophages proliferate in the presence of the chemical regulator but proliferation is inhibited by a crowding response, which also inhibits tumour cell proliferation.

Owen and Sherratt (1997) extended the model to consider spatial heterogeneity using partial differential equations (PDEs) in one dimension. In this model, tumour-macrophage interactions are still defined predominantly though macrophage lysis of tumour cells. There is also spatial crowding which inhibits proliferation of both cell populations. In this model, chemotaxis is implemented using the classical form derived by Keller and Segel (1971) describing the movement of cell populations up a chemical gradient. The distribution of the chemical regulator is defined using a reaction-diffusion equation, where the regulator is produced by tumour cells, decays at a constant rate and diffuses through the domain. Their model produces travelling wave solutions, and the authors show that although the rate of macrophage lysis of cells does not impact the overall rate of tumour growth it has a significant effect on tumour composition behind the wavefront. Their model also exhibits Turing patterns due to the diffusion of the regulator. Rapid chemical diffusion causes macrophage extravasation and activation in areas close to places with high tumour cell density, suppressing the growth of tumour cells. This behaviour is also seen in two spatial dimensions (Owen and Sherratt, 1999).

One weakness of the complementary models presented by Owen and Sherratt (1998, 1997, 1999) is the assumption that tumour composition is heterogeneous. Although small tumours may reasonably be assumed to have a constant oxygen concentration throughout, the composition of larger avascular and vascular tumours is influenced by the non-uniform concentration of oxygen throughout the tumour described in Section 1.2. Greenspan (1972, 1976) presented one of the earliest models of tumour spheroid growth limited by oxygen diffusion, which uses PDEs to describe radially symmetric spheroid growth. This model distinguishes three tumour compartments defined by thresholds of oxygen concentration at which cell behaviour changes: a proliferative outer shell, a layer of viable but non-proliferating cells, and a central necrotic core. Kelly et al. (2002) use this model as a basis to explore macrophage infiltration into tumour spheroids. They propose a radially symmetric PDE model motivated by experiments by Leek et al. (1996, 1999). In these experiments, macrophage infiltration into spheroids made from two different cell lines is measured. Hypoxic cells in the Hepa-1 cell line produce macrophage chemoattractants (VEGF-A) while those in the C4 cell line do not. The model in (Kelly et al., 2002) reproduces key results observed in (Leek, 1999), in particular that chemoattractant production is not required to cause macrophage infiltration, but that it does increase the speed at which macrophages are recruited to the tumour spheroid. The authors reject the explanation that the significant size difference between smaller C4 spheroids and larger Hepa-1 spheroids may be a cause of this difference, although they note that as the spheroid size increases the spatial structure of the tumour varies. In their model, the width of the proliferative rim and hypoxic region tend to a constant value, while spheroid size increase is accounted for by the increasing radius of the necrotic core.

Byrne et al. (2004) develop a complementary ODE model to (Owen and Sherratt, 1998) in which genetically modified macrophages can cause the immediate death of tumour cells. Their model also includes normal, non-cancerous, cells. In (Owen et al., 2004), a spatially resolved version of this model is constructed based on the principles of multiphase models of avascular spheroid growth described by Ward and King (1997, 1999). Like (Kelly et al., 2002), this model simulates macrophage infiltration into avascular spheroids. The use of a multiphase framework allows the authors to relax the implicit assumption in (Kelly et al., 2002) that macrophages do not occupy volume. This model successfully reproduces the infiltration dynamics described by Leek (1999), with macrophages accumulating in the hypoxic/quiescent, chemoattractant producing, region of the spheroid. However, spheroids which express a chemoattractant in their model grow significantly larger than those which do not, which the authors attribute to macrophages occupying space in the hypoxic region.

Webb et al. (2007) develop a model of avascular tumour spheroid response to two types of treatment which use macrophages. In the first, genetically engineered macrophages deliver an enzyme to hypoxic spheroid regions which activates a prodrug to cause cell death. In the second, macrophages are used to directly administer a cytotoxic factor to the hypoxic regions. Their model uses PDEs to describe the tumour spheroid as a multiphase mixture of tumour cells, macrophages and extracellular fluid, with a moving boundary representing tumour spheroid growth. Macrophage movement includes movement up the gradient of a chemokine which is produced by hypoxic tumour cells.

Owen et al. (2011) developed a 2D lattice-based hybrid model to study the infiltration of genetically engineered macrophages into hypoxic tumour regions. Their model combines an agent-based framework in which each cell is treated separately with ordinary and partial differential equations (ODEs and PDEs) submodels. The ODEs describe the dynamics of subcellular species while PDEs describe the concentrations of oxygen and VEGF-A at the tissue scale. Each individual cell has a specified cell type, such as “hypoxic tumour cell”, “normoxic tumour cell” or “macrophage”, and is restricted to occupy sites on a lattice. These cells can move to neighbouring grid sites, proliferate, undergo apoptosis and interact with the diffusible species described by the PDE models. This model also includes vasculature which grows in response to the local cellular VEGF-A concentration and acts as a source or sink for different diffusible species. A novel treatment is considered in which a magnetic field is used to cause macrophages, which have previously consumed magnetic nanoparticles, to extravasate into specific regions of the tumour.

Knútsdóttir et al. (2014) propose a PDE model and a hybrid model of breast cancer cells and macrophages. They focus on a “paracrine loop” in which macrophages are attracted to CSF-1 released by tumour cells, while tumour cells are attracted to endothelial growth factor (EGF) secreted by the macrophages. The authors attribute the aggregation of macrophages and tumour cells, and their subsequent group migration and metastasis, to this loop. Further development of the agent-based model indicated that tumour cell migration towards blood vessels is highly sensitive to the fraction of macrophages infiltrating the tumour (Knútsdóttir et al., 2016). These models provide a good framework for the chemoattraction of macrophages to tumour cells mediated by the EGF / CSF-1 paracrine loop, even though they do not distinguish between macrophage phenotypes. This EGF / CSF-1 loop is also modelled by Elitas and Zeinali (2016) to describe macrophage infiltration into gliomas, and a recent PDE model by Zheng et al. (2018) considers the effect of resistance to immunotherapy targeted at the CSF-1 receptor in gliomas.

Norton et al. (2018) consider the effect of macrophage infiltration on tumour growth. They use an off-lattice 3D agent-based model to simulate angiogenesis and macrophage invasion of a growing tumour. Macrophages in their model increase the migration and proliferation rates of tumour cells, but have no negative effects on tumour growth. The authors find that, under these assumptions, the presence of macrophages increases overall tumour growth but that this is not strongly affected by the spatial location of macrophages either inside or outside the tumour mass. They also find that at excessive rates the presence of macrophages can stifle tumour growth due to crowding effects.

An ODE model by den Breems and Eftimie (2016) considers the effect of macrophage phenotype on tumour growth. The model includes a mutually reinforcing feedback loop between M_1 macrophages and a Th1 immune response and between M_2 macrophages and a Th2 immune response. It also divides the tumour population into two groups, one "immunogenic" population which is recognised by the anti-tumour immune components and triggers an increase in M_1 macrophages, and another "non-immunogenic" population which evades immune recognition. They conclude that tumour escape is ultimately due to the dominance of the tumour population which avoids immune surveillance. This conclusion is undermined by the assumption that the pro-tumour effect of M_2 cells only benefits the non-immunogenic subpopulation rather than providing a global benefit through increased vascularisation that M_2 cells stimulate. Macrophages can change between the M_1 and M_2 phenotypes, but the rate at which they do so depends on two constant parameters which are independent of the tumour environment. Since the model neglects spatial effects and does not account for the effect of tumour cell-derived cytokines on macrophage phenotype, the rate of phenotype switching is independent of the macrophage location in the tumour and only mediated by interactions between M_1 and M_2 macrophages.

Other models involving phenotypic switching of macrophages in cancer include (Carmona-Fontaine et al., 2017), in which agent-based modelling is used to demonstrate that gradients of nutrients produced by tumour cells can lead to local clustering of macrophages with the same phenotype. Wells et al. (2015) incorporate phenotypic switching between M_1 and M_2 macrophages into a spatial model of tumour growth. They use cellular automata to demonstrate that as the macrophage population becomes polarized from anti-tumour to pro-tumour, the probability of controlling tumour growth decreases. (Mahlbacher et al., 2018; Leonard et al., 2017) uses a lattice based model to model the effect of macrophage phenotype of tumour growth. The authors show that a high number of M_2 macrophages is correlated with larger tumour growth.

Jacobsen et al. (2015) consider the effect of blocking macrophage chemoattraction to gliomas in the context of oncolytic virus (OV) treatment. OV treatment uses a genetically engineered virus to kill tumour cells. The authors use a PDE model to demonstrate that blocking chemoattraction reduces macrophage killing of tumour cells and leads to the wider spread of the oncolytic virus, although they do not consider the possibility that macrophages express a pro-tumour rather than anti-tumour phenotype.

Other mathematical models featuring macrophages in the context of cancer include (Chen et al., 2014, 2012), which consider the effect of tumour macrophage hypoxia inducible factors (HIFs) in the context of chemotherapy. In an extension to their PDE model of melanoma invasion, Eikenberry et al. (2009) consider the effect of a cytotoxic immune response from macrophages on preventing melanoma metastasis.

1.4.2 Overview of agent-based modelling methodologies

"Agent-based modelling" (or "individual-based modelling") refers to a wide range of techniques in which each cell is modelled as a distinct object. Each agent-based modelling framework has its own benefits and drawbacks for modelling a particular problem. In this Section, we review some commonly used agent-based models and discuss their suitability for in simulating the infiltration of macrophages into tumour spheroids.

Agent-based models are commonly divided into "on-lattice" models, in which individual cells are restricted to sites in a grid (Figueredo et al., 2013), and "off-lattice" models in which cells may move freely across a continuous space (Osborne et al., 2017; Byrne and Drasdo, 2009). Each framework discussed in this Section can be implemented in either 2D or 3D, although the computational cost of moving to higher dimensions can vary dramatically between frameworks and their implementations. On-lattice models are typically easier to implement and run faster than off-lattice models, while off-lattice models can more readily account for mechanical interactions between cells. The choice of modelling framework therefore has significant impact on how quickly and easily results can be obtained and on the types of interactions that can be incorporated into the mathematical model.

In this Section, we distinguish between modelling "frameworks" and "implementations". We use the word "framework" to refer to the type of model which is used, and the word "implementation" to denote the computer code used to simulate the model. We make this distinction due to the existence of several software tools which are used by the mathematical community to implement different agent-based modelling frameworks. These software tools are generally open source and supported by a range of developers who add functionality to these implementations to enable the efficient creation of new models that capture increasingly complex biological behaviour. Throughout this thesis, we use the Chaste (Cancer, heart and soft tissue engine) software implementation (Mirams et al., 2013; Pitt-Francis et al., 2009) as a foundation for developing our agent-based modelling code, and will ultimately add our contributions to the open source framework. Models implemented using Chaste are programmed in C++, and an object-oriented approach is used in which each component of a mathematical model is encapsulated individually. This means that model components can be added from a library of pre-existing files, and newly developed code can be easily integrated to extend existing models. This approach encourages careful testing of modules and the incremental development of models in which each subcomponent can be individually added, removed or tested for correct implementation. It is possible to implement models in a wide range of frameworks in Chaste, including all of those shown in Figure 1.3.

In addition to Chaste, several other software implementations can be used to simulate different types of agent-based models. Physicell (Ghaffarizadeh et al., 2018) is an open source physics based cell simulator. It implements overlapping spheres models and utilises BioFVM (Ghaffarizadeh et al., 2016) to obtain efficient simulations involving large numbers of diffusing substrates such as oxygen or CSF-1. Morpheus (Starruß et al., 2014) focusses on user-friendliness, with a GUI designed to bypass many of the coding challenges mathematicians face when developing agent-based models. Like Chaste, it can implement models using a range of on-lattice or off-lattice frameworks (Starruß et al., 2014). CompuCell3D (Swat et al., 2012) provides an intuitive way for non-programmers to implement models using the cellular Potts framework (Graner and Glazier, 1992). Other software tools that implement agent-based models include CellSys (Hoehme and Drasdo, 2010), Biocellion (Kang et al., 2014), HAL (Bravo et al., 2018) and Timothy (Cytowski and Szymanska, 2015, 2014). In addition to these resources, many groups use their own software implementations, although this approach increases the risk of undetected coding errors generating incorrect results and reduces code transparency and reuse.

Although many of these software implementations restrict users to a single modelling framework, others including Chaste permit agent-based models to be constructed according to user preference based on the demands of the mathematical model. Figure 1.3 provides schematic diagrams for five different types of agent-based model which we discuss below: cellular automata (CA), cellular Potts models (CP), overlapping spheres (OS), Voronoi tesselation (VT) and vertex-based models (VB) (Osborne et al., 2017).

Cellular automata

Cellular automata (CA) (Van Liedekerke et al., 2015) are one of the oldest and most commonly used types of agent-based model due to their relative simplicity and the ease of implementing them without using external software. The domain is decomposed into a regular lattice, and in general each lattice site is occupied by a single cellular automaton which is able to move around the grid, proliferate into neighbouring lattice sites and occupy space. While their simplicity explains much

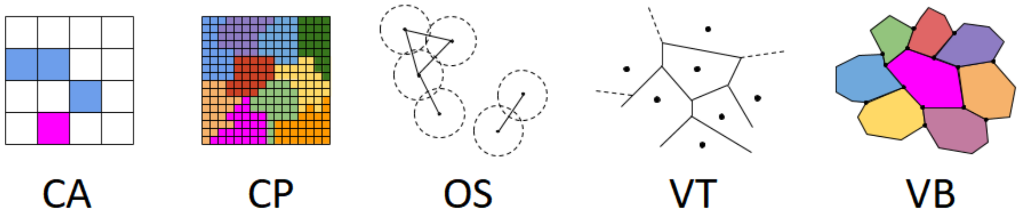

Figure 1.3: Series of schematic diagrams for five different agent-based model frameworks. From left to right: cellular automata (CA), cellular Potts (CP), overlapping spheres (OS), Voronoi tesselation (VT), vertex-based model (VB).
For CA, CP and VB different cells are coloured differently. For OS, approximate cell sizes are given as dashed circles. For VT, lines represent boundaries between neighbouring cells.

of the popularity of CA models, they are unable to represent mechanics associated with morphological changes to a cell and cannot capture nuances of cell movement.

Cellular Potts

Cellular Potts (CP) models (Graner and Glazier, 1992) are lattice-based models in which each cell is represented by a collection of lattice sites. This permits more realistic cell morphologies to be simulated than with a cellular automata, since increasing the number of lattice sites represented by a given cell increases the resolution of the cell boundary and permits more nuanced changes in cell volume than a CA.

Overlapping spheres

"Overlapping spheres" (OS) models are off-lattice models which use a series of nodes to represent cells, with each node corresponding to the centre of a single cell. These points can move around a continuous space according to rules defined by the mathematical model. Cells will interact if they are within a certain distance of each other (Pathmanathan et al., 2009), with neighbouring cells being connected by springs with a specified resting spring length that determine the equations of motion for each cell.

Voronoi tesselation

In the Voronoi tesselation (VT) framework, cell centres are again defined as lattice-free nodes (Pathmanathan et al., 2009). Based on the concept of a Voronoi diagram, each point of the plane (or space) is associated with its nearest node (Meineke et al., 2001). In a VT model, this defines the shape and volume of a cell, and the edges of the cells are always in adhesive contact.

Vertex-based models

A Vertex-Based model (VB) again uses a set of points to represent cells, but instead of each node corresponding to a cell centre they instead define the vertices of a cell (Fletcher et al., 2013). Each cell vertex moves freely according to the rules defined in the mathematical model, which typically includes rules governing cell compressibility and volume. VB models permit the most morphological detail of the models discussed here as an arbitrary number of vertices can be included to increase the resolution of the cell boundary, but they are typically also the hardest of these frameworks to implement and incur greater computational cost since this typically scales with the number of nodes rather than the number of cells.

1.5 Discussion

In this Chapter, we have introduced and reviewed many concepts that will be used throughout this thesis. As this thesis bridges the fields of image analysis, spatial statistics and mathematical modelling, comprehensively reviewing each of these areas is infeasible. As such, we have focussed on reviewing the key ideas that inform subsequent Chapters. Although our thesis integrates these fields, the focus of each Chapter can be broadly categorised as belonging to one of these areas. Chapters 4 and 5 are concerned primarily with mathematical modelling, and we use the model developed in these Chapters to investigate the mechanistic causes of macrophage infiltration into solid tumours. Chapters 2 and 3 introduce and test our image analysis pipeline, and demonstrate how it can be used to generate spatial statistics which can describe macrophage patterns in histology data.

Chapter 2

Development of an image analysis pipeline

2.1 Overview

Our motivation for studying macrophage distributions in solid tumours is the observation that immune cell locations correlate with patient outcome *in vivo* (Chanmee et al., 2014; Pyonteck et al., 2012; Sica et al., 2015; Scholl et al., 1994). Variations in the patterns of macrophage infiltration can be observed in histological images, which provide high-resolution information about the locations of individual cells within a tumour. Despite recent advances in the field of "digital pathology", which aims to develop computational techniques for the analysis of histological images, most analysis of immune cell patterns from histological images is based on qualitative observations made by pathologists. While these observations are based on skilled analysis and understanding of underlying biological processes, automated image analysis provides an opportunity to quantify immune cell patterns, characterise different distributions automatically and provide data which can be used to inform agent-based mechanistic models.

In this Chapter, we develop a versatile pipeline that can process large histology images to reliably obtain cell locations and generate summary statistics for the observed cell distributions. In Section 1.3.1 we described existing software which can be used for whole slide image analysis, and highlighted several features that are required in pathology which these tools are currently unable to meet. Our aim in this Chapter is to develop a tool which is flexible and powerful while remaining user-friendly and accessible to pathologists who may be unfamiliar with programming. Our pipeline can process whole slide images (WSIs) which are several gigabytes in size, and adapts easily to different histological stains and/or intensities. We compare the performance of our tool against the current best practice in pathology laboratories and benchmark its performance in cell counting against counts made by trained pathologists. We use our pipeline in Chapter 3 to obtain detailed spatial summary statistics describing macrophage distribution patterns, and Chapters 2 and 3 together motivate the computational modelling which we undertake in Chapters 4 and 5. The pipeline described in this Chapter is designed to integrate histological images with agent-based models, which we demonstrate in Chapter 5.

The remainder of this Chapter is structured as follows. In Section 2.2 we describe the type of data which is available to us and which our pipeline aims to process, as well as some of the key challenges involved in processing histological slides. This Section introduces two datasets which we refer to throughout this thesis: the Roche Immunohistochemical Tissue Atlas (RITA) data set, and images made available to us by the Pugh Lab at the University of Oxford. Section 2.3 presents the overall structure of our pipeline. It also includes several key algorithms and our rationale for choosing these algorithms in preference to other computer vision techniques described in Section 1.3.2. Section 2.4 describes steps we have taken to ensure that our algorithm correctly identifies immune cells from stained images. We validate the classifiers that we have trained and compare the performance of our pipeline against cell identification by trained pathologists. We then present our procedure for optimising the parameter values used in our cell counting algorithms. Finally, Section 2.5 discusses ways in which our pipeline can be extended to allow deployment by pathologists with no programming experience in laboratory settings.

2.2 Histopathology datasets

In this Section, we introduce the two datasets that we use throughout this Chapter and the remainder of this thesis. Both datasets contain histological sections of *in vivo* tumours resected from human tissue, but are drawn from different tumour indications (types of tumour) and vary in the immunohistochemical markers used for staining. The pipeline described in this Chapter can identify a wide variety of stained cells based on a range of markers, and we use these datasets to demonstrate its broad applicability.

One of the strengths of histological images is the level of detail they contain. High-resolution imaging of tumour slices allows a pathologist to interrogate a tumour at every scale from the micrometre to the millimetre, with intricate details at the scale of cell nuclei available in the context of the surrounding tissue. Such data poses two critical challenges for the automated processing of histological images. Firstly, such images generate files that can be several gigabytes, orders of magnitude larger than a typical photograph. Secondly, a pathologist may be overwhelmed by the quantity of data available in each image. Instead of exploiting all of the detail contained in the image, they must focus on subsections of the tumour which are identified as regions of interest, or summarise the cell-level information via approximate qualitative or semi-quantitative statistics. The growing field of "digital pathology" aims to provide solutions to this second problem, by reducing the level of qualitative human judgement used to assess histological images (Griffin and Treanor, 2017; Williams et al., 2017; Nawaz and Yuan, 2016). Our pipeline provides an automated means of generating quantitative statistics from high-resolution images, and thereby aims to reduce the errors introduced by using human-estimated values for cell counts and area approximations.

2.2.1 Roche Immunohistochemical Tissue Atlas

The Roche Immunohistochemical Tissue Atlas (RITA) dataset was collected by Roche in collaboration with the Bavarian government. The aim was to create a comprehensive resource for understanding immune cell infiltration of tumours. The dataset contains tumours from 240 patients with 4 different indications - breast, non-small cell lung, colorectal, and prostate cancer - with 60 resected tumours for each indication. Each tumour was embedded in a paraffin block and divided into slices as described in Section 1.2.4. Some specimens were pulverised to obtain RNA expression data for the tumour, while the majority were stained with different immune cell markers.

One marker common to large parts of the dataset is the colony stimulating factor-1 receptor (CSF-1R), a member of the receptor protein tyrosine kinase (rPTK) family of growth factor receptors which is expressed by macrophages and monocytes (Rovida and Sbarba, 2015; Ries et al., 2014). CSF-1 mediates the survival, proliferation and differentiation of monocytes into macrophages, and is commonly used as a generic macrophage marker (O'Brien et al., 2010; Stanley et al., 1997). Multiple slides stained using an anti-CSF-1R antibody are present in the datasets for breast, colorectal and prostate cancer, resulting in approximately 540 images across these three indications. Each image is between 3GB and 5GB in size and has a resolution of approximately 0.2 μm per pixel, resulting in images on the order of 10^{10} pixels. Figure 2.1 shows a typical image taken from the RITA dataset. Figure 2.1a is a whole slide image taken from a colorectal cancer sample. It is stained to show macrophages which express CSF-1R, hereafter denoted "CSF1-R^{+}" macrophages, (brown) and counterstained to show all cell nuclei (blue). Tissue scale features, such as colorectal crypts and tumour nests, can be identified from images at this resolution. Figure 2.1b shows a magnified detail from the same sample in which individual macrophages and cell nuclei can be resolved.

Working with the RITA dataset has several significant benefits for the work presented in this thesis. It provides a large sample of data, allowing us to compare macrophage distributions across different patients with the same type of tumour and between groups of patients with different indications. The quantity of data is also beneficial when considering machine learning approaches to image analysis as described in Section 1.3.2, as these techniques often require a large quantity of training data to generate classifiers which perform well. Another advantage is the use of the CSF-1 receptor as the macrophage marker. While a range of markers are known to be expressed by most macrophages, including CD68 and CD63/CD161 (Mantovani et al., 2017), CSF-1R is a target for treatment with emactuzumab, an anti-CSF-1R antibody currently in phase I trials at Roche (Fujimura et al., 2018; Pradel et al., 2016; Cassier et al., 2015; Pradel and Ries, 2015) as described in Section 1.2.3. The small subset of macrophages not identified by the CSF-1R stain in the RITA data will therefore not respond to emactuzumab treatment. Thus the RITA data identifies macrophages that may respond to treatment with emactuzumab.

On the other hand, the RITA dataset poses several challenges. CSF-1R is expressed by macrophages with a range of phenotypes, and so cannot be used to distinguish between different phenotypes. Anti-CSF-1R therapies such as emactuzumab are designed to bias macrophage populations towards an anti-tumour, M_1, phenotype. However, it is likely that the response of a particular macrophage to treatment will depend on its phenotype.

A related disadvantage is that the RITA data does not contain information about the oxygen levels in the tissue. Since macrophage phenotype is heavily influenced by the local tumour microenvironment, it may be possible to infer macrophage phenotype from their location within tumours. However, without knowledge of either the oxygen environment or macrophage phenotype markers, we can only identify CSF-1R$^+$ macrophages from the RITA data.

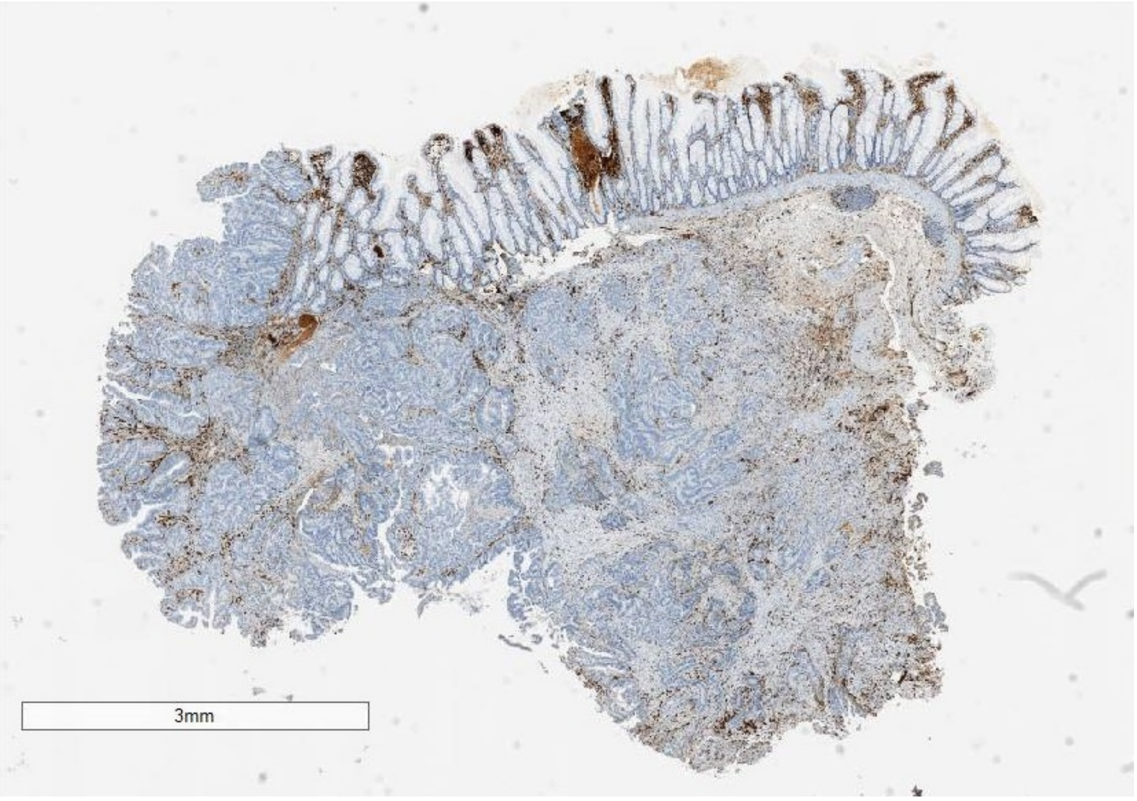

(a) Whole slide image

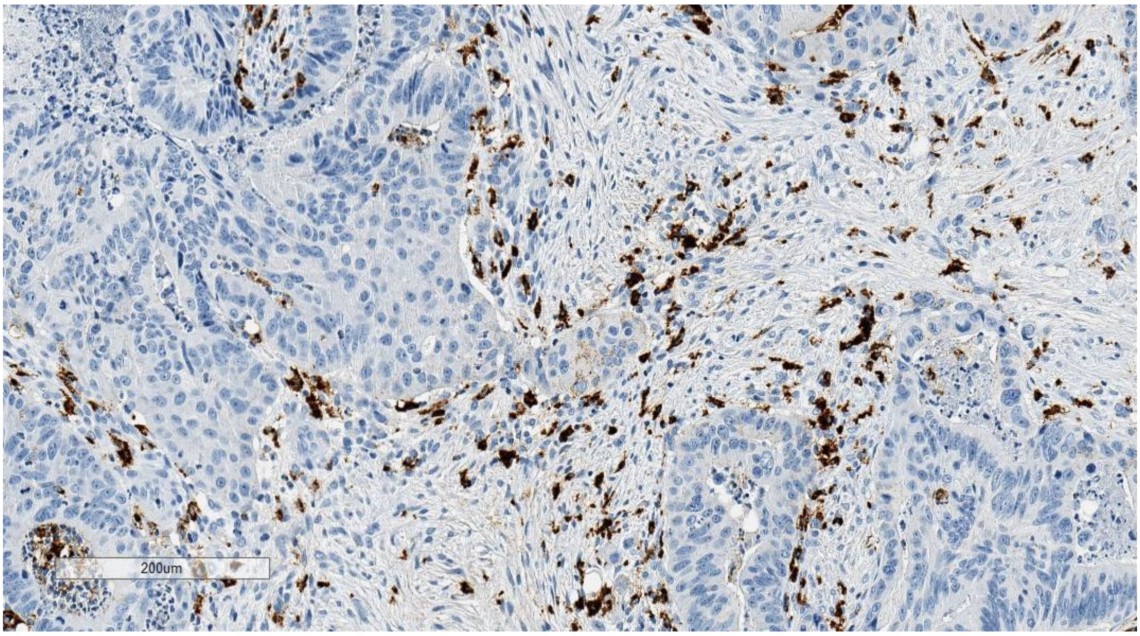

(b) Magnification (20 times)

Figure 2.1: Whole slide (top) and detail (bottom) from a colorectal cancer sample from the RITA dataset. The brown stain shows CSF-$1R^+$ macrophages, with blue counterstain showing cell nuclei.

2.2.2 Pugh Lab, Nuffield Department of Medicine (Oxford)

A second dataset is available through collaboration with Dr Philip Macklin and Professor Christopher Pugh in the Nuffield Department of Medicine at the University of Oxford. This dataset consists of 16 resected tumours from cases of head and neck cancer. A stack of approximately 16 consecutive images is available for each tumour, with stains for five different immune cell subtypes (CD4$^+$ T-cells, CD8$^+$ T-cells, CD68$^+$ macrophages, FoxP3$^+$ T-cells and CD19 B-cells). The dataset also contains stains for pancytokeratin (PanCK), an epithelial cell stain which, in the context of head and neck cancer, can be used to distinguish tumour regions from stroma.

An important feature of the Pugh lab dataset is the presence of multiple markers for hypoxia. Several slides from each tumour are stained to show expression of carbonic anhydrase IX (CAIX), a transmembrane protein expressed by cells in environments with low oxygen concentrations which is an endogenous marker of hypoxia (Olive et al., 2001). The dataset also contains a second hypoxia marker, pimonidazole, which is often considered to be a more accurate marker of extreme hypoxia than CAIX (Rademakers et al., 2011; Varia et al., 1998). CAIX is more commonly used in practice since it is an endogenous marker for hypoxia, while pimonidazole must be injected intraveneously into the patient before the tumour is surgically removed. Pimonidazole binds with thiol-containing proteins specifically within hypoxic cells, and can be detected immunohistochemically using monoclonal antibodies through the same mechanisms as other histological stains described in Section 1.2.4 (Varia et al., 1998). The difference in localisation between pimonidazole and CAIX can be seen in Figure 2.2, which shows detail from one of the head and neck tumours in this dataset. The images come from two consecutive slides, one stained to show pimonidazole (Figure 2.2a) and one to show CAIX (Figure 2.2b).

One major advantage of the Pugh lab dataset is that although each slide is stained to show a single marker only, for each tumour the dataset contains a stack

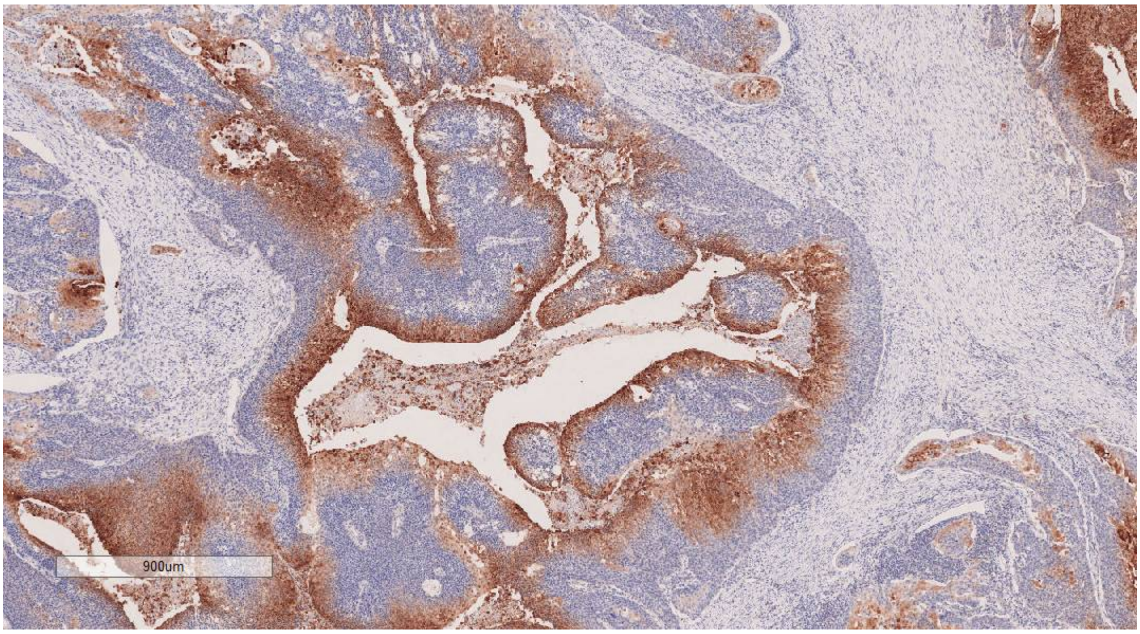

(a) Pimonidazole expression in a head and neck tumour

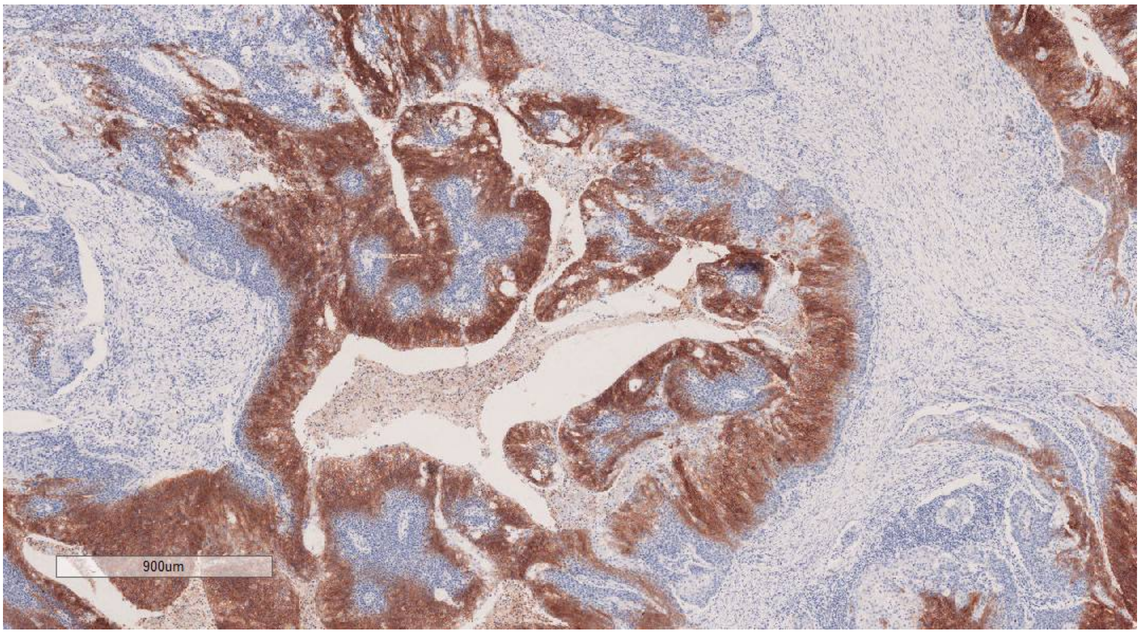

(b) CAIX expression in a head and neck tumour

Figure 2.2: Expression of hypoxia markers (pimonidazole and CAIX, brown) around a necrotic region of a head and neck tumour, taken from consecutive tumour slices. CAIX is expressed in more cells than pimonidazole, although both stains are concentrated around the edge of the necrotic tissue.

Dataset	**RITA**	**Pugh Lab**
Indication	Breast, colorectal, non-small cell lung, prostate	Head and neck
Macrophage staining	CSF-1R	CD68
Hypoxia staining	None	Pimonidazole, CAIX
Consecutive slides?	Unavailable	Available
Number of tumours	60 (per indication)	16

Table 2.1: Comparison between the RITA dataset and the Pugh Lab dataset.

of adjacent slides which are separated by a distance of around $4\,\mu$m. Consecutive slides can be registered to form composite images containing immune cells, hypoxia markers and tumour/stroma markers. We can therefore visualise macrophages in the context of the surrounding tumour microenvironment in a more detailed way than the RITA data permits, as the RITA dataset does not contain consecutive slides which can be coregistered to investigate the colocalisation of different stains.

Table 2.1 contains a brief comparison of features of the two datasets.

2.3 Image analysis pipeline

In this Section, we describe our image analysis pipeline. The pipeline takes as input an image, which can be a region of interest selected by a pathologist or an entire pathology slide, and returns the coordinates of identified cell centres as well as simple summary statistics generated from the image, which can be used for subsequent further analysis.

Figure 2.3 shows the basic structure of our pipeline. We first compress the image while retaining cell boundaries by generating superpixels, a process which we describe in Section 2.3.1. We then assemble a collection of features associated with each superpixel and pass them to an appropriately trained classifier. The classification process described in Section 2.3.2 uses support vector machines, although

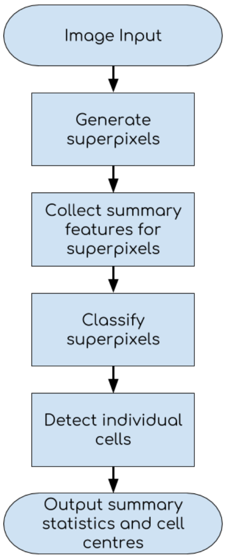

Figure 2.3: Flowchart showing how our pipeline can be used to generate summary statistics and cell locations from a histological slide.

other classification techniques described in Section 1.3.2 can be used instead. Once each superpixel has been classified, the resulting map is passed to a watershedding algorithm which detects individual cells and is described in Section 2.3.3. Finally, we output summary statistics generated from the classification and cell detection data.

2.3.1 Superpixel generation

Although high resolution is necessary to identify individual cells in histological images, it presents challenges in image processing due to the large file size (see Section 2.2). In Section 1.3.2 we introduced the concept of superpixellation as a means of compressing images while simultaneously segmenting boundaries, and described different algorithms for generating superpixels. Rather than trying to identify whether individual pixels belong to a cell or not, we instead use superpixels to reduce the information content of each image while retaining important information about object boundaries and colours. We then assign classifications to superpixels rather than their constituent pixels.

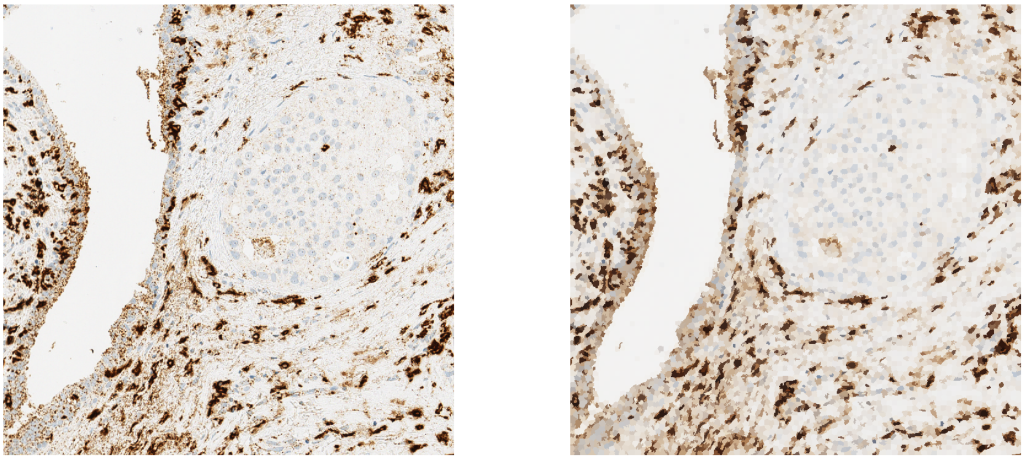

(a) Original Image

(b) SLIC superpixel segmentation

Figure 2.4: Application of SLIC superpixel algorithm to a 1100 × 1100 pixel tile from the RITA data, with number of superpixels chosen to generate an average superpixel size of 150 pixels. Figure 2.4a contains 1,210,000 pixels while Figure 2.4b consists of 8,083 superpixels, each coloured according to the average colour of its constituent pixels. The image is from a breast cancer sample stained for CSF-1R, showing macrophages in brown.

Our pipeline is designed to be fully configurable to use any superpixellation algorithm to generate superpixels. However, in this thesis we use the SLIC (Simple Linear Iterative Clustering) superpixel algorithm (Achanta et al., 2012) due to a range of benefits it exhibits over other algorithms. In their comparison of superpixellation methods, Stutz et al. (2018) recommend SLIC and highlight it as being an especially fast algorithm for which runtime is broadly independent of the number of superpixels. It also provides good control over the size of superpixels and generates superpixels which adhere well to the boundaries of objects in images. In Figure 2.4 we compare a representative result of the output of the SLIC algorithm, in which each superpixel has been coloured according to the mean colour value of the pixels it contains, with the original image. The image used was a randomly chosen section of a breast cancer slide taken from the RITA dataset, with brown staining marking the CSF-1 receptor found on macrophages.

When applying the SLIC algorithm, the relative weighting of distance and colour can be varied to alter the degree of compactness of the superpixels, giving control over how elongated superpixels will become before a thin shape is split into two superpixels. The number of superpixels n within an image is also specified as an input. This leads to superpixels of the same average size - if there are P pixels in the image, each superpixel will contain approximately $^P/_n$ pixels. The ability of SLIC to control the approximate size of superpixels is a useful property for segmenting histology data, since the features of interest (the cells) are typically similarly sized. Furthermore, the degree of compactness can be controlled making SLIC well-suited for generating rounded superpixels which mimic the circular shapes formed by most cells in 2D histology images. Based on the average approximate size of a cell we use $^P/_n = 150$. In Figure 2.4 $P = 1,210,000$, and so we specify $n = 8,066$ initial seeds for the superpixellation. After 10 iterations of SLIC, the resulting superpixel image contains 8,083 superpixels. The edges of individual macrophages are still evident in Figure 2.4b, yet by applying the classifier to these superpixels we can classify 150 times fewer objects than a classifier that is applied to the original image in Figure 2.4a.

Superpixellation can be viewed as an oversegmentation process (Stutz et al., 2018). Since SLIC superpixellation generally preserves object boundaries, we use SLIC to oversegment the image without distinguishing individual cells. We do not rely on superpixellation to identify the boundary between two cells which are touching, but to ensure that stained regions and unstained regions are not included in the same superpixel. While correctly classifying superpixels provides a good first approximation for identifying cell boundaries, as long as we have classified superpixels to identify positively stained regions then we can identify individual cells at a later stage (described in Section 2.3.3).

Although the length of time required to run the SLIC algorithm scales linearly with the number of pixels, high-resolution histology slides are several orders of magnitude larger than the images against which superpixellation algorithms are typically benchmarked (Stutz et al., 2018). The resolution of a histological slide scan is typically 0.2205 microns per pixel, with a tumour section of just 3.5 cm × 2 cm resulting in an image containing approximately 1.4×10^{10} pixels. Even with superpixellation, a large number of objects must still be classified. To increase the speed of superpixel generation for whole slide images, we exploit the local nature of SLIC described in Section 1.3.2 to parallelise the process. We decompose each slide into "tiles" of approximately 1000×1000 pixels which can be segmented in parallel using multiple cores.

The superpixellation process can be further accelerated by using lower-resolution layers of the histology slides to generate the superpixels. While the maximum resolution of scans is 0.2205 microns per pixel, whole slide images are stored at multiple resolutions. This enables slide viewing software to display the whole slide without having to store it in memory at full resolution. When a user zooms in on a portion of an image, the relevant section can be returned at an appropriately chosen resolution to minimise data storage. Where present, we use a lower-resolution layer of the image file to generate superpixels (0.882 microns per pixel), resulting in a further 16-fold decrease in the number of pixels required to generate a superpixellated image. If no layer of the image is stored at this resolution then we rescale the next highest resolution image to reach a new resolution of 0.882 microns per pixel.

2.3.2 Superpixel classification with support vector machines (SVM)

Machine Learning is fast becoming the most popular technique for computer vision. In Section 1.3 we discussed a range of techniques which have been used to identify cells in images. As with the choice of superpixellation algorithm, the flexible nature of our pipeline means that the user can specify their own classifier trained using any of these approaches. In this thesis, we present results generated from classifiers which we have trained using support vector machines (SVMs) (see Section 1.3.2).

While popular techniques such as deep learning can produce exceptional results (Hernández et al., 2018; Janowczyk and Madabhushi, 2016; Sirinukunwattana et al., 2016), we note that SVM approaches can outperform such approaches at detecting cells while being more flexible, less computationally expensive and faster to train (Romero Castro et al., 2018).

SVMs are a form of supervised learning, in which labelled training data is used to generate binary classifiers. This makes them well suited to cell identification, in which each superpixel will belong to only one class (stained or unstained). They rely on pre-specified features to be defined to make predictions. Selecting these features is straightforward in our application as the difference between stained and unstained cells is determined by their colour.

During SLIC superpixel segmentation, each pixel is iterated over several times before it can be assigned to a final superpixel. We exploit this to collect features describing each superpixel without increasing the number of times over which each pixel must be iterated, which would significantly increase the time taken for the pipeline to process an image. These features can then be used to guide classification. One feature of interest is the mean value of each of the colour channels within each superpixel. The SLIC algorithm generates superpixels using the CIELab colour space, with three colour channels L, a and b representing lightness, the red/green axis and the blue/yellow axis. Calculating the mean and variance of each colour channel within a superpixel allows us to assign a "mean" colour to each superpixel, reducing the impact of noise in the image and providing a simple means of visualising superpixellated images. In this thesis we collect only 18 features for each superpixel:

- Feature 1-3: Mean L, a, b values of the superpixel,
- Feature 4-6: Variance of L, a, b colour channels within the superpixel,
- Feature 7-12: Mean and Variance of L, a, b channels within neighbouring superpixels,
- Feature 13-18: Mean and Variance of L, a, b channels within neighbours of

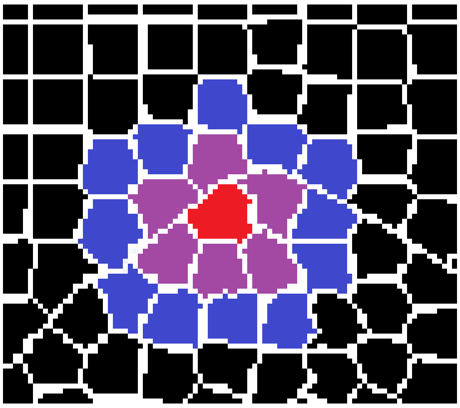

Figure 2.5: The superpixels shown in purple are neighbours of the target superpixel (red). Blue superpixels are neighbours of neighbours for the target superpixel, since they share an edge with a superpixel in the set of neighbours.

neighbouring superpixels.

Using these features ensures that the features of a superpixel account for colour information in the environment of the superpixel as well as the colours of the superpixel itself. We define neighbouring superpixels as being those which share an edge with the target superpixel, and neighbours of neighbours are those sharing an edge with a neighbour (excluding those superpixels in the set of neighbours and the target superpixel itself). An example demonstrating this can be seen in Figure 2.5.

Our list could be extended to include more complex features. Histogram of Oriented Gradients (HOG) features use histograms of local gradient angles to capture the texture of image sections, and have been used in conjuction with SVMs to identify humans in images (Zhu et al., 2006) as well as to identify cells in biomedical images (McCann, 2015; Ergin and Kilic, 2014). SIFT features (Scale Invariant Feature Transform) (Lowe, 2004; Zhang et al., 2007) are also widely used in object identification in medical imaging (Ergin and Kilic, 2014; McCann, 2015). It is likely that the inclusion of such additional features could improve our classification, and we postpone this investigation for future work. We have chosen to use only features directly linked with the colour channels of the histology images as we expect colour to be the most important factor in identifying stained cells, the features are easy to conceptualise for non-programming pathologists and so can be justified easily, and our results indicate that the performance of classifiers trained solely on these features is sufficient for fast and accurate classification of most cells. Since a pathologist can identify cells based on their colour from an image we predict that features based on colour should be sufficient to identify the relevant cells.

A strength of using a preset collection of features is that our SVM classifiers can be trained using far fewer training examples than would be required to train, for example, a deep neural network. Furthermore, a piece of training data takes the form of a single superpixel, and we require only the 18 features specified above in order to train the SVM using that point. In contrast, while a deep neural network may be more powerful, a typical piece of training data would be a section of the image taken around the relevant superpixel from which the network would determine key features.

We therefore require high-quality annotated datasets on which to train and test classifiers. As we wish to classify superpixels rather than pixels, we must first annotate a wide range of superpixels covering a range of tumour cases. To facilitate this process, our pipeline uses a Graphical User Interface (GUI) which can be used to rapidly generate training data from a set of images. Figure 2.6 shows a typical screenshot of our GUI in MATLAB. While MATLAB provides an accessible user interface for selecting superpixels, the underlying image manipulation (including superpixel generation and feature collection) is implemented in C and linked with MATLAB via MEX files. MEX files provide an interface between MATLAB and C, which can increase the processing speed by significant amounts (Andrews, 2012).

Our GUI allows new SVM classifiers to be generated quickly by non-specialist users. For example, in a 15 minute session, a trained pathologist can label up to 1000 superpixels - testing has shown that acceptable classifiers can be generated using far less data than this, although further testing is required to determine the minimum amount of data needed to train an SVM. This means that using our pipeline, pathologists who wish to analyse images using newly developed staining

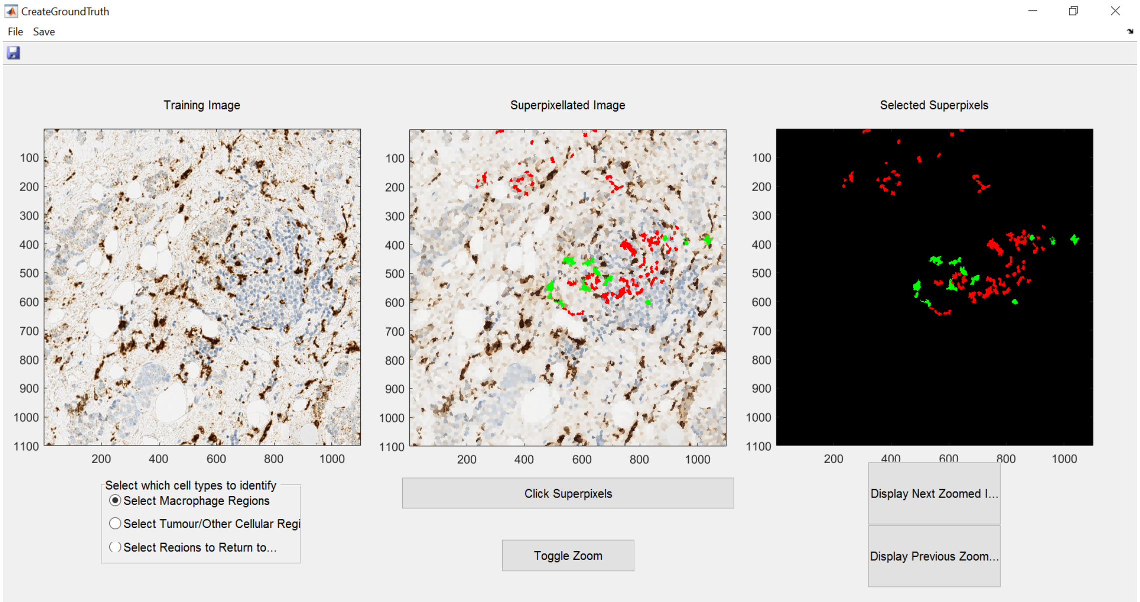

Figure 2.6: Graphical User Interface (GUI) designed in MATLAB for fast generation of "ground truth" data for training or testing of classifiers. Users can see the original image section alongside the superpixellated version, and are able to select either image pixels or superpixels to be labelled as different stains. In this example, green pixels have been labelled as CD68+ staining and red pixels have been labelled as counterstained cell nuclei.

protocols can quickly create their own classifiers and analyse images.

2.3.3 Watershedding for cell identification

After classification of superpixels, we are left with a "mask" identifying positively stained pixels in the image. In our pipeline, we use an average superpixel size of 30 pixels (23.34 μm^2). Human macrophages have a cross-sectional area of approximately 400 μm^2 (Krombach et al., 1997), meaning that each cell will be divided into multiple superpixels, a process called oversegmentation. This increases the confidence that superpixellation will not distort cell boundaries, but means that further processing is required to identify individual cells from the superpixels identified as positively stained. We use watershedding (Meyer, 1994) to divide this classification mask into individual cells and determine the approximate centres of each cell. Watershedding takes its name from the geographical definition of a watershed as a ridge of land that marks the divide at which water will flow into different valleys.

By applying a distance transform, we consider the classification mask as a landscape in which the depth at a given pixel inside the positive classification is proportional to the euclidean distance of that pixel from the edge of the classified region. Cell boundaries can then be identified as the ridges in this landscape marking the divides between basins, and the cell centres are classified as the deepest point in each basin. Watershedding has been successfully used to distinguish between clustered cell nuclei in fluorescence microscopy (Malpica et al., 1997), for the automated counting of red blood cells (Chourasiya and Rani, 2014) and for the detection of lymphocytes in Hematoxylin and Eosin (H&E) images (Romero Castro et al., 2018).

We extend our watershedding algorithm by introducing two parameters, the *opening size* and the *threshold*. The minimum size threshold determines the minimum area of a detected cell. Any cells identified which have an area less than this threshold are assumed to be noise and discarded. The opening size is a parameter used during the process of morphological opening, a process which uses a predefined *structural element* to reshape the boundary of a region (Gonzalez and Woods, 2008). By choosing the structural element to be a disc, morphological opening causes thin or pointed regions to be flattened or removed from the region. We apply morphological opening to the classification mask before watershedding to help separate cells which are touching. The parameter *opening size* is the radius of the disc used as a structural element, so a larger opening size will result in smoother boundaries and result in more of the region being removed. In the following Section, we investigate the sensitivity of our watershedding algorithm to these parameters and identify optimal parameters for accurate cell counting.

Figure 2.7 shows key stages of this procedure applied to a portion of a $CD8^+$ stained histology slide.

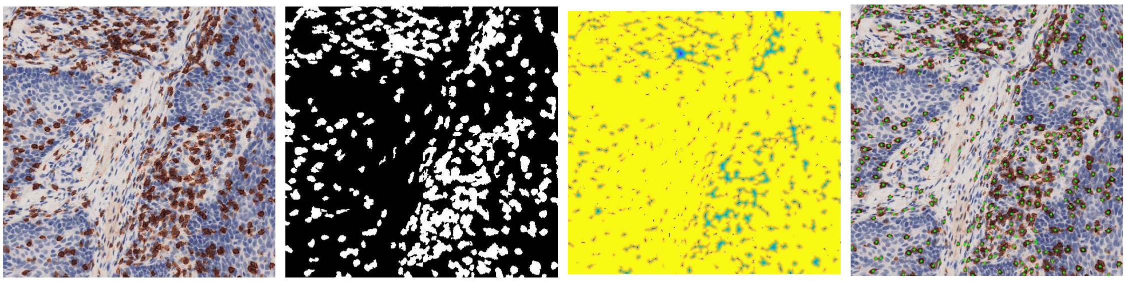
(a) Section of initial image. CD8+ cells are shown in brown.

(b) CD8+ staining identified by a classifier.

(c) Distance transform for watershedding-based segmentation of CD8+ cells.

(d) Predicted CD8+ cell centres (minima of watershedding catchment basins), shown as green dots.

Figure 2.7: Sequence of images illustrating how the watershedding process works. (a) Section of a slide stained for $CD8^+$ T cells; (b) pixels identified as positively stained (white); (c) distances from the edge of the classified boundary used for watershedding. Blue areas indicate staining that is more distant from the stain boundary, and yellow indicates unstained areas; (d) cell centres identified via watershedding.

2.4 Pipeline validation

The process of model validation tests how well a mathematical or computational model captures reality (Cooper et al., 2015; Babuska and Oden, 2004; Peirce, 2008). We must validate our image analysis pipeline to ensure that objects identified as positively stained cells by our algorithms are real cells. In this Section, we demonstrate that our pipeline accurately identifies labelled cells from histological images (true positives) and returns few false positives, background pixels misidentified as being cells. We further demonstrate that our pipeline performs at least as well as trained pathologists at counting cells in images, while being faster and able to generate reproducible results.

2.4.1 Performance of classifiers

In Section 2.3.2 we highlighted the ability of our pipeline to quickly generate new classifiers from a set of initially unlabelled images. The performance of each classifier will depend on the quality of the training data and the diligence of the human who labels the data. As such, the performance of each new classifier must be evaluated when it is created. In this Section, we apply several techniques to evaluate the performance of one specific classifier, but note that these techniques can be easily applied to assess the performance of new classifiers.

The results presented in this Section apply to a classifier trained by Dr Philip Macklin from the Nuffield Department of Medicine, University of Oxford, on images of CD68$^+$ cells. Both CD68 and the CSF-1 receptor are transmembrane proteins expressed on the cell surface of macrophages, and classifiers respond similarly to both stains due to their similarity.

Figure 2.8 shows the Receiver Operating Characteristic (ROC) curve for this classifier (Fawcett, 2006). The ROC provides a way of visualising the relationship between the *sensitivity* and *specificity* of a classifier. The sensitivity of a classifier is the true positive rate, (TP), and measures the proportion of labelled pixels which have been correctly identified as labelled:

$$\text{sensitivity} = TP = \frac{\text{Number of pixels correctly identified as stained}}{\text{Number of stained pixels}}. \tag{2.1}$$

The specificity is the true negative rate, (TN), which measures the proportion of unlabelled pixels which are correctly classified as unlabelled:

$$\text{specificity} = TN = \frac{\text{Number of pixels correctly identified as unstained}}{\text{Number of unstained pixels}}. \tag{2.2}$$

ROC curves are generally shown in terms of the true positive rate, defined in Equation (2.1), and the false positive rate, FP, where:

$$FP = 1 - \text{specificity} = \frac{\text{Number of unstained pixels identified as stained}}{\text{Number of unstained pixels}}. \tag{2.3}$$

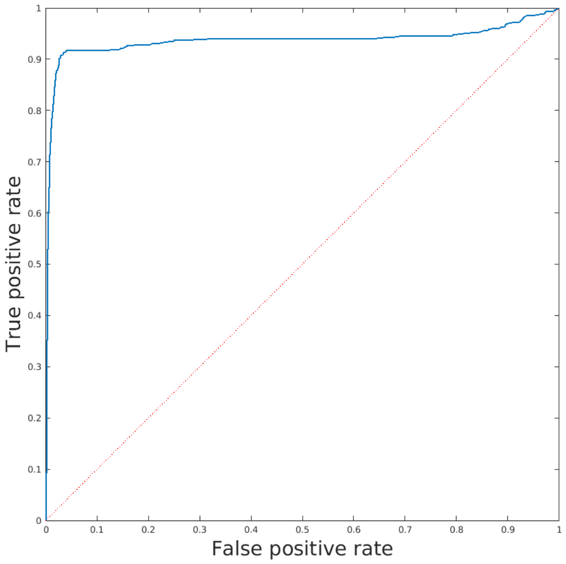

Figure 2.8: Receiver Operating Characteristic (ROC) for our macrophage classifier (blue), showing the true and false positive rates as the threshold parameter varies. The red dotted line shows the ROC for a classifier which randomly allocates superpixels to either class. The ROC for a perfect classifier would pass through (0,1), with no false positives and all true positives correctly identified.

For a binary classifier, the ROC curve is parametrised by a threshold of acceptance. A classifier assigns each pixel (or superpixel) in an image a score which can be thought of as the confidence of the classification. At one extreme, we may require that the classifier is 100% confident that it has detected a labelled pixel before it assigns it to the positive class. This condition will be so stringent that every pixel is classed as a negative. In this case, we will have no false positives, but our true positive rate will also be zero. At the other extreme, we may relax our tolerance so far that every pixel is accepted as labelled by the classifier. In this case we will correctly identify every labelled pixel and have a true positive rate of 1, but our false positive rate will also be 1. A perfect classifier would find every positive pixel and none of the negative ones, returning a true positive rate of 1 and a false positive rate of 0. A classifier which randomly accepts pixels as either positive or negative should identify false positives and true positives at the same rate.

As we vary our threshold parameter from the extreme of rejecting all pixels to the extreme of accepting them all, we obtain a ROC curve such as the one in Figure

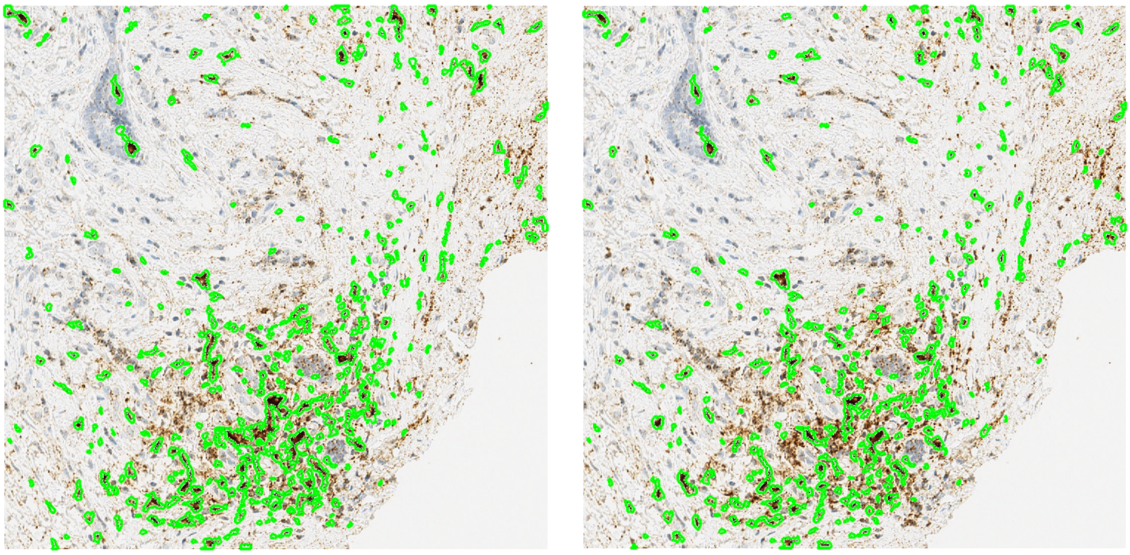

(a) Ground truth labelling of a tile used as validation data.

(b) Prediction of superpixel classifications made by an SVM trained with our pipeline.

Figure 2.9: Comparison of ground truth classification with the SVM predicted classification. Green lines mark the boundaries of superpixels classed as macrophages.

2.8. If the ROC curve is above the red dashed line $y = x$, then the classifier is performing better than random chance for that value of the threshold parameter. The optimal choice of thresholding parameter will be the point on the ROC closest to the upper left corner, $(0, 1)$, and this is the value that should be chosen for classification purposes. For a random classifier, the area under the ROC curve (AUC) will be 0.5, while a perfect classifier will have AUC = 1.

The classifier considered in Figure 2.8 performs far better than random. Indeed, without optimisation our classifier has a true positive rate of 0.749 and a false positive rate of only 0.01. At the optimal point on the ROC, this classifier achieves a true positive rate of 0.917 with a false positive rate of 0.041. These values indicate that our classifier can accurately identify the majority of pixels which are part of a macrophage while incorrectly misclassifying a minimal number of non-macrophage pixels.

In Figure 2.9 we compare the hand-labelled ground truth classification (Figure 2.9a) and the prediction of superpixel classifications obtained from our binary SVM

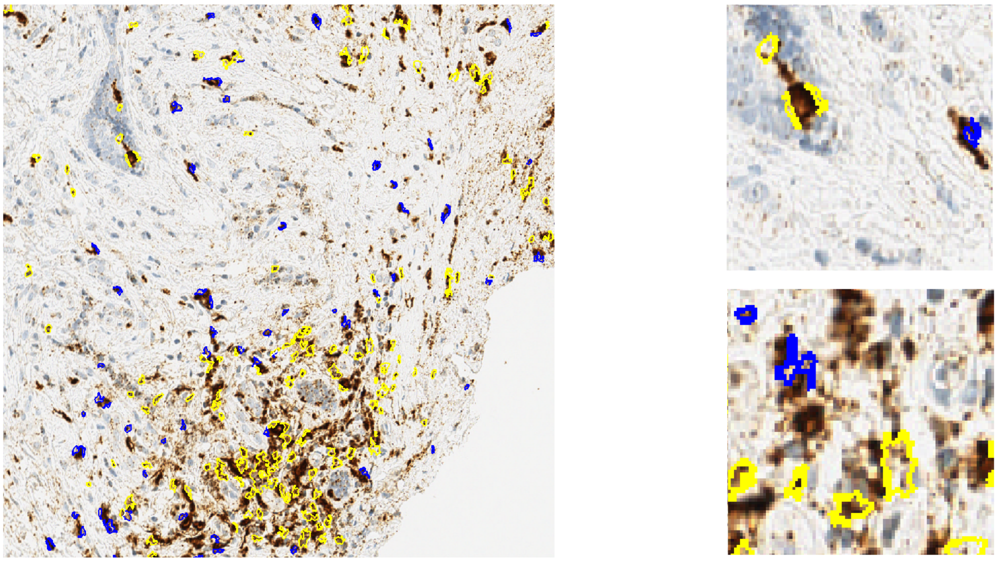

(a) False Negatives (Blue) and False Positives (Yellow) for macrophage classification

(b) Higher magnification details from Figure 2.10a

Figure 2.10: Classification errors for the classification of macrophages (detail of two areas shown on the right). Areas highlighted in yellow have been incorrectly classified as macrophages, while areas in blue contain macrophages but have been misclassified as background.

classifier (Figure 2.9b), in which the "green" labelling is the classifier described by the ROC in Figure 2.8. While these classifications appear similar to the eye, in Figure 2.10 we highlight the false positives and false negatives (labelled pixels which have been missed by the classifier).

Figure 2.10b shows enlarged details of two areas of Figure 2.10a which reveal that several of the misclassifications made by the classifier may in fact be flaws in the manual ground truth classification. Many of the misclassified false positive regions occur along the edges of macrophages, and are frequently long, thin superpixels that may have been misclassified when generating the data. Similarly, some of the false negative regions are small groups of pixels at the edges of macrophages which may easily be misclassifications in the ground truth.

Regardless of problems which may be caused by human error, the classifier appears to perform well at identifying stained macrophages in the sample image presented here.

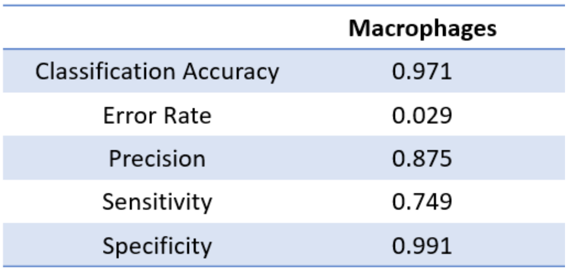

	Macrophages
Classification Accuracy	0.971
Error Rate	0.029
Precision	0.875
Sensitivity	0.749
Specificity	0.991

Table 2.2: Summary statistics for the nonlinear SVM classifier produced by our pipeline without optimisation. Each statistic is measured on a scale from 0 to 1, with 1 being the highest. An ideal classifier would have an accuracy, precision, sensitivity and specificity of 1 and an error rate of 0.

In Table 2.2 we present summary statistics which demonstrate the ability of our classifier to accurately identify macrophages. The sensitivity and specificity are defined in Equations (2.1) and (2.2) respectively. The other metrics shown here are:

$$\text{precision} = \frac{\text{Number of true positives}}{\text{Number of true positives} + \text{Number of false positives}} \tag{2.4}$$

which is also known as the positive predictive value. An ideal classifier would have precision = 1, indicating that very few false positives are returned. The error rate:

$$\text{error rate} = \frac{\text{Number of false positives} + \text{Number of false negatives}}{\text{Total number of pixels}} \tag{2.5}$$

is the proportion of pixels which are misclassified. An ideal classifier would have an error rate of 0.

2.4.2 Optimising cell identification parameters

As described in Section 2.3.3, our watershedding algorithm relies on two parameters, the opening size and a minimum size threshold. In this Section we investigate the sensitivity of the watershedding algorithm to variation in these parameters.

We use a dataset containing 20 images of stained cells not used to train our classifier. Each image was shown to three pathologists, who manually counted the number of stained immune cells visible to them without consultation. This provided three independent collections of cell counts for comparison with the results from our classifier. We investigated opening sizes in the range [0, 20] pixels and minimum size thresholds in the range [0,100] pixels. For each pair of parameters, we counted the number of cells identified in the 20 images. The total squared error between the automatic count and the average of the three pathologists' counts was then calculated and taken as a measure of the goodness of fit.

Figure 2.11 shows how this error changes as the two parameters are varied. In Figure 2.11a we vary the minimum size threshold in increments of 5 and the opening size in increments of 2, giving a coarse view of how they affect the error. We observe that in some regions of parameter space the predicted number of cells differs greatly from the human counts, but that for a large range of values the count is insensitive to small changes in these parameters. To obtain a more refined estimate of the optimal parameter set, we repeated the test using the same dataset for more refined increments of the parameters on a smaller range of values based on the region around the best fit of Figure 2.11a, shown in Figure 2.11b (opening size: [0,8] pixels in increments of 1, minimum size threshold: [25,60] pixels in increments of 1). The error remained similar in size throughout this parameter sweep, again indicating the insensitivity of our watershedding algorithm to these parameters, but indicated that the optimum parameter values correspond to an opening size of 5 pixels and a minimum size threshold of 30 pixels.

Repeating this test for different stains revealed that for all stains there is a range of parameter values for which the fit was comparable with that of the optimum, and hence that the cell counting algorithm is insensitive to these parameters. However, the range of the optimal parameter values depended on the type of stain used. For nuclear stains, the minimum size threshold was approximately half that of cells stained with a membranous marker (38 pixels vs 72 pixels, or $28\,\mu\text{m}^2$ vs $63\,\mu\text{m}^2$).

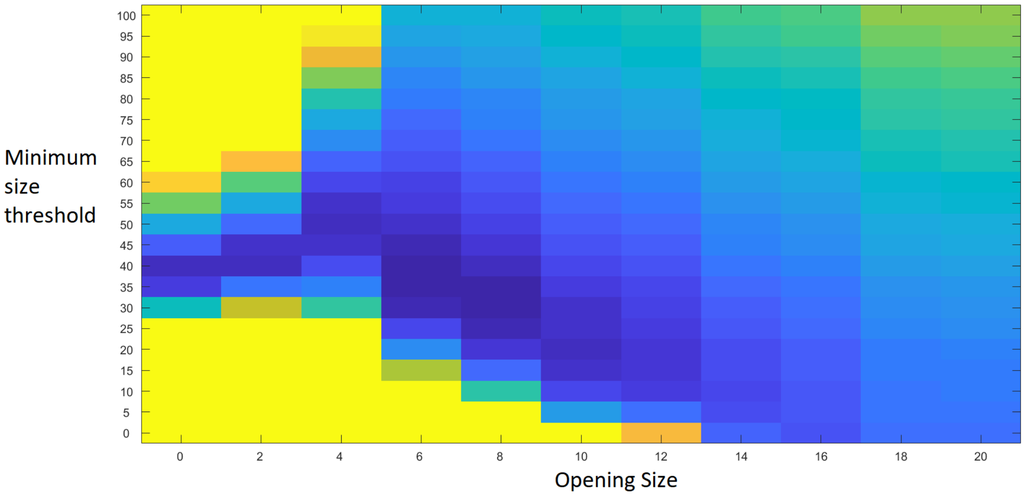

(a) Exploration of parameter space for the two watershedding parameters: the opening size and the size threshold below which candidate cells are discounted as noise.

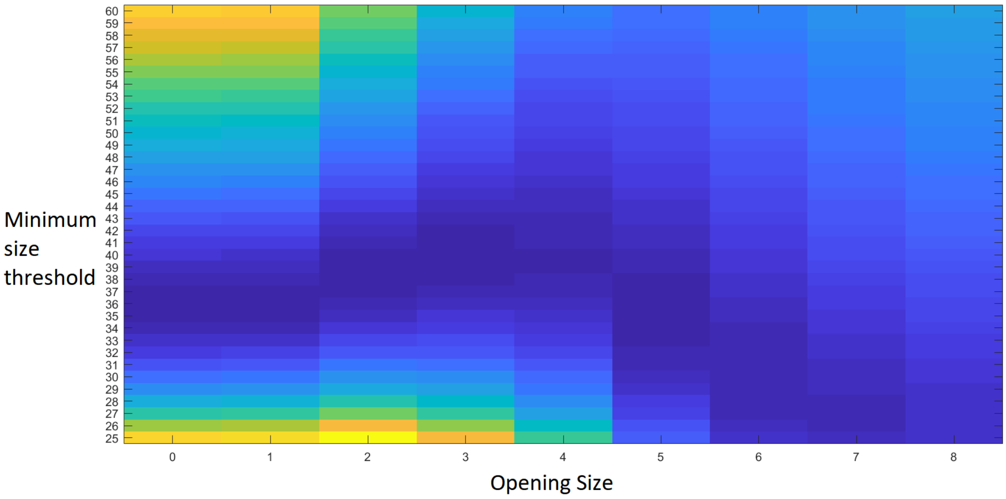

(b) More thorough exploration of parameter space for the watershedding parameters.

Figure 2.11: The colour represents the total squared error between human counts and algorithm counts for 20 images with cell counts provided by three pathologists, with blue representing low error and yellow high. The total predicted cell count is not very sensitive to the choice of parameters in the watershedding algorithm.

We obtained optimal values for CD8$^+$, CD68$^+$ and FoxP3$^+$ stains which we use as default values for watershedding images stained for nuclear, cell membrane and cytoplasmic stains respectively.

2.4.3 Comparison of pipeline performance against trained pathologists

We have also validated the accuracy of our pipeline by comparing the results for cell counts with cell counts obtained by trained pathologists. Cell counting is a common requirement in pathology (Veta et al., 2014), as many clinically relevant metrics rely on properties such as the density of cells per mm^2 of tissue. However, accurate cell counts require a pathologist to examine tissue in a great level of detail and are therefore extremely time consuming and prone to error. Errors can arise because the pathologist has lost count or double counted cells, or may be due to disagreements between pathologists about whether a mark is an artefact or part of a cell.

We used a sample of forty ($1\,\text{mm} \times 1\,\text{mm}$) regions taken from samples of head and neck tumour tissue to test the cell counting ability of our pipeline against three trained pathologists. Each region contained up to several hundred immune cells, and the pathologists independently counted them using an ImageJ app to prevent double counting and to ensure that no cells were overlooked. We applied our pipeline to the same images using our standard CD8$^+$ classifier and the watershedding parameters obtained using the procedure described in Section 2.4.2.

In Figure 2.12 we compare the cell counts obtained by our pipeline (solid black line) and the range of one standard deviation around the mean of the pathologist counts. The images are ordered according to the amount of variation between the pathologists counts. In general, there was more discrepancy between the pathologists' counts as the number of immune cells to be counted increased. The pathologists also reported that for images with a large number of cells the analysis could

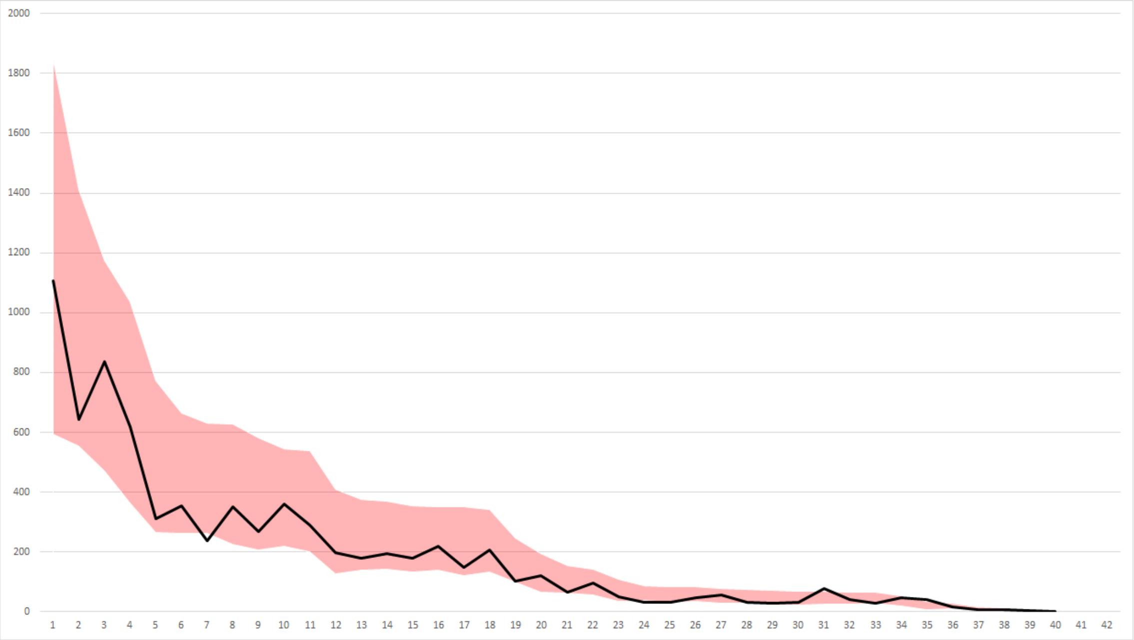

Figure 2.12: Comparison between CD8$^+$ cell counts from the image analysis pipeline and three independent pathologists. The number of CD8$^+$ cells were counted in 40 different images. The red area shows one standard deviation around the mean of the human counts. The black line shows the count by our algorithm. Regions are ordered according to the amount of variation between the pathologists' counts.

take up to an hour. Our algorithm returned the results for all 40 images in less than one minute, marking a significant improvement in speed. In most cases the prediction of the algorithm was within one standard deviation of the mean count of the pathologists.

Automated cell counting has the advantage that it is reproducible. On the other hand, when humans are tasked with such a repetitive task their performance is likely to be affected by factors such as tiredness, boredom, and stress. Our algorithm aims to alleviate these factors rather than exacerbate them.

2.5 Discussion

In this Chapter we have presented a suite of image analysis tools capable of the automated analysis of whole slide histological images. We have validated the key steps in our pipeline and demonstrated that it performs well in comparison with pathologists.

Automated image analysis is slowly becoming accepted by pathologists, and digital pathology is now seen as a potential solution to problems including remote working, receiving second opinions from colleagues in different locations, outsourcing and routine diagnosis (Williams et al., 2017, 2018a). However, the current uses of computational pathology tend to be for reinforcing conventional approaches to pathology. The computer is often used as a tool for the pathologist to easily visualise whole slide images and see specific regions in better detail. While this is encouraging, algorithms to analyse these images quantitatively are not in common use. There is demand for robust algorithms which can streamline pathologist workflows, for instance through automated cell counting, quantification of histopathological staining, automated tumour grading or region of interest detection (Williams et al., 2017; Griffin and Treanor, 2017). In order to improve uptake of such algorithms they must be rigorously validated and shown to be reliable, be user-friendly for non-programming pathologists, and provide an advantage to pathologists over their current routines. Our pipeline is capable of performing these tasks, and future work will include streamlining this process to make use by non-programmers easier.

A key aim of our pipeline is to provide a flexible framework which can combine techniques from across digital pathology and computer vision. While the classifiers described in this Chapter provide a way to accurately identify immune cell locations, the modular design of our pipeline permits new classifiers to be trained quickly by people with no knowledge of the mathematical or computational foundations of computer vision. In this Chapter we have trained classifiers using SVMs, but our pipeline is designed to permit the use of classifiers trained with other methodologies and enable users to incorporate their preferred techniques into the pipeline. The pipeline developed in this Chapter can provide a way of closing the gap between novel techniques in computer vision and their adoption by pathologists with no knowledge of the computer vision literature.

In this thesis we focus on the ability of our pipeline to analyse histological images. However, the techniques which we have used are not unique to this imaging modality. Our pipeline could be applied to images from any methodology in which the objects in the image are distinguished by colour. Extending the list of features in the SVMs removes this requirement, as morphological features could be included to identify cells from their shape. We predict that our colour-based classifiers could be made more robust by the use of stain normalisation techniques which preprocess slides to remove variations in stain intensity before classification (Khan et al., 2014). Future work will involve applying our pipeline to multiplexed immunohistochemistry slides (Stack et al., 2014; Blom et al., 2017). Current multiplexing techniques allow up to eight stains to be visualised robustly on a single slide, and forthcoming techniques may permit an order of magnitude more (Blom et al., 2017).

Metrics which are commonly used in pathology, such as the density of per mm^2 of tissue, were developed before cell quantification techniques existed and are often designed to provide qualitative summaries of cell locations. As such, digital pathology techniques commonly recreate these approximate metrics with more precision, for example by providing more accurate cell counts and tumour areas for the calculation of cell densities. While providing these metrics is a crucial part of digital pathology, having knowledge of cell locations across a tumour provides a range of opportunities that have previously been unavailable to pathologists. Image analysis software such as our pipeline can summarise spatial data using more complex spatial techniques that would be infeasible for a human to calculate and are therefore not in common use. In Chapter 3, we present several applications of our pipeline which leverage such techniques to provide novel quantitative summaries of cell locations within a tumour.

Chapter 3

Applications of the image analysis pipeline

3.1 Overview

In Chapter 2 we introduced an image analysis pipeline designed for the automated analysis of whole slide images (WSIs) of histology data. We described its structure and validated several of the algorithms used in the pipeline. We also benchmarked the performance of our pipeline against trained human pathologists, and demonstrated that it can perform tasks such as cell counting at least as accurately as a trained human in a fraction of the time. In this Chapter, we demonstrate the types of analysis beyond cell counting that can be conducted with our pipeline, and show its flexibility and strength in a range of applications.

In Section 3.2 we briefly describe the methods used in this Chapter. We review existing methods for the analysis of histological data, and discuss how these methods could be extended by applying techniques from the field of spatial statistics to data extracted using our pipeline (Ripley, 1976; Mignardi et al., 2016; Nawaz et al., 2015; Nawaz and Yuan, 2016). We discuss these techniques and additional information that they can provide in comparison with existing clinical metrics, and emphasise the need for automated image analysis to calculate them.

In Section 3.3 we apply our pipeline and spatial statistics to the RITA dataset described in Section 2.2.1. We compare the density of macrophage infiltration between two of the indications present in this dataset, breast cancer and colorectal cancer. We conduct further analysis of images from colorectal cancer using a range of spatial statistics, and consider how each statistic varies between different patterns of macrophage infiltration. These spatial statistics are then used to automatically cluster tumour regions based on the locations of macrophages in the images, and we demonstrate that this process groups the images into clusters with qualitatively similar macrophage distributions. In Section 3.4 we use our pipeline on images from the Pugh lab dataset of head and neck tumour samples described in Section 2.2.2. This dataset contains consecutive tumour slices stained to show the expression of different markers, and we coregister these slides to obtain more information about the oxygen environment of infiltrating immune cells. We obtain immune cell densities for different oxygen environments based on an analysis of randomly selected regions of the tumour. We then use these regions as “artificial biopsies” to investigate how accurately measurements of hypoxia based on tumour samples are able to accurately estimate the degree of hypoxia in the tumour as a whole. In Section 3.5 we demonstrate the flexibility of our pipeline by applying it to *in vitro* experiments on the formation of atherosclerotic plaques conducted by Mr. Hugh Ford at the University of Sydney. In these *in vitro* experiments, we use our pipeline as a high-throughput tool to count the number of plastic microbeads consumed by macrophages.

We end this Chapter in Section 3.6 with a discussion of how the quantitative techniques presented in this Chapter could be more closely integrated with imaging software to make them more accessible to pathologists.

3.2 Methodology

3.2.1 Current techniques used for analysis of whole slide images

Existing methods for the analysis of histological images are primarily qualitative. Historically, pathologists analysing whole slide images had to summarise their findings manually, without the aid of a computer. This has led to semi-qualitative metrics becoming the norm in histological image analysis. Scoring/grading systems provide a series of rules which allow pathologists to assign numeric labels to tissue samples (Cross, 1998). In general, these systems are based on rules for dividing samples into a small number of categories which provide meaningful information about treatment response. For example, tumours may be assigned a grade ranging from 1 to 4 based on site-specific observations that describe tumour invasion which communicate meaningful information about tumour growth to other medical professionals (Edge and Compton, 2010; Cross, 1998). While this type of quantification allows standardisation for the purposes of diagnosis and communication between doctors, scoring may be subjective. In one study, pathologists asked to diagnose a breast biopsy reached the same conclusions only 75% of the time (Elmore et al., 2015). A second problem with pathological scoring is that while such scores are useful for classifying tumours, a whole slide image contains more quantitative information that is not represented by these grades.

More detailed scores have been developed to represent this information. In the case of immune cell infiltration into tumours, a commonly used metric is the density of cells per mm^2 of tissue (Klauschen et al., 2018). In the absence of computational techniques, these densities rely on accurate estimation of both the number of immune cells observed and the area of tissue. Errors due to human miscounting and estimation can therefore prevent these metrics from being reproducible. In Section 2.4.3, we demonstrated that when there are a large number of immune cells present, different pathologists will reach different conclusions as to the number of cells.

As digital pathology becomes more widely accepted by the medical community, computer aided counting techniques are becoming more common in pathology. Our pipeline, and alternative software packages such as Definiens, Visiopharm and QuPath (Bankhead et al., 2017; Baatz et al., 2009; Bueno et al., 2016), permit more accurate cell counts and area calculations. As a result, metrics such as "number of immune cells per mm^2 of tissue" are becoming easier to calculate reliably. At present, digital pathology is largely confined to making these existing metrics more accurate. However, digital pathology permits the generation of more complex summary statistics than cell densities. Whilst spatial statistics are commonly used in other fields, particularly in Geographic Information Systems (GIS) (Unwin, 1996), their use in pathology has been limited (Mignardi et al., 2016). Kather et al. (2017, 2018) conduct spatial quantitative analyses of histological images in which they automatically classify tumours as being "hot" or "cold", qualitative terms which are used to describe the degree of immune infiltration in a tumour. Their results compare mathematical modelling with tumour types defined based on the localisation of proliferation, but do not use more complex spatial statistics.

Yuan (2014) and Nawaz et al. (2015) propose the use of spatial metrics to identify hotspots of tumour-infiltrating lymphocytes in estrogen receptor-negative breast cancer. Nawaz et al. (2015) use spatial statistics for hotspot detection to show that co-localisation of cancer and immune hotspots is a prognostic marker for patient survival, and Yuan (2014) proposes a quantitative intratumour lymphocyte ratio (ITLR) which can only be calculated from digital pathology analyses. Heindl et al. (2015) propose Ripley's K-statistic as a means of mapping spatial heterogeneity. Spatial and density descriptors have also been used to predict likelihood of recurrence in non-small cell lung cancer Corredor et al. (2018).

In Section 3.2.2, we introduce a number of spatial statistics which can be applied to the types of cell location data that our pipeline extracts from histological images, which will be applied in Sections 3.3 and 3.4.

3.2.2 Spatial statistics of point clouds

Cell densities are useful for measuring summary statistics such as the proportion of an image which is covered by macrophages or tumour cells, but more complex measures of spatial distribution can be generated from point cloud data. There are several methods for extracting descriptors of spatial distribution from point datasets. These include *second-order statistics* (Møller and Waagepetersen, 2004) such as Ripley's K-function (Ripley, 1976; Dixon, 2002; Mattfeldt et al., 2009; Kay, 2014) and its successors, the L-function and H-function (Kay, 2014), the pair correlation function (PCF) (Fozard et al., 2011; Mattfeldt et al., 2009; Stoyan and Stoyan, 1994; Møller and Waagepetersen, 2004) and the J-function (van Lieshout and Baddeley, 1996). While second-order statistics focus on relationships between pairs of points, summary statistics for bivariate point data can also be obtained using Quadrat Histograms (Fozard et al., 2011; Buyong, 2007). Bivariate statistics can be used to compare distributions of two different cell types, such as macrophages and tumour cells or macrophages and T cells.

Many of these techniques rely on comparing the point cloud with random point data generated using a Poisson process, in which the location of each point in the point cloud is independent of all other points. Consider, for example, the variance/mean ratio test (Buyong, 2007), which relies on the observation that if point data is generated randomly from a Poisson process then the distribution of points in the squares of a regular grid overlaid onto the domain should have the same mean and variance. If the variance of the observed point distribution is significantly larger than the mean then the data is unlikely to be random and is likely to be grouped into clusters, while if the variance is significantly less than the mean the pattern is instead non-random but scattered or dispersed (Buyong, 2007; O'Sullivan and Unwin, 2010).

Although second-order statistics can provide useful information, they have limitations. Baddeley and Silverman (1984) provide an algorithm for generating point data which appears non-random but has identical second-order statistics to data generated randomly from Poisson processes. It is therefore important to use as diverse a range of summary statistics as possible to ensure that genuine patterns are detected.

Pair correlation function

The pair correlation function (PCF) (Møller and Waagepetersen, 2004; Mattfeldt et al., 2009; Stoyan and Stoyan, 1994) can be used to obtain information about the length scales of emergent patterns in point data (Fozard et al., 2011). To generate the PCF, an annulus of width δr is placed at a distance r around a given point in the point distribution. The number of points within this annulus is calculated, and then normalised with respect to the area of the annulus and the number of points that would be expected to fall inside this ring if the points were distributed according to a random Poisson process. This quantity is calculated for all points within the point cloud and then the average calculated. The process is repeated for a range of values of r to generate the pair correlation function, $g(r)$. $g(r)$ provides a measure of how likely there are to be two points separated by distance r compared to a random point distribution. The value of $g(r)$ indicates the level of clustering at distance r: $g(r) > 1$ indicates that pairs of points separated by distance r occur frequently, while $g(r) < 1$ indicates that pairs of points are less likely to be separated by distance r than would be expected if they were distributed randomly. For a point field generated randomly through a Poisson process, the pair correlation function $g(r) \equiv 1$ (Stoyan and Stoyan, 1994). Importantly, it should be noted that $g(r) = 1$ does not necessarily mean that the points are generated from a random process, just that they share the same second-order characteristics as the Poisson field. This is demonstrated in a cautionary example by Baddeley and Silverman (1984).

Fozard et al. (2011) show how a multivariate version of the pair correlation function can be used to identify the scale of clustering patterns in labelled point data originating from a bipartite population of cells. This is achieved by calculating the PCF for all cells and the PCF for each individual cell population, where the latter PCFs are generated by ignoring all points which do not belong to the population of interest. Gavagnin et al. (2018) propose extensions to the PCF for point data in discrete domains, such as data generated by on-lattice agent-based models.

Ripley's K-function

Ripley's K-function (Ripley, 1976; Dixon, 2002; Mattfeldt et al., 2009; Kay, 2014) and its successors, the L-function (Stoyan and Stoyan, 1994) and H-function (Kay, 2014), are second-order statistics which, like the PCF, relate pairs of points. The construction of the K-function is similar to that of the PCF, with the difference that while the PCF uses an annulus of width δr placed a distance r from a reference point the K-function simply considers a disc of radius r. This leads to a similar second order measure of correlation, except that the K-function is cumulative.

For a random Poisson process, the K function can be written in closed form as

$$K(r) = \pi r^2 \tag{3.1}$$

where r is a measure of distance. This result motivates the definition of the L-function:

$$L(r) = \sqrt{\frac{K(r)}{\pi}} \tag{3.2}$$

which grows linearly with r for a Poisson process. The L-function can be normalised again to give the H-function, which has an expected value of 0 for a Poisson process. The H-function is defined as:

$$H(r) = L(r) - r. \tag{3.3}$$

A value of $H(r) > 0$ indicates clustering on the length scale r, while $H(r) < 0$ indicates dispersion. The H-function and its derivative have been proposed (Kiskowski et al., 2009) as useful measures of the length scale of a pattern, with the H function more useful for predicting this value for experimental data sets with a high density of points and $H'(r)$ more informative at low densities.

The J-function

The J-function is a non-parametric test for identifying clustering and dispersal in point patterns proposed by van Lieshout and Baddeley (1996). The J-function uses two similar cumulative density distributions, the nearest-neighbour distribution and the empty-space function (or spherical-distribution function).

The nearest-neighbour distribution function, $G(r)$, is the cumulative density function (CDF) of the histogram of nearest-neighbours in the point cloud: for each point, the distance to its nearest neighbour is calculated and logged. For a given radius r, $G(r)$ is the number of points whose nearest neighbour is within distance r divided by the total number of points:

$$G(r) = \frac{\text{Number of points within } r \text{ of another point}}{\text{Number of points}}. \tag{3.4}$$

The empty-space function, $F(r)$, is calculated in a similar manner, but instead of calculating the distance between points in the point cloud, randomly chosen reference points are selected from the domain. For each of these random points, we calculate the distance to its nearest neighbour in the point cloud. The CDF is then calculated as before:

$$F(r) = \frac{\text{Number of randomly chosen reference points within } r \text{ of a target point}}{\text{Number of randomly chosen points}}. \tag{3.5}$$

The J-function is defined to be the ratio of these CDFs (van Lieshout and Baddeley, 1996):

$$J(r) = \frac{1 - G(r)}{1 - F(r)}. \tag{3.6}$$

For randomly distributed points, we expect $G(r) \approx F(r)$, in which case $J(r) = 1$. If $J(r) > 1$ then more points from the point cloud can be found in a circle of radius r placed randomly in the domain than in a circle centred at one of the points, indicating dispersal of points within the point cloud. If $J(r) < 1$ then a disc of radius r centred at an element of the point cloud contains more points than a randomly placed disc, and we conclude that the points are clustered (van Lieshout and Baddeley, 1996).

Quadrat Histograms

Quadrat Histograms (QH) are generated by tiling the domain with square "quadrats" of edge length r, where r is much smaller than the domain width and height. They can be applied to standard point clouds as well as bivariate point clouds in which more than one population of cells is identified. In the univariate case, the number of points in each quadrat is counted and the distribution calculated. If points are scattered, then few quadrats will contain a large number of points, while if points are clustered then the distribution will contain many quadrats containing multiple points. Buyong (2007) suggests using a variance/mean ratio (VMR) test on this distribution, since randomly generated points will follow a Poisson distribution in which the mean and variance are equal. A VMR of 1 indicates randomness, while VMR < 1 and VMR > 1 suggest dispersal and clustering respectively.

In the bivariate case, for each quadrat, the number of points of each type is summed, and the proportion of points belonging to the first cell population, p_1, is calculated. Quadrats which contain large clusters of tumour cells will have $p_1 \approx 1$, while those which contain predominantly other cell types will have $p_1 \approx 0$. In a random distribution of two cell populations, we expect $p_1 \approx 0.5$.

The existence of many quadrats with extreme values of p_1 indicates that cells from population one are generally found clustered together. This is measured by creating a histogram of the value of p_1 over all quadrats. Quadrat Histograms and PCFs were used in (Fozard et al., 2011) to distinguish between "noisy patterns" and randomness in labelled point data. They suggest a conservative criterion for the existence of patterns which requires that more than 10% of quadrats have $p_1 < 0.02$ or $p_1 > 0.98$.

If a pattern is detected for a given value of r, it indicates the presence of a pattern with a typical length greater than or equal to r. For example, if a pattern consists of clusters $100\,\mu m$ in diameter then a quadrat histogram for edge length $r = 10\,\mu m$ will detect the existence of a pattern since a large number of quadrats will be entirely within the cluster.

We propose an extension of the quadrat histogram methodology to detect characteristic length scales of clustering or dispersal. The standard quadrat histogram method described above can detect patterns based on only one length scale, defined by the edge length of the quadrats. By repeating the process for a range of quadrat edge sizes $0 \leq r \leq R$, the VMR can be calculated at a range of scales. The shape of this graph will indicate quadrat edge lengths r for which a pattern exists, and any values of r for which the VMR $= 1$ indicate length scales at which clustering or dispersal appear or disappear. We call this graph a "variance/mean ratio curve", or "VMR curve".

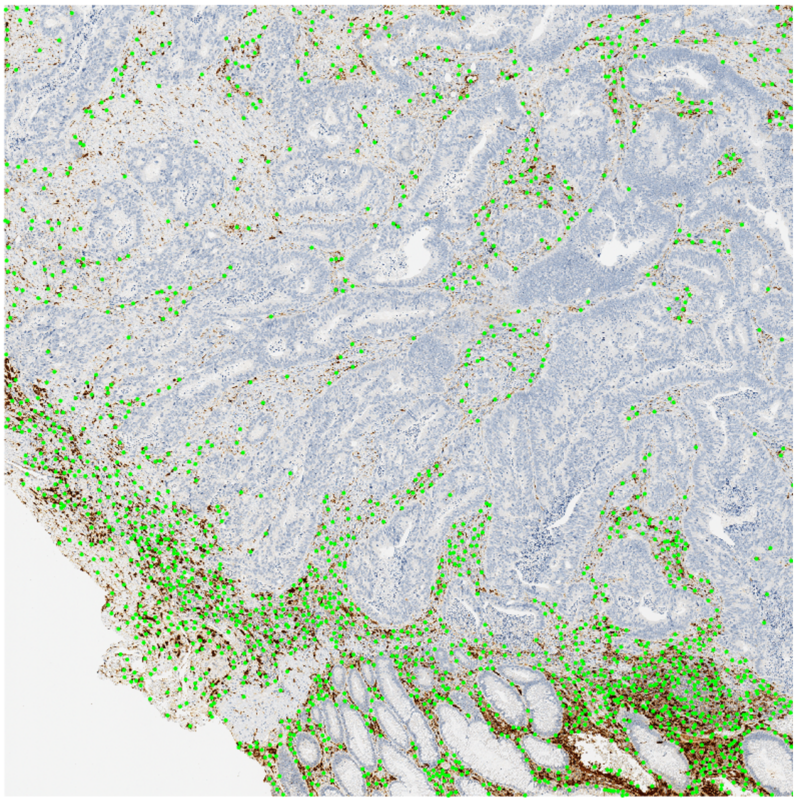

Figure 3.1: An example of macrophage cell centres identified in a 3 mm × 3 mm section from a colorectal cancer tumour in the RITA dataset. Macrophages are marked with a green point. The brown stain is for CSF-1R$^+$ macrophages, with blue counterstain for all other cell nuclei. In this image, macrophages are predominantly in the stroma, and are excluded from tumour cell clusters and colorectal crypts.

3.3 Application of pipeline to RITA data

In this Section we apply our image analysis pipeline to CSF-1R stained images from the RITA dataset described in Section 2.2.1. We focus on the spatial distribution of macrophages in colorectal and breast cancer. For each whole slide image, we randomly select non-overlapping (3 mm × 3 mm) regions of interest from within the tumour region until no more regions can be chosen without overlap. We apply our cell counting process to each region to obtain the coordinates of identified macrophage centres (see Figure 3.1.) These statistics can be used to generate immune cell densities for each tumour or to calculate spatial statistics of macrophages.

3.3.1 Macrophage densities

Figure 3.2 shows a comparison between the observed macrophage densities in breast and colorectal cancer samples, a commonly used measurement for determining the extent of immune cell infiltration. For each region of interest we calculate the number of macrophages per mm^2. Figure 3.2b shows the probability density function obtained by fitting normal distributions to the histograms shown in Figure 3.2a.

Regions from breast cancer samples contain more macrophages on average than those from colorectal cancer, with a mean of 543 macrophages per mm^2 in breast cancer samples compared to 254 macrophages per mm^2 in colorectal cancer. The breast cancer samples show a much wider range of macrophage densities than the colorectal cancer, with a standard deviation of 222 macrophages per mm^2 and an interquartile range of 348 macrophages per mm^2. For the colorectal cancer regions, these values are 114 macrophages per mm^2 and 164 macrophages per mm^2 respectively.

It can be seen from the distributions that breast cancer samples contain more macrophages on average than colorectal cancer samples, and that there is more variability in the density of macrophages in breast cancer samples than in regions from colorectal tumours. While this type of measurement is a common metric for pathological analysis, in the following Sections we demonstrate that applying a range of more complex spatial statistics to the same dataset can provide additional information that this analysis overlooks.

3.3.2 Pair correlation function

In this subsection we apply the pair correlation function (PCF), $g(r)$, to the ($3\,\mathrm{mm} \times 3\,\mathrm{mm}$) regions of interest from the breast and colorectal cancer images. As observed in Section 3.2.2, $g(r) = 1$ indicates that the distribution of macrophages appears random at a distance r from a typical macrophage, while $g(r) > 1$ and $g(r) < 1$ indicate clustering and dispersal at distance r respectively. All PCFs applied to our

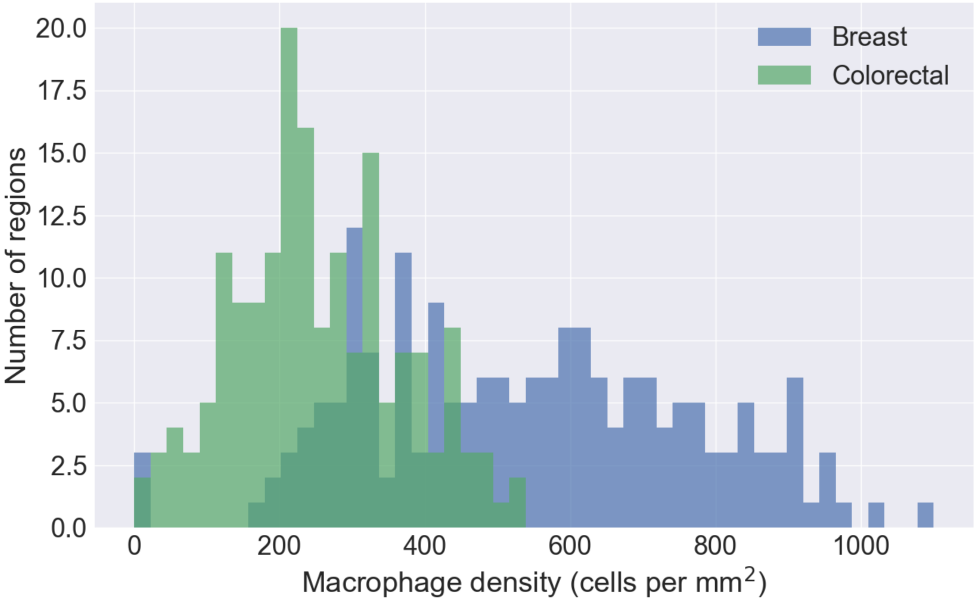

(a) Histogram of macrophage densities for regions of interest in breast and colorectal cancer tumours

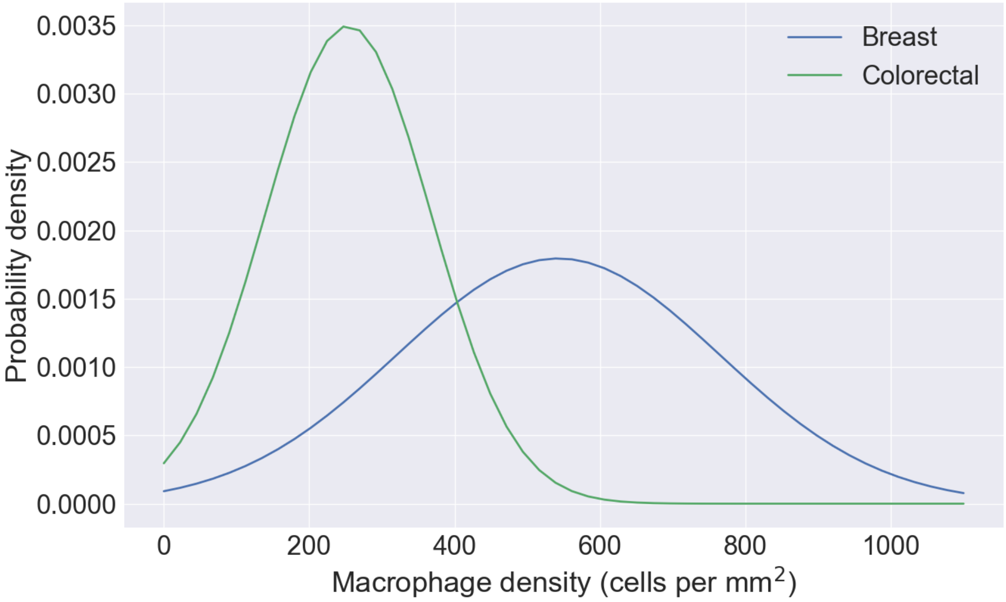

(b) Probability density function for macrophage density in breast and colorectal tumours

Figure 3.2: Differences in the macrophage density (macrophages per mm^2) between breast cancer (blue) and colorectal cancer (green). Figure 3.2a shows the density of macrophages in ($3\,\mathrm{mm} \times 3\,\mathrm{mm}$) regions of interest from breast ($n = 188$ regions, 28 tumours) and colorectal ($n = 173$ regions, 30 tumours) tumours. Figure 3.2b shows the probability density function obtained by fitting a normal distribution to the data.

histology data display a sharp peak for small r which rises sharply from 0 when $r \approx 0.02\,\mathrm{mm}$. This distance is approximately one cell radius, which indicates that our cell identification algorithm generally places macrophage centres at least one cell radius apart. When calculating the PCFs in this Section, we use the central $2\,\mathrm{mm}^2$ of each region of interest and calculate the PCF up to a maximum radius of $r = 1\,\mathrm{mm}$. This eliminates potential boundary effects caused by macrophages close to the edge of the region of interest whose annuli would extend beyond the boundary.

In Figure 3.3 we illustrate some of the characteristic PCFs that we observe, alongside the image from which they were generated. Each image in this demonstration comes from the colorectal cancer dataset. For each PCF, we have chosen the annulus width $\delta r = 0.005\,\mathrm{mm}$. This value ensures that we can measure changes in the PCF that happen on a length scale smaller than that of a cell, which are commonly observed for small r. In order to more easily characterise PCFs, we smooth the PCF using a Savitzky-Golay filter (Savitzky and Golay, 1964) (shown in red overlay in Figure 3.3). This filter uses a fixed-width sliding window to fit a cubic polynomial to a subset of the data. This type of filter produces a smooth curve which closely follows the form of the PCF.

The first example, Figures 3.3a and 3.3b, shows short-range clustering and long range dispersal of macrophages. For small radii, approximately $0.02\,\mathrm{mm} \leq r \leq 0.1\,\mathrm{mm}$, we have $g(r) > 1$. We interpret this as showing that a macrophage is more likely to be within $0.1\,\mathrm{mm}$ of another macrophage than if macrophages were randomly distributed, and that therefore macrophages are clustering on this length scale. For $0.1\,\mathrm{mm} \leq r \leq 0.75\,\mathrm{mm}$, $g(r) < 1$, which we interpret as showing that macrophages are more dispersed on these length scales than would be expected of a random distribution. At $r \approx 0.75\,\mathrm{mm}$ the PCF remains close to 1 and suggests that at this length scale macrophages exhibit no apparent structure. This description matches the distribution shown in Figure 3.3b, in which macrophages are clustered

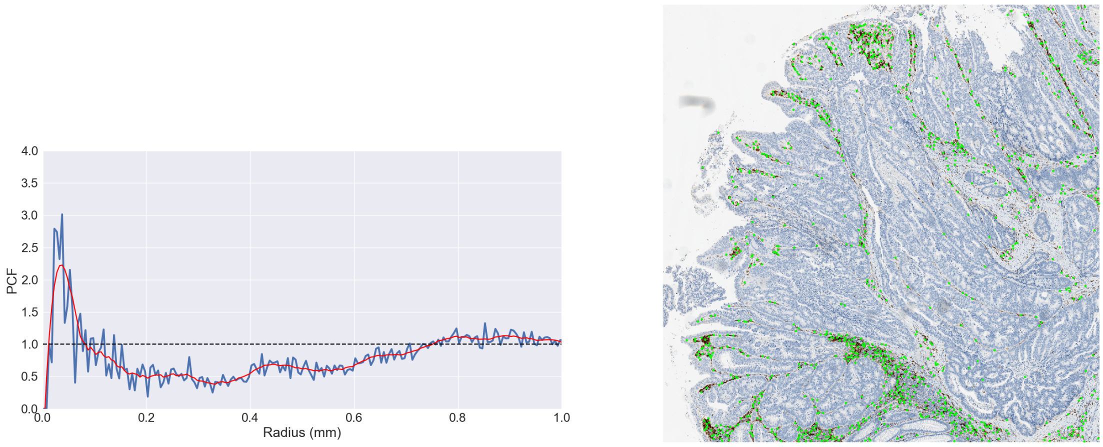

(a) PCF showing short-range clustering and long-range dispersal of macrophages

(b) Image used to generate the PCF in Figure 3.3a.

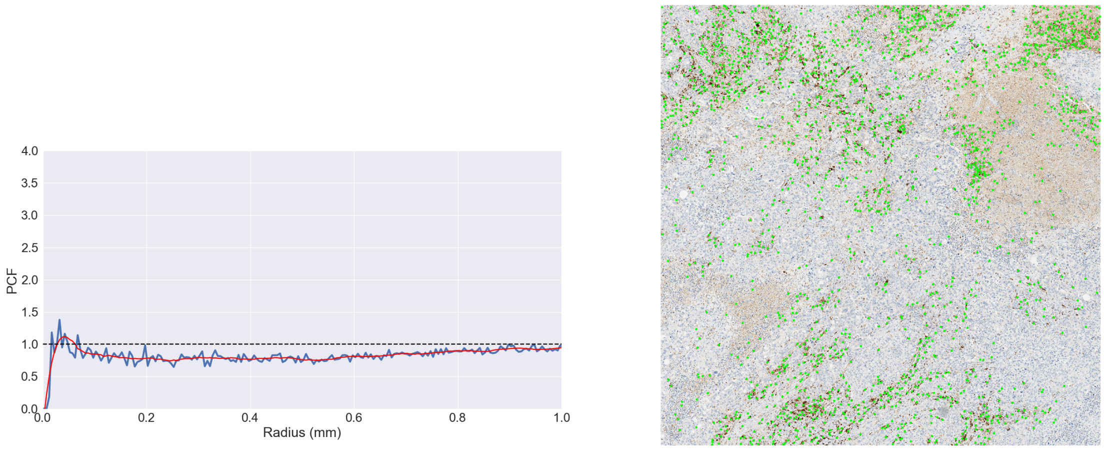

(c) PCF showing unstructured macrophage distribution

(d) Image used to generate the PCF in Figure 3.3c.

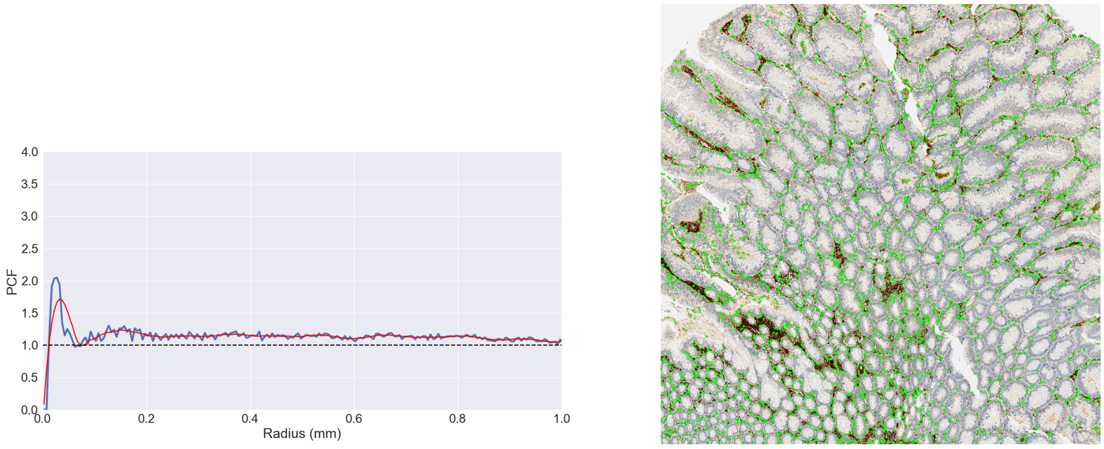

(e) PCF showing macrophage clustering

(f) Image used to generate the PCF in Figure 3.3e.

Figure 3.3: Characteristic behaviours of the PCF. The red curve is the result of smoothing using a Savitzky-Golay filter (Savitzky and Golay, 1964) with a third-order polynomial and a window size of 21 data points.

in groups between tumour cell structures from which macrophages are broadly excluded. The width of these macrophage clusters, and the distances between clusters, appear to qualitatively correspond to the length scales identified by the PCF.

Figures 3.3c and 3.3d show a different characteristic type of PCF, in which $0.7 \leq g(r) \leq 1$ almost everywhere. This indicates that macrophages are slightly more dispersed than would be expected of randomly placed points throughout the region of interest. Figure 3.3d matches this qualitative description, as there are no obvious clusters of macrophages present and no clear structure in the distribution of cells.

Figures 3.3e and 3.3f provide a final example of a type of PCF commonly seen in our analysis. After a peak showing macrophage clustering at small length scales $r < 0.1$, $g(r)$ remains above one for most values of r. This indicates that macrophages are more commonly represented at every length scale than would be expected if they were randomly distributed. While this is at first counterintuitive, Figure 3.3f demonstrates how such a distribution arises. Macrophages are clustered around the edges of colorectal crypts (in cross-section), causing them to appear in approximately regular circles. In this image the crypts vary in size, with diameters ranging from approximately 0.1 mm in the lower left of the image to almost 1 mm at the top of the image, and these varying structure sizes are represented by the range of length scales for which $g(r) > 1$.

3.3.3 J-Function

While the PCF is able to identify spatial structures in the macrophage distribution on a large scale, the J-function can be used to identify smaller scale spatial structures. For almost every region of interest analysed, for both breast and colorectal cancer, $J(r)$ is greater than one for $0\,\mu m \leq r \leq 25\,\mu m$. This indicates dispersal of points on these length scales, and correlates with the approximate radius of a macrophage. Points being dispersed on this length scale agrees with the results from the PCF at this scale, and we again interpret this as an indication that our cell detection algorithm is not placing two points unrealistically close to each other.

After this initial peak, we typically observe relatively few types of behaviour. Two of these types are represented in Figure 3.4, where we calculate the J-function for the regions which are shown in Figure 3.3. Figures 3.4a and 3.4c are both examples of a J-function which shows short-range dispersal at length-scales of approximately one cell radius, followed by clustering on longer length scales. Comparing Figure 3.4a with the corresponding PCF in Figure 3.3a demonstrates the importance of using multiple spatial statistics to describe macrophage distributions. The PCF for this region suggests the macrophages are dispersed from one another for $r \geq 0.1\,\mathrm{mm}$, while the J-function suggests clustering at scales from $0.02\,\mathrm{mm} \leq r \leq 0.3\,\mathrm{mm}$. The reason for this apparent contradiction can be seen in Figure 3.5, which shows the two cumulative density functions from which the J-function is generated.

Figure 3.5 shows that for the J-function presented in Figure 3.4a, the nearest-neighbour distribution function reaches its asymptote much faster than the spherical-contact distribution function. Once the nearest-neighbour function $G(R) = 1$ for some radius R, the J-function $J(r) = 0$ for all $r \geq R$. However, the nearest-neighbour distribution is dependent only on short-range contact between macrophages, rather than long distance structures. For this reason, we suggest that the J-function is informative for short length scales, but does not contain critical information about perceived patterns which exist on scales larger than that of individual macrophage clusters. The PCF provides useful information at this scale, but small-scale information in the range $0\,\mathrm{mm} \leq r \leq 0.1\,\mathrm{mm}$ is likely to be obscured by the need to choose a sufficiently large value δr for the annulus width.

The J-function in Figure 3.4c is a similar shape to that in Figure 3.4a due to the short-range nature of the J-function. While the two distributions are qualitatively different when perceived as a region on the scale of millimeters, at the cell scale

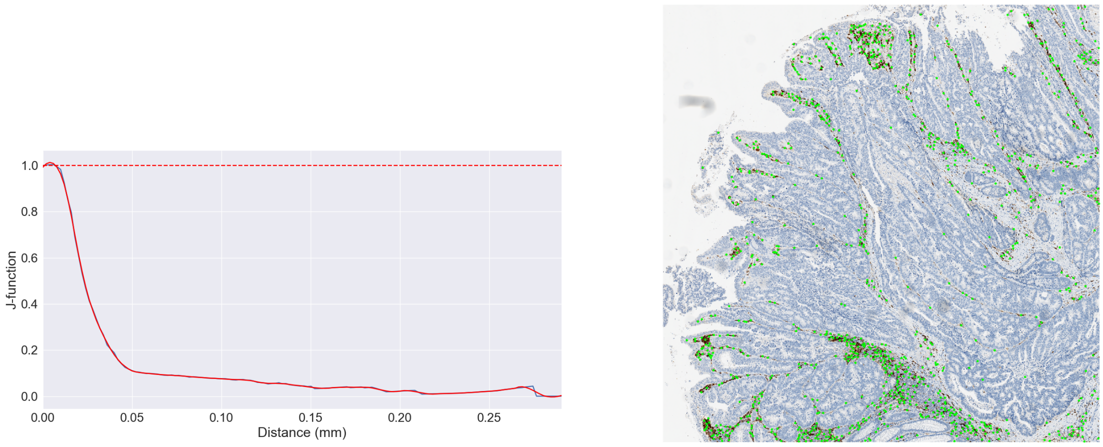

(a) J-function showing clustering of macrophages up to lengths of 0.3 mm

(b) Image used to generate the J-function in Figure 3.4a.

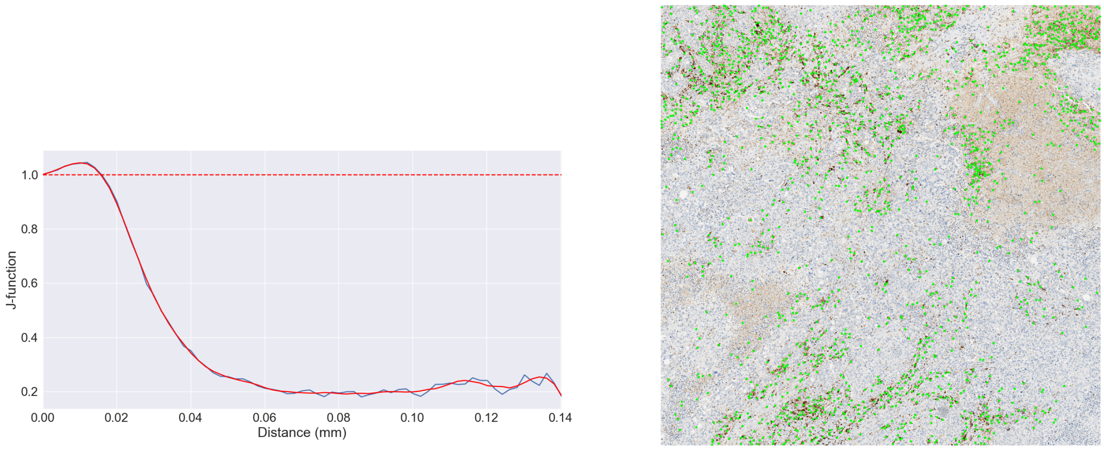

(c) J-function showing clustering of macrophages up to lengths of 0.14 mm

(d) Image used to generate the J-function in Figure 3.4c.

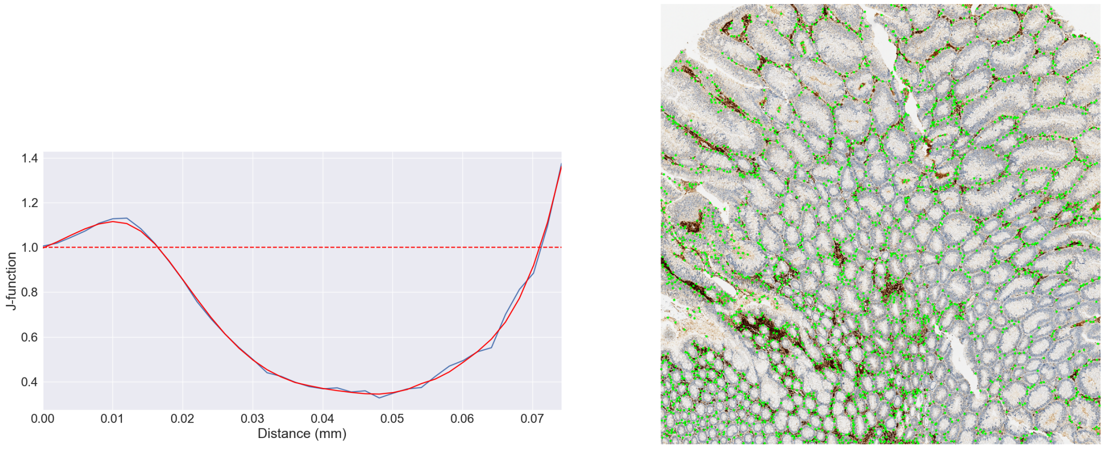

(e) J-function showing short range clustering and longer range dispersal

(f) Image used to generate the J-function in Figure 3.4e.

Figure 3.4: Characteristic behaviours of the J-function. The red curve is the result of smoothing using a Savitzky-Golay filter (Savitzky and Golay, 1964) with a third-order polynomial and a window size of 11 data points.

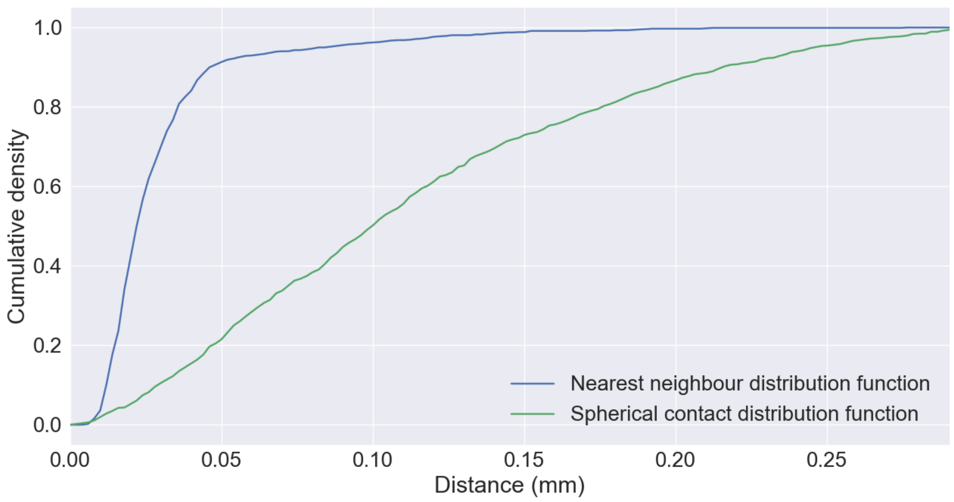

Figure 3.5: Nearest-neighbour density function ($G(r)$) and spherical-contact distribution function ($F(r)$) for the region shown in Figure 3.4b.

relationships between macrophages are similar. Each image contains clusters of macrophages which are identified by the J-function, but to perceive the differences in the structure of these clusters we require the PCF.

In contrast, the J-function in Figure 3.4e displays a different characteristic shape to those in Figures 3.4a and 3.4c. In this case, we see dispersal on the scale of a single cell ($0\,\text{mm} \leq r \leq 0.016\,\text{mm}$), short-range clustering for distances of 3 or 4 cell widths ($0.016\,\text{mm} \leq r \leq 0.07\,\text{mm}$) and then dispersal again for $r \geq 0.07\,\text{mm}$. 0.07 mm is approximately the width of colorectal crypt in this image, and the J-function crossing the line $J = 1$ at this point indicates the presence of these regular, small structures.

3.3.4 Grouping regions of interest based on second-order statistics

We use a combination of three second-order statistics to group tumour regions based on their macrophage distributions. For each tumour region, we calculate the PCF, the J-function and the VMR curve (example VMR curves are shown in Appendix

D). We then summarise these curves with a range of metrics, such as the highest value obtained by the curve, the range of r values for which the curve is within a small distance ϵ of one, or the number of times the curve changes from below one to above. Using this range of features to describe each region as a point in a high dimensional space, we use K-means clustering (Hartigan and Wong, 1979) to identify clusters of similar images. K-means clustering identifies K distinct clusters from a multidimensional dataset by assigning all points to one of K centres. Each centre is then moved to the average position in the high dimensional space of all points assigned to it. The process is then repeated iteratively, with each point being assigned to the nearest centre and the centres moved, until the centres are stationary. K-means clustering relies on knowing how many clusters are present in the data, as the parameter K must be prespecified. There is no definite criterion for knowing the value to choose, and here we select $K = 8$ based on an automatic silhouette analysis optimisation using Scikit-learn, the Python library we used to conduct this analysis (Pedregosa et al., 2012).

Based on the second-order statistics described in this Section, the K-means algorithm separates the data into perceptually meaningful clusters. Figure 3.6 shows examples taken from three of these groupings. Qualitatively, the three groups represent three different patterns of macrophage infiltration. In the top three images of Figure 3.6, macrophages are spaced in circles around the edge of colorectal crypts. In the middle set of three images macrophages are in dense clusters at the edge of regions of tumour tissue from which macrophages are almost completely excluded. Finally, in the bottom set of three images in Figure 3.6 macrophages are spread throughout the tissue with little visible structure.

Although each of the eight clusters identified through our K-means analysis contain regions which are perceptually similar, they do not appear to represent eight qualitatively different patterns of infiltration. For example, two of the clusters appear to contain only regular colorectal crypt structures like those in the top set of

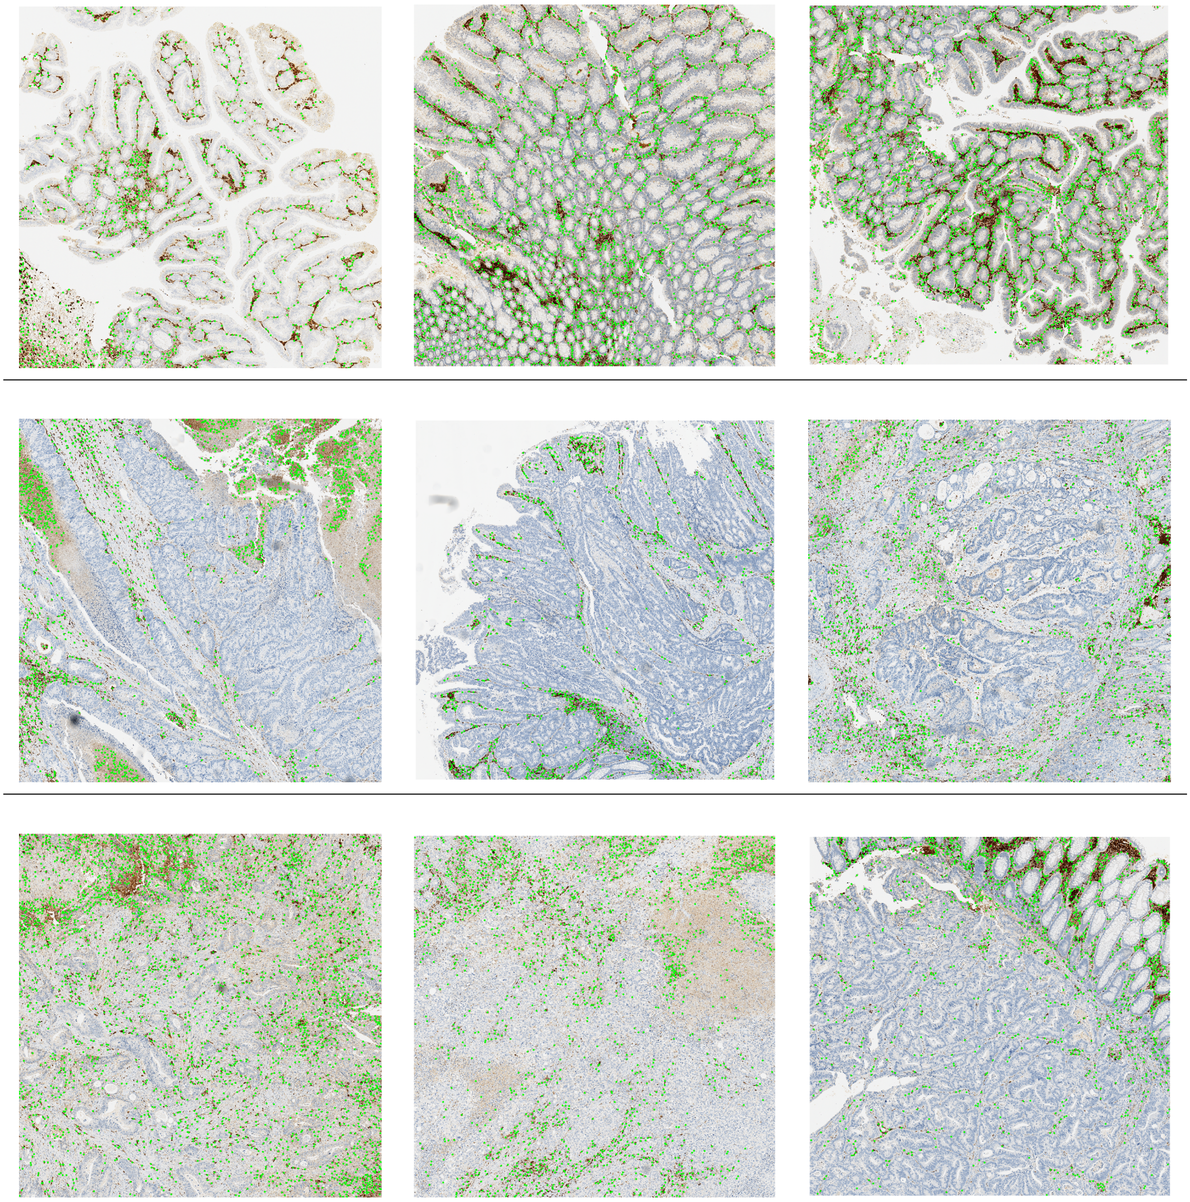

Figure 3.6: Three sets of images which have been grouped as having similar macrophage distributions according to K-means clustering based on second-order statistics.
Top: regular colorectal crypt structures with macrophages in the surrounding stroma.
Middle: dense clusters of macrophages with large areas of near-total macrophage exclusions.
Bottom: dispersed infiltrating macrophages scattered throughout the tissue.

Figure 3.6. This indicates that either fewer clusters could be chosen, or that more statistics are needed to ensure that perceptually similar clusters remain together.

We conclude that it is possible to use second-order statistics to identify perceptually meaningful patterns in macrophage distributions automatically. Further work should include matching these distributions to patient prognoses to determine whether macrophages adopting these patterns has an impact on patient outcome. We also note that these statistics could be applied to cell locations in agent-based models of the kind described in Section 1.4.2. As such, these metrics could provide a useful means of comparing the outputs of agent-based models with histological data.

3.4 Application to head and neck tumours

In this Section, we apply our image analysis process to a sample of 16 human head and neck tumours. As described in Section 2.2.2, this dataset contains stacks of consecutive slides stained to show different oxygen environments and immune cell subtypes. This enables us to investigate the spatial distributions of immune cells in the context of their oxygen environment. In this Section, we use our imaging pipeline to obtain quantitative density counts of three different immune cell subtypes, $CD8^+$ T cells, $CD68^+$ macrophages and $FoxP3^+$ T regulatory cells. Three other stains are present on consecutive slides which we register with the original slide to obtain contextual information about the tumour microenvironment: PanCK, an epithelial cell stain which here distinguishes tumour from stroma; carbonic anhydrase IX (CAIX), a stain which identifies hypoxic cells; and pimonidazole, which stains severely hypoxic cells. The work in this Section was conducted in collaboration with Dr Philip Macklin and Professor Chris Pugh of the Nuffield Department of Medicine, University of Oxford.

3.4.1 Protocol for analysing images

Figure 3.7 illustrates how our image analysis protocol operates for extracting information from a stack of consecutive histology slides. For each slide stained to show immune cells, we identify adjacent PanCK$^+$ and pimonidazole$^+$ slides to form a set of images for registration. Using annotations provided by a pathologist to delineate the tumour region, we saturate the portion of the PanCK$^+$ stained slide containing the tumour with non-overlapping $1.5\,\mathrm{mm} \times 1.5\,\mathrm{mm}$ regions of interest (ROI). These regions are selected by randomly placing $1.5\,\mathrm{mm} \times 1.5\,\mathrm{mm}$ squares into the annotated region, ensuring that each candidate square does not overlap with any previously chosen region or with the edge of the annotation. Once no more squares can be placed without overlap, we then register each ROI with the same region on the other slides in the triplet, one for pimonidazole and one for the immune marker of interest. This registration is semi-automated, as a pathologist uses a GUI integrated into our pipeline to identify reference points in the candidate regions of interest. Our pipeline uses these reference points to calculate a transformation which aligns the images using only translation and rotation of the region. This ensures that images are not distorted by the registration process. Typically only around three points are required for selection, and Figure 3.8 shows a schematic view of how these points can be chosen based on distinctive features of tumour regions such as necrotic boundaries. Using the pipeline described in Chapter 2, we then identify the stained cells in each image and combine the consecutive slides to form a composite image for each region of interest. The composite images are supplemented with annotations for necrotic voids and any artefacts, for instance due to bubbles or scratches on the slides. We then identify cell centres for the immune cells, and distinguish between the numbers of immune cells which have infiltrated the tumour stroma, normoxic tumour areas, hypoxic tumour areas and regions of necrosis.

Some slides were stained for carbonic anhydrase IX (CAIX), which can be coregistered with the pimonidazole, PanCK and immune stained slides to form a set of 4

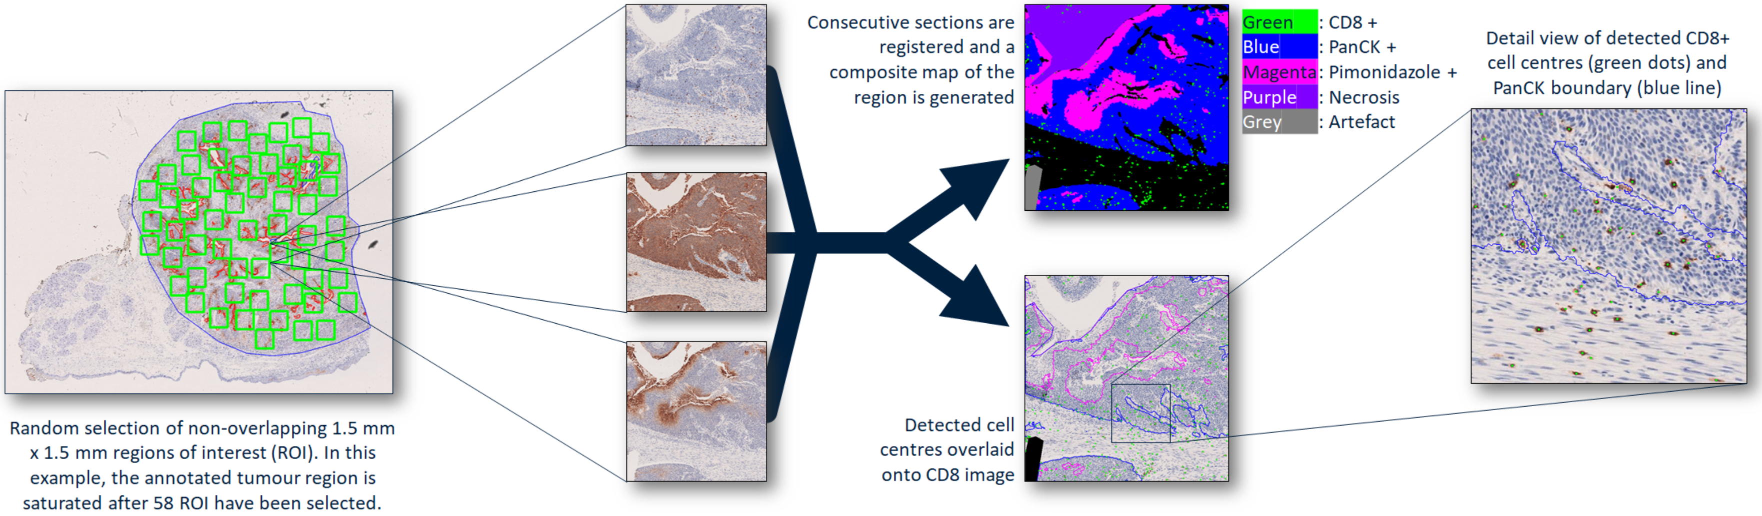

Random selection of non-overlapping 1.5 mm x 1.5 mm regions of interest (ROI). In this example, the annotated tumour region is saturated after 58 ROI have been selected.

Consecutive sections are registered and a composite map of the region is generated

Detected cell centres overlaid onto CD8 image

Detail view of detected CD8+ cell centres (green dots) and PanCK boundary (blue line)

Figure 3.7: Image analysis protocol for the analyses conducted in Section 3.4.

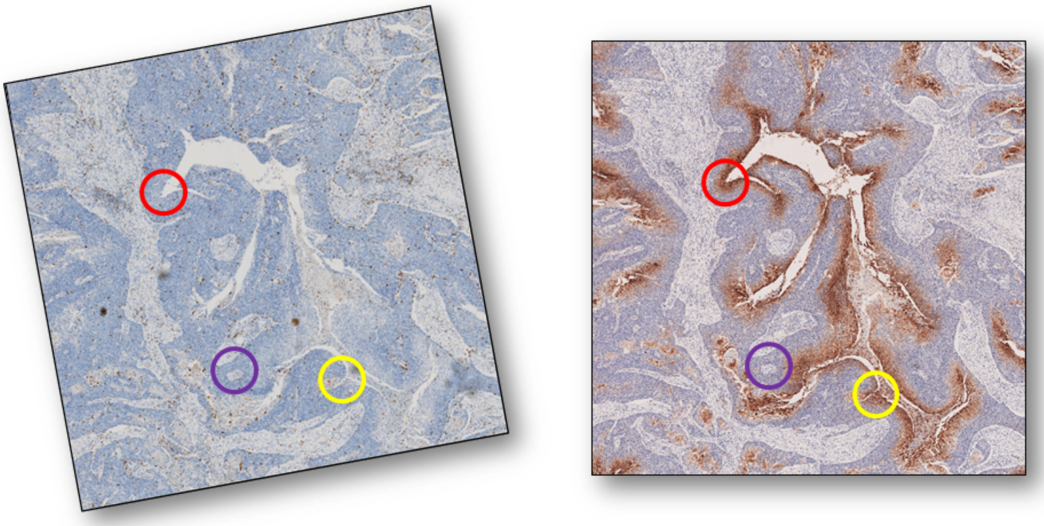

Figure 3.8: Schematic showing the manual selection of three reference points on two consecutive regions of interest.

adjacent slides. Figure 3.9 illustrates this process applied to one region of interest. In Figures 3.9a - 3.9d we identify a (3 mm × 3 mm) region, consisting of the chosen (1.5 mm × 1.5 mm) region of interest together with a 0.75 mm border on each side. This border is included as during registration the rotation and translation of slides causes parts of the slide boundaries to no longer overlap. By registering images with a larger border than required, the registered images can be cropped to the desired 1.5 mm × 1.5 mm size with no boundary effects on the region of interest. The false colour image shown in Figure 3.9e is shown at the correct size, after the boundary regions have been cropped away. In this false colour image, the black region is stroma, blue is PanCK$^+$ pimonidazole$^-$ CAIX$^-$ tumour, purple shows PanCK$^+$ pimonidazole$^-$ CAIX$^+$, magenta is PanCK$^+$ pimonidazole$^+$ CAIX$^+$, and the necrosis annotation is shown in red. In this example, the green cells are CD8$^+$ T cells.

3.4.2 Immune cell in different oxygen environments

Our pipeline allows us to quantify the density of immune cells in each different oxygen environment. We classify each cell as belonging to either the tumour stroma,

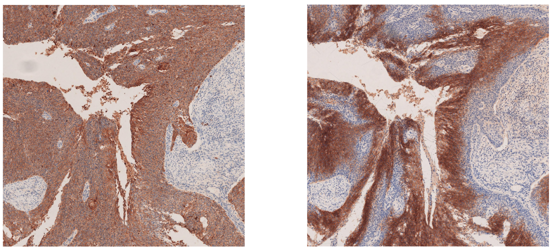

(a) Registered PanCK$^+$ image

(b) Registered CA IX$^+$ image

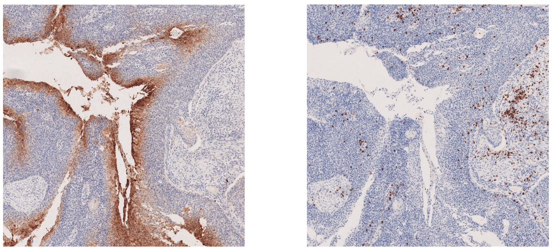

(c) Registered Pimo$^+$ image

(d) Registered CD8$^+$ image

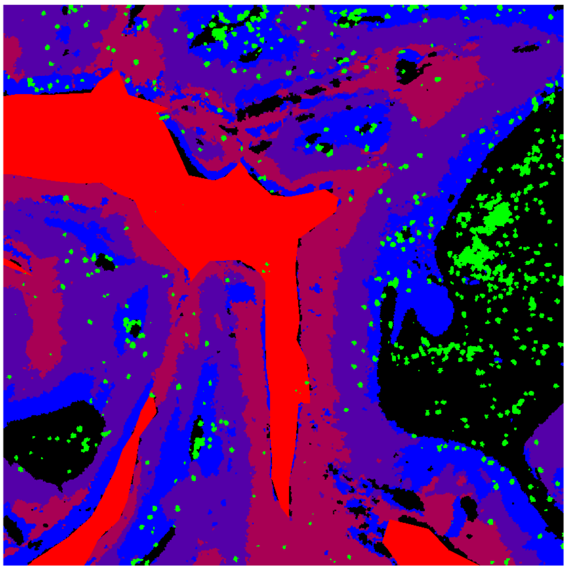

(e) Combined classifications based on Figures 3.9a to 3.9d. In this image, the bright red area is necrosis annotated by a pathologist. Pink/red shows pimonidazole$^+$ stain, purple is CA IX$^+$, blue is PanCK$^+$ and black shows stromal tissue which is negative for all stains. CD8$^+$ cells are shown in green.

Figure 3.9: Original registered images and combined classifications for a randomly chosen 1.5 mm × 1.5 mm region of interest from a head and neck tumour. The regions are registered as 3 mm × 3 mm squares, and then cropped to the central 1.5 mm × 1.5 mm to ensure that artefacts are not introduced at the edge of the regions.

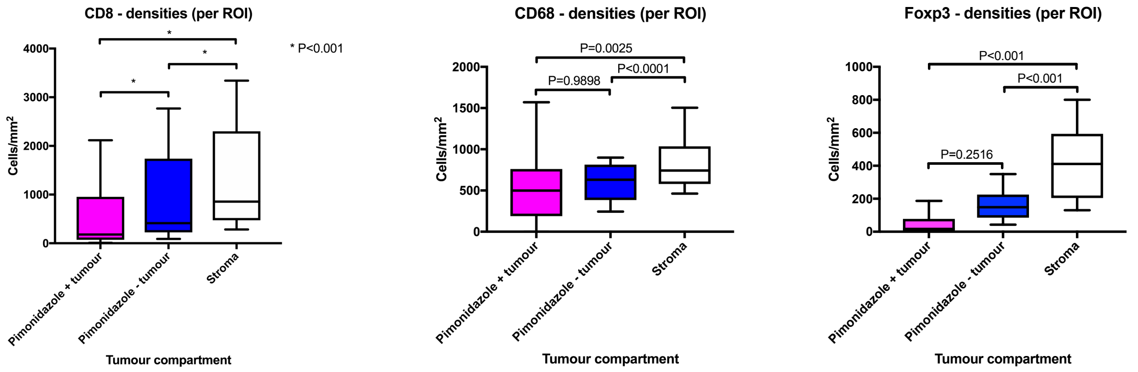

(a) Density of CD8$^+$ cells in each of 258 ROI across 9 tumours

(b) Density of CD68$^+$ cells in each of 128 ROI across 7 tumours

(c) Density of FoxP3$^+$ cells in each of 126 ROI across 7 tumours

Figure 3.10: Densities of three different immune cell subtypes in different tumour environments automatically extracted from randomly placed 1.5 mm × 1.5 mm regions of interest (ROIs) in head and neck tumours. (Number of regions of interest varies due to tumour sizes and ROIs rejected due to excessive artefacts.)

to the pimonidazole$^-$ PanCK$^+$ region ("pimonidazole$^-$ tumour") or to the hypoxic pimonidazole$^+$ PanCK$^+$ region ("pimonidazole $^+$ tumour") for CD8$^+$ cells, CD68$^+$ cells and FoxP3$^+$ cells. We discount the sections of each ROI in which any one of the slides contains an artefact as annotated by a pathologist, and we also discount the necrotic regions of each slide in this analysis due to the degradation of cells in these regions causing counts based on histological data to become unreliable. For each region, we calculate the number of cells per mm^2. These cell densities are presented in Figure 3.10.

The results presented in Figure 3.10 suggest that the density of immune cells within a head and neck tumour depends on the oxygen environment. CD8$^+$ T cells are present in higher densities in the tumour stroma than in normoxic tumour regions, and are present in low densities in hypoxic regions. The same pattern is observed for FoxP3$^+$ regulatory T cells, with lower numbers of regulatory T cells inside the tumour than in the tumour associated stroma. CD68$^+$ macrophages do not follow this trend. While there is a significant difference between the density of macrophages in the stroma and in the tumour, the densities of macrophages in

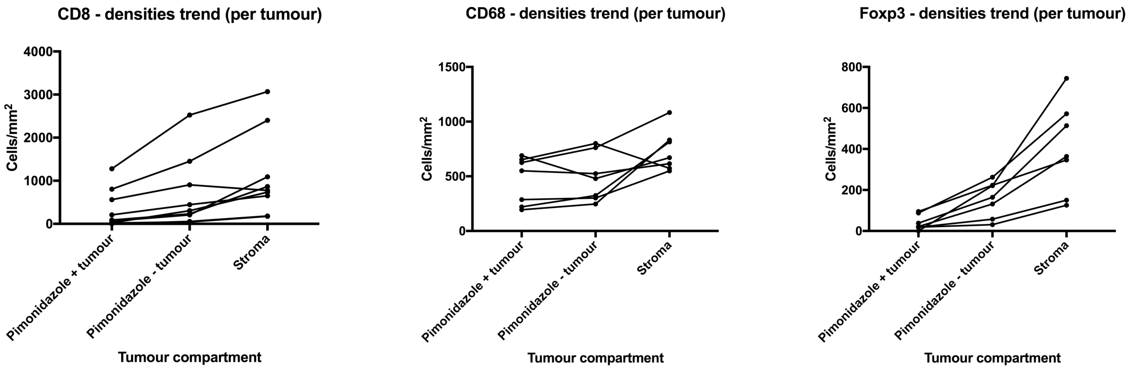

(a) Density of CD8$^+$ cells across 9 tumours

(b) Density of CD68$^+$ cells across 7 tumours

(c) Density of FoxP3$^+$ cells across 7 tumours

Figure 3.11: Trend in densities of three different immune cell subtypes in different tumour environments when regions from the same tumours are linked. Tumours which have high densities of immune cells in the stroma are likely to have high densities of immune cells within the tumour, but the same trends are present in different tumours regardless of the total number of immune cells present.

hypoxic and normoxic tumour regions are very similar. This indicates that while T cells preferentially avoid hypoxia, macrophages within a tumour nest are present at equal densities in both hypoxic and normoxic tumour regions.

This interpretation is reinforced by visualising the results from each tumour individually. In Figure 3.11 we plot the average cell density within each tumour for cells in each oxygen environment by averaging the results from each region of interest within that tumour. For each tumour, the density of cells in the stroma, pimonidazole$^-$ and pimonidazole$^+$ tumour regions is connected with a line in Figure 3.11, showing the change in cell density between tumour compartments. By visualising the data in this way, we see that tumours with a high density of immune cells in one region also have a high density of immune cells in other regions. For each tumour, the pattern is very similar: CD8$^+$ and FoxP3$^+$ cells are densely packed in the stroma, less dense in the normoxic tumour and least dense in the hypoxic tumour. CD68$^+$ macrophages are generally densely packed in the tumour stroma, but are present at similar levels in both the normoxic and hypoxic tumour. Although these observations are based on a small number of tumours, our pipeline provides a

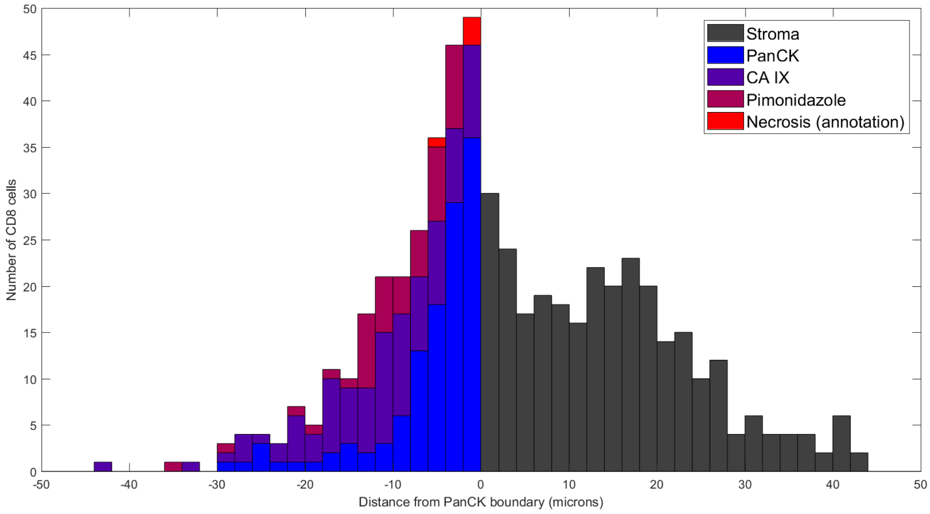

Figure 3.12: Distance of CD8 cells from the PanCK boundary, for cells in the head and neck tumour region of interest shown in Figure 3.9.

simple way to extend this type of study to studies containing many more tumours, and future work will include conducting similar experiments on a larger sample size of head and neck tumours.

Using our quantitative pipeline provides opportunities for innovative data visualisation techniques to describe immune cell distributions. For example, Figure 3.12 shows the distance of each CD8$^+$ cell from the PanCK boundary in the region shown in Figure 3.9. Bars are coloured according to the number of CD8 cells in each oxygen environment. As an alternative to grouping all CD8$^+$ cells in the same oxygen environment together as a cell density, visualising their distances from the tumour/stroma boundary as a histogram reveals that of those T cells observed in hypoxic regions of the tumour, almost all are within 20 μm of the stroma. While a cell density count which indicates that CD8 cells are within a pimonidazole stained region may imply that CD8 cells have moved a long way from the stroma edge, the visualisation used for Figure 3.12 shows that instead the presence of CD8 cells in hypoxic regions is more likely because in some places hypoxia extends almost to the tumour/stroma boundary.

In Figure 3.13 we plot the distribution histograms for CD8$^+$ and CD68$^+$ cells averaged across 71 regions of interest from a head and neck tumour. These distributions clearly show that $CD68^+$ macrophages infiltrate further into the tumour than $CD8^+$ T cells, which cluster in the stroma close to the tumour boundary. More macrophages are observed in the hypoxic and necrotic regions of the tumour than T cells, regardless of how close to the stroma those regions are.

3.4.3 Artificial biopsies

Saturating a tumour with non-overlapping regions of interest enables us to quantify intratumour heterogeneity. In this subsection, we demonstrate that summary statistics which may be naively calculated from a single region of interest in a tumour, such as a biopsy, are not necessarily representative of the true value of that statistic. We use the hypoxic volume fraction, or the proportion of the tumour which is hypoxic, as an example statistic.

To demonstrate this, we use 65 regions of interest taken from one whole slide image to estimate the proportion of the tumour which is stained for pimonidazole. The sample is fully saturated with 65 ($1.5\,\text{mm} \times 1.5\,\text{mm}$) regions of interest. Combining the pimonidazole^+ and PanCK^+ areas from all 65 tumour regions, we calculate the proportion of the tumour slide which is hypoxic via the following formula:

$$\text{Proportion of the tumour which is hypoxic} = \frac{\text{PanCK}^+ \text{ pimonidazole}^+ \text{ area}}{\text{PanCK}^+ \text{ area}} = 0.1503. \tag{3.7}$$

Equation (3.7) says that approximately 15% of the tumour in our example is hypoxic. We attempt to reproduce this result based on “artificial biopsies” from our dataset, in which we randomly select n of the 65 regions and calculate the proportion of the tumour which is hypoxic based on that sample alone. By repeatedly making this prediction for different values of n, and different randomly chosen regions, we ask what the minimum number of ROIs required to recover the “true” value of

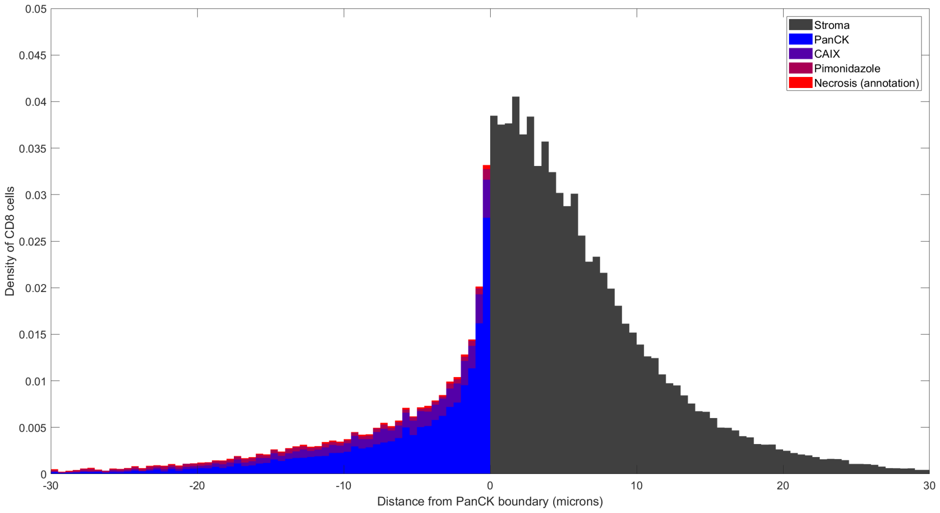

(a) $\mathrm{CD8}^{+}$ cells

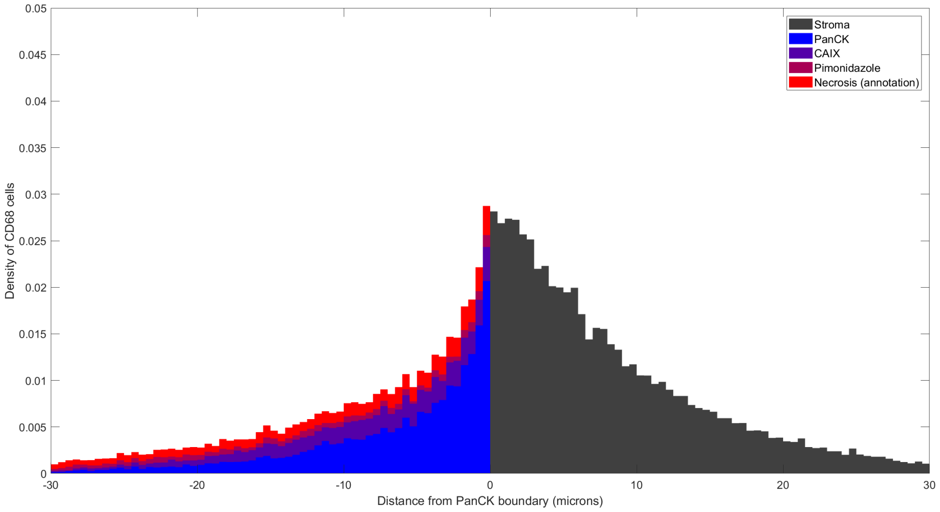

(b) $\mathrm{CD68}^{+}$ cells

Figure 3.13: Distributions of $\mathrm{CD68}^{+}$ macrophages and $\mathrm{CD8}^{+}$ T cells, averaged across 71 (1.5 mm × 1.5 mm) ROIs in one head and neck tumour.

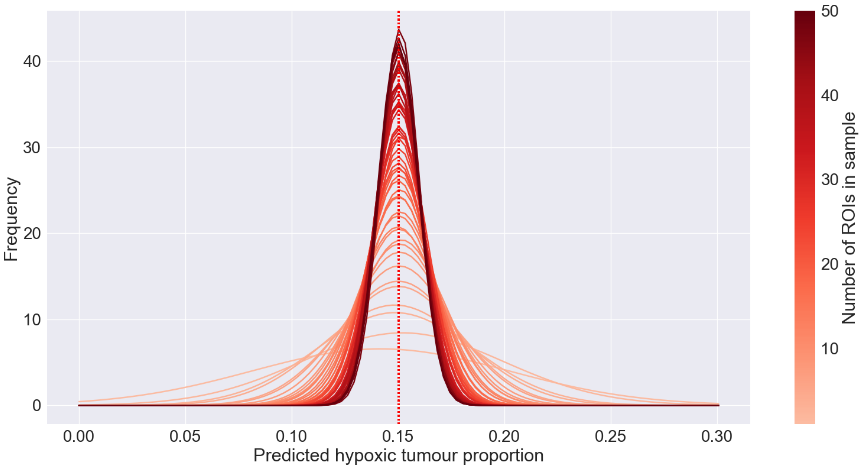

Figure 3.14: Distributions of the predicted hypoxic tumour proportion for a head and neck tumour. The true hypoxic tumour proportion based on all 65 regions of interest needed to saturate the tumour, 0.1503, is shown as a dashed red line. To make a prediction, we draw n samples from the 65 regions, for n between 1 and 50. For each value of n, the distribution shown is generated by randomly choosing n of the 65 regions and calculating the predicted hypoxic tumour proportion based on this subsample, and repeating the process 1000 times.

0.1503 is.

There are 2080 ways of choosing two regions from 65, and over 3×10^{18} for $n = 33$. Hence for each number of sample ROIs, n, we draw n samples 1000 times and calculate the hypoxic volume fraction of the tumour contained in that sample. These 1000 samples give a distribution for the predicted hypoxic tumour proportion. Figure 3.14 shows these distributions for $n = 1$ to 50.

The distributions shown in Figure 3.14 converge towards the true value as the number of samples increases. To determine whether n samples suffice to accurately estimate the proportion of the tumour that is hypoxic, we investigated how many regions of interest must be sampled to return a prediction within a given distance of the true value.

Figure 3.15 indicates that our tolerance of error determines whether our approximation is acceptably accurate. If we only count a prediction as "accurate" if it is with 5% of the true value, then sampling as many as 50 regions of interest re-

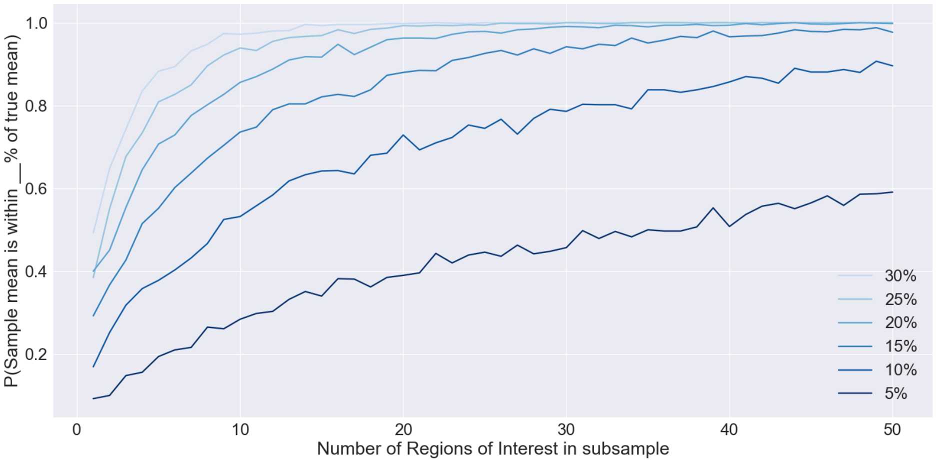

Figure 3.15: Effect of increasing n, the number of ROIs sampled, on the predicted value of the hypoxic volume fraction of the tumour. Taking 50 ROIs in a sample is sufficient to return a prediction within 5% of the true value in only 60% of samples. However, increasing the tolerance to accept predictions within 30% of the true value returns an acceptable prediction with fewer than 10 samples in over 95% of samples.

turned an accurate prediction for only 60% of the 1000 samples. However, accepting predictions within 10% of the true value dramatically increases the number of subsamples of the 1000 iterations which return an acceptably accurate value, with only 13 ROIs needed to reach 60% accuracy. Sampling only 2 regions of interest yields a predicted value within 30% of the true hypoxic proportion for 60% of the samples. We observe similar results when applying this technique to predict the true value of a range of summary statistics, including infiltrating immune cell densities and tumour to stroma ratios within the tumour mass.

Figures 3.14 and 3.15 have implications about the predictions that can be made from clinical biopsies. Our results indicate that it is difficult to accurately predict tumour-wide summary statistics from summary statistics of an individual region of interest. Figure 3.15 suggests that predicting summary statistics from a single (1.5 mm × 1.5 mm) tissue section will yield results within 30% of the true value for only 50% of samples. However, sampling 2 or 3 regions from across the tumour increases this acceptance rate to almost 80%. Our simulations therefore suggest that biopsies should consider samples from different tumour regions in order to make accurate predictions, rather than relying on one portion of the tumour to provide data representative of the whole tumour. The number of samples required to achieve acceptable accuracy will vary depending on the size of the tumour and on how much uncertainty can be tolerated in the predicted values.

3.5 Adapting the pipeline to count microbeads

Our pipeline is sufficiently flexible that it can be used to analyse a variety of different images. In this Section, we show how our pipeline has been used for microbead counting in experiments by Mr Hugh Ford, a PhD candidate at the University of Sydney. Mr Ford studies the role of macrophages in atherosclerosis, a disease in which fatty plaques build up inside arteries (Murray and Wynn, 2011).

Mr Ford's experiments require counting the number of coloured plastic microbeads consumed by macrophages over time. Macrophages are cultured *in vitro* and a large number of coloured plastic microbeads are added to the surface of the substrate. At regular intervals, an image is taken showing the number of beads inside each macrophage, which must then be counted. As macrophages die, their remains are consumed by other macrophages, and the microbeads accumulate inside small numbers of macrophages. The experiment aims to track the distribution of microbeads within cells, and compare this with the predictions of a mathematical model developed by Mr Ford to determine the dynamics of macrophage particle consumption. Figure 3.16 shows the type of images generated during this experiment. Red and blue microbeads are used, and they appear in the images as regularly sized circles.

We have adapted our pipeline for the automatic counting of microbeads inside these macrophages. As there are only macrophages and beads in the images, our pipeline must identify three types of objects, macrophages and both red and blue

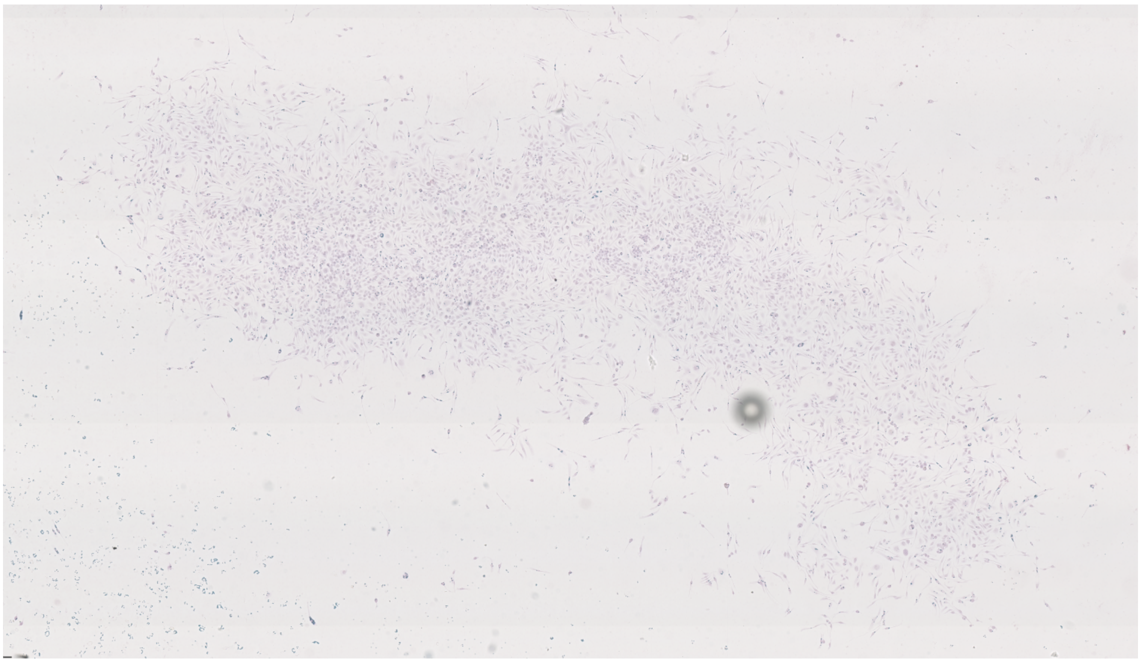

(a) Low-power image showing all macrophages in one experiment.

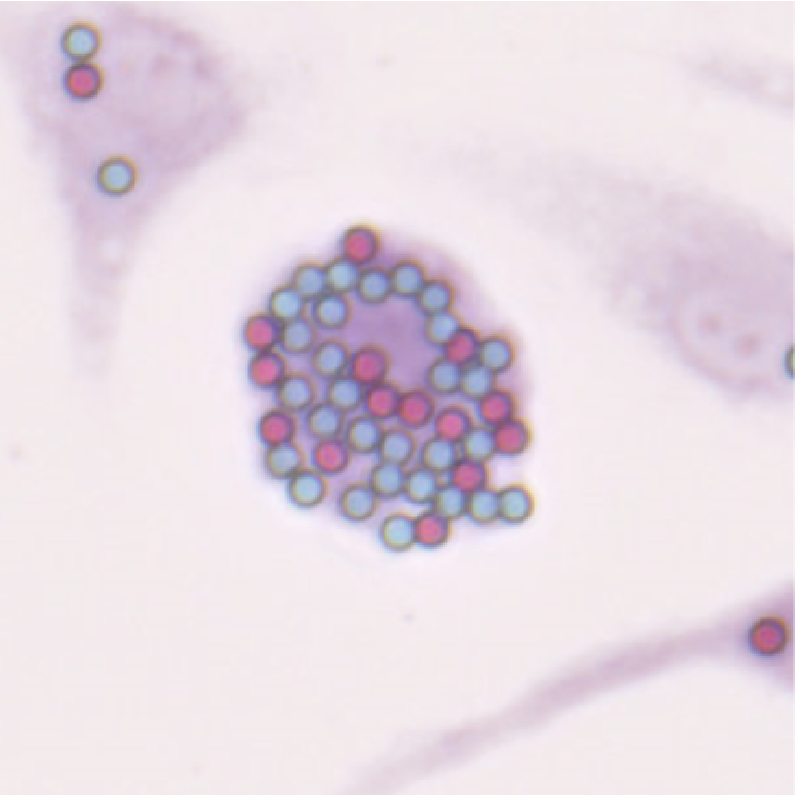

(b) High-power image showing blue and red coloured beads inside macrophages.

Figure 3.16: Experiments showing microbead consumption by macrophages.

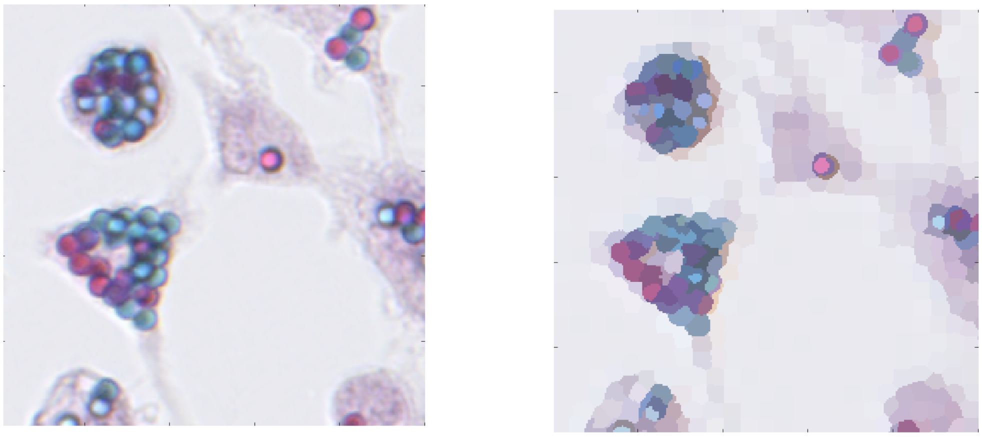

Figure 3.17: Comparison between detail from original image and the superpixellated version used for classification.

microbeads. Individual macrophages must be distinguished, and the number of microbeads within each macrophage counted. It typically takes several hours to analyse a single image manually, but our pipeline can process each image in minutes. In Figure 3.17 we illustrate the superpixellation stage for a section of an image. The microbeads and macrophages differ in size, and so the superpixel sizes are chosen to match the smaller area of the microbeads. As a result, individual microbeads are generally represented by only one superpixel, while macrophages are oversegmented. Microbeads are easily identified in these images due to their regular size and the black "halo" that can be observed around each bead (see Figure 3.16b) which provides clear context information for our SVM classifiers. On the other hand, macrophage boundaries can be difficult to identify due to the often low contrast between the macrophage and the background.

We overcome the problem of macrophage misclassification by varying the parameters for cell identification from those used to analyse the histological images. The initial classification, shown in Figure 3.18a, contains several small superpixels which have been misidentified as macrophages. We can exploit the fact that these isolated superpixels are much smaller than a typical macrophage to exclude them from our count of macrophages. Choosing a large opening size successfully separates adjacent macrophages, and the final classification mask which we use to obtain statistics is shown in Figure 3.18b.

The number of microbeads in each macrophage can be obtained from Figure 3.18b by identifying individual macrophages using basic edge detection methods. Once a macrophage is identified, the total area classified as “microbead” can be compared with the known area of a microbead in cross-section to obtain the statistics for comparison with Mr Ford’s mathematical model.

3.6 Discussion

In this Chapter we have demonstrated a variety of applications of the image analysis pipeline presented in Chapter 2. We have discussed techniques which can be used for the analysis of point data extracted from WSIs. Using a range of spatial statistics, which would be infeasible to calculate without automated cell identification tools, we characterised samples of histology images from the RITA dataset based on their different patterns of macrophage infiltration. We have also demonstrated that these statistics can be used to cluster randomly chosen tumour samples into groups with qualitatively similar patterns of macrophage infiltration. This approach to classifying immune cell infiltration has multiple possible applications, and we have demonstrated that spatial statistics can provide a basis to describe qualitative patterns in immune cell infiltration. Widespread adoption of these statistics could permit clearer comparison of the extent of immune infiltration between different tumours, which has been identified as an important factor in predicting patient response to immunotherapeutic treatments (Mantovani and Allavena, 2015; Chanmee et al., 2014).

Using data from head and neck tumours, we have also used our pipeline to combine information from multiple consecutive histological slides to characterise the oxygen environment of individual immune cells. By characterising whole slide

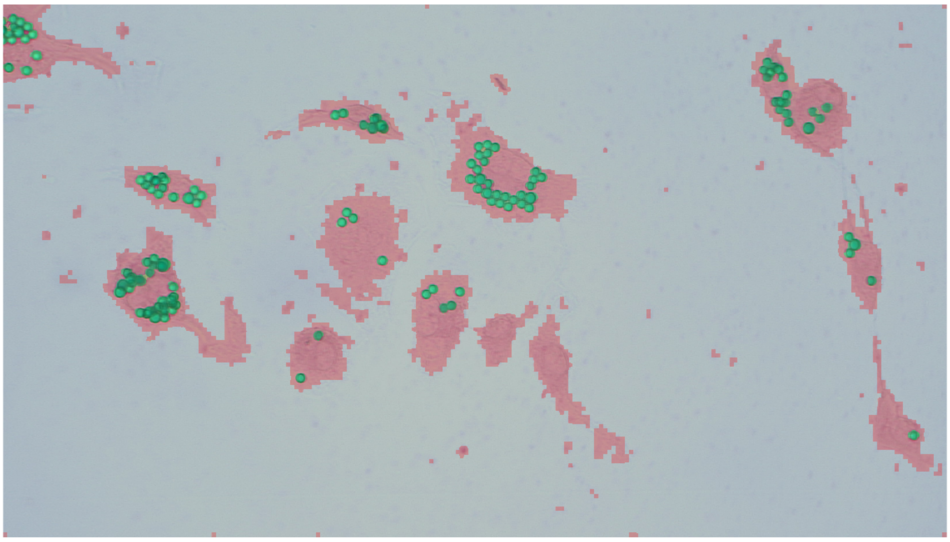

(a) Initial classification overlaid on the image.

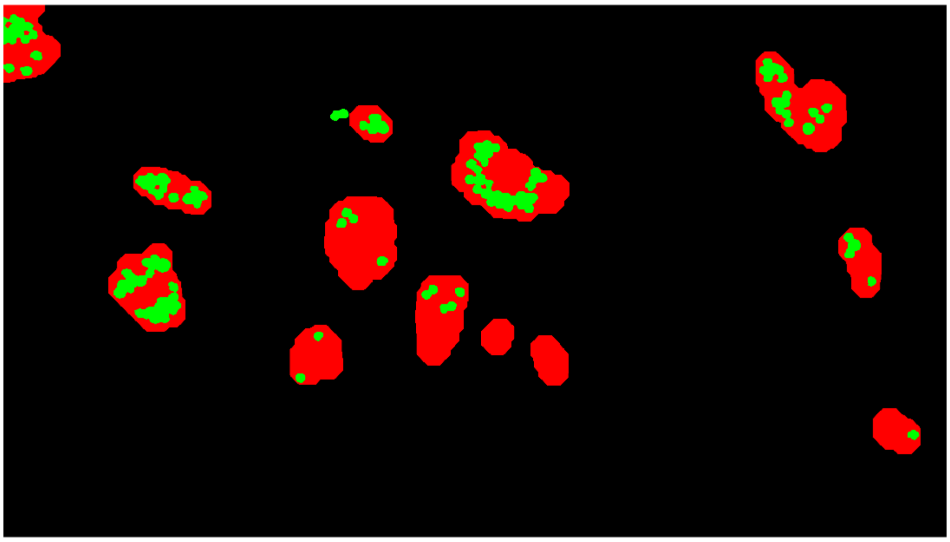

(b) Refined classification after postprocessing used to generate statistics.

Figure 3.18: Effect of postprocessing a classification based on morphological opening and superpixel size.

images, we used our pipeline to assess the accuracy of predicting summary statistics based on only part of the image. Our results show that biopsies which consider only a small part of the tumour are unlikely to accurately represent the whole tumour, and that using more regions when making predictions about the whole tumour improves the accuracy of the prediction. Future work in this area will involve applying this methodology of "artificial biopsies" to a larger range of tumour types and statistics. This approach could be used to quantify the uncertainty around approximations made from patient biopsies, where the amount of tissue which can be extracted is restricted. It could also be used to suggest how large a biopsy should be taken from a patient in order for the pathologist to have the best chance of making correct predictions about the remaining tumour.

As we develop the pipeline described in Chapter 2, we intend to integrate the collection of spatial statistics into the user interface. For instance, automatically identifying regions of a WSI which display macrophage distributions associated with positive or negative patient outcomes and highlighting them to the user can be achieved without requiring any knowledge of spatial statistics on the part of the pathologist. As in Section 2.5, we once again emphasise that widespread uptake of quantitative techniques for computational pathology can only be achieved if they are made accessible to pathologists. Many pathologists have, sadly, not studied mathematics or statistics at a high level and are rightly wary of using techniques they do not understand to make decisions which will determine the treatment of a patient. Mathematicians must demonstrate that these techniques are valuable and robust, and the first step to familiarising pathologists with computational techniques is to demonstrate that these methods work in practice as well as in theory.

As well as providing a valuable tool for pathologists, we also suggest that second-order statistics could provide a means to integrate *in vivo* images with spatial mathematical models. While summary statistics can be easily extracted from many agent-based models, spatial statistics can be used to describe emergent patterns in models which have a complicated spatial structure (Fozard et al., 2011). Spatial statistics applied to point data which arises from the agents in an individual-based model can be directly compared with the same statistics from *in vivo* data. This approach provides the potential for mechanistic models to be compared with histological datasets, providing insight into how the different distributions of macrophages in *in vivo* images arise.

In Chapter 4 we present an agent-based model to investigate the mechanisms which cause these distributions. The model simulates *in vitro* experiments which describe the infiltration of inert microbeads into tumour spheroids, and provides a foundation for Chapter 5 of this thesis, in which we model macrophage infiltration into both *in vitro* spheroids and *in vivo* tumour geometries extracted from histology images.

Chapter 4

Agent-based modelling of *in vitro* tumour spheroids

4.1 Overview

In Chapter 1, we discussed the importance of macrophage infiltration in solid tumours and the value of understanding what causes the distinct patterns of macrophages observed *in vivo*. We discussed several different biological phenomena which may regulate the distribution of macrophages within the tumour environment, including tumour production of chemoattractants such as colony stimulating factor-1 (CSF-1) and chemokine (C-C motif) ligand 2 (CCL2), random macrophage motility and force generated through unchecked proliferation of tumour cells and their interactions with the tumour microenvironment. In this Chapter, we investigate the advective flow of cells created by the localisation of tumour cell proliferation at the spheroid edge and hypoxia-induced cell death in the spheroid centre. We use tumour spheroids (Folkman and Hochberg, 1973) as a model system to explore the movement of tumour cells within the spheroid, and take as our starting point experimental work by Dorie et al. (1982) on the movement of inert microbeads into these spheroids. Subsequently, we develop an agent-based model, implemented within the

Chaste modelling framework (Mirams et al., 2013), to simulate microbead infiltration within tumour spheroids. We conduct a parameter sweep over several model parameters to determine their influence on the microbead infiltration dynamics, and conclude that knowledge of the spatial structure of a spheroid induced by oxygen gradients is vital for characterising the dynamics of this type of "passive" infiltration. Our use of an agent-based model differs from previous mathematical models of (Dorie et al., 1982), and we show that our model can explain the dynamics of microbead infiltration into tumour spheroids in more detail than continuum models. Dorie et al. (1982) also report results of the infiltration of ^{3}H-labelled cells into tumour spheroids, and we show that our model can explain the differences in observed distributions of microbeads and ^{3}H-labelled cells. We conclude that our understanding of immune cell dynamics in tumours must be placed in the context of tumours as a dynamic environment, and that the advective flow which causes the infiltration of inert particles must be accounted for when modelling immune infiltration of solid tumours. This motivates the work in Chapter 5, in which we will discuss the role of chemotaxis in driving macrophage infiltration into tumours, using the results from this Chapter to decompose macrophage movement into advective and chemotactic components.

The remainder of this Chapter is structured as follows. In Section 4.2 we describe experiments by Dorie et al. (1982) in which inert polystyrene microspheres and cells labelled with tritiated thymidine infiltrate multicellular tumour spheroids, and we review previous mathematical models of this system. In Section 4.3 we introduce our new agent-based model of spheroid growth implemented within the Chaste computational environment (Mirams et al., 2013). In Sections 4.4 and 4.5 we present results for spheroid growth dynamics and the infiltration of the microspheres. Results from simulations of the labelled cell experiments are presented in Section 4.6. Finally, Section 4.7 contains a brief summary of our results.

4.2 Experimental basis for the mathematical model

Tumour spheroids are widely used *in vitro* to study the growth and response to treatment of solid tumours. We begin this Section with a short description of how spheroids are grown *in vitro* and the characteristic features of spheroids which make them a useful model of *in vivo* tumours. We then discuss the data described by Dorie et al. (1982) which we use as the basis for our mathematical model in this Chapter. We briefly discuss previous mathematical models of this data in Section 4.2.1 and highlight the differences between our approach and previous approaches.

Tumour spheroids are a good model system for the study of immune infiltration of tumours as they provide a much simpler environment than an *in vivo* tumour. The tumour cells grow *in vitro* in response to local levels of vital nutrients (e.g., oxygen and glucose) which diffuse through the spheroid from the surrounding culture medium. Intercellular adhesion causes the cells to clump in an approximately spherical shape, and all cells in the spheroid continue to proliferate as long as they have sufficient levels of nutrients and oxygen. As nutrients penetrate the spheroid they are taken up by the cells. The associated reduction in nutrient levels gives tumour spheroids a well-defined and reproducible structure. A nutrient-rich proliferative rim surrounds a central nutrient-deprived necrotic core where the concentration of oxygen and other nutrients is too low to sustain the cells. Figure 4.1 shows a schematic diagram of this structure. Typically, a spherical shell of quiescent tumour cells can be found between the proliferative rim and the necrotic core. These quiescent cells have sufficient oxygen to survive, but not enough to proliferate. If their oxygen tension drops too low, quiescent cells may undergo cell death. We refer to these viable cells which are at risk of death as “hypoxic” cells, although experimentally they cannot easily be distinguished from the quiescent cells with sufficient oxygen to survive. If oxygen levels increase sufficiently, hypoxic cells can resume proliferation.

Multicellular tumour spheroids typically grow exponentially in their early stages, before following a logistic growth trajectory until they reach a maximal carrying ca-

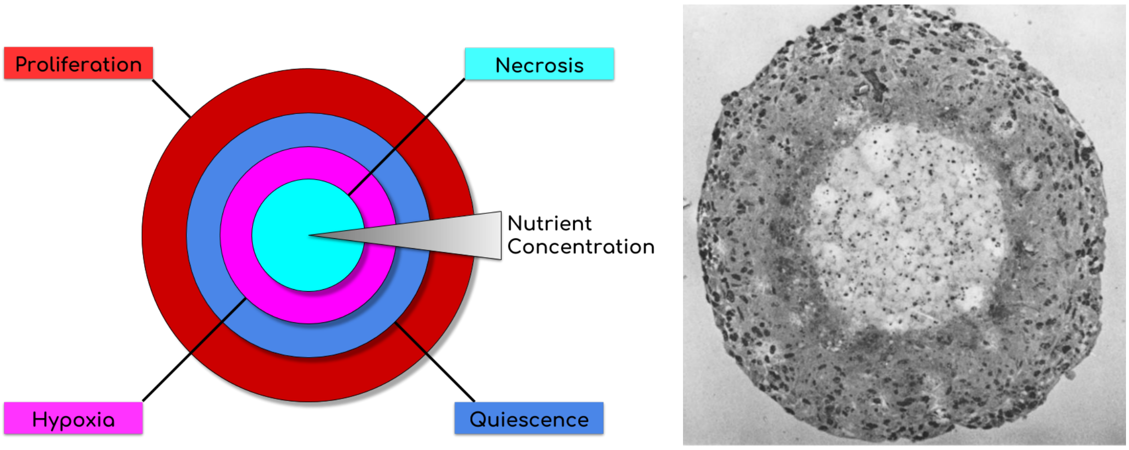

(a) (b)

Figure 4.1: A schematic view of a tumour spheroid (left), together with an image of a tumour spheroid (right) reproduced from (Folkman and Hochberg, 1973).

pacity which they cannot grow beyond (Folkman and Hochberg, 1973). Tumour spheroids are avascular, and generally are not used as a model for complex interactions between tumour cells and the extracellular matrix, although assays investigating these interactions exist (Nath and Devi, 2016; Herter et al., 2017). Experiments have been conducted examining the infiltration of macrophages into tumour spheroids in the presence and absence of chemoattractants, and tumour spheroids are considered a good resource for studying immune infiltration into *in vivo* tumours (Leek, 1999).

In this Chapter, we present a mathematical model based on spheroid experiments by Dorie et al. (1982, 1986) which investigate the movement of tumour cells due to proliferation and hypoxia-induced death. In these experiments, the internalisation of polystyrene microbeads within a tumour spheroid was monitored. The multicellular tumour spheroids Dorie et al. (1982) described are from the EMT6 (mouse, breast mammary carcinoma) and RIF-1 (mouse, sarcoma) cell lines. They were grown *in vitro* in spinner flasks while suspended in a nutrient rich medium, with oxygen diffusing to the spheroid from a gas phase at constant oxygen concentration (95% air, 5% CO_2) (Dorie et al., 1982). The proliferative band of tumour cells (red in Figure 4.1a) was observed to have an approximately constant width of 180 μm, with

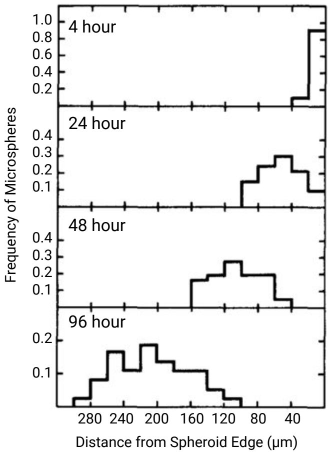

Figure 4.2: Distribution of microspheres in EMT6 spheroids, adapted from Figure 7 in (Dorie et al., 1982). Note the unusual x-axis, which we have reflected in order to keep the spheroid edge on the right hand side of the Figure and maintain consistency throughout this thesis.

the increase in spheroid size during spheroid growth being accounted for by the increased radius of the necrotic region (cyan in Figure 4.1a). The proliferative band moved radially outwards at a rate of approximately $20\,\mu$m per day.

Dorie et al. (1982) used inert polystyrene microbeads to measure the internal dynamics of the spheroids. These microbeads were added to the surface of the spheroids, and their infiltration was tracked over time by sacrificing spheroids 4, 24, 48 and 96 hours after the beads were added. Radial distributions of microbead infiltration were obtained as a time series of frequency histograms; an example is shown in Figure 4.2 (adapted from Figure 7 in (Dorie et al., 1982)). The microbeads form a moving distribution which becomes more dispersed as they travel towards the centre of the spheroid.

In separate experiments reported by Dorie et al. (1982), the dynamics of the spheroids were examined by tracking the progression of labelled tumour cells rather than microbeads. Spheroids grown from both the EMT6 and RIF-1 cell lines were continuously exposed to [^{3}H]thymidine to radioactively label the cells. These spheroids were disaggregated and then added to the edges of either the EMT6 and RIF-1 spheroids following the same protocol as the microspheres. The labelled cells were introduced in both homotypic combinations (^{3}H-labelled EMT6 cells in EMT6 spheroids, ^{3}H-labelled RIF-1 cells in RIF-1 spheroids) and in heterotypic ones (^{3}H-labelled EMT6 cells in RIF-1 spheroids, ^{3}H-labelled RIF-1 cells in EMT6 spheroids). The distribution of these labelled cells was tracked in the same manner as the microspheres.

The results of these experiments indicated that the infiltration pattern of ^{3}H-labelled tumour cells is different to that of microbeads. Although the cells still move in a wave towards the centre of the spheroid which disperses over time, there is a peak which remains close to the edge of the spheroid. Specifically, this appears to represent a population of ^{3}H-labelled cells which are unaffected by the advective flow and remain clustered on the spheroid boundary.

4.2.1 Previous mathematical models of data from Dorie et al. (1982)

Microbeads

The data described by Dorie et al. (1982) has motivated a series of mathematical models, which we describe in this Section. McElwain and Pettet (1993) formulated a continuum partial differential equation (PDE) model and showed that a possible cause of bead internalisation is pressure gradients caused by differential cell proliferation and death between the proliferative rim and the necrotic core of the spheroid. Their model captures the observed movement and dispersal of the microbead distribution. To reproduce the dynamics of the ^{3}H-labelled cells, the authors assumed that the tumour cells move via chemotaxis towards the higher oxygen levels on the spheroid boundary. The model does not distinguish between proliferative, quiescent and hypoxic cells, and so is unable to capture common spheroid structures such as the proliferative rim and necrotic core. The authors argued that the observed inward motion of beads is "passive" and due to an advective flow, and that while tumour cells may attempt to migrate via chemotaxis up the nutrient gradient they are in fact “swimming against the tide” (McElwain and Pettet, 1993).

Thompson and Byrne (1999) developed an alternative PDE model of this system which reproduces the data in the absence of chemotaxis. They assumed that multicellular tumour spheroids are incompressible and introduced a more realistic proliferation rule. They argue that different internalisation patterns of microbeads are due solely to nonuniform death and proliferation in the spheroids. However, this model assumes for simplicity that the tumour spheroids do not contain a necrotic core, and dead cells do not occupy space. Notably, their model is unable to reproduce predicted infiltration patterns unless the microsphere wave begins at a distance of 80% of the spheroid radius, i.e. inside the tumour rather than from the tumour boundary.

Pettet et al. (2001) modify the model presented in (McElwain and Pettet, 1993) to distinguish between proliferative and quiescent cells. This adaptation is more successful at reproducing the experimental results, but they still retain the assumption that the dead cells are immediately removed from the spheroid and do not occupy space, so that the model does not capture the formation of a necrotic core at the spheroid centre. They assume that proliferative and quiescent cells react differently to the chemotactic gradient which attracts them towards the high nutrient concentrations at the spheroid boundary. The authors argue that proliferative cells, which are expending energy for mitosis, are less motile than quiescent ones, and that the quiescent cells will sense the nutrient gradient and follow it towards the spheroid edge. As such, queiscent cells in their model attempt to move via chemotaxis while proliferative cells do not.

We also note that there are similarities between these models and several of the models of macrophage infiltration described in Section 1.4.1, in particular (Kelly et al., 2002), (Owen et al., 2004) and (Webb et al., 2007).

^{3}H-labelled cells

Reproducing the dynamics of the ^{3}H-labelled cells has proven problematic to previous continuum models. McElwain and Pettet (1993) reproduce similar behaviour to that observed experimentally by including a chemotactic component to the movement of tumour cells. All tumour cells in their model are attracted up the nutrient/oxygen gradient towards the spheroid boundary, which provides an outward force. By choosing parameters appropriately, the authors are able to generate an advective flow moving outwards from the spheroid centre which causes labelled cells to remain at the spheroid boundary. The assumption of chemotaxis drawing cells towards the boundary appears questionable, and the authors do not provide experimental evidence for this type of chemotactic movement. Their model also fails to include proliferation of the ^{3}H-labelled cells, something which our modelling indicates plays a key role in the observed distributions of cells.

Thompson and Byrne (1999) do not include such a chemotactic rule in their model, and provide a more successful explanation for the observed clustering at the spheroid boundary by using the assumption that the spheroid is incompressible and implementing a proliferation rule for the ^{3}H-labelled cells. The labelled cells in their model do not cluster in the spheroid centre, but the experimentally observed peak at the spheroid boundary is still absent. Pettet et al. (2001) distinguish between proliferative and quiescent subtypes of cells, but dead cells are instantaneously removed from the simulation and as such their model cannot reproduce a necrotic core. By varying chemotaxis coefficients between the subpopulations, the authors are able to reproduce the "classical" expected infiltration patterns. The authors note that in order to reproduce the experimental behaviour, both proliferative and quiescent cells in their model always have consistently negative velocity profiles. This means that they will always be advected towards the spheroid centre.

In this Chapter, we propose an agent-based model of spheroid growth motivated by the data presented by Dorie et al. (1982). Using a cell-scale model to simulate the experiments has several advantages over continuum models. First, each microbead can be tracked individually through the simulation. This enables us to see how the movement of individual microbeads causes the dispersing distribution observed in the experimental data. Continuum models cannot distinguish between dispersion caused by beads moving at different speeds and dispersion caused by a delay in bead infiltration, which we discuss in detail in Section 4.5.2. Agent-based modelling also extends naturally to the asymmetric geometries observed in *in vivo* tumours, permitting the extension of this model presented in Chapter 5. We present this model in Section 4.3.

4.3 Model development

The models reviewed in Section 4.2.1 reproduce several important features of "passive" migration, but they are limited by describing the evolution of populations of cells and beads instead of resolving individual cells or microbeads. In this Section we present an off-lattice agent-based model of tumour spheroid growth and then use it to study the passive internalisation of microbeads. Our model is implemented within the cancer, heart and soft tissue environment (Chaste) modelling framework (Mirams et al., 2013; Pitt-Francis et al., 2009) described in Section 1.4.2, which provides an object-oriented means of implementing agent-based models of tumours.

4.3.1 Choice of modelling framework

Any of the frameworks described in Section 1.4.2 could be used to simulate the experiments described in (Dorie et al., 1982), and we choose to use an overlapping spheres framework for our model. Overlapping spheres is a centre-based framework, which permits individual beads and cells to be represented by single nodes, avoiding the additional computational complexity that a vertex-based framework would entail by modelling cells as collections of vertices. Centre-based methods also allow us to track cells over time by following the cell centre, while the changing shapes of cells in a vertex model make it less intuitive to track cells over time. While a Voronoi tessellation framework could also be appropriate, the irregular size and shape of cells in this framework would make the addition of microspheres more complex than in an overlapping spheres framework.

In our model, the centre of each tumour cell is modelled as a point which is connected to its neighbouring cells, defined as those which lie within a "radius of interaction", R_{int}. If two cells are within this proximity, their centres are connected with a spring and the spring resting length therefore represents an equilibrium cell diameter. Cell movement is determined via a force balance applied to each cell, accounting for intercellular forces exerted on cells by their neighbours and stochastic fluctuations that represent random cell movement. The following Sections contain the equations solved at each timestep in our model and a more detailed description of the forces present in our model.

4.3.2 Oxygen distribution

While we model cells and microbeads discretely, we assume that the distribution of oxygen throughout the domain is continuous and that its spatio-temporal evolution can be described via a PDE. At the beginning of each timestep, the oxygen concentration at time t and point $\mathbf{x}$, $\omega(\mathbf{x}, t)$, is updated by solving a reaction-diffusion equation, with point sinks at the centre of viable cells modelling oxygen consumption by live tumour cells.

The equation governing the spatio-temporal evolution of the oxygen concentration at position $\mathbf{x}$ is thus

$$\frac{\partial \omega}{\partial t} = D_\omega \nabla^2 \omega - \kappa \omega \sum_i \delta(\mathbf{x} - \mathbf{x}_i). \tag{4.1}$$

where $\mathbf{x}_i$ is the location of cell i, the parameter κ controls the oxygen consumption rate of tumour cells and the parameter D_ω is the diffusion coefficient of oxygen. $\delta(\mathbf{x} - \mathbf{x}_i)$ is a delta function which equals 1 when the argument is zero, and is otherwise equal to zero. We rescale Equation (4.1) using the following scales:

$$\tilde{\omega} = \frac{\omega}{\omega_0}; \qquad \tilde{t} = \frac{t}{t_0}; \qquad \tilde{x} = \frac{x}{x_0}; \qquad \tilde{y} = \frac{y}{x_0} \tag{4.2}$$

to obtain a non-dimensional form of Equation (4.1):

$$\frac{\partial \tilde{\omega}}{\partial \tilde{t}} = \tilde{D_\omega} \nabla^2 \tilde{\omega} - \tilde{\kappa} \tilde{\omega} \sum_i \delta(x_0(\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_i)). \tag{4.3}$$

where

$$\tilde{D_\omega} = \frac{D_\omega t_0}{x_0^2} \qquad \text{and} \qquad \tilde{\kappa} = \kappa t_0. \tag{4.4}$$

We choose $t_0 = 1$ hour as the timescale and $x_0 = 1$ cell diameter as the lengthscale. Dropping the tildes for convenience, the non-dimensional equation describing the spatio-temporal evolution of the oxygen concentration at position $\mathbf{x}$ is thus

$$\frac{\partial \omega}{\partial t} = D_\omega \nabla^2 \omega - \kappa \omega \sum_i \delta(\mathbf{x} - \mathbf{x}_i) \tag{4.5}$$

where $\mathbf{x}_i$ is the location of cell i, and values of the dimensionless parameters κ and D_ω coincide with their dimensional equivalents due to our choice of timescales and lengthscales in Equation (4.4). We note that for our choices of lengthscale and timescale the non-dimensional Equation (4.5) and the dimensional Equation (4.1) are identical, up to a rescaling of the oxygen concentration. Throughout this Chapter we will generally, for convenience, refer to time in hours and distance in cell widths, but we note that due to our choice of rescaling these values are the same as their dimensionless equivalents.

As tumour spheroids are avascular and obtain their oxygen through diffusion, we assume that oxygen diffuses in from the edge of the spheroid, where its concentration is constant. Each cell in our model is represented by a single point, so the boundary of the spheroid must be carefully defined in order to apply boundary conditions. We use the concept of an α-shape (Edelsbrunner et al., 1983) to provide our definition of which cells belong to the spheroid boundary in a simulation, which we describe in Appendix A.

We use the oxygen concentration on the boundary, which we assume is a constant ω_∞, to define the scaling factor ω_0 used for non-dimensionalisation. Setting $\omega_0 = \omega_\infty$, our boundary condition becomes that $\omega(\mathbf{x}, t) = 1$ on the spheroid boundary. We apply Dirichlet boundary conditions, $\omega(\mathbf{x}, t) = 1$ to the nodes in the α-shape. Since the timescale for oxygen diffusion is much faster than that of cell movement, we use a quasi-steady state approximation to solve Equation (4.5):

$$0 = D_\omega \nabla^2 \omega - \kappa \omega \sum_i \delta(\mathbf{x} - \mathbf{x}_i). \tag{4.6}$$

We solve Equation (4.6) on a regular tetrahedral finite element mesh. This mesh is defined over a square domain large enough to contain all the cell centres. Any nodes of the mesh outside the α-shape are assumed to have concentration $\omega = 1$.

4.3.3 Cell phenotypes and cell cycle model

Cell phenotypes

To capture oxygen limited spheroid growth we distinguish three phenotypes of tumour cells based on the local oxygen concentration, ω. We also simulate dead cells, which do not have a cell cycle but occupy space in the model. The relationships between these compartments are described in Figure 4.3.

- A tumour cell is **proliferative** if its local oxygen concentration exceeds a threshold value, ω_q.
- For $\omega \leq \omega_\mathrm{q}$, the cell becomes **quiescent** and stops proliferating immediately. If its oxygen level increases above ω_q it will immediately resume its cell cycle.

- There is a second threshold, $\omega_\mathrm{h} \leq \omega_\mathrm{q}$, below which the cell is deemed to be **hypoxic**. A hypoxic cell can re-enter the quiescent compartment if its oxygen level rises back above ω_h. If a cell's oxygen concentration remains below this threshold for sufficiently long (τ_ρ hours) then it will undergo apoptosis and die (Greijer and Van Der Wall, 2004).

We also distinguish two other types of agent in our model:

- **Microbeads** are treated as a separate cell phenotype. They have no cell cycle, do not consume oxygen and are unaffected by changes in the local oxygen concentration.
- **Labelled cells** are not treated as a separate cell phenotype, and instead progress through the above phenotypes according to the same rules as normal cells. ^{3}H-labelled are indistinguishable from other tumour cells, but will pass their label on to any daughter cells.

A dead cell will continue to occupy space for τ_apop hours before being removed from the simulation. While these dead, or **necrotic** cells do not have a cell cycle, they form a fourth cellular compartment in our model. Figure 4.3a shows a snapshot of a 2D simulation with cells coloured according to compartment, and Figure 4.3b shows the relationships between the phenotypes schematically.

Cell cycle model

Cell cycle progress depends on the local oxygen concentration $\omega(\mathbf{x}, t)$. At birth, each cell is assigned a cell cycle duration, τ_η, drawn from a uniform distribution $U(0.75\,\eta, 1.25\,\eta)$ where the parameter η controls the average cell cycle length.

If the age of a cell at position $\mathbf{x}$ and time t is given by $\gamma_\mathrm{cyc}(t)$, then we have that

$$\frac{\mathrm{d}\gamma_\mathrm{cyc}}{\mathrm{d}t} = \mathcal{H}(\omega(\mathbf{x}, t) - \omega_\mathrm{q}) \tag{4.7}$$

where $\omega(\mathbf{x}, t)$ is the local oxygen concentration at time t and location $\mathbf{x}$, $\gamma_\mathrm{cyc} = 0$ when a cell is born and cell division occurs when $\gamma_\mathrm{cyc} = \tau_\eta$. $\mathcal{H}(\omega(\mathbf{x}, t) - \omega_\mathrm{q})$ is the

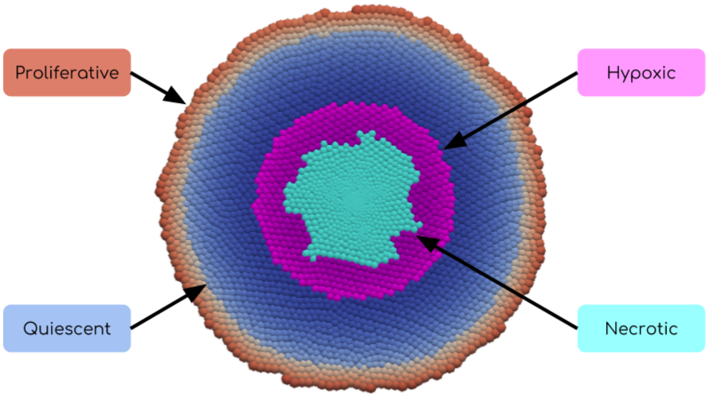

(a) Snapshot from a simulation showing the four different compartments within the model. There are three compartments of living cells with cell cycles, showing proliferating cells, viable quiescent cells and viable hypoxic cells at risk of death. The fourth compartment contains necrotic cells, which are subject to intercellular forces and fill space but no longer progress through the cell cycle.

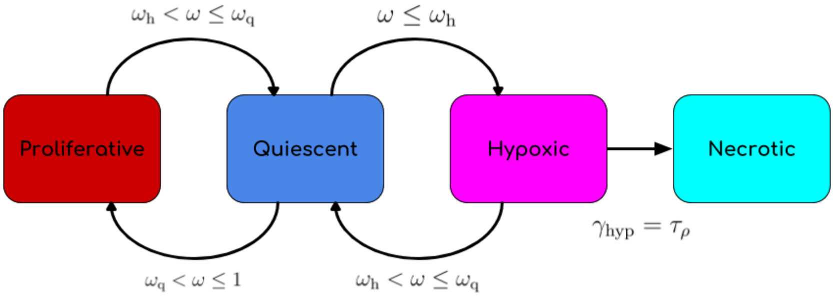

(b) Schematic showing how cells move between phenotypes.

Figure 4.3

Heaviside function, which is equal to zero when the argument is negative and one when it is positive. When $\omega < \omega_{\mathrm{q}}$, the cell cycle pauses and the cell remains dormant until either ω returns to above the threshold and the cell cycle continues or ω drops below ω_{h} and the cell is at risk of hypoxia induced death depending on the time the oxygen concentration remains below ω_{h}. When a cell reaches the end of its cell cycle, a daughter cell is placed half a cell diameter away in a randomly chosen direction and both cells are assigned new cell cycle lengths. For both new cells, $\gamma_{\mathrm{cyc}} = 0$ and γ_{cyc} again evolves according to Equation (4.7). At the start of each simulation, all cells are randomly assigned a value for γ_{cyc} from a uniform distribution $U(0, 0.75\,\eta)$.

Cell death

Cell death occurs when a cell experiences hypoxia, $\omega_{\mathrm{h}} \leq \omega$, for longer than a threshold time τ_ρ, drawn from a uniform distribution $U(0.75\,\rho, 1.25\,\rho)$ where the parameter ρ controls the average time between a cell becoming hypoxic and the onset of cell death.

As such, we assign to each cell an internal hypoxia time, γ_{hyp}, which evolves as follows:

$$\frac{\mathrm{d}\gamma_{\mathrm{hyp}}}{\mathrm{d}t} = \mathcal{H}(\omega_{\mathrm{h}} - \omega(\mathbf{x}, t)) \tag{4.8}$$

where $\gamma_{\mathrm{hyp}} = 0$ if $\omega \leq \omega_{\mathrm{h}}$, to model hypoxic cells once again receiving sufficient nutrient to prevent cell death. A cell becomes apoptotic when $\gamma_{\mathrm{hyp}} = \tau_\rho$. It is then irreversibly marked for cell death. It remains intact for a period of τ_{apop} hours, during which time it is subject to the force laws described in Section 4.3.4.

For each individual cell, the cell cycle model therefore depends on two internal clocks: the “cell cycle time” denoted by γ_{cyc}, and the “hypoxia time” denoted by γ_{hyp}. Pseudocode describing the update implementation for the cell cycle is given in Algorithm 4.1. The timers γ_{cyc} and γ_{hyp} are updated based on the oxygen concentration.

```
Input: All cells, viable or necrotic
for All cells do
    if Cell is alive then
        if $\omega_q < \omega \leq 1$ then
            // Cell is proliferative
            // Move cell through cell cycle by one timestep
            Set $\gamma_{\text{cyc}} = \gamma_{\text{cyc}} + dt$;
            // Ensure hypoxia timer is unset
            Set $\gamma_{\text{hyp}} = 0$;
            // If cell is less than one hour old, increase the cell
               radius
            if $\gamma_{cyc} < 1$ then
                Set $s_i = s_i + R_{\text{Cell}} dt$;
            end
            // If cell is at end of cell cycle, proliferate
            if $\gamma_{cyc} = \tau_\eta$ then
                Choose random location within $R_{\text{int}}$ of cell $i$;
                Place daughter cell $j$ in selected location;
                Set $s_i = \frac{R_{\text{Cell}}}{2}$;
                Set $s_j = \frac{R_{\text{Cell}}}{2}$;
                Set $\gamma_{\text{cyc}} = 0$ for cells $i$ and $j$;
                Choose new cell cycle durations $\tau_\eta$ for cells $i$ and $j$;
            end
        else if $\omega_h < \omega \leq \omega_q$ then
            // Cell is quiescent
            // Ensure hypoxia timer is unset
            Set $\gamma_{\text{hyp}} = 0$;
        else if $\omega \leq \omega_h$ then
            // Cell is hypoxic
            // Increment hypoxia timer by one timestep
            Set $\gamma_{\text{hyp}} = \gamma_{\text{hyp}} + dt$;
            // Check for cell death
            if $\gamma_{hyp} = \tau_\rho$ then
                Mark cell as dead;
            end
        end
    else
        // Cell is necrotic
        // Reduce necrotic cell radius linearly over $\tau_{\text{apop}}$ hours to
           model decay
        Set $s_i = s_i - \frac{R_{\text{Cell}}\,dt}{\tau_{\text{apop}}}$;
        if $s_i$ = 0 then
            Remove cell from simulation;
        end
    end
end
```

Algorithm 4.1: Update procedure for the cell cycle.

4.3.4 Force balance

Cells are subject to forces arising from several different mechanisms:

- Mechanical forces, caused by intercellular adhesion and volume exclusion,
- Random cell motility,
- Surface tension.

At each timestep, we calculate the net force acting at each node and update the node position accordingly. Newton's second law in the over-damped limit, and neglecting inertial effects, gives the following force balance for cell i:

$$\nu \frac{\mathrm{d}\mathbf{x_i}}{\mathrm{d}t} = \mathbf{F}_i^m + \mathbf{F}_i^r + \mathbf{F}_i^s \tag{4.9}$$

where ν is a damping coefficient representing drag between cells and the culture in which the spheroid is growing. We denote the mechanical force by $\mathbf{F}_i^m$, the random motility by $\mathbf{F}_i^r$ and the surface tension by $\mathbf{F}_i^s$. We describe the form of each of these forces in the following sections.

Using the forward Euler method to approximate the time derivative, we obtain an update rule for each cell location, $\mathbf{x_i}$, as follows:

$$\mathbf{x}_i^{\text{new}} = \mathbf{x}_i^{\text{old}} + \frac{\mathbf{F}_i dt}{\nu} \tag{4.10}$$

where $\mathbf{F}_i = \mathbf{F}_i^m + \mathbf{F}_i^r + \mathbf{F}_i^s$, and dt is the timestep.

Mechanical Forces

Due to a cell i interacting with its neighbours, it will experience a mechanical force, $\mathbf{F}_i^m$, which includes both volume exclusion and intercellular adhesion. Intercellular adhesion forces arise when neighbouring cells form bonds between their cell membranes using adhesion molecules such as ICAM-1 (intercellular adhesion molecule-1) to pull each other closer (Usami et al., 2013; Turner and Sherratt, 2002). We assume that two cells interact if their cell centres are within radius R_{int} of each other. We then connect the centres of these interacting cells i and j with a spring if $|\mathbf{x}_i - \mathbf{x}_j| \leq R_{\text{int}}$ (Meineke et al., 2001). The resting length of this spring, $s_{i,j}$, is the sum of the equilibrium springs for each cell s_i and s_j, such that $s_{i,j} = s_i + s_j$. With the exception of newly divided cells and necrotic cells, we assume that $s_i = R_{\text{Cell}}$, a constant parameter. For newly divided cells, $s_i = R_{\text{Cell}} \,/\, 2$ and increases linearly over the course of one hour until $s_i = R_{\text{Cell}}$. For necrotic cells, s_i decreases linearly from R_{Cell} to 0 over the course of τ_{apop} hours, leading to a reduction in cell size until the cell is removed from the simulation. To ensure that this does not generate artificially high forces by "stretching" apoptotic cells, the associated spring constant is reduced linearly at the same rate representing a weakening intercellular force between an apoptotic cell and any other cell. These reductions in cell size are reflected in the pseudocode in Algorithm 4.1.

The net force acting on cell i due to mechanical forces is the sum of the contributions over all cells j within R_{int}:

$$\mathbf{F}_i^m = \sum_{\{j \mid |\mathbf{x}_i - \mathbf{x}_j| \leq R_{\text{int}}\}} \mathbf{F}_{i,j}^m. \tag{4.11}$$

where $\mathbf{F}_{i,j}^m$ is the mechanical force between cells i and j. This force always points in the direction of the vector between the cells. The magnitude of $\mathbf{F}_{i,j}^m$ is given by:

$$|\mathbf{F}_{i,j}^m| = \begin{cases} \mu s_{i,j} \log\left(1 + \frac{x}{s_{i,j}}\right) \text{ if } x < 0 \\ \mu x s_{i,j} \exp\left(-\delta \frac{x}{s_{i,j}}\right) \text{ if } x \geq 0 \end{cases} \tag{4.12}$$

where $x = |\mathbf{x}_i - \mathbf{x}_j| - s_{i,j}$ is the overlap between cells, μ is a coefficient representing the spring stiffness and δ is a parameter determining the strength of intercellular adhesion on neighbouring cells. The magnitude of this force as the distance between cell centres varies is shown in Figure 4.4. The force law in Equation (4.12) assumes that once cell centres are within the radius of interaction, R_{int}, forces of adhesion pull them together. This adhesive force grows stronger as the cell centres draw closer,

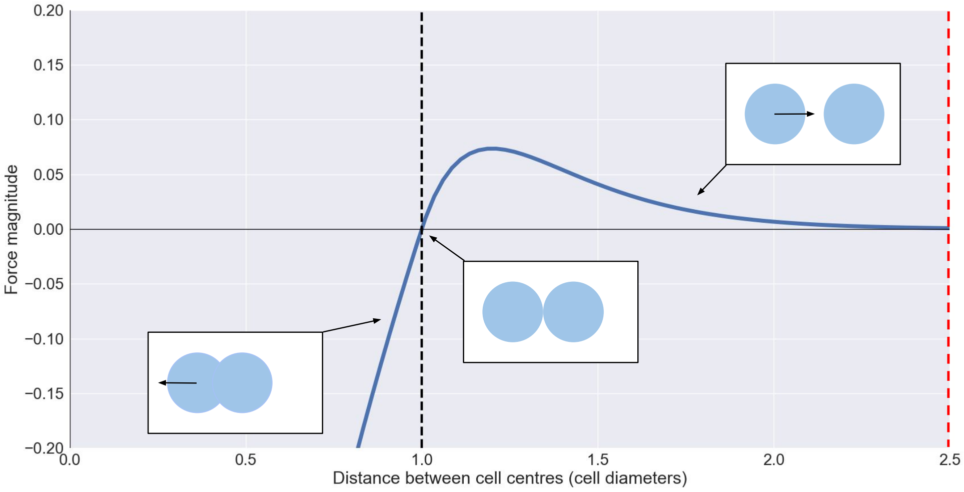

Figure 4.4: Magnitude of the force between two neighbouring cells i and j as a function of the distance between those cells. If the distance is greater than $s_{i,j}$, the cells experience an attractive force. If the distance is less than $s_{i,j}$ cells experience a repulsive force which becomes exponentially stronger as cells are compressed. The black dashed line marks the choice in this image of $s_{i,j} = 1$, and in this example the parameters $\mu = 1$ and $\delta = 5$. The radius of interaction $R_{\text{int}} = 2.5$ is marked with a red dashed line.

since more intercellular bonds form when the surface area on contact between the cells increases (Usami et al., 2013; Drasdo and Höhme, 2005). The form of Equation (4.12) for $x < 0$ is chosen to model volume exclusion, as cells which are closer than the resting spring length $s_{i,j}$ are compressed and repel each other.

Random Motility

We include a random force in our model which simulates cell motility. The random force, $\mathbf{F}_i^r = (F_x^r, F_y^r, F_z^r)$, has a contribution at each timestep in $d \in \{x, y, z\}$ of

$$F_d^r = \frac{\nu}{dt}\sqrt{2Ddt}\,\xi \tag{4.13}$$

where ν is the damping coefficient (see Equation (4.9)), ξ is a random variable drawn from a standard normal distribution and D is a diffusion coefficient determining the intensity of movement. We include a prefactor of $\nu\,/\,dt$ in Equation (4.13) to ensure that F_d has the dimensions of a force and that D has the dimensions of a diffusion coefficient.

Choosing this form for $\mathbf{F}_i^r$ ensures that in each direction the movement of a cell due to random motility alone is normally distributed with mean 0 and variance $2Ddt$. This is Brownian motion, which is a commonly used model of random cell motility (Macfarlane et al., 2018). Appendix B.1 contains results demonstrating that our random motion model is correctly implemented.

Surface tension force

A constant force, $\mathbf{F}_i^s$, is applied to the spheroid boundary nodes in the direction of the spheroid centroid with the form

$$\mathbf{F}_i^s = \begin{cases} -\beta\hat{\mathbf{x}}_i & \text{if cell } i \text{ is on the boundary,} \\ \mathbf{0} & \text{if cell } i \text{ is not on the boundary,} \end{cases} \tag{4.14}$$

where $\hat{\mathbf{x}}_i$ is a unit vector in the direction of the spheroid centroid for a boundary node i and β is a parameter determining the strength of the surface tension force.

Although it acts only on the spheroid boundary, this force resists the formation of voids caused by necrotic cell decomposition within the tumour spheroid and maintains spheroid shape. Strong surface tension effects cause spheroids to grow as compact masses (Byrne and Drasdo, 2009; Greenspan, 1972), since voids within the spheroid are filled with cells that collapse under the force due to surface tension. Landman and Please (2001) note that in the absence of surface tension on the outer surface of a spheroid, the tumour will grow without bound. We observe that in the necrotic core of some spheroids, loss of volume in dead cells due to fluid leakage can cause "cracking", or fluid filled voids, when coupled with intercellular adhesion (Ma et al., 2012; Ghosh et al., 2007). While this behaviour has been reported in existing agent-based models of tumour spheroids (Ghaffarizadeh et al., 2018; Drasdo and Höhme, 2005), we note that in the absence of a compressive force resisting tumour growth, our attempts at reproducing existing agent-based models featuring this "void" effect frequently produce travelling wave solutions which do not reach a steady state, as predicted by Landman and Please (2001). We have observed this effect across a range of spheroid models implemented in different agent-based and continuum frameworks, and note that while tumour spheroids may contain a void *in vitro* these "spherical shell" simulations in which a spheroid with constantly increasing radius exists as a shell of fixed width are not observed experimentally.

Addition of microbeads

Microbeads are implemented as another type of cell without a cell cycle and are identified by the positions of their centres. Like other cells, they interact with neighbouring cells via the force balance described in Equation (4.9). Since microbeads do not adhere to each other we modify the mechanical force described in Equation (4.12) as follows:

$$|\mathbf{F}^m_{i,j}| = \begin{cases} \mu_{\text{bead}} s_{i,j} \log\left(1 + \frac{x}{s_{i,j}}\right) & \text{if } x < 0 \\ 0 & \text{if } x \geq 0 \end{cases} \tag{4.15}$$

where $x = s_{i,j} - |\mathbf{x}_i - \mathbf{x}_j|$ is the overlap between the beads. Hence if the centres of the beads are beyond $s_{i,j} = 2R_{\text{Bead}}$ for bead radius R_{Bead}, the beads do not exert a force on each other. Beads therefore resist being compressed, but do not adhere to other beads. Microbeads are subject to random motion.

4.3.5 Model initialisation and parametrisation

We initialise our model by placing 300 tumour cells in contact with each other in a circle. Each cell is assigned a division time τ_η as described in Section 4.3.3, and an age which is chosen at random from a uniform distribution $U(0, \tau_\eta)$. We allow the simulation to run for 300 hours in the absence of microbeads or labelled cells. For most parameter sets this is sufficient for simulations to reach a steady state. At

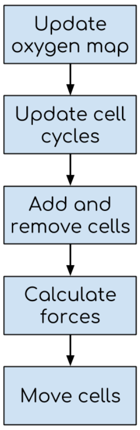

Figure 4.5: Flowchart describing the simulation level process undertaken at each timestep to update the model.

$t = 300$ hours, 100 microbeads are placed randomly in a circular shell at the edge of the spheroid or 50 boundary cells are labelled randomly.

Where possible, model parameters are estimated from the literature. Where experimental data is lacking, we choose parameter values that produce simulations which replicate the *in vitro* growth rates described in Section 4.2. The parameter values we have used can be seen in Table C.1 in Appendix C, with ranges of values where relevant.

While our model can simulate spheroid growth in two or three spatial dimensions, we focus on two-dimensional simulations unless stated otherwise in order to reduce computational running time.

The flowchart in Figure 4.5 illustrates the order in which the different steps of our computational model are executed at a simulation level. Cell cycles are updated according to the process described in Algorithm 4.1.

4.4 Generation of spheroids with different compositions

We first demonstrate that our model can reproduce tumour spheroid dynamics. The volume of a tumour spheroid over time typically undergoes logistic growth (Folkman and Hochberg, 1973), with exponential growth at early times when all cells are nutrient rich and proliferation is dominant. The growth rate decreases as lack of nutrient in the spheroid core triggers the onset of hypoxia and eventually necrosis, and ultimately plateaus when the net rates of cell proliferation and cell death balance. In this Section, we demonstrate that our model generates logistic growth curves, and present example parameter sets demonstrating that similar growth dynamics can be obtained from parameter sets in which the thresholds for quiescence and hypoxia, ω_{q} and ω_{h}, vary substantially. In subsequent Subsections, we show that although the dynamics of these growth curves and their associated steady state volumes are similar for different parameter sets, the infiltration patterns of the microbeads vary dramatically.

Figure 4.6 shows typical simulation results at three timepoints for three parameter sets that generate spheroids of approximately equal size. By varying parameters associated with cell proliferation and death (ω_{q}, ω_{h}, η and τ_ρ), we demonstrate that spheroids with different internal dynamics can exhibit similar growth dynamics. Growth curves averaged over 30 realisations for each parameter set can be seen in Figures 4.7 and 4.8. Figure 4.7 shows the spatial composition of each spheroid in two different ways, first by measuring the approximate radius of each oxygen environment inside the tumour - normoxic, quiescent, hypoxic or necrotic - over time, and then by showing how the number of tumour cells in each oxygen compartment changes over time (stacked plot). Figure 4.8 shows how the total area of each 2D spheroid changes over time, with the spheroid area being defined as the region bounded by the α-shape of the spheroid.

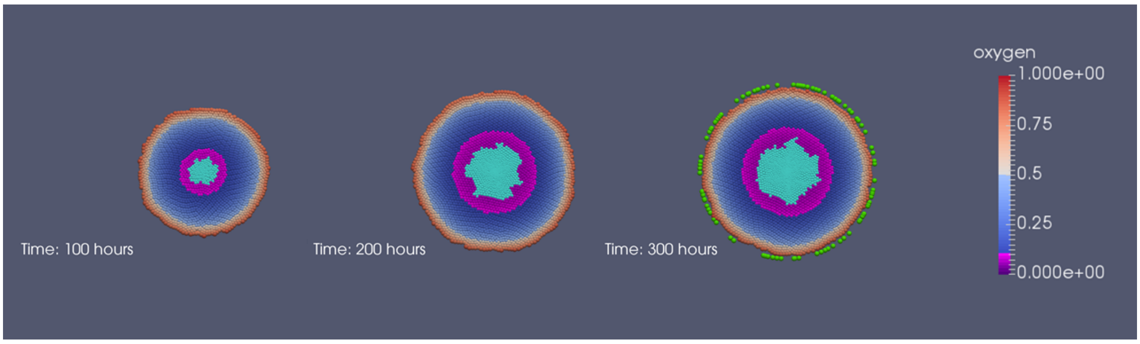

(a) Spheroid A: $\omega_\mathrm{q} = 0.5$, $\omega_\mathrm{h} = 0.1$, $\eta = 16$, $\tau_\rho = 16$

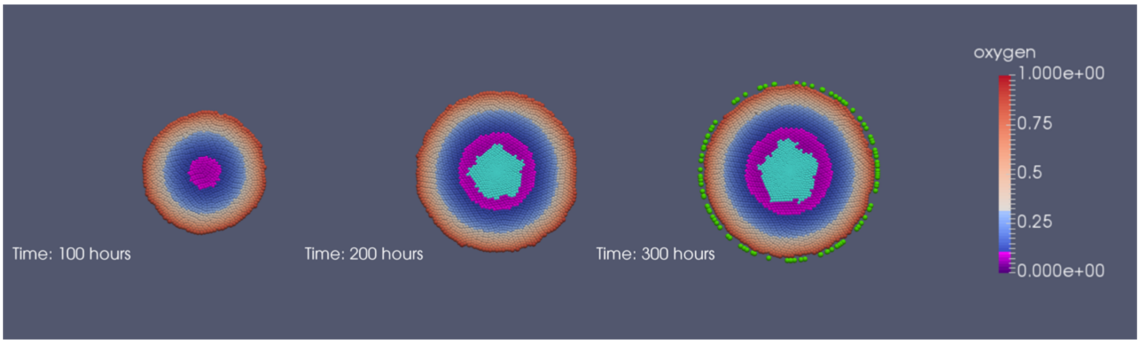

(b) Spheroid B: $\omega_\mathrm{q} = 0.3$, $\omega_\mathrm{h} = 0.1$, $\eta = 32$, $\tau_\rho = 16$

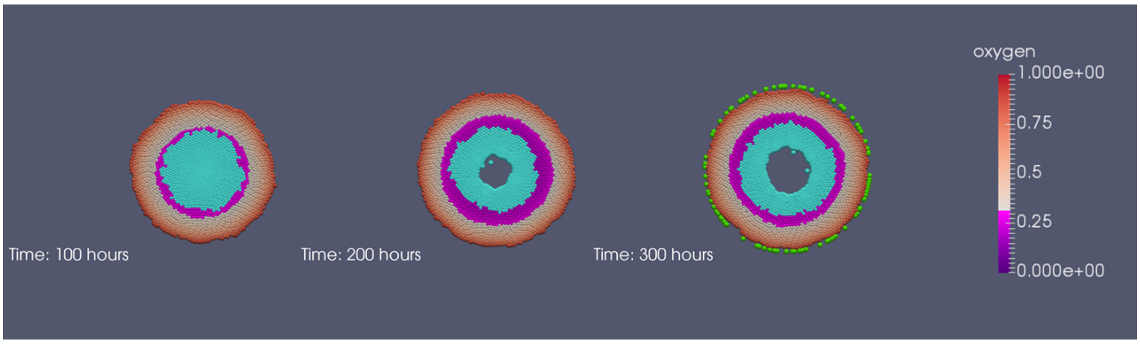

(c) Spheroid C: $\omega_\mathrm{q} = 0.3$, $\omega_\mathrm{h} = 0.3$, $\eta = 16$, $\tau_\rho = 8$

Figure 4.6: Series of simulation results showing how the tumour spheroids evolve over time for three different parameter sets, A (top), B (middle) and C (bottom). Cells are coloured according to their oxygen concentration, red cells are proliferative ($\omega > \omega_\mathrm{q}$), blue cells are quiescent ($\omega_\mathrm{q} \geq \omega > \omega_\mathrm{h}$), magenta cells are hypoxic ($\omega_\mathrm{h} \geq \omega$) and cyan cells are necrotic. The ring of green cells shown around the largest spheroids is the initial configuration of microspheres, which are randomly placed in a circle at the spheroid edge after 300 hours.
The only differences between the parameter sets are the thresholds for the onset of hypoxia and quiescence, ω_h and ω_q, the average tumour cell cycle length, η, and the length of time that it takes for a hypoxic cell to die, τ_ρ

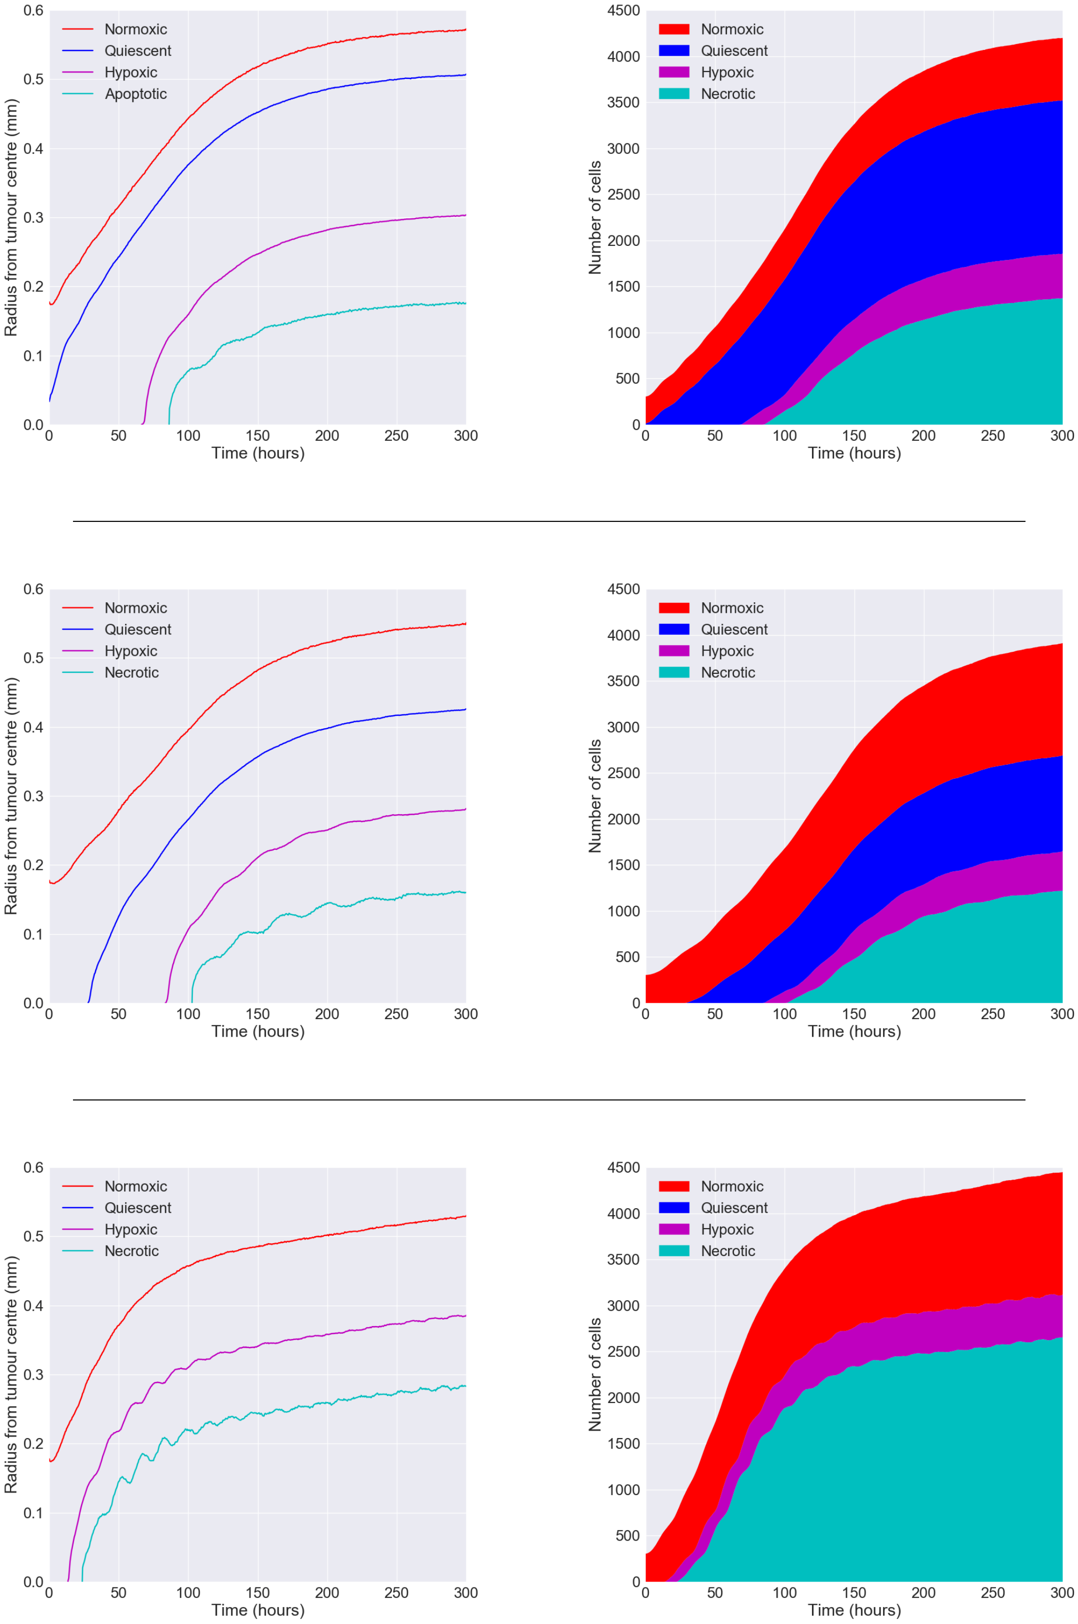

Figure 4.7: Averaged growth dynamics from 30 realisations of simulated spheroids for the three parameter sets described in Figure 4.6.

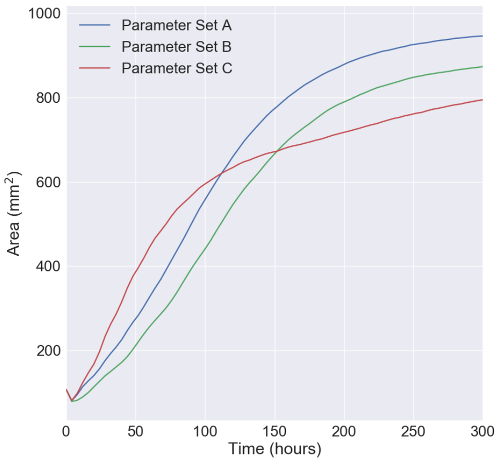

Figure 4.8: Spheroid area as a function of time, averaged over 40 realisations of simulated spheroids for the three specimen parameter sets described in Figure 4.6. After an initial transient, the growth rate of all three spheroids is similar and they reach comparably sized steady states.

These three parameter sets have been chosen as examples since they characterise distinct points on a spectrum of tumour behaviours. Spheroid A is a relatively “quiet” spheroid, in which a low rate of growth at the spheroid periphery is balanced by a comparatively small core of necrotic material caused by a low choice of ω_{h}. Spheroid B shares the same values of ω_{h} and τ_ρ as Spheroid A, and consequently cells die at a similar rate. By reducing the value for ω_{q} a much thicker proliferative rim has formed. In order to achieve a similar steady state size, the proliferation rate η has been slowed down. Finally, Spheroid C is more “energetic” than the previous spheroids, with a large proliferative rim, no quiescent compartment and low η being balanced by a substantial core of necrotic and hypoxic material which is removed quickly due to a low τ_ρ. In Figure 4.6c, cells are removed sufficiently quickly that the necrotic core pulls away to form a void. This type of behaviour is observed *in vitro* (Ma et al., 2012; Ghosh et al., 2007). This phenomenon has been observed in other computational models when volume loss in necrotic cells and intercellular adhesion cause “cracking”, or the formation of fluid-filled voids in the necrotic core (Ghaffarizadeh et al., 2018; Drasdo and Höhme, 2005).

Over a period of approximately 10 days, the radius of each spheroid grows by approximately 16 cell diameters (32 μm per day), which is comparable to the rate of 20 μm per day reported by Dorie et al. (1982) for spheroid growth rate. Figure 4.7 also shows that after an initial transient the size of the proliferative rim of each spheroid remains approximately constant, the growth in spheroid volume being caused by increased numbers of quiescent, hypoxic and necrotic cells. Each simulation follows the logistic growth curves observed *in vitro* and *in silico* for oxygen limited tumour spheroid growth (Benzekry et al., 2014; Spratt et al., 1993; Folkman and Hochberg, 1973; Greenspan, 1972).

While these sample parameter sets generate spheroids with similar growth rates and equilibrium sizes, their internal structures are very different. Parameter set A results in a spheroid with a significant quiescent region, comprising approximately a third of the total tumour radius. The width of the proliferative rim in Spheroid A is also much smaller than that of the hypoxic region, in contrast to Spheroid C where the width of the proliferative rim is comparable with that of the neighbouring hypoxic region. Spheroids A and B have similar structures, since the much slower proliferation rate in Spheroid B is balanced by a correspondingly lower threshold for quiescence, although the proliferative rim in Spheroid B is approximately twice as thick as that in Spheroid A. Spheroid C has an even more markedly different structure, with the most evident difference exhibited in the size of the central necrotic core - Spheroid A has a necrotic core approximately 8 cell diameters in radius, compared with almost double this in Spheroid C.

Figure 4.9 shows the effect of varying ω_{q} and η on the radius of a tumour spheroid and the width of the proliferative rim after 400 hours of growth. Spheroids with the same approximate radius can be generated with different values of ω_{q} and η, but Figure 4.9b indicates that the proliferative rim in these spheroids will vary substantially in thickness.

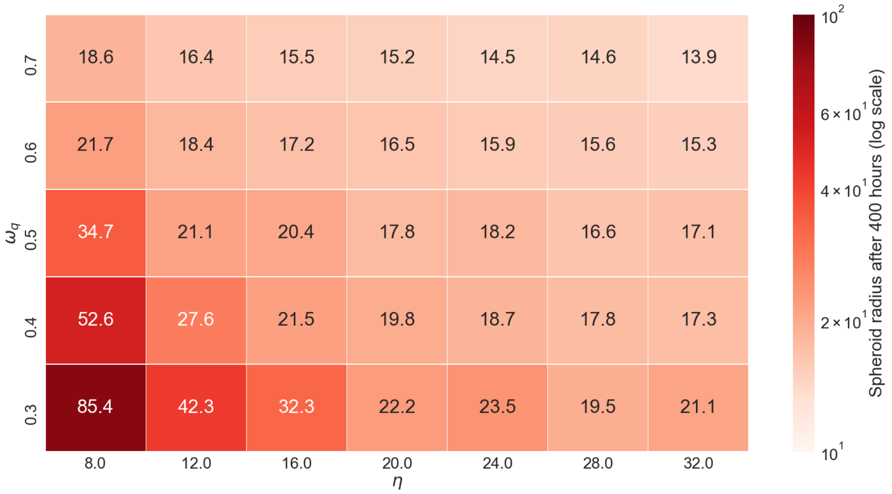

(a) Effect of ω_q and η on steady state spheroid radius (in cell widths).

(b) Effect of ω_q and η on the width of the proliferative rim (in cell widths).

Figure 4.9: Effect of varying ω_q and η on the radius of tumour spheroids after 400 hours of growth (Figure 4.9a) and the corresponding width of the proliferative rim (Figure 4.9b).

4.5 Effect of spheroid composition on microsphere infiltration

4.5.1 Reproduction of distributions from (Dorie et al., 1982)

We randomly distribute 100 microspheres around the boundary of each tumour spheroid after 300 hours of simulation time, by which time most spheroids have reached a steady state. In Figure 4.10 we present histograms of the distributions of the internalised microspheres. Each bar shows the proportion of infiltrating microbeads within an annulus of fixed width at a given radius from the spheroid centre at a fixed timepoint, with the approximate spheroid radius marked by a vertical dashed line. As the spheroids are not radially symmetric due to stochastic fluctuations at the boundary, the approximate radius is that of the largest circle centred at the spheroid centroid which can fully enclose every tumour cell. The results shown are aggregated over 30 realisations for each parameter set, and the spheroid radius shown is the average of the approximate radii over all realisations.

The histograms show the proportion of microbeads found a given distance from the spheroid centroid, averaged over 40 iterations for each parameter set. Figures 4.10a - 4.10c each show a distribution moving towards the tumour centre, but the distributions indicate different internal dynamics between the simulations. The peak in Figure 4.10c is lower than those in Figures 4.10a and 4.10b, indicating a more dispersed distribution in this parameter regime. The wave also reaches the spheroid centre faster, with cells accumulating at the edge of the central void by the end of the simulation. Figures 4.10a and 4.10b also differ, with similar initial peaks dispersing at different rates. In each case, the distribution of beads becomes more dispersed as the simulation progresses as observed in (Dorie et al., 1982). However, these histograms indicate that the size of the spheroid alone is insufficient to predict the dynamics of the microspheres.

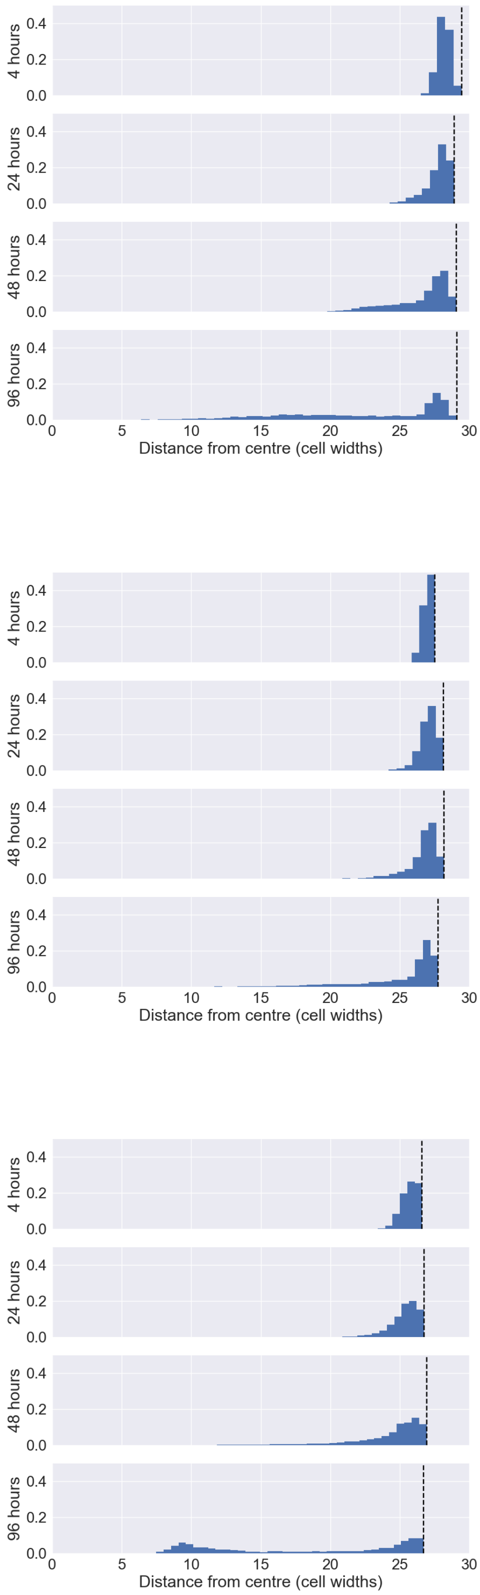

Figure 4.10: Frequency histograms showing the distribution of microbeads within Spheroids A, B and C over time. In contrast to Figure 4.2, the x-axis indicates distance from the spheroid centroid with the vertical dashed line denoting the approximate tumour radius.

4.5.2 Tracking individual bead trajectories

An advantage of our agent-based model compared to existing PDE models of the infiltration experiments in (Dorie et al., 1982) is that it enables us to track trajectories of individual microbeads over time. Previous continuum models of the experiments in (Dorie et al., 1982) are unable to distinguish individual microbeads, instead modelling the density of microbeads.

Based on the simulation results presented in Figure 4.10, we now propose two alternative hypotheses to explain the observed patterns of microbead infiltration and demonstrate that by tracking the movement of individual beads we are able to distinguish between the cases.

As shown in Figure 4.2, the microbeads form a moving distribution which disperses over time. This observation is consistent with either of the following hypotheses:

Hypothesis 1. *Microbeads move steadily towards the centre of the tumour, but with different average velocities.*

Hypothesis 2. *Microbeads move with similar average velocities towards the centre of the tumour, but at first remain trapped at the spheroid edge for different lengths of time due to random cell proliferation.*

Under the second hypothesis, microbead distributions disperse because some beads are trapped in the proliferative rim by the Brownian motion induced by high proliferation rates there; these may push beads either radially inwards or outwards. Once a bead is pushed sufficiently far from the outer spheroid boundary, advection caused by the balance of birth and death causes beads to infiltrate at approximately the same speed. In contrast, under Hypothesis 1 the same dispersal of distributions is explained without distinguishing between these regimes. When referring to Hypothesis 2, we will occasionally describe microbeads as starting their infiltration at different times. Although technically all beads begin infiltrating when they cross the outer spheroid boundary, this terminology will be used to distinguish between beads which remain in the regime dominated by Brownian motion and those which have “started their infiltration” by leaving this regime and entering the advection dominated region.

Previous continuum models of bead infiltration are unable to distinguish between these hypotheses since they do not resolve the dynamics of individual cells, although some previous models do suggest a transient period before beads start to infiltrate (Thompson and Byrne, 1999). However, it is important to distinguish between Hypotheses 1 and 2 when considering the effect of passive infiltration on a single microbead or immune cell.

Figure 4.11 shows how the radial distance from the spheroid centroid changes over time for 10 randomly chosen beads from two parameter sets ($\omega_{\mathrm{q}} = 0.5$ or 0.3, $\omega_{\mathrm{h}} = 0.3$, $\eta = 16$, $\tau_{\rho} = 8$). Trajectories are coloured according to the oxygen concentration of the bead at each timepoint. We see that, although the beads begin to infiltrate at different times, their trajectories can be divided into three different regimes. After a transient period close to the edge of the spheroid, microbeads pass through the proliferative rim with approximately constant radial velocities until they stop at the spheroid centre. This indicates that beads can be “trapped” inside the proliferative rim before reaching the quiescent compartment and suggests that Hypothesis 2 more accurately describes the way in which the beads infiltrate the spheroids. Under this hypothesis, the average radial velocity profile at a given distance from the spheroid centroid is unchanging and the dispersion of the microbead distribution is due to differences in the times at which beads leave the Brownian motion dominated regime near the spheroid edge. This behaviour is consistent with all the parameter sets we have analysed.

In order to understand how the rate of bead infiltration varies between parameter sets, for each simulation we fit a straight line to the linear portion of the trajectory of each bead. As beads in the same simulation may stay near the spheroid rim

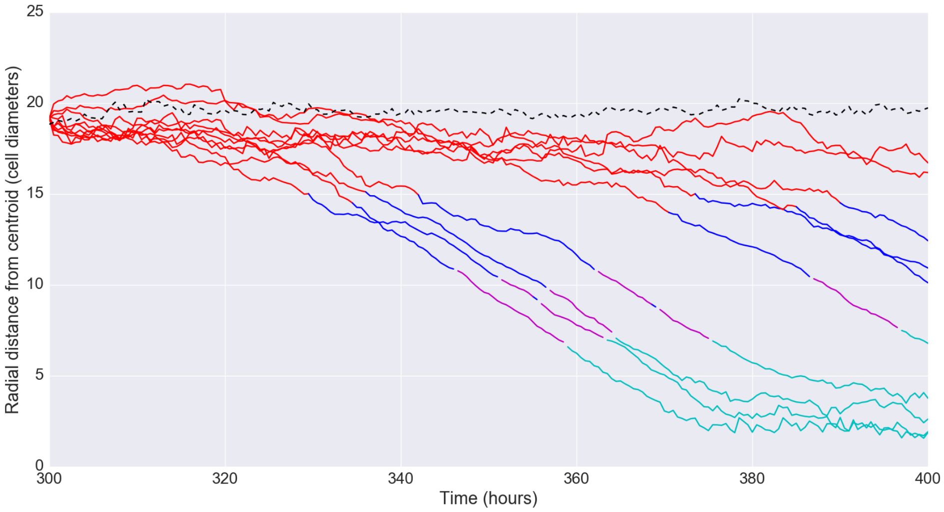

(a) Trajectories from 10 random beads from a simulation with $\omega_\mathrm{q} = 0.5$, $\omega_\mathrm{h} = 0.3$, $\eta = 16$, $\tau_\rho = 8$

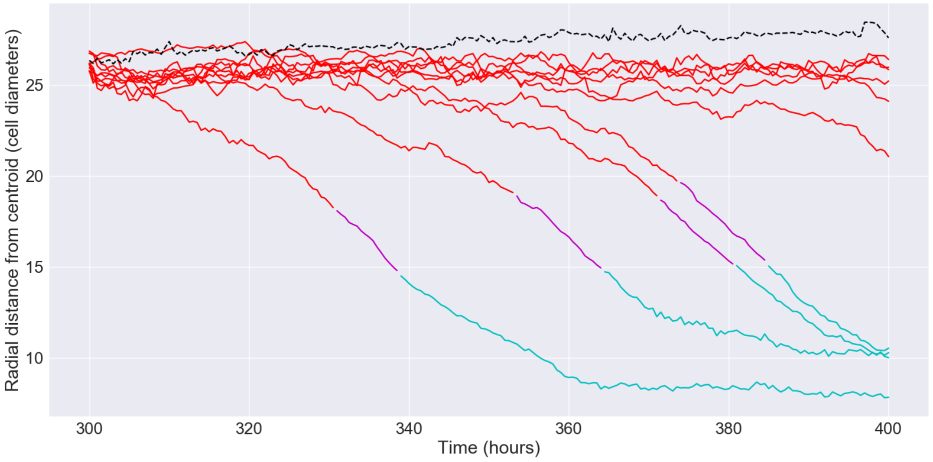

(b) Trajectories from 10 random beads from a simulation with $\omega_\mathrm{q} = 0.3$, $\omega_\mathrm{h} = 0.3$, $\eta = 16$, $\tau_\rho = 8$

Figure 4.11: Trajectories of 10 randomly chosen beads from spheroids generated using two different parameter sets. Each bead is coloured according to its oxygen environment using the same colour scheme as in Figure 4.6. The black dashed line shows the approximate radius of the spheroid for this simulation. The parameters used in Figure 4.11b are the same as those in Figure 4.11a, but with a lower value of ω_q, causing a thicker proliferative rim.

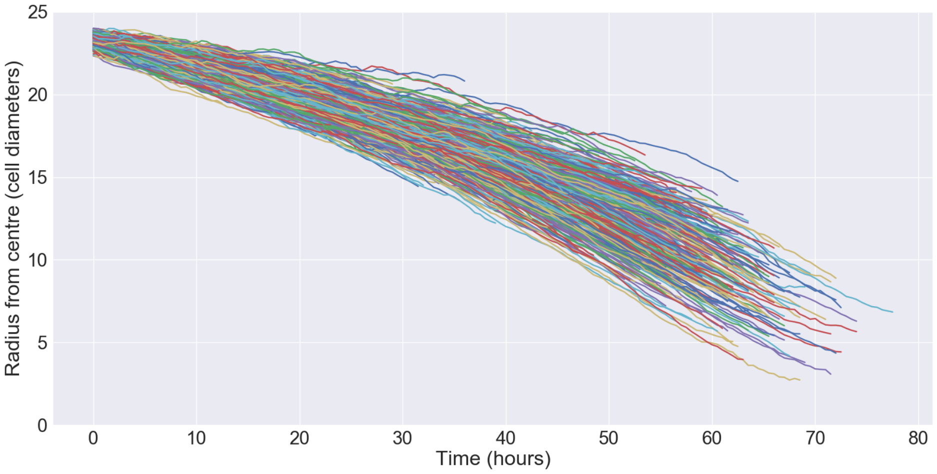

Figure 4.12: Radial trajectory of 1950 microbeads (30 realisations using one parameter set, $\omega_\mathrm{q} = 0.5$, $\omega_\mathrm{h} = 0.1$, $\eta = 16$, $\tau_\rho = 16$), aligned according to the time at which they reach a distance of 80% of the tumour radius.

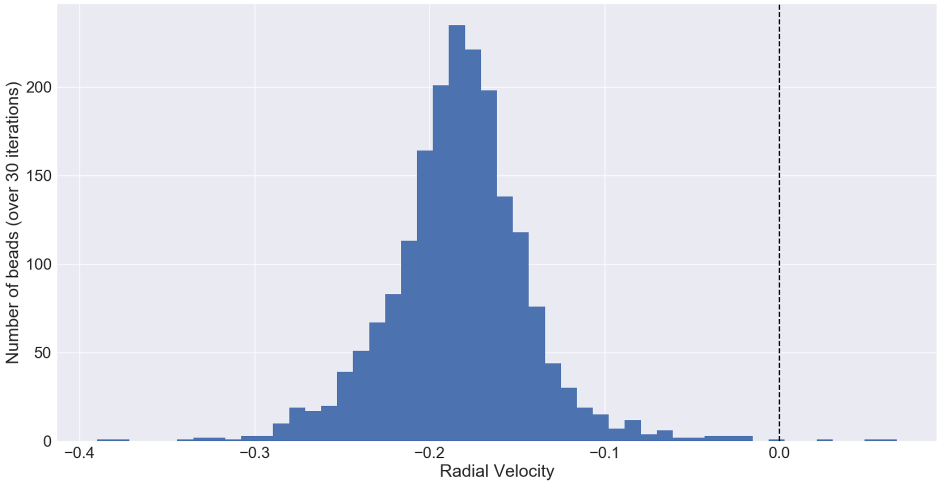

Figure 4.13: Average radial velocity (cell diameters per hour) of each infiltrating microbead within 30 repeats of a representative parameter set ($\omega_\mathrm{q} = 0.5$, $\omega_\mathrm{h} = 0.1$, $\eta = 16$, $\tau_\rho = 16$).

for different lengths of time before beginning their infiltration due to advection, we fit the line from the time at which an individual bead crosses a threshold of 80% of the spheroid radius. In Figure 4.12 we align the infiltrating portion of each microbead trajectory based on the time that the bead leaves the spheroid boundary (for one representative parameter set). For each trajectory shown in Figure 4.12 we calculate the bead radial velocity by fitting a straight line to the trajectory and calculating its gradient. We generate a histogram of these velocities to understand their distribution. Figure 4.13 shows the distribution of the velocities for all microbeads from 30 realisations of a particular parameter set (here, parameter set A). The histogram shows an approximately normal distribution of bead velocities, with a mean radial velocity of approximately -0.18 cell diameters per hour and a standard deviation of 0.04. In the following Sections, we use this mean radial velocity as a summary statistic to compare the speed of infiltration of microbeads in spheroids generated using different parameter sets.

4.5.3 Impact of η and ω_{q} on infiltration velocity

In Figure 4.9b, we showed that decreasing the oxygen threshold for quiescence, ω_{q}, generates spheroids with thicker proliferative rims. The width of the proliferative rim also plays an important role in determining the speed of passive infiltration of microbeads.

Figure 4.14 demonstrates the impact of decreasing the threshold for quiescence, ω_{q}, on the speed of microbead infiltration when all other parameters are held constant. The three histograms shown in Figure 4.14 are generated using the same procedure as for Figure 4.13. Each bead velocity is calculated by tracking the bead from the time at which it reaches 80% of the spheroid radius, fitting a straight line to this trajectory (as in Figure 4.12) and calculating the gradient. The histograms in Figure 4.14 each show the frequency of these bead velocities obtained from 30 realisations of the same parameter set. All parameters are held constant between

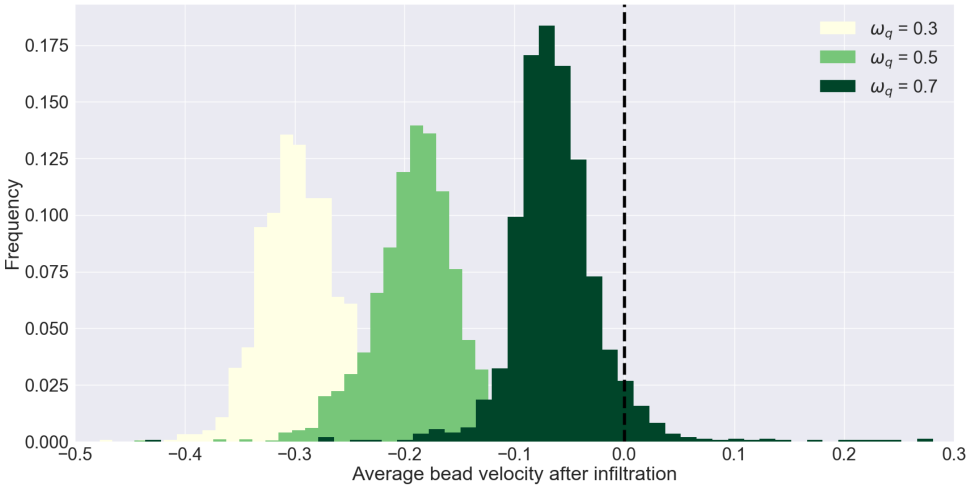

Figure 4.14: Histograms showing that decreasing ω_{q}, and so decreasing the width of the proliferative rim, causes microbeads to infiltrate at a slower velocity. Darker histograms show higher values of ω_{q}, and all other parameters are held constant ($\omega_{\mathrm{h}} = 0.1$, $\eta = 16$, $\tau_\rho = 8$).

the three parameter regimes shown except for ω_{q}.

Figure 4.14 demonstrates that by lowering ω_{q}, and hence increasing the thickness of the proliferative rim, the speed of passive infiltration increases and infiltrating microbeads are carried at a higher velocity than in spheroids with thin proliferative rims. The relationship between ω_{q} and infiltration velocity is further explored in Figure 4.15, which shows the effect of ω_{q} and η on the mean velocity of infiltrating microbeads. The contours show combinations of ω_{q} and η which result in the same infiltration velocity.

As well as demonstrating that a thicker proliferative rim causes faster internalisation of microbeads, Figure 4.15 also shows that η, the average cell cycle length of a tumour cell, has a strong impact on infiltration velocity. Microbeads infiltrate faster into spheroids containing rapidly dividing cells (low η).

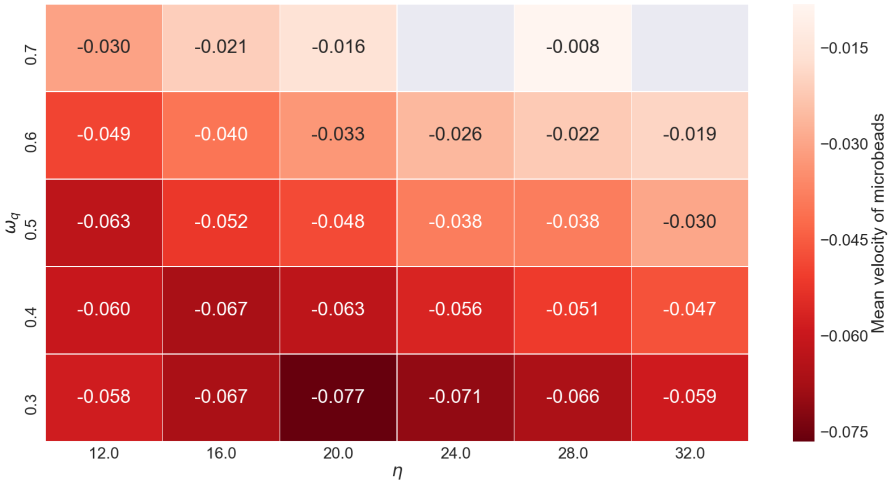

Figure 4.15: Mean velocity of infiltrating microbeads (cell widths per hour) as the threshold for quiescence, ω_q, and the average cell cycle length, η, are varied. Darker shades of red indicate faster movement towards the spheroid centre. Empty boxes show parameter sets for which no beads infiltrated beyond 80% of the spheroid radius.

4.5.4 Impact of η and ω_q on bead waiting time

We have shown that the balance of proliferation and death in a tumour spheroid determines the bead velocity once a bead has begun infiltrating the spheroid. In Figures 4.11 and 4.12 we showed that beads can become trapped in the proliferative rim before beginning their infiltration. Here, we demonstrate that the average duration for which a bead is "stuck" in the proliferative rim, another important facet of "passive infiltration", is also dependent on η and ω_q.

In Figure 4.16 we examine the effect of varying η and ω_q on the average time that it takes a bead to cross the proliferative rim, which we term the bead "waiting time" for a spheroid. In terms of the bead trajectories shown in Figure 4.11, the "bead velocity" refers to the average gradient of the inward trajectory, while the "waiting time" is the average time that it takes a bead to begin this inward movement.

Where ω_q is high, the proliferative rim forms a narrower band (Figure 4.9b) and beads are able to cross it quickly, resulting in a low average waiting time for beads to leave the proliferative compartment. This is due to the proliferation of cells in

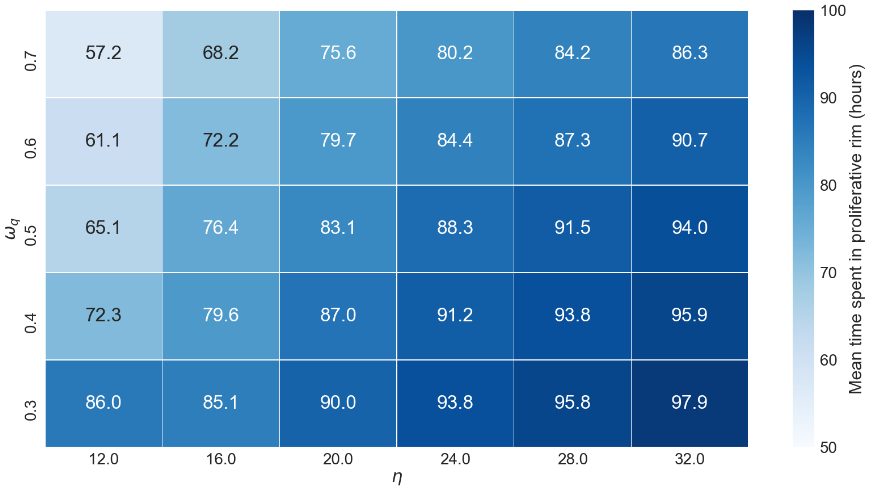

Figure 4.16: Average time taken for a microbead to cross the proliferative rim. (Darker shades of blue indicate a longer time in the proliferative rim). Beads take longer to cross the proliferative rim when it is wide, but they are also inhibited by increased cell cycle length. Where cells proliferate slowly, beads are trapped in the proliferative rim for longer.

the spheroid rim being equally likely to push beads toward the spheroid boundary as to push them towards the centre. A thicker proliferative rim causes microbeads to be exposed to this environment for longer, increasing the likelihood that in any given timestep they are pushed away from the spheroid centre. When a bead is near the spheroid edge, most proliferation events will push it away from the core. In contrast, most proliferation events will push a bead near the boundary between the quiescent and proliferative regions towards the centre, making it more likely to reach the linear inwards section of its trajectory. Hence lower values of ω_q which increase the proliferative rim width increase the waiting time by increasing the likelihood of a bead being pushed away from the centre for longer.

This interpretation is supported by the effect of increasing η, which also causes much longer average waiting times. Passive infiltration moves beads towards the centre because over long timescales proliferation events which push beads inwards are more common than those which push them outwards. Increasing η reduces the frequency of proliferation, and so a bead must wait longer to accumulate the same number of inward “nudges” to cross the proliferative rim. Rather than providing an obstacle to beads by producing an outward force at the spheroid boundary, in the long term increased proliferation instead drives beads through the proliferative rim and into the linear regime observed in Figure 4.11.

4.6 Infiltration of ^{3}H-labelled cells

In order to simulate ^{3}H-labelled cells in our agent-based model, we track the lineage of each cell over time. When a target cell divides, both of the daughter cells are assigned the same label as the original.

Figure 4.17 shows an example of this lineage tracing. Viable, living cells (proliferative, quiescent or hypoxic) are coloured dark blue while necrotic cells are coloured cyan. In this example, two randomly selected proliferative cells were selected at an early timepoint to be labelled. One cell was labelled yellow, and another in pink. Each time one of these cells divided, the progeny were assigned the same colour label.

Many of these labelled cells are caught in the advective flow moving towards the spheroid centre, and ultimately succumb to hypoxia. Necrotic labelled cells lose their labelling in our simulations, and so cannot be traced into the necrotic core. To simulate the addition of ^{3}H-labelled cells to the edge of our tumour spheroids, we randomly tag 50% of the boundary cells with the same label and track the distribution of labelled cells (and their progeny) in the same manner as we track microbeads.

In our model, we note that the pressure gradient driving advective flow moves from the proliferative rim towards the necrotic core in the same way described in (Pettet et al., 2001). However, the stochastic nature of proliferation in our model allows for local movement away from the spheroid centroid which is not observed in their model. This phenomenon is due to cells near the edge of the spheroid giving

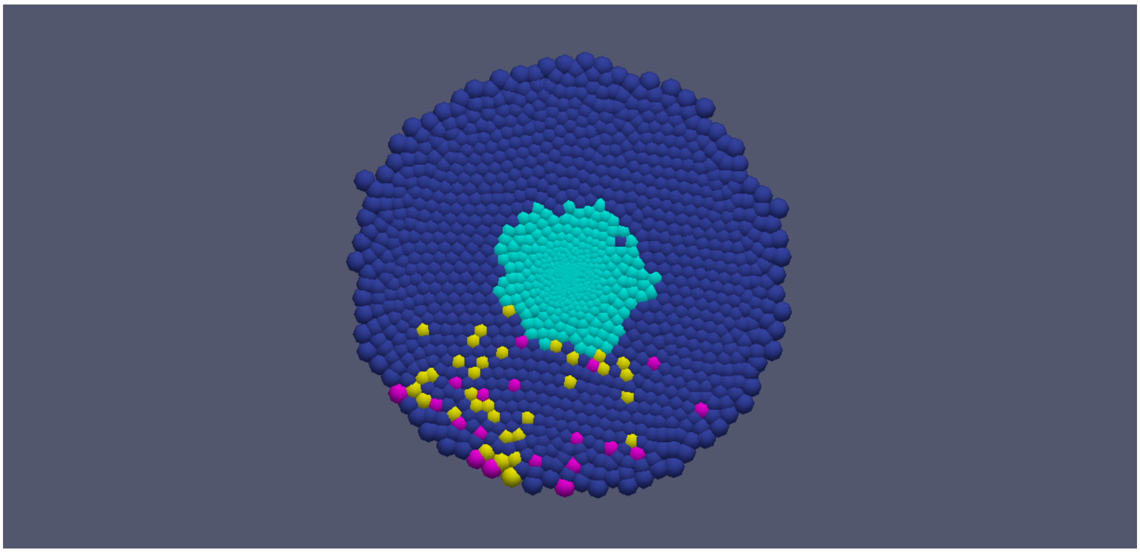

Figure 4.17: Result of labelling two randomly selected proliferative cells at the start of the simulation. At $t = 0$ two cells in the proliferative compartment of the tumour were randomly selected and coloured (pink, yellow). Each time one of the labelled cells divided, the two daughter cells were coloured with the same label. All other living cells are coloured dark blue regardless of their state (proliferative, quiescent or hypoxic), with cyan colouring indicating dead cells (including dead labelled cells).

rise to daughter cells on or near the spheroid boundary, causing the ^{3}H-labelled cell density to appear to move outwards. In contrast to the movement of microbeads, which all follow the advective flow towards the spheroid centre, we see an increase in the number of ^{3}H-labelled cells at the spheroid boundary due to proliferation.

This phenomenon is a consequence of cell proliferation combined with the way the proliferative rim "traps" particles described in Section 4.5.4. We have observed in Sections 4.5.2 and 4.5.4 that microbeads may remain in the proliferative rim for a significant length of time before beginning their inward movement due to advection. A similar delay applies to ^{3}H-labelled cells in the proliferative rim, which we have observed may remain there for several days of simulation time before ultimately being pushed into the quiescent compartment due to the pressure gradient. During this time, a ^{3}H-labelled cell may proliferate several times, producing labelled daughter cells which may be pushed either towards the spheroid boundary or inwards towards the quiescent region. This causes a significant growth in the number of ^{3}H-labelled in the simulation, with the new cells being added in the proliferative rim.

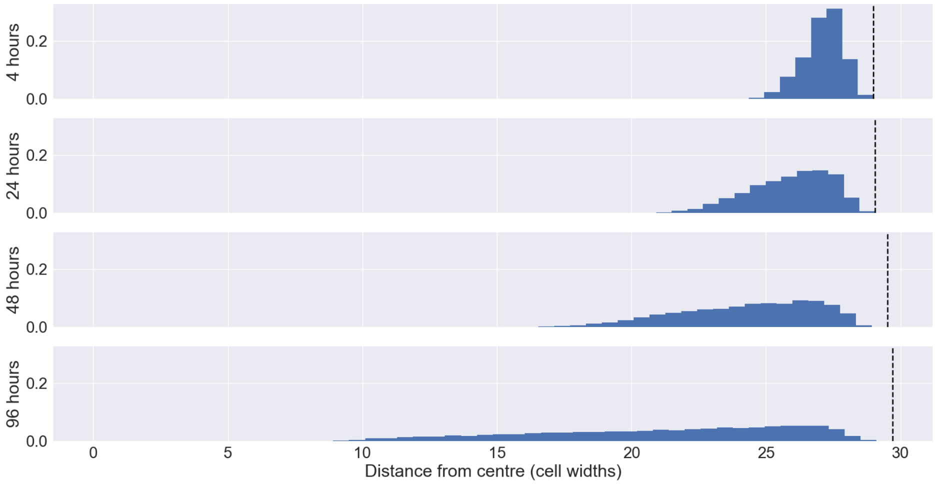

Figure 4.18: Frequency of ^{3}H-labelled tumour cells at four timepoints over 30 iterations of one example parameter set ($\eta = 16, \omega_\mathrm{q} = 0.5, \omega_\mathrm{h} = 0.1, \tau_\rho = 8$). The wave is more heavily skewed towards the tumour boundary due to the proliferation of ^{3}H-labelled cells in the proliferative rim.

Figures 4.18 and 4.19 demonstrate this process at work for one representative parameter set ($\eta = 16, \omega_\mathrm{q} = 0.5, \omega_\mathrm{h} = 0.1, \tau_\rho = 8$). In Figure 4.18 we present a histogram showing how the frequency of labelled cells changes over time in the same manner as Figure 4.10 demonstrates for microbeads. Although the dynamics of our simulated spheroid are identical to those of the spheroids described in the microbead simulations, the profile appears substantially different. The wave disperses more slowly than in the bead simulations and becomes flatter, with the peak of the wave always remaining near the spheroid boundary.

Figure 4.19 shows the same data set, but this time without scaling to show frequency. We show the total number of cells at each location rather than the proportion, which highlights the increase in cell number due to proliferation. Each bar in the histogram is coloured according to the proportion of cells counted in that bar which fall in to each of the three viable cell compartments: red for proliferative, blue for quiescent and magenta for hypoxic. In this Figure, the behaviour we describe above becomes clear. While the distribution is moving towards the spheroid centre, the height of the peak at the spheroid edge remains approximately constant. The

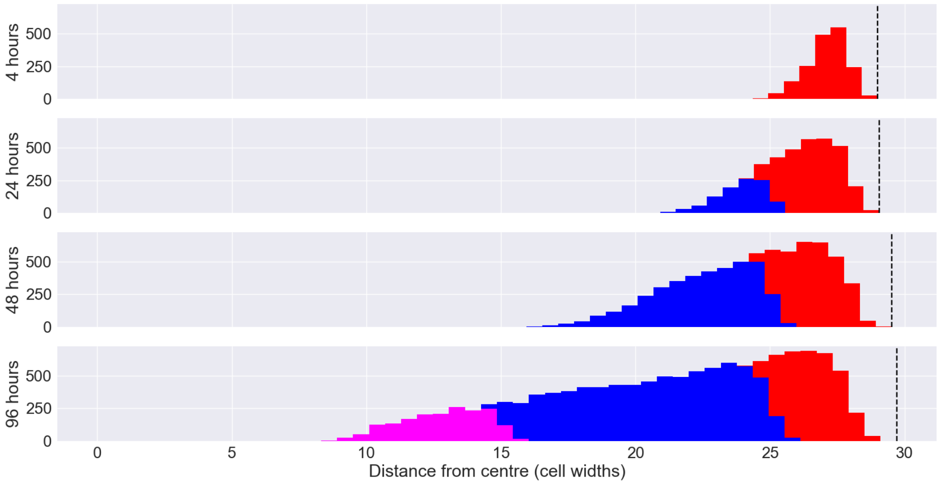

Figure 4.19: Total number of ^{3}H-labelled tumour cells at four timepoints over 30 iterations of one example parameter set ($\eta = 16, \omega_{\mathrm{q}} = 0.5, \omega_{\mathrm{h}} = 0.1, \tau_\rho = 8$). Each bar is coloured according to the corresponding cell compartments, with red indicating the number of proliferating labelled cells, blue denoting the number of quiescent labelled cells and magenta representing the number of hypoxic labelled cells.

increase in the number of cells is due to additional cells in the quiescent compartment and, at later timesteps, in the hypoxic compartment.

The distribution in Figure 4.19 could be due to a small population of “immortal” labelled cells residing permanently in the proliferative rim and sending their daughter cells inwards, but the most likely explanation is that any cell in the proliferative rim will eventually transit to the quiescent compartment. When a ^{3}H-labelled cell divides in the proliferative rim, it is equally likely that the mother or the daughter is closer to the spheroid centroid. However, since both cells are labelled identically the effect on the ^{3}H-labelled cell distribution is identical in either case. The effect is that the number of ^{3}H-labelled cells in the proliferative rim remains approximately constant over time, with the dispersion of the distribution due to cells in the proliferative rim dividing before being caught in the advective flow.

4.7 Discussion

In this Chapter, we have developed and presented a new agent-based hybrid model of diffusion limited tumour spheroid growth implemented in Chaste using an overlapping spheres framework. This model reproduces the classic logistic growth dynamics of multicellular tumour spheroids and can generate spheroids with similar growth curves but different proportions of proliferative, quiescent, and hypoxic cells and necrotic debris.

We have used our model to investigate the dynamics of "passive" movement within spheroids due to the advective flux generated between high pressure areas at the proliferative rim and low pressure areas of necrosis. Our model is successful at describing the experimentally observed movement of polystyrene microspheres added to such spheroids in experiments by Dorie et al. (1982, 1986), and provides a new mathematical framework for investigating these experiments. While several continuum models of this system have been proposed, our individual-based model is capable of providing an explanation for why the microbead distribution is a moving distribution which disperses over time. By tracking the progress of individual beads in this wave, we demonstrate that infiltrating particles all travel at the same speed but that the dispersion is caused by variations in the duration of time each spends trapped in the proliferative rim.

Our model indicates that the thickness of the proliferative rim and the balance between proliferation and death in a spheroid are crucial factors in both the speed of the advective flow towards necrotic regions and in the average time taken for a particle to cross the proliferative rim. These are key components of passive migration, and indicate that an understanding of how quickly tumour cells are proliferating relative to their rate of death by hypoxia is important in understanding immune cell penetration of tumours *in vivo*.

When extended to simulate the infiltration of ^{3}H-labelled cells into the spheroids, our model shows more success than previous continuum models. Most importantly, we do not need to postulate additional forces to cause labelled cells to remain at the spheroid boundary. Instead, the stochastic nature of proliferation near the spheroid edge is sufficient to cause a population of ^{3}H-labelled cells to remain at the spheroid boundary while the rest of the labelled population is carried with the advective flow.

Critically, our findings suggest that the immune invasion of solid tumours must be understood against the background of an advective flow driving cells towards necrotic regions. In the case of immune cells such as FoxP3$^+$ cells which rarely penetrate into tumours ((García-Martínez et al., 2014), see also Section 3.4.2) or of patients in which the bulk of the macrophage population resides in the stroma, we note that immune cells may need to actively resist this flow to remain on the tumour boundary. This is in contrast to many mathematical models of immune cell invasion, which often view the tumour as broadly static and infiltration as a purely active process (Ghaffarizadeh et al., 2018). We suggest that any future mechanical models involving movement into tumours where proliferation and death are localised should take the advective flow between these regions into account.

The model we have described in this Chapter provides a foundation for describing the passive advective movement of cells and particles within a spheroid. This foundation can be extended to account for macrophage invasion of spheroids by the addition of "active" properties to our microbead population. McElwain and Pettet (1993) note that "if internalisation is an active process, the most obvious candidate is chemotaxis". By adding macrophages which respond to chemotactic gradient produced by spheroids generated using our model, we can directly compare the effect that chemotaxis has on macrophage migration against a null model provided by the microbead simulations. In the following Chapter, we will describe extensions to our model which implement chemotaxis and compare the effects of these submodels against the null model provided in this Chapter.

Chapter 5

Macrophage infiltration: from *in vitro* to *in vivo*

5.1 Overview

In this Chapter, we extend the *in silico* model introduced in Chapter 4 to simulate the effect of chemotaxis on macrophage infiltration into tumour spheroids. In Chapter 4 we developed an agent-based model to investigate the cellular dynamics inside a tumour spheroid which are induced by the localisation of cell proliferation at the spheroid edge and cell death in the spheroid core. We use this work as a “null” model against which we compare predictions of a model that includes active migration by macrophages. We focus on the role of chemotaxis in promoting macrophage migration, with macrophages moving up a gradient of colony stimulating factor-1 (CSF-1) produced by hypoxic tumour cells.

The structure of this Chapter is as follows. In Section 5.1.1 we briefly motivate the extension of our agent-based model to include chemotactic movement of macrophages. Section 5.2 describes the model modifications required to simulate chemotaxis, including the production of CSF-1 by hypoxic tumour cells, the introduction of macrophages and extensions to the force laws used in Chapter 4 to accommodate chemotaxis. Section 5.3 contains results which demonstrate the impact of chemotaxis on the average radial velocities and waiting times of infiltrating macrophages, and in Section 5.4 we consider the effect of chemotaxis on the shape of the distribution formed by the radial locations of macrophages over time. In Section 5.5 we integrate our model directly with the imaging data analysed in Chapter 3 to simulate *in vivo* macrophage infiltration in realistic tumour geometries. Finally, Section 5.6 contains a discussion of the implications of chemotaxis for macrophage infiltration.

In Appendix B.2 we validate our chemotaxis model by comparing a simplified model of a macrophage subject to a chemotactic force with analytically calculated predictions for the time taken to cross a domain of fixed length subject to a chemotactic gradient. We also discuss assumptions made in our numerical solution for the CSF-1 gradient which could affect this implementation.

5.1.1 Biological motivation

In Section 1.2.2 we explained how the tumour microenvironment can determine macrophage behaviour. For example, tumour expressed factors such as CSF-1, IL-4, IL-10 and TGF-β polarize macrophages towards a pro-tumour phenotype, stimulating macrophage production of growth factors including IL-6 and TNFα (Sica et al., 2015). In addition to affecting macrophage phenotype, there are several cytokines, known as chemokines, which induce chemotaxis in cells that possess the relevant receptors.

Many cytokines present in the tumour microenvironment are macrophage chemoattractants, including colony stimulating factor-1 (CSF-1), CCL2, CCL5 and vascular endotheial growth factor-A (VEGF-A) (Harwood et al., 2012; Murdoch et al., 2004). Some of these cytokines, in particular VEGF-A, are primarily associated with angiogenesis and are produced by both tumour associated macrophages and hypoxic tumour cells (Harwood et al., 2012). During wound healing in disease-free tissue, CSF-1 is produced by hypoxic cells to recruit macrophages to the wound site from the surrounding tissue and vasculature. While recruited macrophages will typically be highly activated M_1 macrophages, particularly when a wound has broken the skin, the presence of factors such as CSF-1 eventually biases the phenotype of macrophages towards a pro-tumour phenotype. This stimulates macrophage production of vascular growth factors such as VEGF-A, which cause angiogenesis and revascularisation of the damaged tissue (Sidibe et al., 2018; Murdoch et al., 2004; Wynn et al., 2013; Wynn and Vannella, 2016). In tumours, CSF-1 is produced by hypoxic tumour cells, which are often in the centre of “tumour nests”, clusters of tumour cells which group together in structures that resemble tumour spheroids. In these nests, the absence of nearby blood vessels can cause hypoxia if insufficient levels of nutrients are available to diffuse through the densely packed tumour cells. This CSF-1 gradient stimulates macrophage recruitment to the tumour and their infiltration into tumour nests. This explains the high density of macrophages in hypoxic regions in comparison to other immune cells ((Bindea et al., 2013), see also Section 3.4.2).

In this Chapter, we extend the model described in Chapter 4 to account for macrophage chemotaxis towards cytokines expressed by hypoxic tumour cells. Although a range of macrophage chemokines are produced by hypoxic tumour cells, including CSF-1, VEGF-A, CCL2 and CCL5, for simplicity we consider only CSF-1 and its effect on macrophages and assume that this is representative of the response of macrophage to other chemokines expressed by hypoxic tumour cells. This choice reflects interest in the CSF-1 receptor as a target for immunotherapy: emactuzumab, an anti-CSF-1R antibody, is currently in Phase I trials at Roche (Pradel et al., 2018, 2016). Treatment with CSF-1R inhibitors is theorised to bias macrophages towards an anti-tumour phenotype, as well as preventing tumour recruitment of macrophages through expression of CSF-1.

Previous mathematical models of macrophage chemotaxis within avascular tumour spheroids and vascular tumours also model only one chemoattractant. A series of models by Owen and Sherratt (1998, 1997, 1999) focus on macrophage lysis of tumour cells in small, avascular tumours. They consider a single "chemical regulator" which acts as a macrophage chemoattractant and promotes macrophage proliferation. Kelly et al. (2002) use PDEs to model macrophage infiltration into tumour spheroids which express a generic chemoattractant, based on data from (Leek et al., 1999) of tumour spheroids which express VEGF-A, a macrophage chemoattractant. Webb et al. (2007) use a multiphase model to simulate macrophage infiltration into radially symmetric avascular tumour spheroids in response to a single chemokine expressed by hypoxic tumour cells. In addition to chemokines produced by hypoxic tumour cells, recent models have focussed on the CSF-1/EGF (epidermal growth factor) paracrine loop (Goswami et al., 2005). In this loop, macrophages are attracted to CSF-1 produced by tumour cells which are in turn attracted to EGF produced by the macrophages. This loop has been modelled using PDEs and agent-based models by Knútsdóttir et al. (2014, 2016), and is implicated in tumour metastasis. More detail about these models can be found in Section 1.4.1, where we reviewed previous mathematical models of macrophages in cancer.

5.2 Extensions to agent-based model to account for macrophage chemotaxis

In this Section, we present modifications to the agent-based model which we developed in Chapter 4. We describe how microbeads and macrophages differ in our model, the rules used to include CSF-1 production by hypoxic tumour cells, and the force laws used to model macrophage chemotaxis.

5.2.1 Introduction of macrophages

The model which we presented in Chapter 4 included two classes of “cell”: tumour cells and microbeads. We model macrophages by modifying the rules which apply to microbeads to account for the differences between inert microbeads and living macrophages. The main difference in the models is that macrophages actively move up the gradient of CSF-1. We implement this as an additional force applied to the cell, which we describe in Section 5.2.3. The motility and malleability of macrophages could be implemented by reducing μ and adjusting D. However, in order to retain as many similarities as possible between our investigation of passive infiltration in Chapter 4 and this investigation of active migration, we opt to keep these parameters the same.

We do not include macrophage phagocytosis of tumour cells, or their production of cell killing factors such as nitric oxide (NO) or reactive oxygen intermediates (ROIs) (Capece et al., 2013; Mantovani et al., 2002). These cell-killing phenotypes are a key part of the models of Owen and Sherratt (1998, 1997, 1999). We note that this behaviour is predominantly associated with the classically-activated, M_1, phenotype. Tumour associated macrophages are generally considered to express a pro-tumour, M_2, phenotype, particularly in the presence of CSF-1 (Kratochvill et al., 2015; Mantovani et al., 2002). While the effect of macrophage phenotype on tumour growth is an important one, our model does not distinguish between different phenotypes of macrophage, and macrophages in our simulation are neither pro-tumour nor anti-tumour. We also do not consider macrophage proliferation or death in our simulations. Macrophages proliferate slowly, and are presumed to have a higher tolerance for hypoxia than other cell types. By choosing not to include macrophage proliferation or death, we can examine the effect of chemotaxis of spheroid infiltration without needing to account for the effects of proliferation and death on the distribution.

As a result of these modelling choices, there is no difference between macrophages which do not respond to chemoattractants in this Chapter and microbeads in Chapter 4. By changing only the response to chemotaxis, we aim to more fully understand the impact that chemotaxis has in the context of the advective flow examined in Chapter 4.

5.2.2 CSF-1 production

In this Chapter we use CSF-1 as a surrogate for all macrophage chemoattractants produced by hypoxic tumour cells, as discussed in Section 5.1.1. As CSF-1 is produced by hypoxic tumour cells (Harwood et al., 2012), we assume that viable tumour cells in the hypoxic compartment - those with oxygen concentration $\omega \leq \omega_{\mathrm{h}}$ - are the only sources of CSF-1. In particular, necrotic cells do not produce CSF-1 in our model.

We use a reaction-diffusion equation to simulate CSF-1 production and diffusion, with hypoxic cells acting as point sources of CSF-1. The CSF-1 concentration at a point $\mathbf{x}$ and time t, $c(\mathbf{x}, t)$, is thus given by:

$$\frac{\partial c}{\partial t} = D_c \nabla^2 c + \kappa_c \sum_i \delta(\mathbf{x} - \mathbf{x}_i) \tag{5.1}$$

where D_c is the diffusion coefficient for CSF-1, which we assume to be constant, κ_c is the production rate of CSF-1 by hypoxic cells and $\mathbf{x}_i$ indexes the hypoxic cells in the simulation. $\delta(\mathbf{x} - \mathbf{x}_i)$ is a delta function which equals 1 when the argument is zero, and is otherwise equal to zero. We assume that diffusion of CSF-1 takes place on a much faster timescale then cell movement, and so as when modelling the oxygen distribution described in Equation (4.6), we make the approximation that ${}^{\partial c}/_{\partial t} = 0$. The values in this Equation are non-dimensionalised using the same scaling regime as in Chapter 4, described in Equation (4.4). At each simulation timestep, we calculate the CSF-1 concentration according to the equation:

$$0 = D_c \nabla^2 c + \kappa_c \sum_i \delta(\mathbf{x} - \mathbf{x}_i). \tag{5.2}$$

To apply boundary conditions, we first specify the domain on which Equation (5.2) is solved by surrounding the spheroid with a square of side length L, where L is much greater than the typical size of a spheroid which reaches growth equilibrium. We assume that the CSF-1 concentration decays to 0 at large distances from the spheroid, and model this with Dirichlet boundary conditions, $c(\mathbf{x}, t) = 0$, on the boundary of the square. As for the oxygen concentration described in Equation (4.6), we solve Equation (5.2) on the nodes of a regularly spaced triangular finite element mesh which spans this square domain.

We note that an alternative implementation of CSF-1 concentration would be to include decay of CSF-1 with zero-flux boundary conditions. We have chosen to use the model described in Equation (5.2) to simplify the implementation of the CSF-1 model. The lack of CSF-1 degradation in our model means that a biologically realistic CSF-1 gradient must be established by applying the far-field Dirichlet boundary conditions described in the previous paragraph. These boundary conditions generate a CSF-1 distribution which smoothly reduces to zero at large distances from the spheroid. In the region of interest, close to the tumour spheroid, our model generates a CSF-1 gradient which follows the approximate form expected biologically. The CSF-1 concentration is highest at the hypoxic boundary, and decreases radially through the quiescent and proliferative portions of the spheroid. The analysis in this Chapter focuses on the impact of varying the sensitivity of macrophages to the chemotactic gradient rather than on the shape of the gradient itself. Equation (5.2) is suitable for this analysis, but care should be taken if our model is used to draw conclusions about the precise shape of the chemotactic gradient within a tumour spheroid.

Figure 5.1a shows an example of the CSF-1 distribution around a spheroid, with half of the spheroid cropped away to show the distribution underneath. This spheroid has reached steady state, and no macrophages have been added to the simulation. The distribution is not visibly distorted by the size of the box, but we note that for tumours with sufficiently large steady state radii this choice of domain shape could distort the radial symmetry of the CSF-1 distribution. In practice, we fix L to be greater than any of the observed steady state sizes (none of the simulations in this Chapter appear to show this type of distortion). Figure 5.1b shows the radial cross section of the CSF-1 concentration for a spheroid at steady state.

5.2.3 Modified force law

Previous agent-based models have implemented macrophage chemotaxis by the addition of an extra force. In their model of macrophage chemotaxis in the EGF/CSF-1 paracrine loop, Knútsdóttir et al. (2016) add a chemotactic force when the gradient of the chemoattractant is above a threshold, although the magnitude of this force is not specified. Following Perfahl et al. (2017), we define the chemotactic force experienced by a single macrophage i as:

$$\mathbf{F}_i^{\chi} = \chi \nabla c(\mathbf{x}, t). \tag{5.3}$$

We extend our previous model by adding this term to the force balance given in Equation (4.9), which gives a new equation of motion for macrophage movement:

$$\nu \frac{\mathrm{d}\mathbf{x_i}}{\mathrm{d}t} = \mathbf{F}_i^m + \mathbf{F}_i^r + \mathbf{F}_i^s + \mathbf{F}_i^{\chi} \tag{5.4}$$

where, as before, ν is the damping coefficient, $\mathbf{F}_i^m$ is mechanical force due to interactions with other cells, $\mathbf{F}_i^r$ is random motility and $\mathbf{F}_i^s$ is the surface tension acting on the spheroid boundary. Physical interactions between macrophages and tumour cells are included in the term $\mathbf{F}_i^m$, but we do not include other interactions between macrophages and tumour cells. Specifically, in this model macrophages do not at-

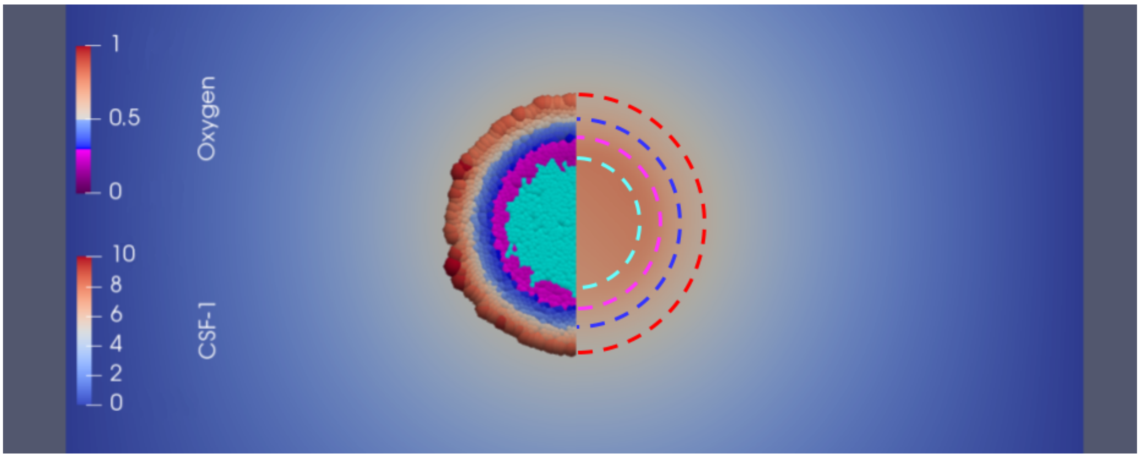

(a) Snapshot showing CSF-1 distribution. Cells are coloured according to oxygen concentration, and the background shows CSF-1 concentration. Half of the spheroid is hidden to make the CSF-1 distribution visible, with dashed lines showing the approximate boundaries of the different tumour compartments. ($\omega_{\mathrm{q}} = 0.5$, $\omega_{\mathrm{h}} = 0.3$, $\eta = 32$, $D_c = 1$, $\kappa_c = 0.01$ and $L = 100$).

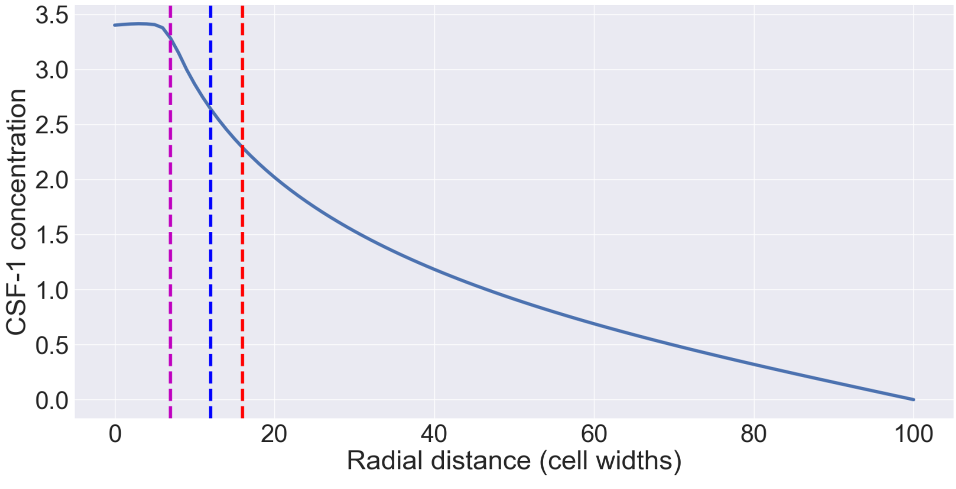

(b) Radial distribution of CSF-1. Dashed lines show approximate boundaries for hypoxic (magenta), quiescent (blue) and proliferative (red) compartments. ($\omega_{\mathrm{q}} = 0.5$, $\omega_{\mathrm{h}} = 0.3$, $\eta = 16$, $D_c = 1$, $\kappa_c = 0.01$ and $L = 100$).

Figure 5.1: Examples of the spatial distribution of CSF-1 associated with spheroids at steady state.

tack tumour cells as in (Owen and Sherratt, 1998, 1997, 1999). Macrophages also do not produce pro-tumour growth factors in our model (Seager et al., 2017; Sousa et al., 2015). Macrophage phenotype is crucial in determining their behaviour, but is dependent on the tumour microenvironment and cannot be inferred from the data available to us. As a result, we choose not to include either pro-tumour or anti-tumour behaviour in this model. Instead of focussing on the impact of the tumour environment on macrophage phenotype, our model considers the effect of the tumour composition on macrophage location. This simplification allows the role of chemotaxis to be evaluated directly by comparing the distribution of macrophages for varying χ, without confounding effects introduced by varying macrophage phenotype.

We introduce one more new parameter to our model, t_{add}, which is the time at which we add macrophages to the spheroid. In Chapter 4 this parameter was fixed at $t_{\text{add}} = 300$ hours, when all spheroids were assumed to have reached their steady state size. When $\chi = 0$ and $t_{\text{add}} = 300$, our parameter choices match those in Chapter 4.

5.3 Impact of chemotaxis on macrophage velocity and waiting time

Prior to adding macrophages, spheroid growth follows the same dynamics as those described in Section 4.4. In Chapter 4, we investigated the effect that varying parameters associated with spheroid composition had on the passive infiltration of inert microbeads. In this Section, we compare the average radial velocity and waiting times that we described for microbeads in Chapter 4 with those of macrophages in response to chemotaxis and consider the effect of varying the chemotactic sensitivity coefficient, χ. Throughout this Chapter, we vary χ while keeping D_c and κ_c constant.

In Figures 5.2 - 5.10 we compare macrophage infiltration into two spheroids generated using the same parameter set ($\eta = 32$, $\omega_{\mathrm{h}} = 0.3$, $\omega_{\mathrm{q}} = 0.5$, $\tau_\rho = 8$, where η is the average cell cycle length, ω_{h} and ω_{q} are the oxygen thresholds for hypoxia and quiescence, and τ_ρ is the critical duration of hypoxia which induces cell death.) Macrophages were added after $t_{\mathrm{add}} = 300$ hours of spheroid growth, and the simulations differ only in the value of χ. In one spheroid, macrophages are insensitive to chemotaxis ($\chi = 0$) while in the other chemotaxis is relatively strong ($\chi = 10$). The results shown in these Figures are representative of a range of parameter values, and we investigate the effect of varying τ_ρ, η and ω_{q} in Figure 5.9.

Figure 5.2 shows snapshots from these simulations when the macrophages are added, and again after 2 days and 4 days. The spheroid on the left has $\chi = 0$, and on the right $\chi = 10$. Two differences in macrophage distribution are apparent from these simulation results. Firstly, when chemotaxis is not present several macrophages move away from the spheroid into the extracellular space. This behaviour was also observed in the microbead experiments in (Dorie et al., 1982), who reported that many of the microbeads would not adhere to the spheroids. In contrast, when $\chi > 0$, all of the macrophages are adhering to the spheroid rim or infiltrating the spheroid. Secondly, when $\chi = 0$ only a few macrophages have reached the edge of the necrotic region after 96 hours. When $\chi > 0$, a large number of macrophages have localised within the necrotic core. Experimentally, macrophages under these severe hypoxic conditions would eventually die themselves, but we do not include macrophage death in our simulations.

In Figure 5.3 we show the proportion of macrophages in each tumour compartment over time, averaged over 20 realisations of the same parameter sets. Proportions are coloured according to the colour scheme previously used for tumour compartments, but we have introduced a green compartment in Figure 5.3 which shows the proportion of macrophages which are in extracellular space and have not infiltrated the tumour. Figure 5.3 shows that when $\chi > 0$ almost all macrophages

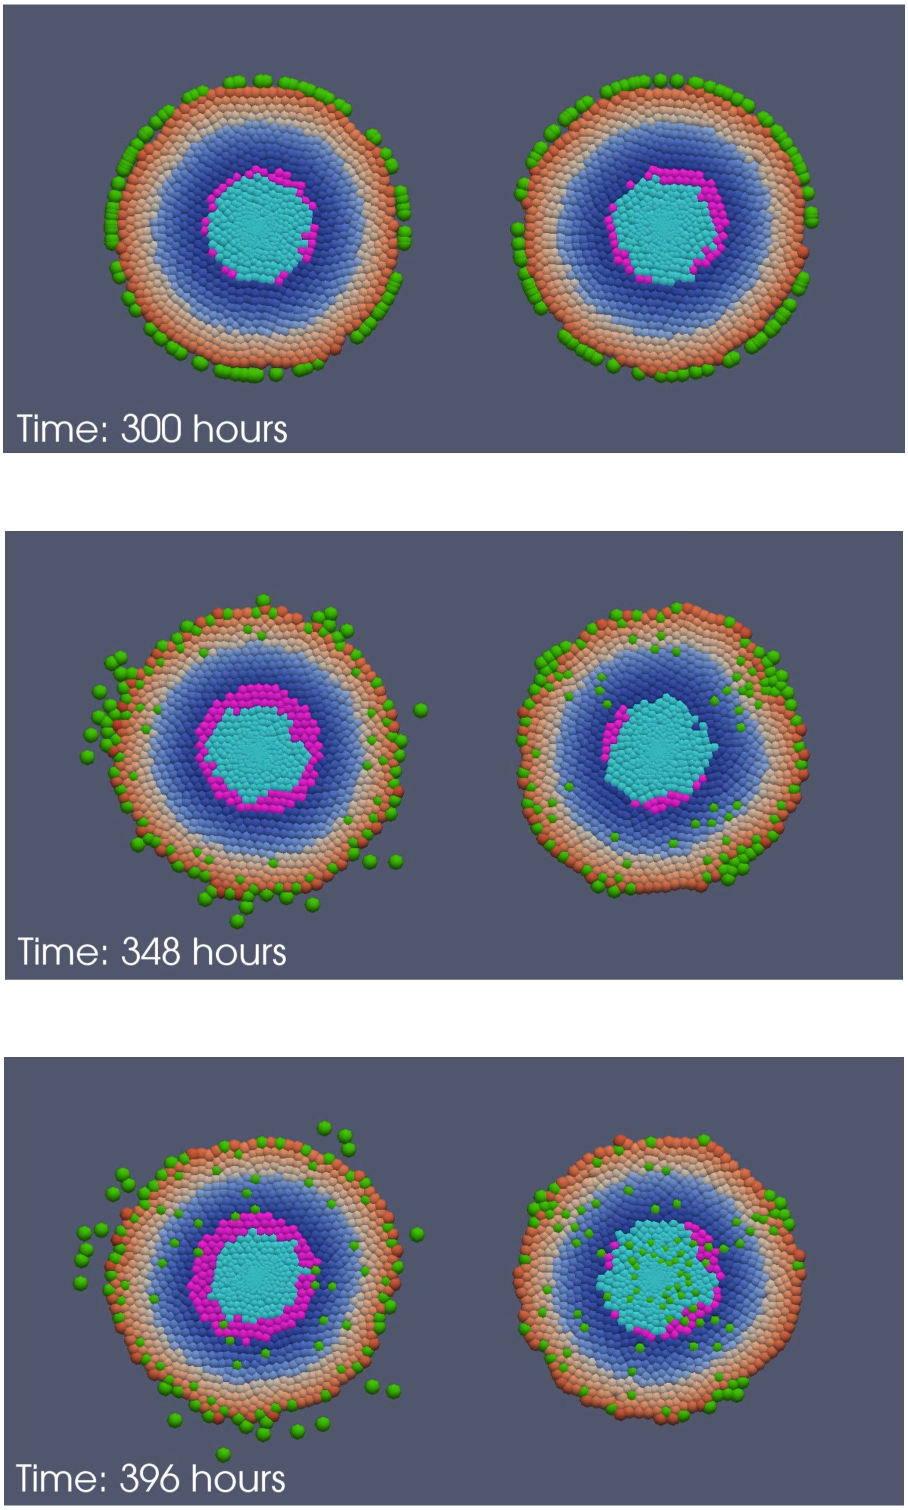

Figure 5.2: Comparison between macrophage infiltration in two identical spheroids, with no chemotaxis on the left ($\chi = 0$) and high chemotaxis on the right ($\chi = 10$). Screenshots are shown when macrophages are added at $t = 300$ hours, after 2 days, and after 4 days. (Parameters for both spheroids: $\eta = 32$, $\omega_{\mathrm{h}} = 0.3$, $\omega_{\mathrm{q}} = 0.5$, $\tau_\rho = 8$)

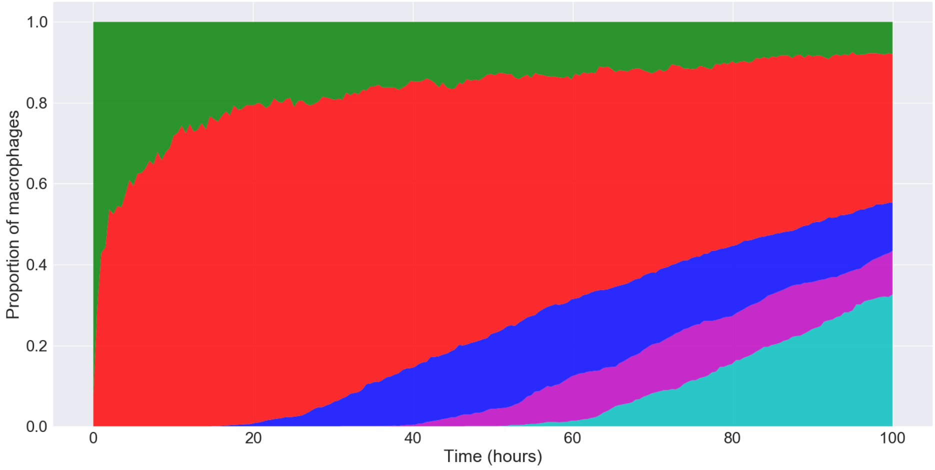

(a) Proportion of macrophages in different tumour compartments over time ($\chi = 0$).

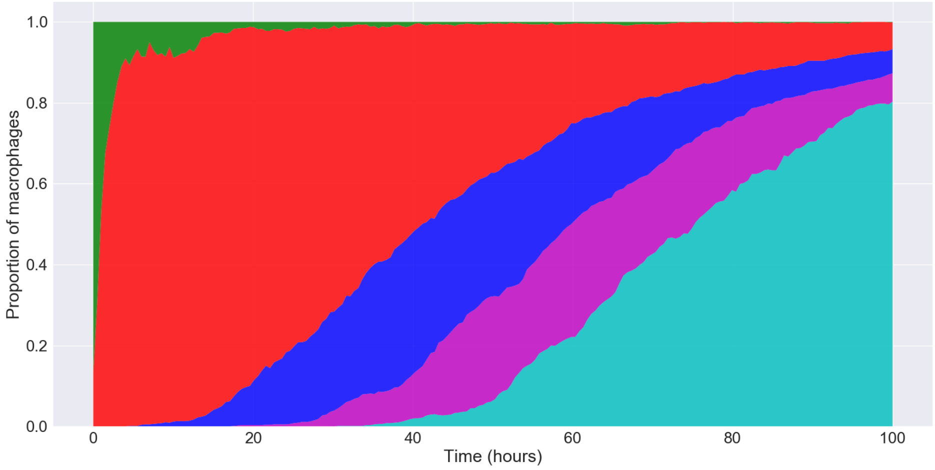

(b) Proportion of macrophages in different tumour compartments over time ($\chi = 10$).

Figure 5.3: Proportion of macrophages in different tumour compartments over time in the presence of absence of chemotaxis (average over 20 realisations). Macrophages are more likely to enter the spheroid when χ is high, and infiltrate to internal compartments faster than when $\chi = 0$. (Parameters for both spheroids: $\eta = 16$, $\omega_{\mathrm{h}} = 0.3$, $\omega_{\mathrm{q}} = 0.5$, $\tau_\rho = 8$)
Cyan = necrotic, magenta = hypoxic, blue = quiescent, red = proliferative, green = outside the spheroid.

eventually infiltrate the spheroid. In contrast, when $\chi = 0$ approximately 10% of macrophages have not begun to infiltrate after 100 hours. Figure 5.3 also shows that when $\chi = 10$ macrophages reach the quiescent, hypoxic and necrotic compartments of the tumour faster than when $\chi = 0$.

From Figures 5.2 and 5.3, we conclude that chemotaxis increases macrophage infiltration. Figure 5.4 demonstrates that although chemotaxis causes macrophages to infiltrate at a faster radial velocity, the increase in velocity varies in different compartments of the tumour. In Figure 5.4a, we show the average radial velocity of a macrophage ($\chi = 10$) and a microbead ($\chi = 0$) as a function of distance from the spheroid centre, averaged over 20 realisations of a representative parameter set. In both cases, average radial velocity is slower in the proliferative rim of the tumour compared with either the hypoxic or quiescent compartments. The slow movement in the proliferative rim is representative of the waiting time, and is caused by proliferation pushing macrophages either towards the spheroid centre or outwards towards the rim. Velocity increases in the quiescent region as the interplay between proliferation and death causes the macrophage to move with the advective flux. In both regions, increasing chemotaxis causes a slight increase in average radial velocity inwards, but as average velocity in the proliferative rim is lower than in the quiescent region, the effect of chemotaxis is more noticeable here.

Figure 5.4a shows that in the necrotic compartment of the tumour, higher chemotaxis causes macrophages to move inwards more slowly than microbeads. This implies that active migration towards CSF-1 is resisting the passive advection of cells in this area, rather than aiding it, although we note that the low number of microbeads which reach the necrotic region makes this observation potentially unreliable. In the hypoxic compartment, the difference between the infiltration velocities for $\chi = 0$ and $\chi = 10$ disappears. This difference can be seen in Figure 5.4b, which shows the difference between the infiltration velocities for $\chi = 0$ and $\chi = 10$. The change in the velocity of macrophages caused by increasing chemotaxis is approximately constant

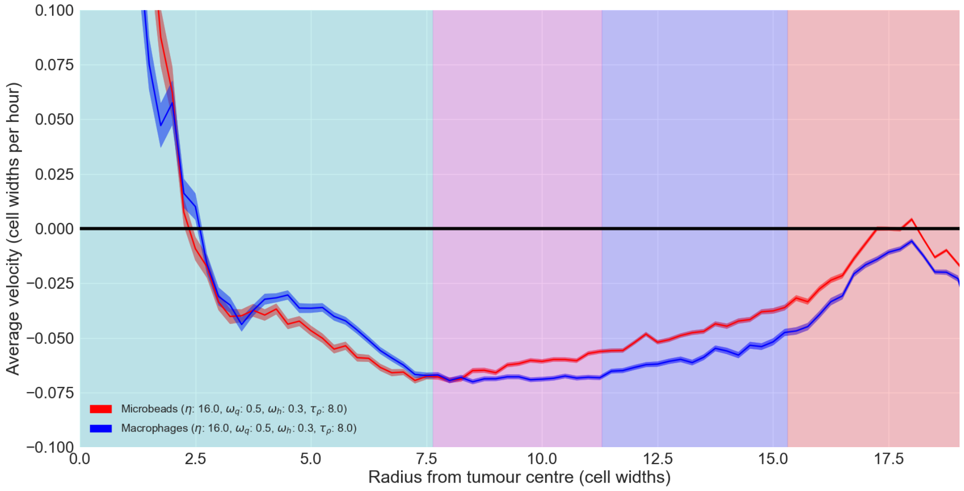

(a) Comparison of the average radial velocity of microbeads ($\chi = 0$) and macrophages ($\chi = 10$) as a function of radial distance from the spheroid centre (± Standard Error).

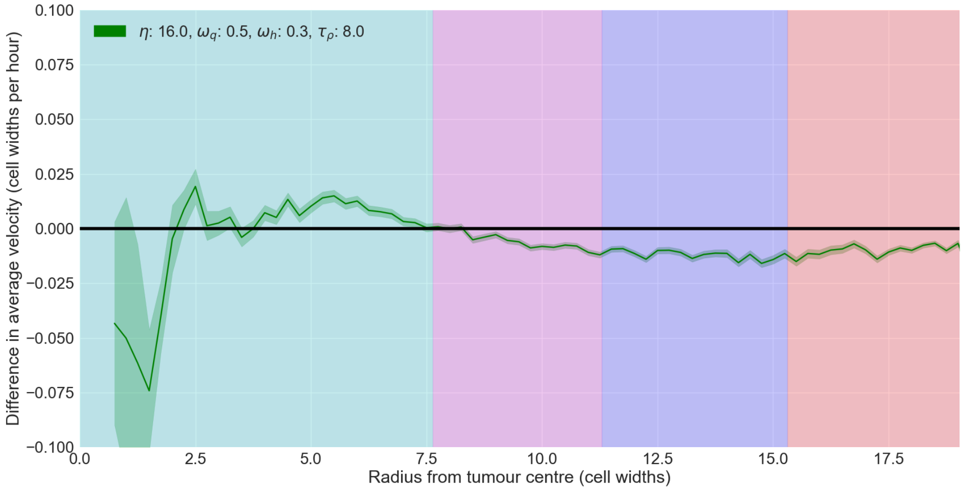

(b) Difference in average radial infiltration velocities between $\chi = 0$ and $\chi = 10$ (± Standard Error). Negative differences show that chemotaxis is causing an increase in velocity towards the spheroid centre, and positive differences show that it is reducing inward radial velocity.

Figure 5.4: Velocity of microbeads or macrophages as a function of radius from the spheroid centre averaged over 20 realisations of spheroids with $\eta = 16$, $\omega_{\mathrm{h}} = 0.3$, $\omega_{\mathrm{q}} = 0.5$, $\tau_\rho = 8$.
The radius of tumour compartments is shown by colour: Cyan = necrotic, magenta = hypoxic, blue = quiescent, red = proliferative.

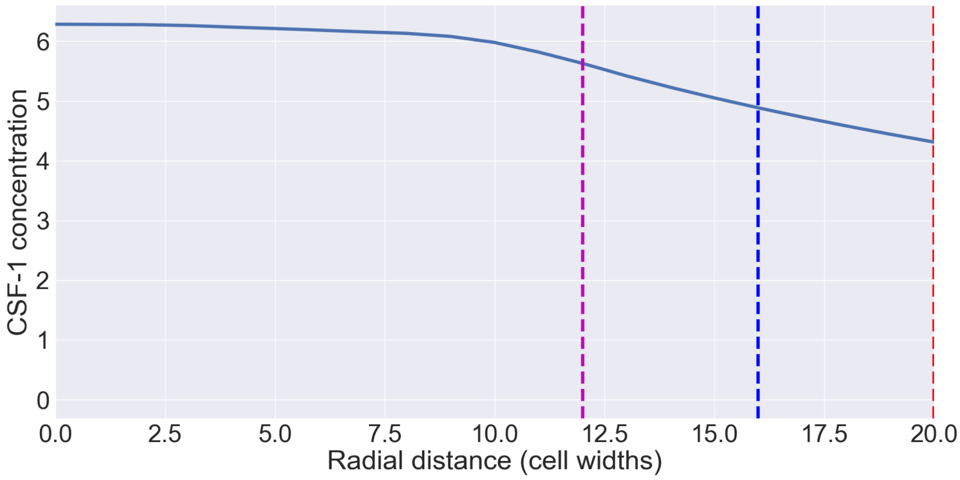

Figure 5.5: CSF-1 concentration associated with the spheroids described in Figure 5.4, with approximate compartment boundaries marked by dashed vertical lines.

in the proliferative and quiescent regions, but increases in the hypoxic region until the difference vanishes at the inner boundary of the hypoxic region, which is where cells stop producing CSF-1.

The results in Figure 5.4 represent an average of 20 realisations of one parameter set. To show that the effect of chemotaxis is to produce similar behaviour in the same spheroid compartments across parameter sets, in Figure 5.6 we consider the average behaviour across a range of parameter combinations. As varying the spheroid growth parameters causes the relative compartment sizes to change, we show the difference in average velocity across each compartment by first normalising the radius of the compartment. If $R(\omega)$ is the steady state spheroid radius at which the average oxygen concentration is ω, then for a given parameter set the proliferative rim is r such that $R(\omega_{\mathrm{q}}) < r \leq R(1)$. For each spheroid, we map this interval onto the interval $(0, 1]$. Figure 5.6 shows a schematic of this process applied to two different parameter sets (top), as well as the average difference in velocities for all parameter sets in the proliferative rim ($n = 36$ parameter sets). Each value of χ causes a small increase in macrophage infiltration velocity close to the spheroid boundary at $R(1)$. The velocity increase in the inner 60% of the proliferative rim is more

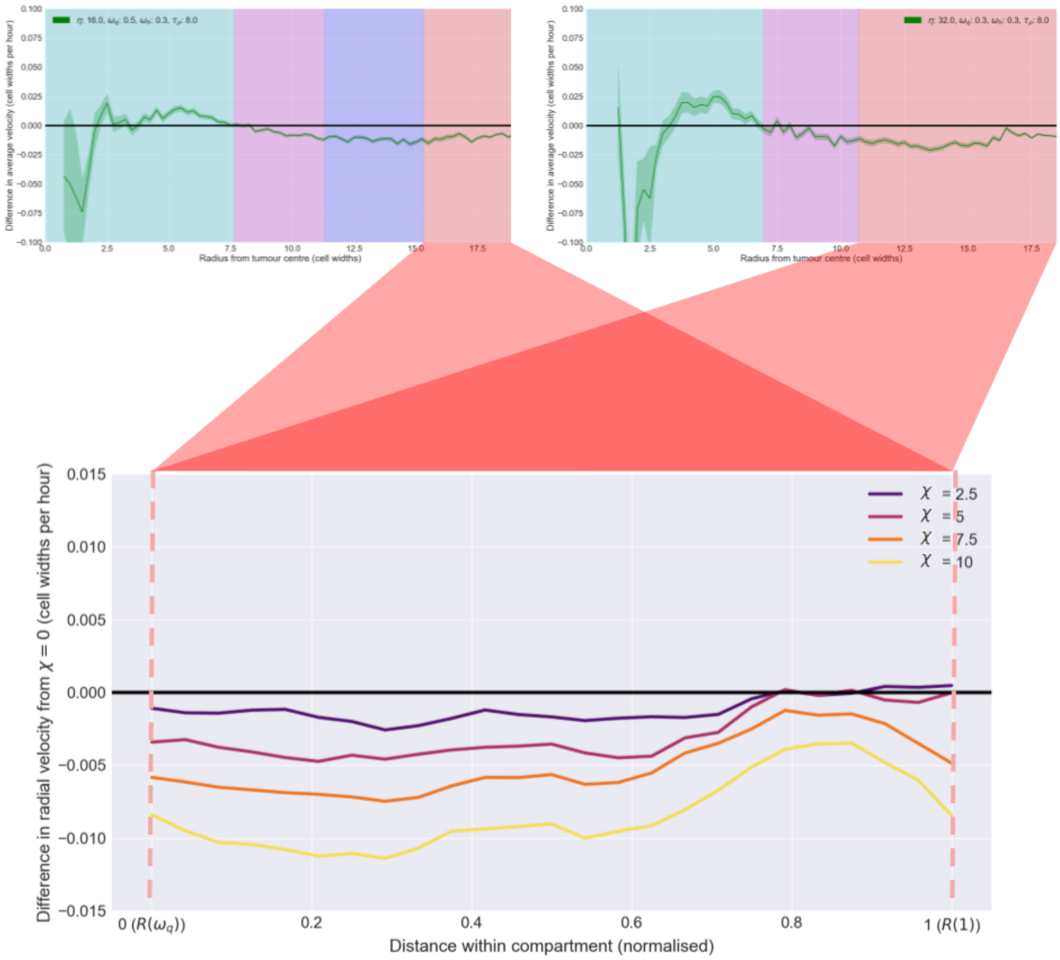

Figure 5.6: Average effect of adding chemotaxis on radial velocity in the proliferative compartment of spheroids with different widths, averaged over 20 realisations each for 36 parameter sets. The upper two graphs demonstrate how the proliferative rim of each parameter set is normalised onto the region (0,1].

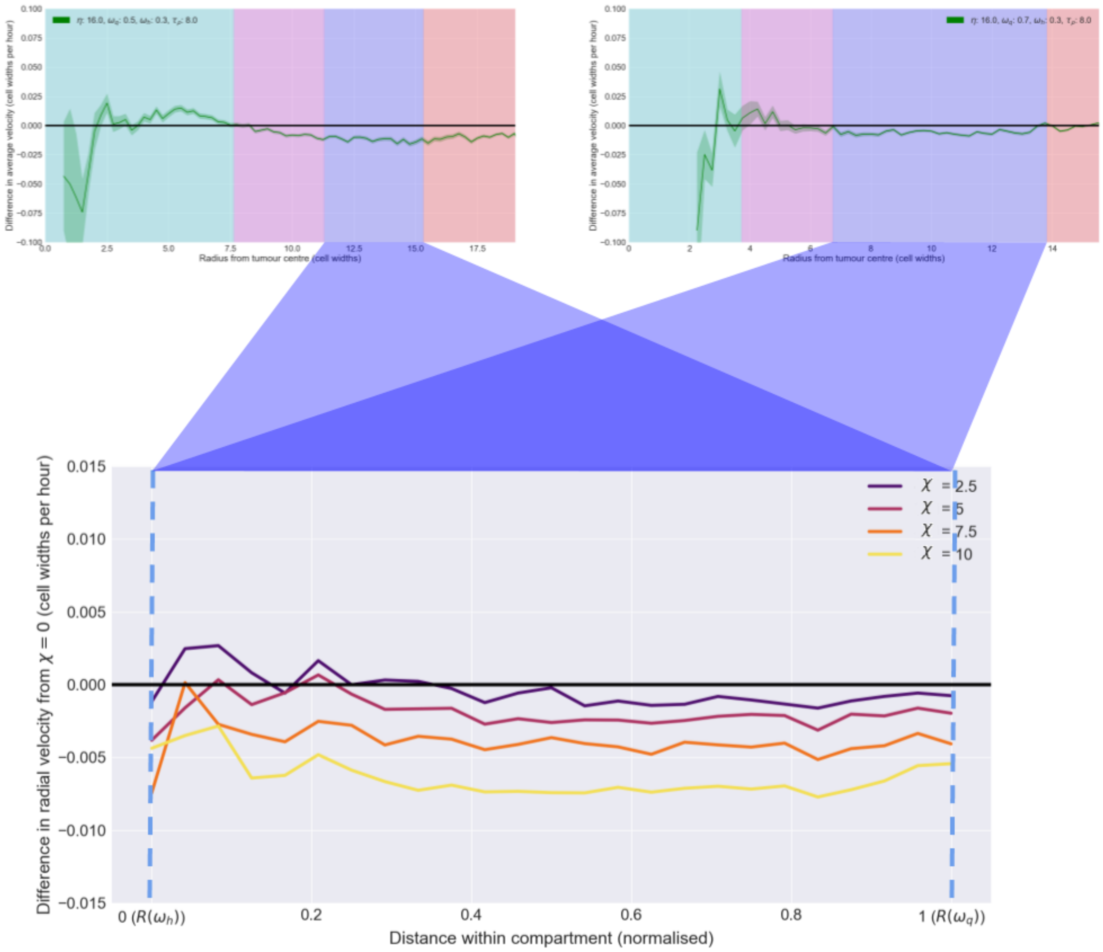

Figure 5.7: Average effect of adding chemotaxis on radial velocity in the quiescent compartment of spheroids with different widths, averaged over 20 realisations each for 36 parameter sets. The upper two graphs demonstrate how the quiescent compartment of each parameter set is normalised onto the region (0,1].

pronounced, and is approximately constant. Figure 5.7 follows the same protocol as Figure 5.6, but for the quiescent compartment. It confirms that increasing χ causes macrophages to move through the quiescent compartment at a higher radial velocity. The difference is approximately constant, although there is a noisy region close to the boundary of the hypoxic compartment at $R(\omega_{\mathrm{h}})$.

In Chapter 4, microbeads passively infiltrated the spheroids with an approximately constant velocity, but with a delay (or "waiting time"), caused by the birth of new cells pushing microbeads away from the spheroid centre. In Figure 5.8 we compare the trajectories taken by microbeads in Figure 4.11a with those of macrophages responding to chemotaxis, and note that macrophages follow similar trajectories, although macrophages appear to spend less time in the initial phase waiting close to the spheroid boundary. In Section 4.5, we assigned each microbead a waiting time and an infiltration velocity based on the time at which they crossed a threshold of 80% of the spheroid radius. We use the same methodology in Figure 5.9a to assign average velocities to macrophages as we vary η, ω_{q}, χ and t_{add}. Figure 5.9a shows that increasing χ from 0 to 10 leads to an increase in average macrophage radial velocity of between 0.02 and 0.1 cell widths per hour. This increase is small compared to the effect of varying η and ω_{q}, which can cause an increase of up to 0.4 cell widths per hour. The time at which macrophages are added, t_{add}, does not have a strong effect on the average macrophage infiltration velocity.

We also showed in Chapter 4 that the microbead "waiting time", the average time taken for a microbead to cross the proliferative rim, depended on η and ω_{q}. We repeat this analysis in Figure 5.9b, where we also vary χ and t_{add}. Although varying t_{add} has no effect on waiting time, increasing χ is associated with a reduced waiting time for every combination of parameters considered.

We note that the effect of increasing χ is more significant in comparison with varying η and ω_{q} for the waiting time than for the average macrophage velocity. Varying any one of χ, η or ω_{q} the range of values tested causes waiting times to

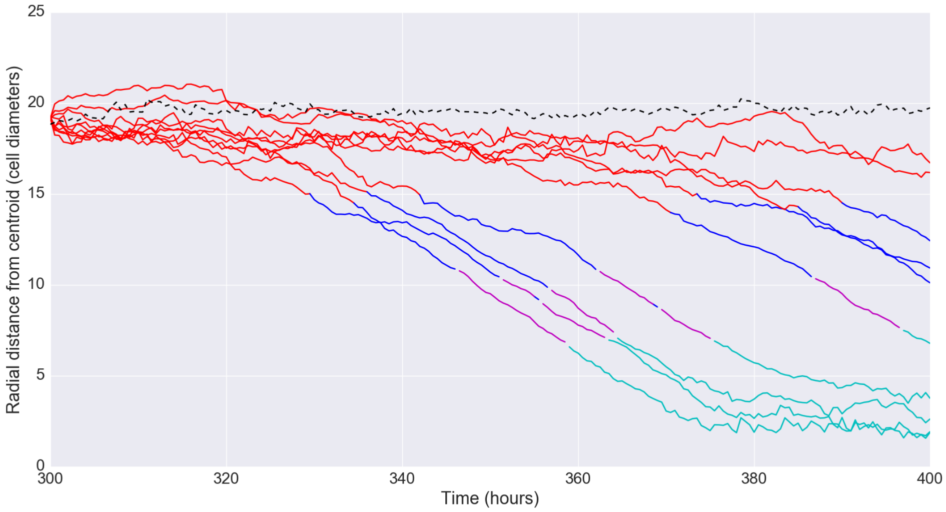

(a) Trajectories from 10 random beads from a simulation with $\omega_{\mathrm{q}} = 0.5$, $\omega_{\mathrm{h}} = 0.3$, η= 16, τ_{ρ}= 8 (reproduced from Figure 4.11a)

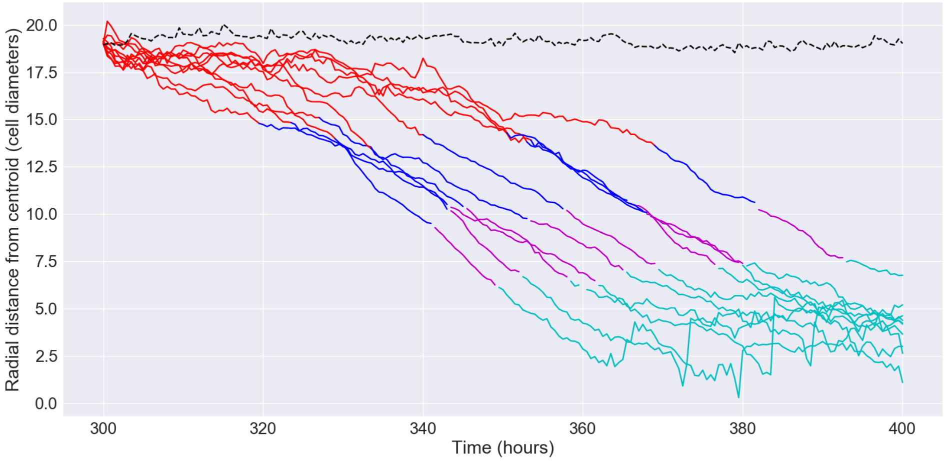

(b) Trajectories from 10 random macrophages from a simulation with the same parameters as Figure 5.8a, but with $\chi = 10$.

Figure 5.8: Trajectories of microbeads (Figure 5.8a) and macrophages (Figure 5.8b) infiltrating a tumour spheroid generated with the same parameters. The most notable difference is that the macrophages responding to chemotaxis spend less time in the Brownian motion dominated regime at the spheroid edge than the microbeads.

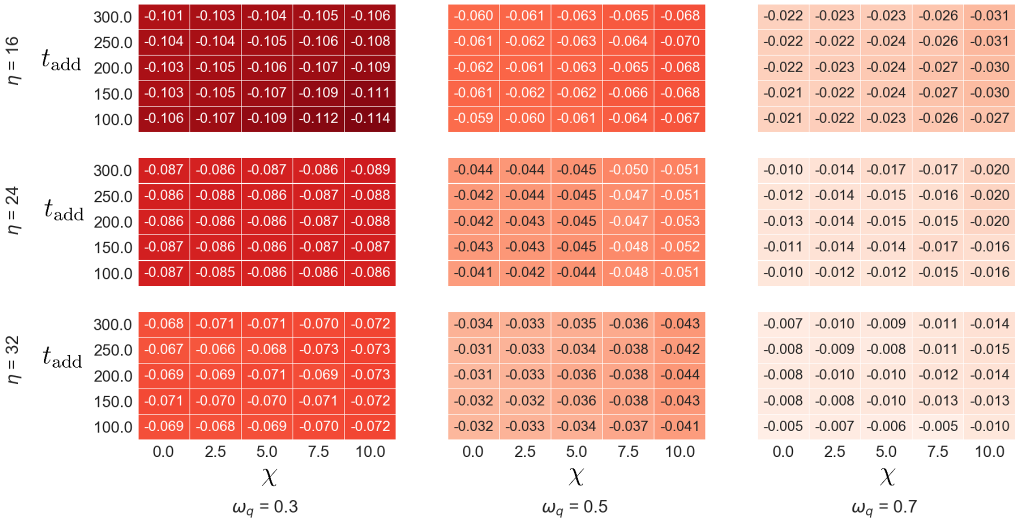

η	t_{add}	ω_q = 0.3					ω_q = 0.5					ω_q = 0.7				
	χ	0.0	2.5	5.0	7.5	10.0	0.0	2.5	5.0	7.5	10.0	0.0	2.5	5.0	7.5	10.0
η = 16	300.0	-0.101	-0.103	-0.104	-0.105	-0.106	-0.060	-0.061	-0.063	-0.065	-0.068	-0.022	-0.023	-0.023	-0.026	-0.031
	250.0	-0.104	-0.104	-0.105	-0.106	-0.108	-0.061	-0.062	-0.063	-0.064	-0.070	-0.022	-0.022	-0.024	-0.026	-0.031
	200.0	-0.103	-0.105	-0.106	-0.107	-0.109	-0.062	-0.061	-0.063	-0.065	-0.068	-0.022	-0.023	-0.024	-0.027	-0.030
	150.0	-0.103	-0.105	-0.107	-0.109	-0.111	-0.061	-0.062	-0.062	-0.066	-0.068	-0.021	-0.022	-0.024	-0.027	-0.030
	100.0	-0.106	-0.107	-0.109	-0.112	-0.114	-0.059	-0.060	-0.061	-0.064	-0.067	-0.021	-0.022	-0.023	-0.026	-0.027
η = 24	300.0	-0.087	-0.086	-0.087	-0.086	-0.089	-0.044	-0.044	-0.045	-0.050	-0.051	-0.010	-0.014	-0.017	-0.017	-0.020
	250.0	-0.086	-0.088	-0.086	-0.087	-0.088	-0.042	-0.044	-0.045	-0.047	-0.051	-0.012	-0.014	-0.015	-0.016	-0.020
	200.0	-0.086	-0.086	-0.086	-0.087	-0.088	-0.042	-0.043	-0.045	-0.047	-0.053	-0.013	-0.014	-0.015	-0.015	-0.020
	150.0	-0.087	-0.086	-0.086	-0.087	-0.087	-0.043	-0.043	-0.045	-0.048	-0.052	-0.011	-0.014	-0.014	-0.017	-0.016
	100.0	-0.087	-0.085	-0.086	-0.086	-0.086	-0.041	-0.042	-0.044	-0.048	-0.051	-0.010	-0.012	-0.012	-0.015	-0.016
η = 32	300.0	-0.068	-0.071	-0.071	-0.070	-0.072	-0.034	-0.033	-0.035	-0.036	-0.043	-0.007	-0.010	-0.009	-0.011	-0.014
	250.0	-0.067	-0.066	-0.068	-0.073	-0.073	-0.031	-0.033	-0.034	-0.038	-0.042	-0.008	-0.009	-0.008	-0.011	-0.015
	200.0	-0.069	-0.069	-0.071	-0.069	-0.073	-0.031	-0.033	-0.036	-0.038	-0.044	-0.008	-0.010	-0.010	-0.012	-0.014
	150.0	-0.071	-0.070	-0.070	-0.071	-0.072	-0.032	-0.032	-0.036	-0.038	-0.043	-0.008	-0.008	-0.010	-0.013	-0.013
	100.0	-0.069	-0.068	-0.069	-0.070	-0.072	-0.032	-0.033	-0.034	-0.037	-0.041	-0.005	-0.007	-0.006	-0.005	-0.010

(a) Average velocity of infiltrating macrophages (cell widths per hour).

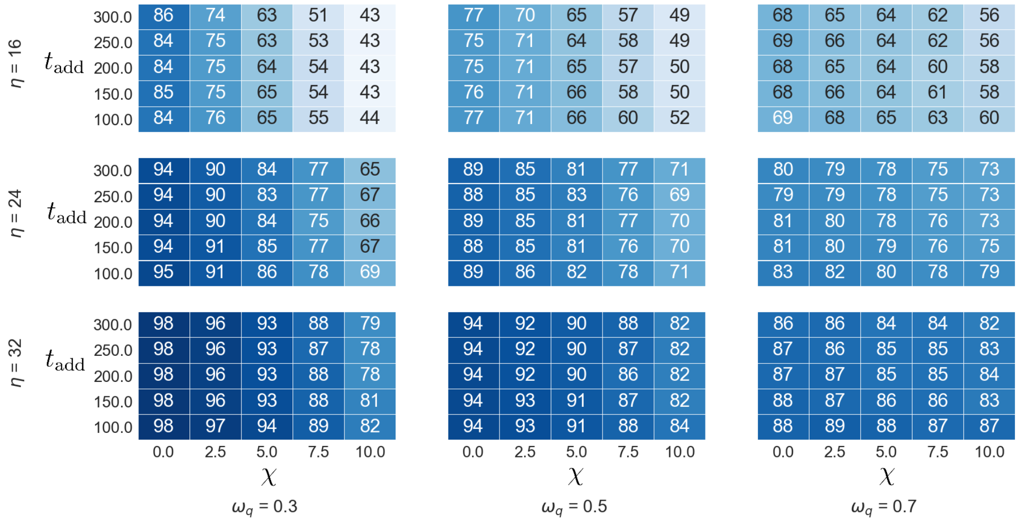

η	t_{add}	ω_q = 0.3					ω_q = 0.5					ω_q = 0.7				
	χ	0.0	2.5	5.0	7.5	10.0	0.0	2.5	5.0	7.5	10.0	0.0	2.5	5.0	7.5	10.0
η = 16	300.0	86	74	63	51	43	77	70	65	57	49	68	65	64	62	56
	250.0	84	75	63	53	43	75	71	64	58	49	69	66	64	62	56
	200.0	84	75	64	54	43	75	71	65	57	50	68	65	64	60	58
	150.0	85	75	65	54	43	76	71	66	58	50	68	66	64	61	58
	100.0	84	76	65	55	44	77	71	66	60	52	69	68	65	63	60
η = 24	300.0	94	90	84	77	65	89	85	81	77	71	80	79	78	75	73
	250.0	94	90	83	77	67	88	85	83	76	69	79	79	78	75	73
	200.0	94	90	84	75	66	89	85	81	77	70	81	80	78	76	73
	150.0	94	91	85	77	67	88	85	81	76	70	81	80	79	76	75
	100.0	95	91	86	78	69	89	86	82	78	71	83	82	80	78	79
η = 32	300.0	98	96	93	88	79	94	92	90	88	82	86	86	84	84	82
	250.0	98	96	93	87	78	94	92	90	87	82	87	86	85	85	83
	200.0	98	96	93	88	78	94	92	90	86	82	87	87	85	85	84
	150.0	98	96	93	88	81	94	93	91	87	82	88	87	86	86	83
	100.0	98	97	94	89	82	94	93	91	88	84	88	89	88	87	87

(b) Average waiting time of infiltrating macrophages (hours).

Figure 5.9: Effect of varying η, ω_{q}, χ and t_{add} on the average velocity and waiting time of infiltrating macrophages. For this parameter sweep, we fix $\omega_{\mathrm{h}} = 0.3$.

change by approximately 10 hours. In contrast, the velocity change caused by varying χ by the same amount is approximately 4 times less than varying η or ω_{q}.

These results indicate that although chemotaxis is associated with a slight increase in the average velocity of macrophages which have crossed the proliferative rim, it has a stronger effect in helping macrophages to cross the proliferative rim. This interpretation explains why the proportion of macrophages reaching the necrotic core shown in Figure 5.3 is much higher for high χ, even though the velocity profiles shown in Figures 5.4, 5.6 and 5.7 differ by only around 0.01 cell widths per hour. The speed of the advective flow from the proliferative rim to the necrotic core is independent of χ, but increasing χ allows macrophages to reach this advective zone faster by helping them cross the proliferative rim.

5.4 Effect of chemotaxis on macrophage distribution

In Section 5.3, we concluded that increasing χ causes increased macrophage infiltration by reducing the time that they spend in the proliferative rim. This allows them to take advantage of the advective flow from the proliferative rim towards the necrotic core. In this Section, we reinforce this interpretation of the role of chemotaxis by considering how the shape of the macrophage distribution changes over time in the presence or absence of chemotaxis.

In Figure 5.10 we present two example sets of histograms which show the distribution of infiltrating macrophages as a function of the radial distance from the spheroid centre. These histograms were obtained by averaging over 20 realisations each of simulations conducted with the parameter set used in Figures 5.2 - 5.4 and with either $\chi = 0$ (Figure 5.10a) or $\chi = 10$ (Figure 5.10b). As time elapses, both distributions travel radially inwards with peak densities which remain close to the spheroid boundary. The profiles disperse over time, as observed in Chapter 4. Even-

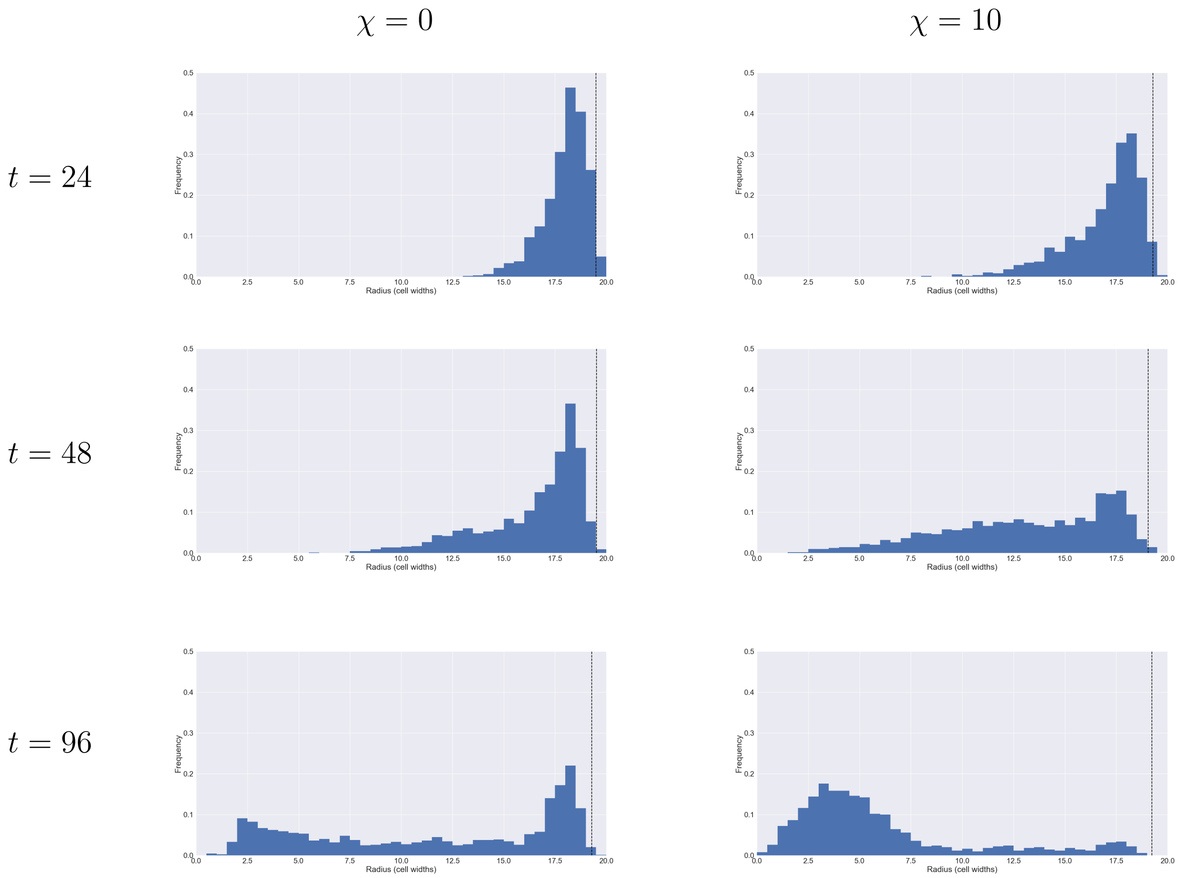

(a) (b)

Figure 5.10: Histograms for low chemotaxis (left: $\chi = 0$) and high chemotaxis (right: χ=10). The black dashed line shows the spheroid boundary.

tually, a second peak forms close to the spheroid centre. Our results from Chapter 4 indicate that this behaviour is caused by the movement of individual macrophages leaving the spheroid boundary and eventually accumulating at the spheroid centre, where they accumulate due to the absence of macrophage death in our model.

In this Section, we use distributions like those pictured in Figure 5.10 to approximate our discrete model, and describe how the shapes of these waves are affected by chemotaxis. For the examples shown, we hold all parameters apart from χ constant ($\eta = 16$, $\omega_\mathrm{h} = 0.3$, $\omega_\mathrm{q} = 0.5$, $\tau_\rho = 8$, $t_\mathrm{add} = 300$) and consider the effect of varying χ on three properties, which we define in Section 5.4.1: the velocity of the wavefront of the distribution, the interquartile range of the distribution and the wave skewness. Our simulations show that the peak of the wave does not move continuously, but instead remains stationary near the spheroid boundary. The magnitude of this peak decreases as individual cells move towards the spheroid centre, and once enough cells have accumulated there the mode of the distribution shifts to the spheroid centre (results not shown). Other parameter sets which we have examined exhibit the same behaviour.

5.4.1 Statistics which describe macrophage distribution shape

In this Section we describe the shape of the macrophage distribution using three summary statistics, which we describe below: the wavefront velocity, the interquartile range and the skewness.

Wavefront velocity

The waves formed by the evolution of the macrophage distribution over time have a leading edge which consists of a small number of macrophages. We define the location of the wavefront of the distribution as the smallest radial distance from the spheroid centre which has non-zero macrophage density. It is likely that for many distributions this wavefront consists of a single macrophage transiting from the proliferative rim to the necrotic core.

Interquartile range

The interquartile range is a commonly used measure of dispersion which is useful when data is not normally distributed (Upton and Cook, 1996). For our distribution, the interquartile range is defined as the distance in cell widths between the 25th and 75th percentile of the macrophage distribution. A low interquartile range indicates that most macrophages are clustered, while a high interquartile range can be interpreted as dispersion of the macrophage distribution.

Skewness

Figure 5.11, reproduced from (Doane and Seward, 2011), shows example distributions with negative skew (left), no skew (middle) or positive skew. There are many ways to define the skewness of a distribution (Doane and Seward, 2011; Upton and Cook, 1996).

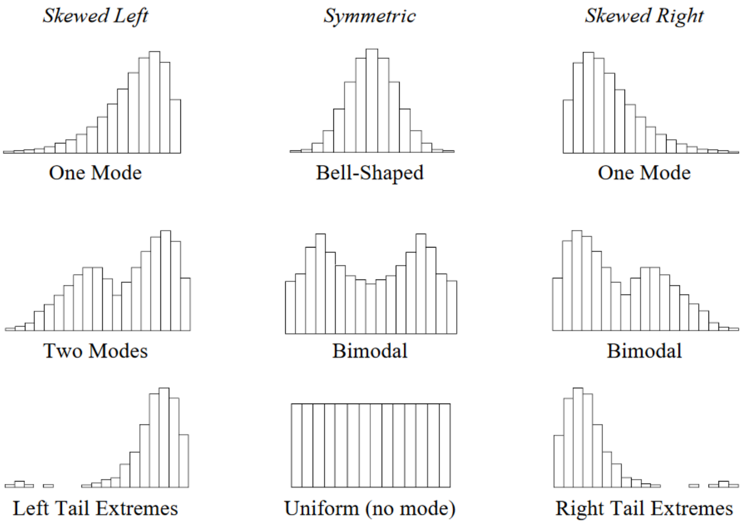

Figure 5.11: Representative histograms demonstrating skewed distributions, reproduced from Figure 2 in (Doane and Seward, 2011). Left: Distributions with negative skew. Middle: Unskewed distributions. Right: Distributions with positive skew.

Pearson's first skewness coefficient, defined as

$$\frac{\text{mean} - \text{mode}}{\text{standard deviation}} \tag{5.5}$$

is perhaps the most commonly used (Upton and Cook, 1996), but it does not work well for bimodal data. From Figure 5.10, we expect the mode of the macrophage distribution to remain near the spheroid boundary until the distribution becomes temporarily bimodal when the number of macrophages at the centre matches the number remaining in the proliferative rim. At this point, we expect the mode to abruptly switch to the spheroid centre. This makes Pearson's first skewness coefficient discontinuous over time, and it is dominated by the distance of the wavefront from the spheroid edge rather than accurately capturing skewness for the observed distributions (results not shown). Instead, we use Pearson's second skewness coefficient:

$$\frac{3(\text{mean} - \text{median})}{\text{standard deviation}} \tag{5.6}$$

as it is independent of the distribution mode. The median of our distribution moves continuously over time, so we expect the second skewness coefficient to be continuous with time. Pearson's second skewness coefficient does not rely on knowledge of the form of the underlying distribution. A skewness of 0 indicates that a distribution is symmetric.

5.4.2 Effect of chemotaxis on wave velocity

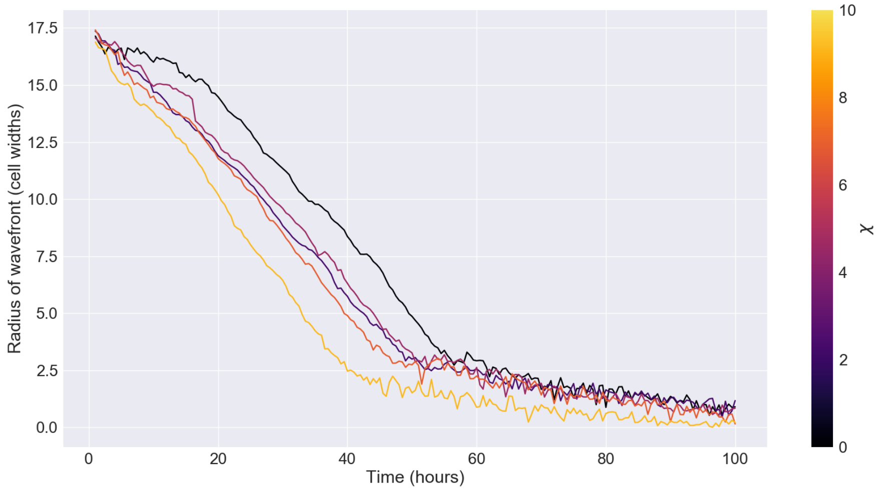

Figure 5.12: Effect of increasing χ on the radius of the distribution front, averaged over 15 realisations for each χ. The wavefront velocity is similar for different values of χ, but moves towards the spheroid centre at earlier times for larger χ.

Figure 5.12 shows the radial distance of the wavefront from the spheroid centre over the 100 hours after the macrophages are added to the simulation. We show the radial distance of the wavefront from the spheroid centroid averaged over 15 realisations of each parameter set, for 5 different values of χ. Figure 5.12 reinforces the conclusions from Section 5.3 regarding the effect of chemotaxis on the average macrophage velocity, as the trajectories follow approximately parallel courses. This indicates that varying χ has little effect on the velocity at which the wavefront is moving. The wavefronts follow similar trajectories to the bead trajectories identified in Figure 4.11 and Figure 5.8, which we interpret as suggesting that the wavefront trajectory is essentially the trajectory of the macrophage with the shortest waiting time. The wavefront trajectories shown in Figure 5.12 are therefore consistent with the conclusions reached in Section 5.3, as they indicate that increased chemotaxis reduces macrophage waiting time but has a relatively small effect on infiltration velocity.

5.4.3 Effect of chemotaxis on wave interquartile range

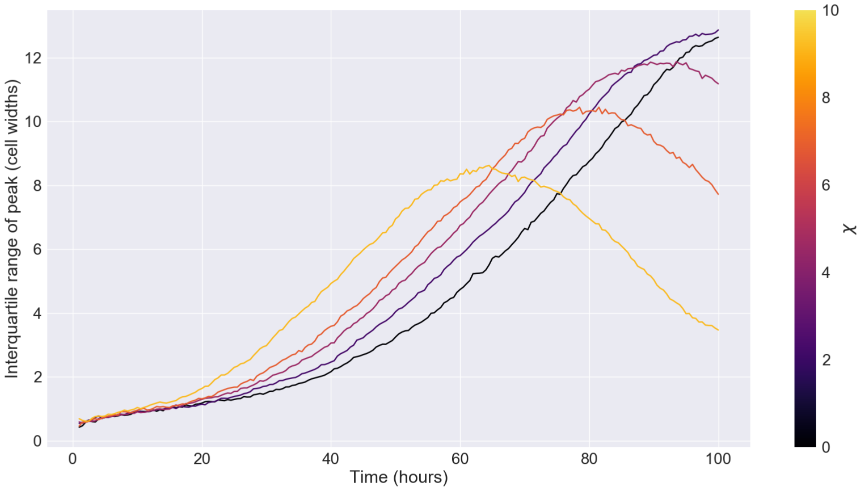

Figure 5.13: Effect of increasing χ on the interquartile range of the macrophage distribution. Distributions become more dispersed over time until reaching a maximum. After this time, the macrophage distribution becomes more clustered as macrophages accumulate in the spheroid centre.

Figure 5.13 shows the interquartile range of the wave over time, which we use to quantify how the spatial dispersion of the distribution changes. For each value of χ the interquartile range initially grows, showing dispersal of the wave. Higher chemotaxis causes a faster increase in the interquartile range, which supports the interpretation that chemotaxis causes macrophages to start moving towards the spheroid centre at earlier times. For $\chi \geq 5$ the interquartile range reaches a maximum and begins to reduce, with the height of this maximum decreasing as macrophages become more sensitive to chemotaxis. The curves for $\chi < 5$ exhibit signs of the same behaviour, but do not reach their maximum value before the end of the simulation.

We interpret the maxima of the interquartile ranges as marking the point at which macrophages are maximally dispersed across the spheroid radius. At first, the interquartile range increases as macrophages leave the proliferative rim and move away from the peak at the spheroid edge. Once a critical mass of macrophages have accumulated in the spheroid centre, the interquartile range begins to reduce as macrophages which leave the spheroid rim are moving towards the mass of macrophages rather than away from it. Simulations with high chemotaxis sensitivity reach this point faster as the waiting time of macrophages in the spheroid rim is shorter. This also explains why the interquartile range is lower in simulations with higher chemotaxis sensitivity, as macrophages leave the spheroid rim in faster succession and are thus less dispersed across the domain when they begin clustering at the spheroid core.

5.4.4 Effect of chemotaxis on wave skewness

In Figure 5.14 we show the effect of varying χ of the skewness of the macrophage distribution. Figure 5.14 shows that while macrophage distributions are initially approximately symmetric, with skewness close to 0, the skewness eventually reaches a minimum. This minimum can be interpreted as representing the shape of the distribution when the median macrophage transits from the spheroid rim to the core. From this minimum, the skewness gradually increases until the distribution is once again symmetric, with the peak at the spheroid boundary approximately balanced by the peak at the spheroid centre. The continuing movement of macrophages towards

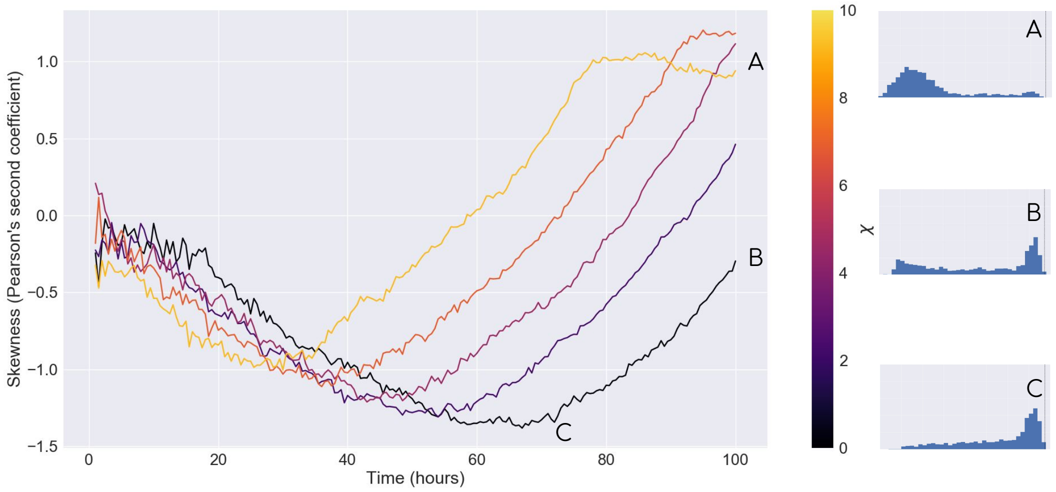

Figure 5.14: Effect of increasing χ on the skewness of the moving distribution. The points marked A, B and C correspond to histograms with positive skew, small skew, and negative skew respectively.

the necrotic core causes the skewness to increase until no macrophages are left in the proliferative rim and the increase in skewness stops. This behaviour happens faster when χ is high, and so the size of the skewness at the stationary points is lower in these cases. This indicates that although the increase in macrophage infiltration velocity and decrease in waiting time cause the distribution to become skewed faster for high χ than in simulations with lower sensitivity to chemotaxis, simulations without chemotaxis exhibit more extreme skewness. In simulations with low waiting times, the peak at the spheroid boundary shrinks rapidly and the dispersal of the distribution has less impact on the overall asymmetry of the distribution.

In Figure 5.15 we plot the skewness of the distribution as a function of both the mean (Figure 5.15a) and the median (Figure 5.15b) macrophage distance from the spheroid centre. In this way, we can compare the infiltration distributions when the distances of the average macrophage from the centre are aligned, rather than at the same timepoints. Figure 5.15 supports our interpretation of Figure 5.14, showing that the transition from negative skewness to positive skewness is due to the average macrophage density switching from being in the proliferative rim to

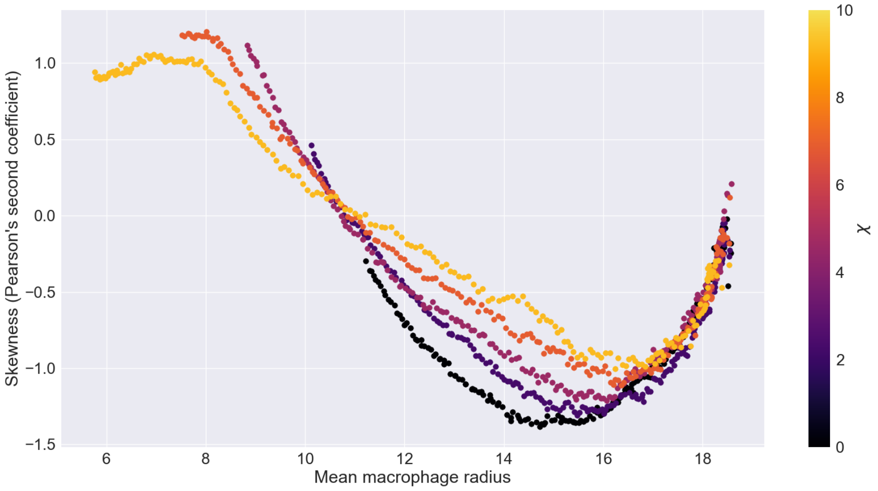

(a) Skewness as a function of mean wavefront radius.

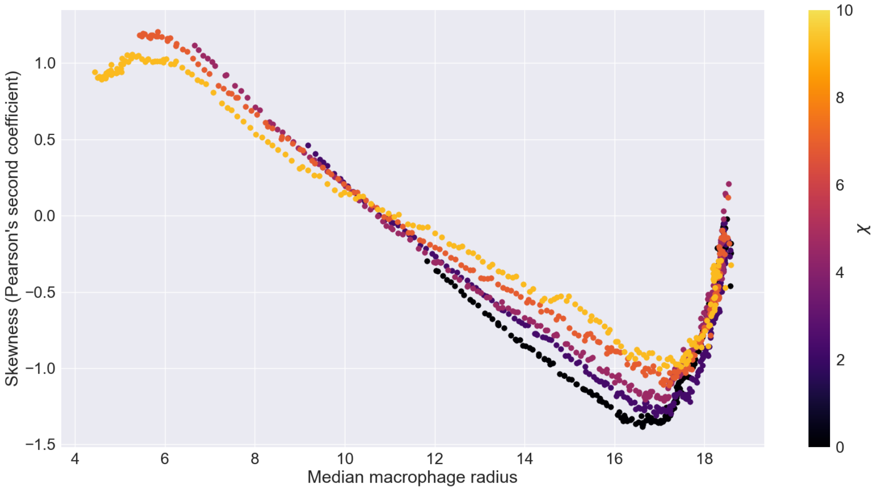

(b) Skewness as a function of median wavefront radius

Figure 5.15: Skewness as a function the average macrophage in the distribution, for both the mean (Figure 5.15a) and median (Figure 5.15b) macrophage.

being in the spheroid centre. As the location of the median/mean of the distribution moves towards the spheroid centre, the profile switches from an increasingly negative skewness to an increasingly positive one. The movement of individual macrophages causes a jump in the location of the wave peak from the boundary to the core, with the transition being triggered by the movement of the "average" macrophage location from the outside of the spheroid to the inside.

5.4.5 Conclusions: effect of chemotaxis on macrophage infiltration

The results presented in Sections 5.3 and 5.4 reach a complementary conclusion based on the velocities and waiting times of individual macrophages and on the shape of the macrophage distribution. This conclusion is that increasing the sensitivity of macrophages to CSF-1 results in increased macrophage infiltration. From our agent-based model, we conclude that increased infiltration is caused by individual macrophages passing through the proliferative rim more easily in the presence of chemotaxis, rather than being due to macrophages moving faster once they have crossed the proliferative rim. Chemotaxis provides a direction to macrophage movement within the proliferative rim, and in the absence of this direction macrophage movement is dominated by the interaction with proliferating cells. Cell proliferation can push macrophages in the proliferative rim either towards the spheroid centre or the spheroid edge, but increased chemotaxis causes macrophages to resist the forces pushing them away from the spheroid core.

Kelly et al. (2002) were able to reproduce in a PDE model experimental results by Leek (1999) which show increased infiltration of macrophages into spheroids that produce chemoattractants. These results match the observations from our agent-based model: in the presence of chemotaxis, more macrophages will have infiltrated a spheroid after a fixed time, and those macrophages which have not infiltrated the spheroid will remain localised near the spheroid edge. One discrepancy between

(Kelly et al., 2002) and (Leek, 1999) is that in the absence of chemotaxis, the peak of the macrophage distribution is observed to be inside the spheroid but close to the edge, while (Kelly et al., 2002) suggests that the peak should be outside the spheroid. The authors conclude from this that a small amount of chemoattractant is likely to be produced in the experiments. However, our simulations indicate that this behaviour arises (see, e.g., Figure 5.10b) *in silico*. Macrophage density in the PDE model of (Kelly et al., 2002) moves only via chemotaxis or random motion. Using an agent-based framework allows us to identify another influence on macrophage motion, namely the intercellular forces produced by the proliferation of tumour cells at the spheroid rim. In the absence of chemotaxis, the forces generated by this proliferation are sufficient to trap macrophages within the tumour, but close to the boundary, as observed by Leek (1999). These forces are still present in spheroids which produce chemoattractant, but increased chemotaxis allows macrophages to escape from this region more quickly than they would otherwise be able to. Hence for spheroids with the same composition, waiting time is heavily influenced by chemotaxis (Figure 5.9b) while the velocity with which infiltrating macrophages move towards the spheroid centre is only slightly increased (Figure 5.9a). These results have implications on what we can learn about *in vitro* spheroids based on infiltration studies, which we discuss in Section 5.6.

5.5 Infiltration in *in vivo* tumour geometries

In this Section, we extend the agent-based model of tumour spheroid infiltration described in Chapters 4 and 5 by combining it with the image analysis pipeline described in Chapters 2 and 3 to describe macrophage infiltration *in vivo*. We present preliminary results based on an example drawn from the Pugh lab dataset of head and neck tumour images described in Section 2.2.2.

The multiple stains for oxygen in the Pugh lab dataset have parallels with the

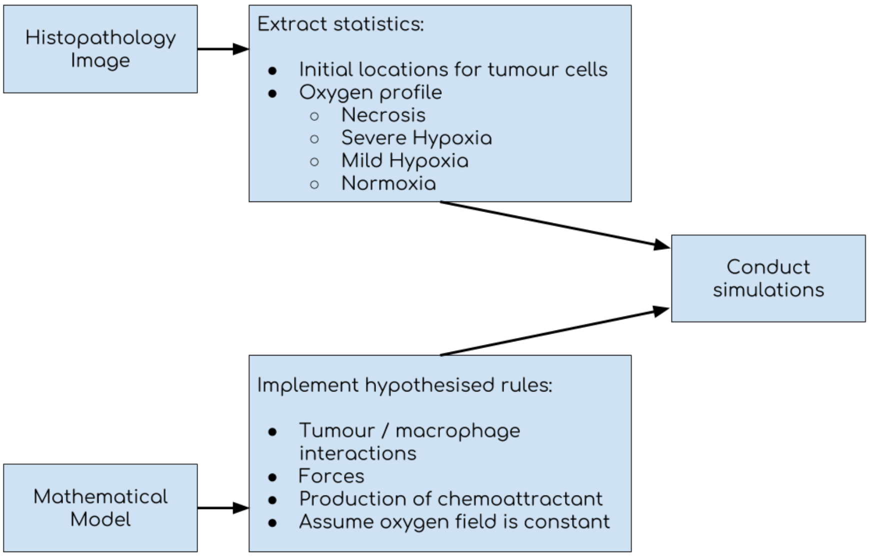

Figure 5.16: Overview of the process by which our image analysis pipeline can be used to prescribe initial conditions for an agent-based model.

oxygen environments used in our agent-based modelling of tumour spheroids. In our spheroid simulations we distinguish between normoxic, quiescent, hypoxic and necrotic tumour regions. As CAIX indicates less severe hypoxia than pimonidazole (Rademakers et al., 2011; Varia et al., 1998), the borders of these stains can be used to define contours of constant oxygen concentration which we can map onto our model compartments.

The contours of the oxygen environments can be obtained directly from our image analysis pipeline. In Figure 5.17 we show the composite mask produced by our image analysis pipeline in which PanCK, CAIX and CD68 stains have been identified and aligned for a 2 mm × 2 mm region of a head and neck tumour. We have chosen this image as it contains a wide range of tumour environments including stroma, normoxia, quiescence, hypoxia and necrosis. We use our image analysis pipeline to estimate the oxygen concentration across the domain, and to prescribe initial conditions for an agent-based simulation. An overview of this process is shown in Figure 5.16. In this Section, we fix the oxygen concentration along the borders of each of these stains (and the stain for pimonidazole, not pictured) to correspond with

Stain	**CAIX$^+$**	**CAIX$^-$**
Pimonidazole$^+$	$0 \leq \omega \leq \omega_\mathrm{h}$	$\omega_\mathrm{h} < \omega \leq \omega_\mathrm{q}$
Pimonidazole$^-$	N/A	$\omega_\mathrm{q} < \omega \leq 1$

Table 5.1: Choices of oxygen threshold based on histological stains for oxygen

ω_h and ω_q (the choices made for this example are shown in Table 5.1.) We assign oxygen concentrations between the stain boundaries by interpolating between 0, ω_h, ω_q and 1 as appropriate for the regions defined in Table 5.1. In necrotic regions we set $\omega = 0$.

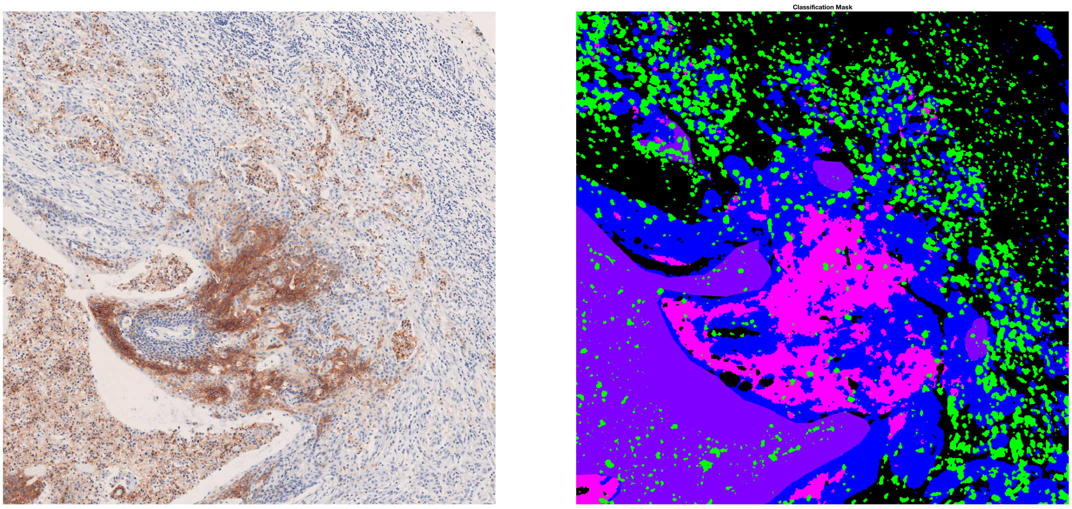

(a) CAIX stained image showing the geometry of the region which we wish to model.

(b) Composite mask produced by our image analysis pipeline.

Figure 5.17: A 2 mm × 2 mm region of interest from one of the head and neck tumour samples in the Pugh lab dataset used to provide a realistic geometry for agent-based modelling.

To initialise an agent-based simulation with this information we downsample the oxygen classification masks in Figure 5.17 so that one pixel corresponds to a cell width squared in our models. For a 2 mm × 2 mm region and an assumed cell width of 20 μm, this results in a composite mask which is 100 × 100 pixels. We assign each one of these new pixels the oxygen environment which dominates the relevant $20\,\mu m \times 20\,\mu m$ region of the full composite mask. We then assign oxygen concentrations to the nodes of a regular triangular finite element mesh which spans the domain. Since in this example we fix the oxygen concentration rather than update it using the PDE, we use the mesh only for the linear interpolation of oxygen values for cells which are in between nodes of the oxygen mesh. This oxygen distribution can now be used as the backdrop for agent-based modelling of the *in vivo* geometry.

5.5.1 Adaptations to the mathematical model

Using our agent-based model to simulate *in vivo* geometry requires some modifications of the mathematical model. In particular, the oxygen concentration throughout the domain is fixed and therefore not updated via the PDE described in Equation (4.6). The cell cycle rules describing tumour cell proliferation and death described in Section 4.3.3 still determine the behaviour of cells, but cells no longer act as oxygen sinks. We also add stromal cells to the model to occupy space, and add boundary conditions to confine cells to the domain. We describe these modifications in the rest of this Section.

Our decision to keep the oxygen concentration profile constant is due to practical limitations in the data available. We are restricted by the lack of knowledge of the position of blood vessels throughout these tumours. As such, we cannot define the location of oxygen sources in the tissue sample. A more sophisticated approach could be to use the oxygen boundaries to predict where the sources of the oxygen should lie in the domain. It must also be noted that histological images provide a 2D slice through the tumour, while *in vivo* the 3D tumour geometry is more complex. Mathematical models of vascular remodelling can be highly intricate, and accurately fitting a vascular model to either 2D or 3D imaging data is a challenging task (Perfahl et al., 2017; Kim et al., 2012). While integrating our model with this type of vascular data would make it more biologically realistic, it is beyond the scope of both our data and our current model. We therefore assume that the vasculature is sufficiently stable on the timescale of the simulation that oxygen sources are constant. We further assume that the movement of tumour cells does not significantly alter this distribution during the simulation. This assumption is partially justified through observation of spheroid models, in which individual cells move through the spheroid but the bulk behaviour at steady state means that the oxygen distribution in the domain remains steady.

Another modification to our model is the inclusion of stromal cells based on the histological data. We introduce stromal cells in order to keep the tumour boundary approximately steady, justifying the assumption that the oxygen distribution in the domain is fixed. These also account for one of the key differences between tumour spheroids and *in vivo* tumours, which is the presence of a surrounding tissue matrix which resists tumour growth. All PanCK$^-$ regions of the domain are therefore populated by stromal cells. The presence of stromal cells removes the need for the surface tension force present in the spheroid models, which is therefore removed from these simulations. Stromal cells in our model are identical to tumour cells in almost all respects: they progress through the cell cycle in exactly the same way as tumour cells, and can proliferate or die through the same mechanisms. The only difference between the cell populations in our model is the proliferation rate: tumour cells are assumed to have a proliferative advantage, which we model by doubling the average cell cycle time for stromal cells in comparison to that of tumour cells.

We must also apply boundary conditions for the cells in our model to prevent cells leaving the domain. We confine cells to the $2\,\mathrm{mm} \times 2\,\mathrm{mm}$ on which the oxygen concentration is defined by placing cells which attempt to leave back on the boundary of the domain at the nearest point. While this may introduce edge effects, we note that the vast majority of cells in our simulations are not close to the region boundary and therefore do not account for edge effects in this preliminary example.

We initialise the model by placing cells in contact on a square lattice separated by $20\,\mu$m. Macrophages are seeded randomly throughout the stroma at a density of 5% of that of the stromal cells. While it would be possible to initialise the model using macrophage locations identified from Figure 5.17b, we choose to initialise them in the stroma to compare the rates of infiltration in simulations with no chemotaxis and high chemotaxis, mirroring the spheroid simulations conducted earlier in this Chapter. This, along with other potential applications of this combination of imaging and modelling, is discussed in Section 5.6.

5.5.2 Macrophage response to chemotaxis: *in vivo* simulations

In this Section we present preliminary results regarding the effect of chemotaxis in *in vivo* tumour geometries. More work is required to establish detailed results, and these results are intended as a preliminary demonstration of how our imaging pipeline and agent-based modelling can be combined.

Figure 5.18 shows a typical snapshot of a simulation initialised using the histological images shown in Figure 5.17. The domain is 2 mm × 2 mm, but necrosis has caused the lower-left portion of the tumour to die away, leaving a void. This effect is observed both *in vivo* and *in vitro* (Ma et al., 2012), and can be seen in Figure 5.17a. Proliferation of tumour and stromal cells causes the tumour boundary to shift and causes tumour and stromal cells along the boundary to mix, blurring the boundary. Additional snapshots of this simulation can be seen in Appendix E, along with simulations based on different tumour geometries.

The comparison presented in this Section is between two simulations conducted in this *in vivo* geometry. Both simulations use the same geometry and initial conditions, and differ only in the chemotaxis sensitivity coefficient, χ. In one simulation $\chi = 0$ ("no chemotaxis"), and in the other $\chi = 10$ ("with chemotaxis").

As we do not initialise macrophages along the tumour boundary, we cannot compare the distributions in terms of distance from the edge of the tumour in the same way as we have previously in our spheroid models. As such, our analysis is based on the velocities of the macrophages. In Figure 5.19 we present the average

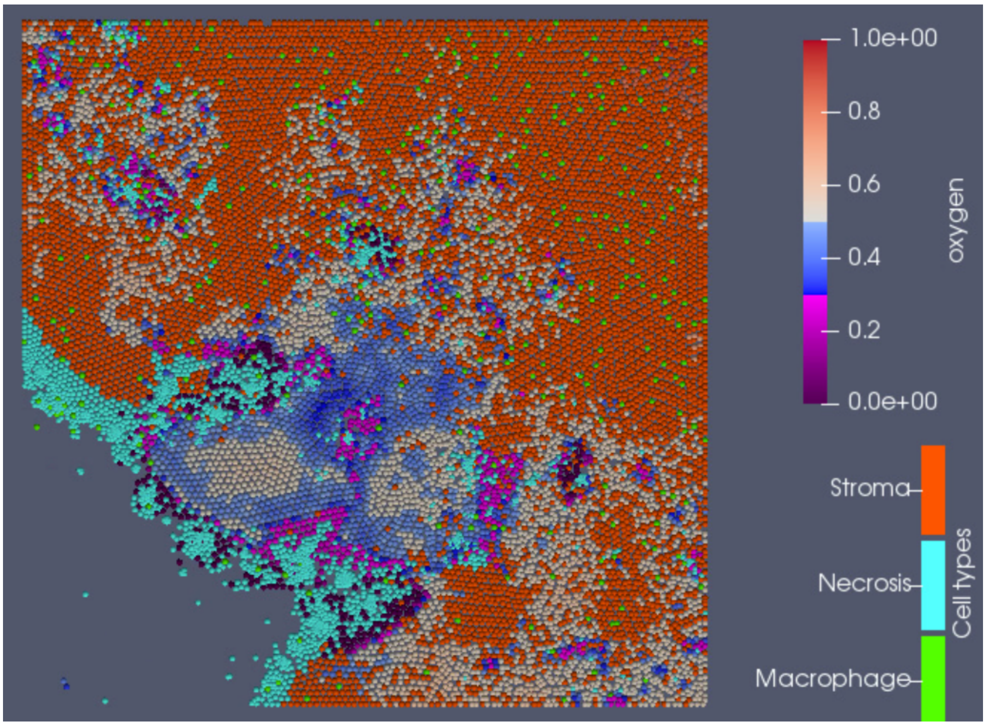

Figure 5.18: Snapshot of agent-based model initalised using the geometry in Figure 5.17. The oxygen concentration at each point is defined by the detected stains shown in Figure 5.17b: Necrotic / pimonidazole$^+$ regions are defined to be hypoxic, CAIX$^+$ regions are quiescent, PanCK$^+$ regions are populated with proliferative tumour cells and PanCK$^-$ regions are populated with stromal cells. Stroma cells are shown in orange, regardless of their oxygen concentration, but their progress through the cell cycle and probability of death are dependent on oxygen concentration in the same way as for the tumour cells. All other cells are coloured according to the same rules as previous simulations. The image shown is after $t = 50$ hours. Snapshots from other timepoints and for other tumour geometries can be seen in Appendix E.

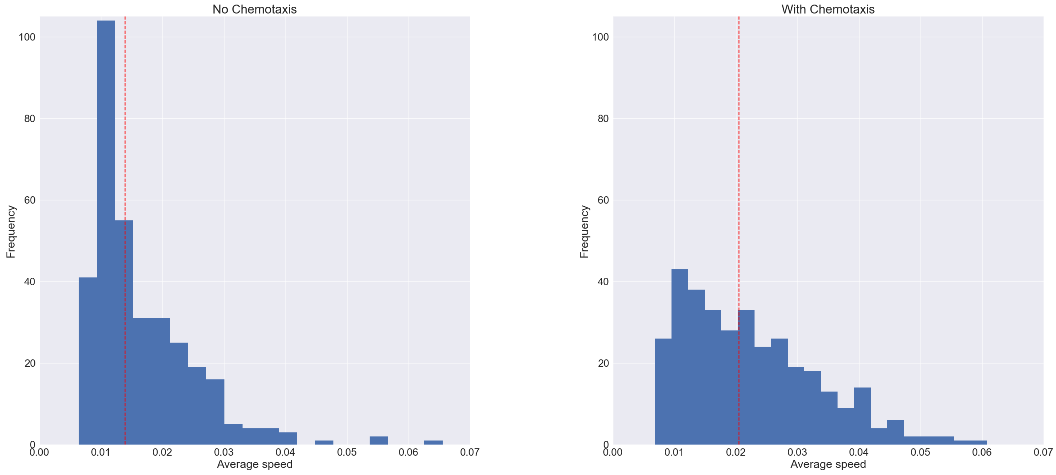

(a) Histogram showing the mean speed of each macrophage in the absence of chemotaxis ($\chi = 0$).

(b) Histogram showing the mean speed of each macrophage in the presence of strong chemotaxis ($\chi = 10$).

Figure 5.19: Histograms showing the mean speed of macrophages in realistic geometry simulations, in the absence and presence of chemotaxis. The dashed red line shows the median average speed of macrophages across the simulation. In the absence of chemotaxis, macrophages appear to move much more slowly than when it is present.

speed of each macrophage across both simulations ($n \approx 580$). The peak of both distributions occurs at low speeds of approximately 0.01 cell widths per hour. However, the distribution of the average speeds of macrophages which move according to chemotaxis has a much larger tail, with many more macrophages moving at higher speeds than in the simulation without macrophages. The dashed red line indicates the median macrophage speed in each simulation; the median speed of macrophages in the presence of chemotaxis is much higher than that in the simulation without chemotaxis.

In Figures 5.20 and 5.21 we show how the average speed of macrophages changes at different locations in the domain. Figure 5.20 visualises the average speed of a macrophage placed at a given point in the domain, and Figure 5.21 shows the trajectories of macrophages starting at different places in the domain. In both simulations, we see that the regions of high velocity coincide with the edges of

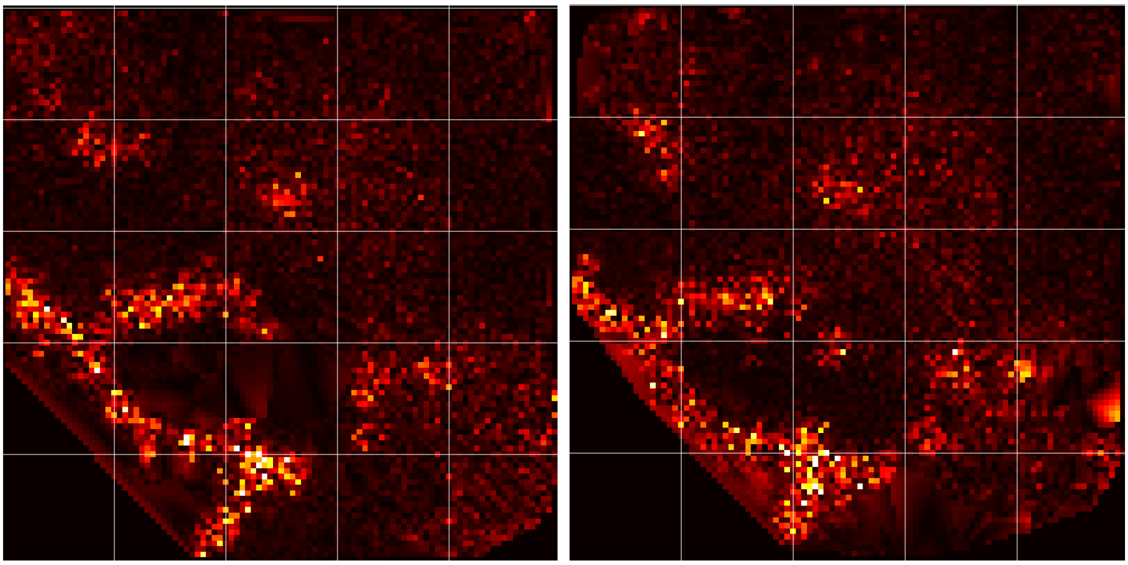

(a) $\chi = 0$ (b) $\chi = 10$

Figure 5.20: Heatmaps showing the average speed of cells throughout the simulation. Cells close to necrotic regions move much faster than those in the stroma.

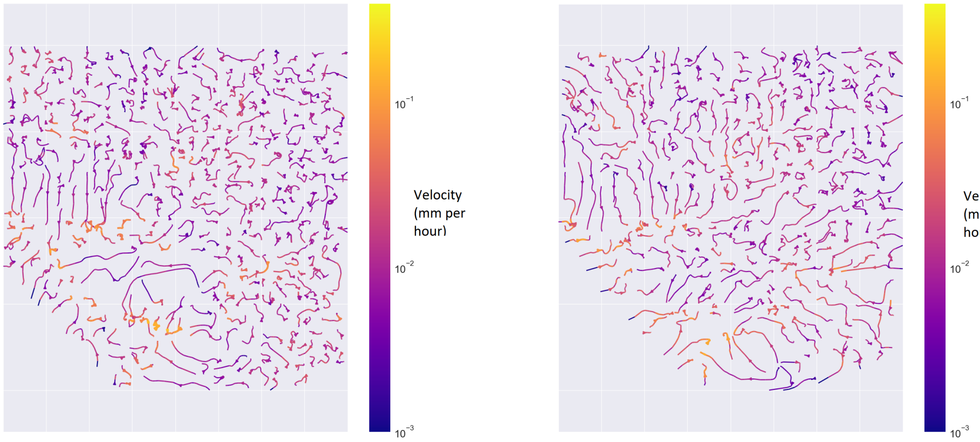

(a) χ= 0 (b) χ= 10

Figure 5.21: Particle paths throughout the simulation. Macrophages move faster close to the necrotic regions, both in the presence and absence of chemotaxis.

necrotic regions. The lowest velocities in the tissue occur in the stromal regions. These regions are the most distant from necrotic regions, and the stromal cells proliferate slowly. Results from Chapters 4 and 5 indicate that large values of η and large proliferative regions are both associated with slower movement. This is because movement in these areas is dominated by random cell proliferation rather than by the advective flux towards necrotic regions, and hence macrophages in a large stromal region with a slow proliferation rate will move slowly.

Although Figure 5.19 indicates that macrophages move faster in the simulation with $\chi = 10$, Figure 5.20 shows that the movement of macrophages at any given point in the domain appears to happen at comparable speeds between the simulations. This indicates that macrophages in simulations with chemotaxis must spend more time in the “faster” parts of the tissue than in the “slower” parts. This reinforces the results from our spheroid simulations, which indicate that chemotaxis causes macrophages to move through the proliferative rim faster, and suggests that chemotaxis plays a key role in drawing macrophages from the stroma towards necrotic regions.

Figure 5.22 supports the interpretation that chemotaxis draws macrophages from the stroma towards necrosis. It shows the time taken for each macrophage to move from their randomly seeded positions in the stroma to the boundary of the quiescent region. The median time taken for a macrophage to reach the quiescent region is 51 hours in the presence of chemotaxis, but 112 hours without it.

Our model suggests that although chemotaxis increases the speed of infiltrating macrophages, its main role is to facilitate the movement of macrophages towards regions where passive migration is strong. Passive movement in our model acts like a conveyor belt, and the speed of the conveyor is the dominant driver of macrophage movement in the model. The role of chemotaxis is to move macrophages onto this conveyor - the increase in infiltration velocity due to chemotaxis is small in comparison with the impact of the tumour cell proliferation/death balance, but

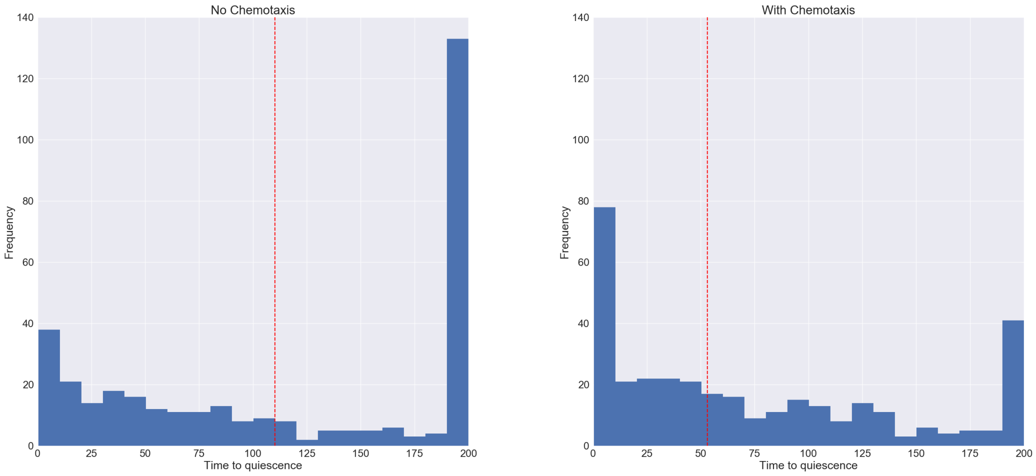

(a) $\chi = 0$

(b) $\chi = 10$

Figure 5.22: Histograms showing the time taken for macrophages in realistic geometry simulations to reach the quiescent (CAIX) boundary, in the absence and presence of chemotaxis. The dashed red line shows the median time taken for macrophages to leave the proliferative zone. Macrophages which remained in the stroma or proliferative parts of the tumour are included in the final bin.

without active recruitment of macrophages by chemotaxis they are less likely to reach the conveyor.

5.6 Discussion

In this Section, we have extended the agent-based model presented in Chapter 4 to include the production of CSF-1 by hypoxic tumour cells, and to include macrophage chemotaxis along the resulting CSF-1 gradient. In our model, chemotaxis is associated with only a small change in the velocity at which macrophages travel from the proliferative rim to the necrotic core. The influence of the advective flow discussed in Chapter 4 is several times stronger, and is the key determinant of macrophage infiltration velocity. We find that the main impact of adding chemotaxis is to increase the speed at which macrophages can cross the proliferative rim of the spheroid. Chemotaxis provides a mechanism through which macrophage movement in the proliferative rim can be directed towards the spheroid core, resulting in more macrophages infiltrating into the spheroid centre. Our preliminary work in Section 5.5 indicates that there is an *in vivo* parallel of this *in vitro* observation, where macrophage movement through the tumour stroma and normoxic tumour regions is similar in nature to movement in the proliferative rim of spheroids.

While our model does not include the phenotypic change that CSF-1 induces in macrophages, our results have implications for treatment using CSF-1R inhibitors. Treatments which block CSF-1 receptors on macrophages are intended to bias the population towards an anti-tumour phenotype (Mantovani et al., 2017; Pradel et al., 2016; Edwards V et al., 2018). Our model suggests that desensitising macrophages to CSF-1 would also have the effect of reducing macrophage infiltration into tumours. As a result, while treatment with CSF-1R inhibitors is likely to bias macrophages towards an anti-tumour phenotype, these macrophages are likely to remain at the periphery of the tumour. The main benefit of this treatment is likely to be in depleting the number of pro-tumour macrophages within tumour nests, rather than in causing tumour infiltrating macrophages to attack the tumour from within. Extending our model to include the effect of CSF-1 on macrophage phenotype could help to confirm this hypothesis. Previous mathematical models have considered the use of genetically modified macrophages as a vehicle for expressing drugs or prodrugs in hypoxic regions of tumours (Webb et al., 2007; Byrne et al., 2004; Owen et al., 2004, 2011). Our model suggests that any such treatment which alters the chemotactic response of these macrophages should be considered carefully, as it could inhibit the macrophages from localising in necrotic regions and reduce the effectiveness of these treatments.

Our *in vitro* simulations could also be used to help interrogate the relationship between macrophages and tumour cell lines in infiltration studies such as those conducted by Leek (1999). Our results indicate that the degree of chemokine production by hypoxic cells could be inferred from the proportion of infiltrating macrophages at the periphery of the tumour. Conducting experiments in which microbeads and macrophages are added simultaneously to tumour spheroids would result in two different infiltration profiles. Our results indicate that by comparing the numbers of infiltrating particles or macrophages in the spheroid core, the degree of dispersion of each profile, and the number of microbeads or macrophages at the spheroid boundary, it should be possible to infer whether migration is dominated by active or passive processes. In the absence of these additional experimental results, further work could conduct this analysis *in silico* by attempting to fit the current model to artificial data for macrophage distributions generated using known parameters.

In this Section we have also developed a method through which the output classifications from our image analysis pipeline can be used to provide realistic geometries for our agent-based model. We have demonstrated this by presenting preliminary results in which we use the oxygen thresholds provided by pimonidazole and CAIX staining to fix the oxygen environments in our model. These results support the conclusion drawn from our spheroid modelling that the addition of chemotaxis causes increased macrophage infiltration by drawing macrophages from the stroma towards regions of the tumour in which passive migration is strong.

Although the results from our simulations of *in vivo* data are preliminary, they provide a framework from which future work can extend our model significantly. The work in this Section provides a means through which our model can be simulated in hundreds of different *in vivo* oxygen environments. The limiting factor in this is the time for a realisation of our agent-based simulation to run, which can be several days depending on size of the domain and complexity of the interactions. Conducting these additional simulations, with enough realisations conducted in each geometry to provide statistically significant results, would produce a powerful modelling dataset for further interrogation. Further adaptations could include initialising macrophages according to their observed locations in the histological images, or finding parameter regimes which cause macrophages to adopt these distributions.

Further work should also consider the implications of other roles of macrophage chemotaxis in the tumour microenvironment. In particular, the paracrine loop formed between EGF and CSF-1 is of particular interest in the movement of macrophages away from tumour nests towards the tumour vasculature (Wyckoff et al., 2007; Goswami et al., 2005; Wyckoff et al., 2004). This mechanism is implicated in tumour metastasis (Lewis et al., 2016a; Laoui et al., 2014; Hughes et al., 2015; Kitamura et al., 2015), and has been modelled by Knútsdóttir et al. (2016, 2014) to replicate the results of *in vitro* experiments. Integrating this system with our current work for *in vivo* geometries could provide insight into which factors determine whether macrophages are predominantly perivascular or in the tumour nest. A comprehensive model to describe the role of chemotaxis on macrophage localisation must also include macrophage phenotype, which is often critically affected by chemokines such as CSF-1 (Van Overmeire et al., 2015).

Chapter 6

Discussion and future directions

6.1 Discussion

The automated analysis of histological images is becoming one of the most important problems in histopathology today (Williams et al., 2018b; Bell, 2017). Traditional methods of analysing slides are being superceded by digital techniques to increase efficiency and save costs (Williams et al., 2018b). As digitising whole slide images becomes cheaper and more commonplace, new techniques which enable pathologists to manage the quantity of data available will be required. Current software for digital pathology permits automatic measurement of basic metrics (Bankhead et al., 2017). In this thesis, we have shown that computational techniques can achieve more than simply providing more accurate assessments of currently used statistics. Instead of using qualitative terms like "hot" or "cold" to describe the immune activity within a tumour, automated image processing provides an opportunity to define quantitative metrics. We advocate the use of mathematical techniques for the analysis of histological images.

Analysing the distribution of macrophages within *in vivo* tumours is a problem which highlights the role these techniques can play. In this thesis, we have combined approaches from mathematical modelling, image analysis and spatial statistics to investigate macrophage infiltration into solid tumours. We have used each of these approaches to engage with different facets of this issue.

In Chapter 2 we described a novel pipeline for the analysis of whole slide images. Our pipeline can process high-resolution images to identify cells stained for the expression of an arbitrary range of markers, and new classifiers can be trained and evaluated in minutes by users with no knowledge of programming. We validated the ability of our pipeline to reproducibly detect macrophages in histology images, and demonstrated that it performs well in comparison with trained pathologists. This pipeline is currently being used for research by pathologists in the Nuffield Department of Medicine, and will be made publicly available following the submission of this thesis.

In Chapter 3 we used our pipeline to interrogate image datasets provided by the Nuffield Department of Medicine and by Roche. We applied spatial statistics to macrophage locations obtained from the RITA dataset via our pipeline, and demonstrated that these can be used for the effective comparison of regions of interest between tumours. Spatial statistics which describe the distribution of macrophages in a tumour region can be used to automatically group tumour samples into qualitatively similar groups. In Chapter 3, we also used our pipeline to combine multiple consecutive histological slides to characterise the oxygen environment of tumour associated macrophages. We obtained statistics describing the density of immune cells in different compartments, but also developed a histogram-based approach through which the extent of infiltration by macrophages in different oxygen environments can be visualised. We also used “artificial biopsies” to quantify how accurately summary statistics can be estimated based on a small amount of tumour tissue. We analysed approximately 120 whole slide images from breast and colorectal cancer samples in the RITA dataset described in Section 2.2.1, and 1536 $1.5\,\mathrm{mm} \times 1.5\,\mathrm{mm}$ regions of interest from the head and neck tumour described in Section 2.2.2.

We developed an agent-based model in Chapter 4 to model experiments by (Dorie et al., 1982) on the infiltration of inert microbeads into tumour spheroids, implemented in the Chaste modelling environment (Mirams et al., 2013; Pitt-Francis et al., 2009). This model let us investigate the causes of passive infiltration that occurs due to the advective flow of tumour cells caused by proliferation in the spheroid rim and death in the core. Our model reproduces experimental observations and the results of previous continuum models of this dataset. Our agent-based model provides insight regarding the causes of the observed distributions of microbeads that previous continuum models have been unable to observe, as by tracking the trajectories of individual beads we identified that infiltration occurred in two stages. In the first stage, individual microbeads are trapped near the spheroid boundary by tumour cell proliferation. After a "waiting time" which depends on the balance of proliferation and death in the spheroid, microbeads move radially towards the spheroid centre with an approximately constant velocity. We investigated the effect of varying parameters associated with proliferation and death on the waiting time and infiltration velocity, and showed that the tumour composition plays an important role in determining these properties. We also used our model to reproduce results from (Dorie et al., 1982) describing the infiltration of ^{3}H-labelled tumour cells into spheroids, and showed that the results can be replicated as a result of proliferation of labelled cells alone without adding chemotactic mechanisms required to reproduce this in previous continuum models of this data (Pettet et al., 2001; McElwain and Pettet, 1993). The results in this Thesis are based on approximately 5760 simulations (3600 simulations of microbead and macrophage infiltration and 2160 of ^{3}H-labelled cell infiltration), which were simulated in parallel on compute clusters in the Mathematical Institute, with each simulation taking between approximately 1 hour and 48 hours. Compute time varied due to the number of cells in each simulation, with simulations over long timescales with high proliferation and low death rates taking the longest. While simple spheroid models can be reproduced from examples on the Chaste website (Mirams et al., 2013; Pitt-Francis et al., 2009), we have substantially added to this code base. This code will be added to the publicly available Chaste source code following the submission of this thesis.

In Chapter 5, we extended the model developed in Chapter 4 to include macrophages and their response to chemotactic gradients of CSF-1 produced by hypoxic tumour cells. We found that chemotaxis led to increased macrophage infiltration across 975 tumour spheroid simulations with varying parameter sets, but that the increase in macrophage infiltration velocity was small compared to its magnitude. We found that the effect of increased chemotaxis was to cause macrophages to cross the proliferative rim more quickly, but that once they had crossed this rim they travelled at a speed determined by the advective flow investigated in Chapter 4. We analysed the shape of the wave formed by the macrophage distribution, and found that the presence of chemotaxis caused less distortion in the shape of this moving distribution. We also presented preliminary work on macrophage infiltration into *in vivo* tumour geometries. We used data obtained from our imaging pipeline to define the boundaries of the tumour and the shapes of contours of the oxygen concentration throughout the tissue. We observed that these preliminary simulations of macrophage infiltration supported the conclusions of our simulations of *in vitro* infiltration. They showed that more macrophages reach quiescent parts of the tumour in the presence of chemoattractants, and that increased macrophage infiltration occurs because chemotaxis draws macrophages towards regions of the tumour where the advective flow of tumour cells is strong rather than because chemotaxis directly increases macrophage velocity.

6.2 Future directions

In this Section, we describe ways to develop and extend the work in this thesis from the perspectives of image analysis, spatial statistics and mathematical modelling. We also discuss how these fields can be further integrated to provide additional benefit to pathologists.

6.2.1 Development of image analysis software

The image analysis pipeline which we developed in Chapter 2 has the potential to provide significant value to the pathology community. We aim to develop this pipeline into a standalone application which can be used by non-programmers. By providing a tool which can be easily used by pathologists with minimal training, we aim not only to improve current workflows but also reveal to pathologists the benefits of quantitative rather than qualitative analyses. This is crucial for more complex mathematical and statistical metrics to gain widespread use, as we advocate in Section 6.2.4. This program should have an intuitive graphical user interface, and integrate tools such as image registration and cell identification alongside tools for basic image manipulation and annotation. We will collaborate with pathologists from both Roche and from the Nuffield Department of Medicine to ensure that the software meets the needs of pathologists, rather than providing an interface for unwanted computational tools. Two tools which should be added to our image analysis toolbox to achieve this are automated image registration and automated region detection to identify, for example, necrotic areas, without requiring annotation.

We will also extend our classifiers to enable the detection of cell types in multiplexed histology images. While most histology slides contain at most three separate stains, with new techniques it is possible to stain for up to 250 different markers on a single histology slide (Stack et al., 2014; Blom et al., 2017). This technique would allow better identification of the phenotypes of individual tumour infiltrating macrophages, as the expression of different cytokines within each macrophage could be identified. While images with a single stain can be qualitatively analysed by a pathologist, slides stained to show tens of markers will rapidly become impossible to analyse without quantitative metrics.

6.2.2 Additional metrics for analysis of point data

The spatial statistics used to describe immune cell distributions in Chapter 3 can be extended using a wide range of techniques. The field of topological data analysis (TDA) provides mathematical descriptions of the topology of point data sets (Ghrist, 2008; Bubenik and Dłotko, 2016; Otter et al., 2017). Preliminary results obtained by applying techniques from TDA to simulations of macrophage infiltration into tumour spheroids (see Chapter 5) show promising ways of describing the evolution of the immune cell distribution. Persistent homology, a tool in TDA, has recently been used to aid segmentation of tumours in histological images (Qaiser et al., 2018), indicating that spatial statistics can be used to aid the image analysis process as well as analyse its output. The approaches that we described in Section 3 to automatically cluster tissue samples using spatial statistics should also be extended to aid automatic segmentation in our pipeline.

We will also aim to develop new techniques for measuring the colocalisation of different stains. In particular, quantifying the differences between the hypoxia stainings pimonidazole and CAIX could provide a way of approximating pimonidazole using CAIX. As pimonidazole is considered a better marker of hypoxia than CAIX but cannot be applied to samples after they have been resected, quantifying this difference could add value to CAIX staining. In Chapter 3 we discussed how spatial statistics can be applied to bivariate labelled point sets to provide measures of colocalisation. Future work should extend these techniques to apply spatial statistics to quantify colocalisation of points in datasets labelled with multiple labels. This could then be used to describe the colocalisation of multiple immune cell subtypes throughout a tumour.

6.2.3 Further development of mathematical models

Mechanistic mathematical models can be used to investigate a wide range of biological processes involving macrophages. In particular, we highlight the need for models which can accurately predict macrophage phenotype in response to environmental factors. We stress that although dividing labels into M_1 and M_2 subtypes is an attractive approach for achieving this, mathematical models of macrophage phenotype should represent the current understanding that macrophages express a spectrum of phenotypes and often cannot be easily classified as M_1 or M_2 (Sica and Mantovani, 2012; Bronte and Murray, 2015; Lewis et al., 2016b). We propose the creation of ODE models which can be used to predict macrophage phenotype in response to the local concentration of cytokines. Such models could be integrated with our agent-based modelling to investigate how the spatial infiltration of macrophages influences their phenotype and tumour prognosis. Including macrophage phenotype in our models would allow simulation of tumour treatment using anti-CSF-1R antibodies. This type of model could eventually be used to predict which patients will respond well to this treatment based on observations of patterns of macrophages within the tumour. Our model can also be extended to simulate other treatment types. Preliminary work (results not shown) on applying radiation treatments to tumour spheroids indicates that the spheroid model we describe in Chapter 4 can help explain how tumour composition affects radiation treatment.

If histological images with a stain for blood vessels such as CD31 or CD34 (Nagatsuka et al., 2005) can be added to those describing tumour hypoxia and immune cell locations, we could use this data to integrate multiple relevant biological processes into our model of infiltration into *in vivo* tumours. These processes include extravasation of macrophages from the vasculature, and the role of the EGF / CSF-1 paracrine loop in migration of tumour cells towards blood vessels to form metastases.

6.2.4 Combining image analysis, statistics and mechanistic modelling

Making improvements in any one of the fields of image analysis, spatial statistics and mathematical modelling can help provide benefit to pathologists and aid our understanding of immune cell infiltration into tumours. However, we advocate that these fields be integrated. In this thesis, we have shown that each field is valuable for understanding different aspects of macrophage infiltration, in the form of quantitative descriptions of spatial patterns or through assessing the roles of active and passive mechanisms of macrophage movement. To ensure that mathematical and statistical models have the maximum impact on biological research, we must demonstrate the importance of these techniques to biological researchers.

One way of doing this is to ensure that these quantitative tools are readily accessible to pathologists. By integrating spatial statistics into our imaging software, pathologists with no previous explicit interest in quantitative methods will benefit, and become more accustomed to the concept of using computationally derived metrics to describe tumour images. In Chapter 5 we also demonstrated how our pipeline can be integrated with agent-based models to better understand the mechanisms which cause patterns observed in histology data. We hope that mathematical modelling tools can ultimately be incorporated into imaging software, to provide a direct interface between imaging and predictive modelling for pathologists.

6.3 Final comments

In this thesis, we have developed a set of tools that can be applied to understand the causes of patterns of macrophages observed in histological data. These tools are generic in nature, and can be applied to a wide range of problems throughout immune biology and histology and beyond. We have discussed an interdisciplinary process through which imaging data can be combined with spatial statistics and mechanistic modelling. Although this process has been motivated in this thesis by patterns of macrophage distribution, these tools can be applied to a broad range of problems which involve spatial data in biology. An example of these wider applications is seen in the discussion of microbead counting in Section 3.5, but we have also applied our pipeline to problems of cell detection in fluorescence microscopy images, tracking immune cell movement through videos of *in vitro* experiments and identifying germinal clusters in histology images. We have shown that integrating techniques from image analysis, spatial statistics and mathematical modelling is a realistic way of gaining deeper understanding of underlying biological processes than any one of these fields alone can provide.

Bibliography

Abels, E. and L. Pantanowitz
2017. Current State of the Regulatory Trajectory for Whole Slide Imaging Devices in the USA. *Journal of Pathology Informatics*, 8(23):1–5.

Achanta, R., A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk
2012. SLIC Superpixels Compared to State-of-the-Art Superpixel Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2281.

Andrews, T.
2012. *Computation Time Comparison Between Matlab and C++ Using Launch Windows.* PhD thesis, California Polytechnic State University San Luis Obispo.

Arteta, C., V. Lempitsky, J. A. Noble, and A. Zisserman
2012. Learning to Detect Cells Using Non-overlapping Extremal Regions. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012*, Pp. 348–356.

Baatz, M., J. Zimmermann, and C. G. Blackmore
2009. Automated analysis and detailed quantification of biomedical images using Definiens Cognition Network Technology. *Combinatorial chemistry & high throughput screening*, 12(9):908–16.

Babuska, I. and J. T. Oden
2004. Verification and validation in computational engineering and science: basic concepts. *Computer methods in applied mechanics and engineering*, 193:4057–66.

Baddeley, A. J. and B. W. Silverman
1984. A Cautionary Example on the Use of Second-Order Methods for Analyzing Point Patterns. 40(4):1089–1093.

Bankhead, P., M. B. Loughrey, J. A. Fernández, Y. Dombrowski, D. G. McArt, P. D. Dunne, S. McQuaid, R. T. Gray, L. J. Murray, H. G. Coleman, J. A. James, M. Salto-Tellez, and P. W. Hamilton
2017. QuPath: Open source software for digital pathology image analysis. *Scientific Reports*, 7(1):16878.

Beck, A. H., I. Espinosa, B. Edris, R. Li, K. Montgomery, S. Zhu, S. Varma, R. J. Marinelli, M. D. Van Rijn, and R. B. West
2009. The macrophage colony-stimulating factor 1 response signature in breast carcinoma. *Clinical Cancer Research*, 15(3):778–787.

Belgiovine, C., M. D'Incalci, P. Allavena, and R. Frapolli
2016. Tumor-associated macrophages and anti-tumor therapies: complex links. *Cellular and Molecular Life Sciences.*

Bell, J.
2017. Life sciences: industrial strategy - A report to the Government from the life sciences sector. Technical report, Life Sciences Industrial Strategy Board.

Benzekry, S., C. Lamont, A. Beheshti, A. Tracz, J. M. L. Ebos, L. Hlatky, and P. Hahnfeldt
2014. Classical Mathematical Models for Description and Prediction of Experimental Tumor Growth. *PLoS Computational Biology*, 10(8).

Bindea, G., B. Mlecnik, M. Tosolini, A. Kirilovsky, M. Waldner, A. C. Obenauf, H. Angell, T. Fredriksen, L. Lafontaine, A. Berger, P. Bruneval, W. H. Fridman, C. Becker, F. Pagès, M. R. Speicher, Z. Trajanoski, and J. Galon
2013. Spatiotemporal dynamics of intratumoral immune cells reveal the immune landscape in human cancer. *Immunity*, 39(4):782–795.

Blom, S., L. Paavolainen, D. Bychkov, R. Turkki, P. Mäki-Teeri, A. Hemmes, K. Välimäki, J. Lundin, O. Kallioniemi, and T. Pellinen
2017. Systems pathology by multiplexed immunohistochemistry and whole-slide digital image analysis. *Scientific Reports*, 7(1):15580.

Bravo, R., M. Robertson-Tessi, and A. R. A. Anderson
2018. Hybrid Automata Library. *Preprint bioRxiv.*

Bronte, V. and P. J. Murray
2015. Understanding Local Macrophage Phenotypes In Disease: Modulating macrophage function to treat cancer. *Nature Medicine*, 21(2):117–119.

Bubenik, P. and P. Dłotko
2016. A persistence landscapes toolbox for topological statistics. *Journal of Symbolic Computation*, 1:1–24.

Buchwalow, I. B. and W. Böcker
2010. *Immunohistochemistry: Basics and Methods.* Heidelberg: Springer-Verlag.

Bueno, G., M. M. Fernández-Carrobles, O. Deniz, and M. García-Rojo
2016. New Trends of Emerging Technologies in Digital Pathology. *Pathobiology*, 83(2-3):61–69.

Buyong, T.
2007. *Spatial Data Analysis for Geographic Information Science.* Skudai: Penerbit UTM.

Byrne, H., S. M. Cox, and C. E. Kelly
2004. Macrophage-Tumour Interactions: in Vivo Dynamics. *Discrete and Continuous Dynamical Systems - Series Band*, 4(1):81–98.

Byrne, H. M. and D. Drasdo
2009. Individual-based and continuum models of growing cell populations: A comparison. *Journal of Mathematical Biology*, 58(4-5):657–687.

Cannarile, M. A., C. H. Ries, S. Hoves, and D. Rüttinger
2014. Targeting tumor-associated macrophages in cancer therapy and understanding their complexity. *Oncoimmunology*, 3(9):e955356.

Capece, D., M. Fischietti, D. Verzella, A. Gaggiano, G. Cicciarelli, A. Tessitore, F. Zazzeroni, and E. Alesse
2013. The Inflammatory Microenvironment in Hepatocellular Carcinoma: A Pivotal Role for Tumor-Associated Macrophages. 2013.

Cappuccio, A., M. Elishmereni, and Z. Agur
2006. Cancer immunotherapy by interleukin-21: Potential treatment strategies evaluated in a mathematical model. *Cancer Research*, 66(14):7293–7300.

Carmona-Fontaine, C., M. Deforet, L. Akkari, C. B. Thompson, J. A. Joyce, and J. B. Xavier
2017. Metabolic origins of spatial organization in the tumor microenvironment. *Proceedings of the National Academy of Sciences*, 114(11):201700600.

Cassier, P. A., A. Italiano, C. A. Gomez-Roca, C. Le Tourneau, M. Toulmonde, M. A. Cannarile, C. Ries, A. Brillouet, C. Müller, A.-M. Jegg, A.-M. Bröske, M. Dembowski, K. Bray-French, C. Freilinger, G. Meneses-Lorente, M. Baehner, R. Harding, J. Ratnayake, K. Abiraj, N. Gass, K. Noh, R. D. Christen, L. Ukarma, E. Bompas, J.-P. Delord, J.-Y. Blay, and D. Rüttinger
2015. CSF1R inhibition with emactuzumab in locally advanced diffuse-type tenosynovial giant cell tumours of the soft tissue: a dose-escalation and dose-expansion phase 1 study. *The Lancet Oncology*, 16(8):949–956.

Chanmee, T., P. Ontong, K. Konno, and N. Itano
2014. Tumor-Associated Macrophages as Major Players in the Tumor Microenvironment. *Cancers*, 6(3):1670–1690.

Chaplot, S., L. M. Patnaik, and N. R. Jagannathan
2006. Classification of magnetic resonance brain images using wavelets as input to support vector machine and neural network. *Biomedical Signal Processing and Control*, 1(1):86–92.

Chen, D., A. A. Bobko, A. C. Gross, R. Evans, C. B. Marsh, V. V. Khramtsov, T. D. Eubank, and A. Friedman
2014. Involvement of Tumor Macrophage HIFs in Chemotherapy Effectiveness: Mathematical Modeling of Oxygen, pH, and Glutathione. *PLoS ONE*, 9(10):e107511.

Chen, D., J. M. Roda, C. B. Marsh, T. D. Eubank, and A. Friedman
2012. Hypoxia Inducible Factors-Mediated Inhibition of Cancer by GM-CSF: A Mathematical Model. *Bulletin of Mathematical Biology*, 74(11):2752–2777.

Chourasiya, S. and G. U. Rani
2014. Automatic Red Blood Cell Counting using Watershed Segmentation. *International Journal of Computer Science and Information Technologies*, 5(4):4834–4838.

Comaniciu, D. and P. Meer
2002. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619.

Cooper, J., J. O. Vik, and D. Waltemath
2015. A call for virtual experiments: Accelerating the scientific process. *Progress in Biophysics and Molecular Biology*, 117(1):99–106.

Corredor, G., X. Wang, Y. Zhou, C. Lu, P. Fu, K. N. Syrigos, D. L. Rimm, M. Yang, E. Romero, K. A. Schalper, V. Velcheti, and A. Madabhushi
2018. Spatial architecture and arrangement of tumor-infiltrating lymphocytes for predicting likelihood of recurrence in early-stage non-small cell lung cancer. *Clinical Cancer Research*, P. clincanres.2013.2018.

Cortes, C. and V. Vapnik
1995. Support-Vector Networks. *Machine Learning*, 20(3):273–297.

Cross, S. S.
1998. Grading and scoring in histopathology. *Histopathology*, 33(2):99–106.

Cytowski, M. and Z. Szymanska
2014. Large Scale Parallel Simulations of 3-D Cell Colony Dynamics. (October):1–15.

Cytowski, M. and Z. Szymanska
2015. Large-Scale Parallel Simulations of 3D Cell Colony Dynamics: The Cellular Environment. *Computing in Science and Engineering*, 17(5):44–48.

de Pillis, L. G., K. R. Fister, W. Gu, C. Collins, M. Daub, D. Gross, J. Moore, and B. Preskill
2009. Mathematical Model Creation for Cancer Chemo-Immunotherapy. *Computational and Mathematical Methods in Medicine*, 10(3):165–184.

de Visser, K. E., A. Eichten, and L. M. Coussens
2006. Paradoxical roles of the immune system during cancer development. *Nature reviews. Cancer*, 6(1):24–37.

den Breems, N. Y. and R. Eftimie
2016. The re-polarisation of M2 and M1 macrophages and its role on cancer outcomes. *Journal of Theoretical Biology*, 390:23–39.

Deroulers, C., D. Ameisen, M. Badoual, C. Gerin, A. Granier, and M. Lartaud
2013. Analyzing huge pathology images with open source software. *Diagnostic pathology*, 8:92.

Dixon, P. M.
2002. Ripley's K function. *Encyclopedia of Environmetrics*, 3(December):1796–1803.

Doane, D. P. and L. E. Seward
2011. Measuring skewness: A forgotten statistic? *Journal of Statistics Education*, 19(2):1–18.

Dorie, M. J., R. F. Kallman, and M. A. Coyne
1986. Effect of cytochalasin B, nocodazole and irradiation on migration and internalization of cells and microspheres in tumor cell spheroids. *Experimental Cell Research*, 166(2):370–378.

Dorie, M. J., R. F. Kallman, D. F. Rapacchietta, D. Van Antwerp, and Y. R. Huang
1982. Migration and internalization of cells and polystyrene microspheres in tumor cell spheroids. *Experimental Cell Research*, 141(1):201–209.

Dranoff, G.
2004. Cytokines in cancer pathogenesis and cancer therapy. *Nature Reviews Cancer*, 4(1):11–22.

Drasdo, D. and S. Höhme
2005. A single-cell-based model of tumor growth in vitro: monolayers and spheroids. *Physical Biology*, 2(3):133–147.

Edelsbrunner, H., D. G. Kirkpatrick, and R. Seidel
1983. On the Shape of a Set of Points. *IEEE Transactions on Information Theory*, 29(4):551–559.

Edge, S. B. and C. C. Compton
2010. The american joint committee on cancer: The 7th edition of the AJCC cancer staging manual and the future of TNM. *Annals of Surgical Oncology*, 17(6):1471–1474.

Edwards V, D. K., D. T. Sweeney, H. Ho, C. A. Eide, A. Rofelty, A. Agarwal, S. Q. Liu, A. V. Danilov, P. Lee, D. Chantry, S. K. McWeeney, B. J. Druker, J. W. Tyner, S. E. Spurgeon, and M. M. Loriaux
2018. Targeting of colony-stimulating factor 1 receptor (CSF1R) in the CLL microenvironment yields antineoplastic activity in primary patient samples. *Oncotarget*, 9(37):24576–24589.

Eftimie, R., J. J. Gillard, and D. A. Cantrell
2016. Mathematical Models for Immunology: Current State of the Art and Future Research Directions. *Bulletin of Mathematical Biology*, 78(10):2091–2134.

Eikenberry, S., C. Thalhauser, and Y. Kuang
2009. Tumor-Immune Interaction, Surgical Treatment, and Cancer Recurrence in a Mathematical Model of Melanoma. *PLoS Computational Biology*, 5(4).

El-Naqa, I., Y. Yang, M. N. Wernick, N. P. Galatsanos, and R. M. Nishikawa
2002. A support vector machine approach for detection of microcalcifications. *IEEE Transactions on Medical Imaging*, 21(12):1552–1563.

Elitas, M. and S. Zeinali
2016. Modeling and Simulation of EGF-CSF-1 pathway to Investigate Glioma - Macrophage Interaction in Brain Tumors. *International Journal of Cancer Studies & Research (IJCR)*, Pp. 1–8.

Elmore, J. G., G. M. Longton, P. A. Carney, B. M. Geller, T. Onega, A. N. Tosteson, H. D. Nelson, M. S. Pepe, K. H. Allison, S. J. Schnitt, F. P. O'Malley, and D. L. Weaver
2015. Diagnostic concordance among pathologists interpreting breast biopsy specimens. *JAMA - Journal of the American Medical Association*, 313(11):1122–1132.

Ergin, S. and O. Kilic
2014. A new feature extraction framework based on wavelets for breast cancer diagnosis. *Computers in Biology and Medicine*, 51:171–182.

Fawcett, T.
2006. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874.

Felzenszwalb, P. F. and D. P. Huttenlocher
2004. Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59(2):167–181.

Figueredo, G. P., T. V. Joshi, J. M. Osborne, H. M. Byrne, and M. R. Owen
2013. On-lattice agent-based simulation of populations of cells within the open-source Chaste framework. *Interface Focus*, 3(2).

Fletcher, A. G., J. M. Osborne, P. K. Maini, and D. J. Gavaghan
2013. Implementing vertex dynamics models of cell populations in biology within a consistent computational framework. *Progress in Biophysics and Molecular Biology*, 113(2):299–326.

Folkman, J. and M. Hochberg
1973. Self-regulation of growth in three dimensions. *The Journal of Experimental Medicine*, 138:745–753.

Fozard, J. A., G. R. Kirkham, L. D. Buttery, J. R. King, O. E. Jensen, and H. M. Byrne
2011. Techniques for analysing pattern formation in populations of stem cells and their progeny. *BMC bioinformatics*, 12:396.

Franklin, R. A., W. Liao, A. Sarkar, M. V. Kim, M. R. Bivona, K. Liu, E. G. Pamer, and M. O. Li
2014. The cellular and molecular origin of tumor-associated macrophages. *Science (New York, N.Y.)*, 344(6186):921–5.

Friedman, A. and K.-L. Liao
2015. The role of the cytokines IL-27 and IL-35 in cancer. *Mathematical Biosciences and Engineering*, 12(6):1203–1217.

Fujimura, T., Y. Kambayashi, Y. Fujisawa, T. Hidaka, and S. Aiba
2018. Tumor-Associated Macrophages: Therapeutic Targets for Skin Cancer. *Frontiers in Oncology*, 8(January):1–6.

Fukushima, K.
1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4):193–202.

García-Martínez, E., G. L. Gil, A. C. Benito, E. González-Billalabeitia, M. A. V. Conesa, T. G. García, E. García-Garre, V. Vicente, and F. A. de la Peña
2014. Tumor-infiltrating immune cell profiles and their change after neoadjuvant chemotherapy predict response and prognosis of breast cancer. *Breast Cancer Research*, 16(6):1–17.

García Rojo, M., G. Bueno García, and M. Carbajo Vicente
2003. Techniques in Pathology Critical Comparison of 31 Commercially Available. *Techniques in Pathology*, 14(4):285–305.

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

Ghaffarizadeh, A., S. H. Friedman, and P. MacKlin
2016. BioFVM: An efficient, parallelized diffusive transport solver for 3-D biological simulations. *Bioinformatics*, 32(8):1256–1258.

Ghaffarizadeh, A., R. Heiland, S. H. Friedman, and M. Shannon
2018. PhysiCell : an Open Source Physics-Based Cell Simulator for 3-D Multicellular Systems. *PLoS Computational Biology*, 14(2):1–34.

Ghosh, S., M. B. Joshi, D. Ivanov, C. Feder-Mengus, G. C. Spagnoli, I. Martin, P. Erne, and T. J. Resink
2007. Use of multicellular tumor spheroids to dissect endothelial cell-tumor cell interactions: A role for T-cadherin in tumor angiogenesis. *FEBS Letters*, 581(23):4523–4528.

Ghrist, R.
2008. Barcodes: The persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75.

Ginhoux, F. and M. Guilliams
2016. Tissue-Resident Macrophage Ontogeny and Homeostasis. *Immunity*, 44(3):439–449.

Gonzalez, R. C. and R. E. Woods
2008. *Digital image processing.*

Goswami, S., E. Sahai, J. B. Wyckoff, C.-s. F. Epidermal, M. Cammer, D. Cox, F. J. Pixley, E. R. Stanley, J. E. Segall, and J. S. Condeelis
2005. Macrophages Promote the Invasion of Breast Carcinoma Cells via a Colony-Stimulating Factor-1 / Epidermal Growth Factor Paracrine Loop Macrophages Promote the Invasion of Breast Carcinoma. (12):5278–5284.

Graner, F. and J. A. Glazier
1992. Simulation of Biological Cell Sorting Using a Two-Dimensional Extended Potts Models. *Physical Review Letters*, 69(13):2013–2017.

Greenspan, H. P.
1972. Models for the Growth of a Solid Tumor by Diffusion.

Greenspan, H. P.
1976. On the growth and stability of cell cultures and solid tumors. *Journal of Theoretical Biology*, 56(1):229–242.

Greijer, A. E. and E. Van Der Wall
2004. The role of hypoxia inducible factor 1 (HIF-1) in hypoxia induced apoptosis. *Journal of Clinical Pathology*, 57(10):1009–1014.

Griffin, J. and D. Treanor
2017. Digital pathology in clinical use: Where are we now and what is holding us back? *Histopathology*, 70(1):134–145.

Hagemann, T., T. Lawrence, I. McNeish, K. A. Charles, H. Kulbe, R. G. Thompson, S. C. Robinson, and F. R. Balkwill
2008. "Re-educating" tumor-associated macrophages by targeting NF-kappaB. *The Journal of experimental medicine*, 205(6):1261–1268.

Hamilton, P. W., P. Bankhead, Y. Wang, R. Hutchinson, D. Kieran, D. G. McArt, J. James, and M. Salto-Tellez
2014. Digital pathology and image analysis in tissue biomarker research. *Methods*, 70(1):59–73.

Hanahan, D. and R. A. Weinberg
2011. Hallmarks of cancer: the next generation. *Cell*, 144(5):646–74.

Harney, A. S., E. N. Arwert, D. Entenberg, Y. Wang, P. Guo, B.-Z. Qian, M. H. Oktay, J. W. Pollard, J. G. Jones, and J. S. Condeelis
2015. Real-Time Imaging Reveals Local, Transient Vascular Permeability, and Tumor Cell Intravasation Stimulated by TIE2hi Macrophage-Derived VEGFA. *Cancer Discovery*, 5(9):932–943.

Hartigan, J. A. and M. A. Wong
1979. Algorithm AS 136: A K-Means Clustering Algorithm. *Applied Statistics*, 28(1):100.

Harwood, R. J., C. E. Lewis, and S. K. Biswas
2012. The role of hypoxia in re-educating macrophages in the tumour environment. *Tumour Microenvironment and Myelomonocytic cells*, Pp. 89–111.

Heindl, A., S. Nawaz, and Y. Yuan
2015. Mapping spatial heterogeneity in the tumor microenvironment: a new era for digital pathology. *Laboratory Investigation*, 95(4):377–384.

Hernández, C. X., M. M. Sultan, and V. S. Pande
2018. Using Deep Learning for Segmentation and Counting within Microscopy Data.

Herter, S., L. Morra, R. Schlenker, J. Sulcova, L. Fahrni, I. Waldhauer, S. Lehmann, T. Reislï¿½nder, I. Agarkova, J. M. Kelm, C. Klein, P. Umana, and M. Bacac
2017. A novel three-dimensional heterotypic spheroid model for the assessment of the activity of cancer immunotherapy agents. *Cancer Immunology, Immunotherapy*, 66(1):129–140.

Hoehme, S. and D. Drasdo
2010. A cell-based simulation software for multi-cellular systems. *Bioinformatics*, 26(20):2641–2642.

Holten-Rossing, H., M. L. M. Talman, A. M. B. Jylling, A. V. Lænkholm, M. Kristensson, and B. Vainer
2017. Application of automated image analysis reduces the workload of manual screening of sentinel lymph node biopsies in breast cancer. *Histopathology*, 71(6):866–873.

Hoos, A.
2016. Development of immuno-oncology drugs - from CTLA4 to PD1 to the next generations. *Nature reviews. Drug discovery*, 15(4):235–47.

Hsu, C.-W. and C.-J. Lin
2002. A comparison of methods for multiclass support vector machines. *IEEE Transactions on Neural Networks*, 13(2):415–425.

Hughes, R., B.-Z. Qian, C. Rowan, M. Muthana, I. Keklikoglou, O. C. Olson, S. Tazzyman, S. Danson, C. Addison, M. Clemons, A. M. Gonzalez-Angulo, J. A. Joyce, M. De Palma, J. W. Pollard, and C. E. Lewis
2015. Perivascular M2 Macrophages Stimulate Tumor Relapse after Chemotherapy. *Cancer research*, Pp. 1–14.

Jacobsen, K., L. Russell, B. Kaur, and A. Friedman
2015. Effects of CCN1 and Macrophage Content on Glioma Virotherapy: A Mathematical Model. *Bulletin of Mathematical Biology*, 77(6):984–1012.

Janowczyk, A. and A. Madabhushi
2016. Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases. *Journal of Pathology Informatics*, 7(1):29.

Kadioglu, E. and M. de Palma
2015. Cancer metastasis: Perivascular macrophages under watch. *Cancer Discovery*, 5(9):906–908.

Kang, S., S. Kahan, J. McDermott, N. Flann, and I. Shmulevich
2014. Biocellion: Accelerating computer simulation of multicellular biological system models. *Bioinformatics*, 30(21):3101–3108.

Kather, J. N., P. Charoentong, M. Suarez-Carmona, E. Herpel, F. Klupp, A. Ulrich, M. Schneider, I. Zoernig, T. Luedde, D. Jaeger, J. Poleszczuk, and N. Halama
2018. High-throughput screening of combinatorial immunotherapies with patient-specific in silico models of metastatic colorectal cancer. *Cancer Research*, P. canres.1126.2018.

Kather, J. N., J. Poleszczuk, M. Suarez-Carmona, J. Krisam, P. Charoentong, N. A. Valous, C. A. Weis, L. Tavernar, F. Leiss, E. Herpel, F. Klupp, A. Ulrich, M. Schneider, A. Marx, D. Jäger, and N. Halama
2017. In silico modeling of immunotherapy and stroma-targeting therapies in human colorectal cancer. *Cancer Research*, 77(22):6442–6452.

Kay, S. K.
2014. *Cell fate mechanisms in colorectal cancer*. PhD thesis, University of Oxford.

Keklikoglou, I. and M. De Palma
2009. Metastasis risk after anti-macrophage therapy. *Nature*, Pp. 10–11.

Keller, E. F. and L. A. Segel
1971. Model for chemotaxis. *Journal of Theoretical Biology*, 30(2):225–234.

Kelly, C. E., R. D. Leek, H. M. Byrne, S. M. Cox, A. L. Harris, and C. E. Lewis
2002. Modelling Macrophage Infiltration into Avascular Tumours. *Journal of Theoretical Medicine*, 4(1):21–38.

Khalil, D. N., E. L. Smith, R. J. Brentjens, and J. D. Wolchok
2016. The future of cancer treatment: immunomodulation, CARs and combination immunotherapy. *Nature reviews. Clinical oncology*, 13(5):273–290.

Khan, A. M., N. Rajpoot, D. Treanor, and D. Magee
2014. A nonlinear mapping approach to stain normalization in digital histopathology images using image-specific color deconvolution. *IEEE Transactions on Biomedical Engineering*, 61(6):1729–1738.

Kim, E., S. Stamatelos, J. Cebulla, Z. M. Bhujwalla, A. S. Popel, and A. P. Pathak
2012. Multiscale imaging and computational modeling of blood flow in the tumor vasculature. *Annals of Biomedical Engineering*, 40(11):2425–2441.

Kiskowski, M. A., J. F. Hancock, and A. K. Kenworthy
2009. On the use of Ripley's K-function and its derivatives to analyze domain size. *Biophysical Journal*, 97(4):1095–1103.

Kitamura, T., B.-Z. Qian, and J. W. Pollard
2015. Immune cell promotion of metastasis. *Nature Reviews Immunology*, 15(2):73–86.

Klauschen, F., K. R. Müller, A. Binder, M. Bockmayr, M. Hägele, P. Seegerer, S. Wienert, G. Pruneri, S. de Maria, S. Badve, S. Michiels, T. O. Nielsen, S. Adams, P. Savas, F. Symmans, S. Willis, T. Gruosso, M. Park, B. Haibe-Kains, B. Gallas, A. M. Thompson, I. Cree, C. Sotiriou, C. Solinas, M. Preusser, S. M. Hewitt, D. Rimm, G. Viale, S. Loi, S. Loibl, R. Salgado, and C. Denkert
2018. Scoring of tumor-infiltrating lymphocytes: From visual estimation to machine learning.

Kloepper, J., L. Riedemann, Z. Amoozgar, G. Seano, K. Susek, V. Yu, N. Dalvie, R. L. Amelung, M. Datta, J. W. Song, V. Askoxylakis, J. W. Taylor, C. Lu-Emerson, A. Batista, N. D. Kirkpatrick, K. Jung, M. Snuderl, A. Muzikansky, K. G. Stubenrauch, O. Krieter, H. Wakimoto, L. Xu, L. L. Munn, D. G. Duda, D. Fukumura, T. T. Batchelor, and R. K. Jain
2016. Ang-2/VEGF bispecific antibody reprograms macrophages and resident microglia to anti-tumor phenotype and prolongs glioblastoma survival. *Proceedings of the National Academy of Sciences*, P. 201525360.

Knútsdóttir, H., J. S. Condeelis, and E. Pálsson
2016. 3-D individual cell based computational modeling of tumor cell–macrophage paracrine signaling mediated by EGF and CSF-1 gradients. *Integrative Biology*, 8:104–119.

Knútsdóttir, H., E. Pálsson, and L. Edelstein-Keshet
2014. Mathematical model of macrophage-facilitated breast cancer cells invasion. *Journal of Theoretical Biology*, 357:184–199.

Kratochvill, F., G. Neale, J. M. Haverkamp, L.-A. Van de Velde, A. M. Smith, D. Kawauchi, J. McEvoy, M. F. Roussel, M. A. Dyer, J. E. Qualls, and P. J. Murray
2015. TNF Counterbalances the Emergence of M2 Tumor Macrophages. *Cell Reports*, 12(11):1902–1914.

Krizhevsky, A., I. Sutskever, and G. Hinton
2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, number 2, Pp. 1097–1105.

Krombach, F., S. Münzing, A.-m. Allmeling, J. T. Gerlach, J. Behr, and M. Dorger
1997. Cell Size of Alveolar Macrophages: An Interspecies Comparison. *Environmental Health Perspectives*, 105:1261–1263.

Lamagna, C., M. Aurrand-Lions, and B. A. Imhof
2006. Dual role of macrophages in tumor growth and angiogenesis. *Journal of leukocyte biology*, 80(4):705–713.

Landman, K. A. and C. P. Please
2001. Tumour dynamics and necrosis: surface tension and stability. *IMA Journal of Mathematics Applied in Medicine and Biology*, 18:131–158.

Laoui, D., E. Van Overmeire, P. De Baetselier, J. A. Van Ginderachter, and G. Raes
2014. Functional Relationship between Tumor-Associated Macrophages and Macrophage Colony-Stimulating Factor as Contributors to Cancer Progression. *Frontiers in Immunology*, 5(October):1–15.

Leek, R. D.
1999. *The Role of Tumour Associated Macrophages in Breast Cancer Angiogenesis*. PhD thesis, Oxford Brookes University.

Leek, R. D., R. J. Landers, A. L. Harris, and C. E. Lewis
1999. Necrosis correlates with high vascular density and focal macrophage infiltration in invasive carcinoma of the breast. *Br J Cancer*, 79(5-6):991–995.

Leek, R. D., C. E. Lewis, R. Whitehouse, M. Greenall, J. Clarke, and A. L. Harris
1996. Association of Macrophage Infiltration with Angiogenesis and Prognosis in Invasive Breast Carcinoma. *Cancer Research*, 56(16):4625–4629.

Leonard, F., L. T. Curtis, M. J. Ware, T. Nosrat, X. Liu, K. Yokoi, H. B. Frieboes, and B. Godin
2017. Macrophage polarization contributes to the anti-tumoral efficacy of mesoporous nanovectors loaded with albumin-bound paclitaxel. *Frontiers in Immunology*, 8(JUN).

Levinshtein, A., A. Stere, K. N. Kutulakos, D. J. Fleet, S. J. Dickinson, and K. Siddiqi
2009. TurboPixels: Fast superpixels using geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2290–2297.

Lewis, C. E., A. S. Harney, and J. W. Pollard
2016a. The Multifaceted Role of Perivascular Macrophages in Tumors.

Lewis, C. E., A. S. Harney, and J. W. Pollard
2016b. The Multifaceted Role of Perivascular Macrophages in Tumors. *Cancer Cell*, 30(1):18–25.

Lewis, C. E. and J. W. Pollard
2006. Distinct role of macrophages in different tumor microenvironments. *Cancer Research*, 66(2):605–612.

Lewis, J. S., J. A. Lee, J. C. E. Underwood, A. L. Harris, and C. E. Lewis
1999. Macrophage responses to hypoxia: relevance to disease mechanisms. *Journal of leukocyte biology*, 66(6):889–900.

Litjens, G., T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. van der Laak, B. van Ginneken, and C. I. Sánchez
2017. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42(December 2012):60–88.

López, Á. G., J. M. Seoane, and M. a. F. Sanjuán
2014. A Validated Mathematical Model of Tumor Growth Including Tumor–Host Interaction, Cell-Mediated Immune Response and Chemotherapy. *Bulletin of Mathematical Biology*, 76(11):2884–2906.

Lowe, D. G.
2004. Distinctive image features from scale invariant keypoints. *International Journal of Computer Vision*, 60:91–11020042.

Lucchi, A., P. Marquez-Neila, C. Becker, Y. Li, K. Smith, G. Knott, and P. Fua
2015. Learning structured models for segmentation of 2-D and 3-D imagery. *IEEE Transactions on Medical Imaging*, 34(5):1096–1110.

Lucchi, A., K. Smith, R. Achanta, V. Lepetit, and P. Fua
2010. A Fully Automated Approach to Segmentation of Irregularly Shaped Cellular Structures in EM Images. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2010 13th International Conference, Beijing, China, September 20-24, 2010, Proceedings, Part II*, T. Jiang, N. Navab, J. Pluim, and M. Viergever, eds., Pp. 463–471, Beijing, China.

Ma, H. L., Q. Jiang, S. Han, Y. Wu, J. C. Tomshine, D. Wang, Y. Gan, G. Zou, and X. J. Liang
2012. Multicellular tumor spheroids as an in vivo-like tumor model for three-dimensional imaging of chemotherapeutic and nano material cellular penetration. *Molecular Imaging*, 11(6):487–498.

Ma, J., L. Liu, G. Che, N. Yu, F. Dai, and Z. You
2010. The M1 form of tumor-associated macrophages in non-small cell lung cancer is positively associated with survival time. *BMC cancer*, 10:112.

Macfarlane, F. R., T. Lorenzi, and M. A. Chaplain
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.

Mahlbacher, G., L. T. Curtis, J. Lowengrub, and H. B. Frieboes
2018. Mathematical modeling of tumor-associated macrophage interactions with the cancer microenvironment. *Journal for ImmunoTherapy of Cancer*, 6(1):1–17.

Malpica, N., C. O. De Solórzano, J. J. Vaquero, A. Santos, I. Vallcorba, J. M. García-Sagredo, and F. Del Pozo
1997. Applying watershed algorithms to the segmentation of clustered nuclei. *Cytometry*, 28(4):289–297.

Mantovani, A. and P. Allavena
2015. The interaction of anticancer therapies with tumor-associated macrophages. *Journal of Experimental Medicine*, 212(4).

Mantovani, A., P. Allavena, and A. Sica
2004. Tumour-associated macrophages as a prototypic type II polarised phagocyte population: Role in tumour progression. *European Journal of Cancer*, 40(11):1660–1667.

Mantovani, A., F. Marchesi, A. Malesci, L. Laghi, and P. Allavena
2017. Tumour-associated macrophages as treatment targets in oncology. *Nature Reviews Clinical Oncology*, 14(7):399–416.

Mantovani, A., S. Sozzani, M. Locati, P. Allavena, and A. Sica
2002. Macrophage polarization: Tumor-associated macrophages as a paradigm for polarized M2 mononuclear phagocytes. *Trends in Immunology*, 23(11):549–555.

Mattfeldt, T., S. Eckel, F. Fleischer, and V. Schmidt
2009. Statistical analysis of labelling patterns of mammary carcinoma cell nuclei on histological sections. *Journal of microscopy*, 235(1):106–18.

McCann, M. T.
2015. *Tools for Automated Histology Image Analysis*. PhD thesis, Carnegie Mellon University.

McElwain, D. L. S. and G. J. Pettet
1993. Cell Migration in Multicell Spheroids: Swimming Against the Tide. *Bulletin of Mathematical Biology*, 55(3):655–674.

Meineke, F. A., C. S. Potten, and M. Loeffler
2001. Cell migration and organization in the intestinal crypt using a lattice-free model. *Cell Proliferation*, 34(4):253–266.

Meyer, F.
1994. Topographic distance and watershed lines. *Signal Processing*, 38(1):113–125.

Mignardi, M., O. Ishaq, X. Qian, and C. Wahlby
2016. Bridging Histology and Bioinformatics—Computational Analysis of Spatially Resolved Transcriptomics. *Proceedings of the IEEE*, 105(3):1–12.

Mills, C. D., K. Kincaid, J. M. Alt, M. J. Heilman, and A. M. Hill
2000. M-1/M-2 Macrophages and the Th1/Th2 Paradigm. *Journal of Immunology*.

Mirams, G. R., C. J. Arthurs, M. O. Bernabeu, R. Bordas, J. Cooper, A. Corrias, Y. Davit, S. J. Dunn, A. G. Fletcher, D. G. Harvey, M. E. Marsh, J. M. Osborne, P. Pathmanathan, J. Pitt-Francis, J. Southern, N. Zemzemi, and D. J. Gavaghan
2013. Chaste: An Open Source C++ Library for Computational Physiology and Biology. *PLoS Computational Biology*, 9(3).

Mohri, M., A. Rostamizadeh, and A. Talwalkar
2012. *Foundations of Machine Learning*. Cambridge, Massachusetts: The MIT Press.

Møller, J. and R. P. Waagepetersen
2004. *Statistical Inference and Simulation for Spatial Point Processes.* Boca Raton: Chapman & Hall/CRC.

Mosser, D. M. and J. P. Edwards
2008. Exploring the full spectrum of macrophage activation. *Nature reviews. Immunology*, 8(12):958–69.

Munn, D. H. and V. Bronte
2016. Immune suppressive mechanisms in the tumor microenvironment. *Current Opinion in Immunology*, 39:1–6.

Murdoch, C., A. Giannoudis, and C. E. Lewis
2004. Mechanisms regulating the recruitment of macrophages into hypoxic areas of tumors and other ischemic tissues. *Blood*, 104(8):2224–2234.

Murray, P. J., J. E. Allen, S. K. Biswas, E. A. Fisher, D. W. Gilroy, S. Goerdt, S. Gordon, J. A. Hamilton, L. B. Ivashkiv, T. Lawrence, M. Locati, A. Mantovani, F. O. Martinez, J.-L. Mege, D. M. Mosser, G. Natoli, J. P. Saeij, J. L. Schultze, K. A. Shirey, A. Sica, J. Suttles, I. Udalova, J. A. van Ginderachter, S. N. Vogel, and T. A. Wynn
2014. Macrophage Activation and Polarization: Nomenclature and Experimental Guidelines. *Immunity*, 41(1):14–20.

Murray, P. J. and T. A. Wynn
2011. Protective and pathogenic functions of macrophage subsets. *Nature Reviews Immunology*, 11(11):723–737.

Mylonas, K. J., M. G. Nair, L. Prieto-Lafuente, D. Paape, and J. E. Allen
2009. Alternatively activated macrophages elicited by helminth infection can be reprogrammed to enable microbial killing. *Journal of immunology (Baltimore, Md. : 1950)*, 182:3084–3094.

Nagatsuka, H., K. Hibi, M. Gunduz, H. Tsujigiwa, R. Tamamura, T. Sugahara, A. Sasaki, and N. Nagai
2005. Various immunostaining patterns of CD31, CD34 and endoglin and their relationship with lymph node metastasis in oral squamous cell carcinomas. *Journal of Oral Pathology and Medicine*, 34(2):70–76.

Nath, S. and G. R. Devi
2016. Three-dimensional culture systems in cancer research: Focus on tumor spheroid model. *Pharmacology and Therapeutics*, 163:94–108.

Nawaz, S., A. Heindl, K. Koelble, and Y. Yuan
2015. Beyond immune density: critical role of spatial heterogeneity in estrogen receptor-negative breast cancer. *Modern Pathology*, 28(6):766–777.

Nawaz, S. and Y. Yuan
2016. Computational pathology: Exploring the spatial dimension of tumor ecology. *Cancer Letters*, 380(1):296–303.

Norton, K. A., K. Jin, and A. S. Popel
2018. Modeling triple-negative breast cancer heterogeneity: Effects of stromal macrophages, fibroblasts and tumor vasculature. *Journal of Theoretical Biology*, 452:56–68.

Noy, R. and J. W. Pollard
2014. Tumor-Associated Macrophages: From Mechanisms to Therapy. *Immunity*, 41(1):49–61.

O'Brien, J., T. Lyons, J. Monks, M. S. Lucia, R. S. Wilson, L. Hines, Y. G. Man, V. Borges, and P. Schedin
2010. Alternatively activated macrophages and collagen remodeling characterize the postpartum involuting mammary gland across species. *American Journal of Pathology*, 176(3):1241–1255.

Olive, P. L., C. Aquino-Parsons, S. H. MacPhail, S. Y. Liao, E. J. Stanbridge, J. A. Raleigh, and M. I. Lerman
2001. Carbonic anhydrase 9 as an endogenous marker for hypoxic cells in cervical cancer. *Cancer Research*, 61(24):8924–8929.

Osborne, J. M., A. G. Fletcher, J. M. Pitt-Francis, P. K. Maini, and D. J. Gavaghan
2017. *Comparing individual-based approaches to modelling the self-organization of multicellular tissues*, volume 13.

O'Sullivan, D. and D. J. Unwin
2010. *Geographic Information Analysis*. Hoboken, New Jersey: Wiley.

Otter, N., M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington
2017. A roadmap for the computation of persistent homology. *EPJ Data Science*, 6(17).

Owen, M. R., H. M. Byrne, and C. E. Lewis
2004. Mathematical modelling of the use of macrophages as vehicles for drug delivery to hypoxic tumour sites. *Journal of Theoretical Biology*, 226(4):377–391.

Owen, M. R. and J. A. Sherratt
1997. Pattern formation and spatiotemporal irregularity in a model for macrophage-tumour interactions. *Journal of theoretical biology*, 189(1):63–80.

Owen, M. R. and J. A. Sherratt
1998. Modelling the macrophage invasion of tumours: Effects on growth and composition. *IMA Journal of Mathematics Applied in Medicine and Biology*, 15:165–185.

Owen, M. R. and J. A. Sherratt
1999. Mathematical modelling of macrophage dynamics in tumours. *Mathematical Models and Methods in Applied Sciences*, 09(04):513–539.

Owen, M. R., I. J. Stamper, M. Muthana, G. W. Richardson, J. Dobson, C. E. Lewis, and H. M. Byrne
2011. Mathematical modeling predicts synergistic antitumor effects of combining a macrophage-based, hypoxia-targeted gene therapy with chemotherapy. *Cancer Research*, 71(8):2826–2837.

Pathmanathan, P., J. Cooper, A. Fletcher, G. Mirams, P. Murray, J. M. Osborne, J. Pitt-Francis, A. Walter, and S. J. Chapman
2009. A computational study of discrete mechanical tissue models. *Physical biology*, 6(3):036001.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay
2012. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Peirce, S. M.
2008. Computational and mathematical modeling of angiogenesis. *Microcirculation*, 15(8):739–751.

Perfahl, H., B. D. Hughes, T. Alarcón, P. K. Maini, M. C. Lloyd, M. Reuss, and H. M. Byrne
2017. 3D hybrid modelling of vascular network formation. *Journal of Theoretical Biology*.

Peterson, T. E., N. D. Kirkpatrick, Y. Huang, C. T. Farrar, K. A. Marijt, M. Datta, Z. Amoozgar, G. Seano, K. Jung, W. S. Kamoun, T. Vardam, M. Snuderl, J. Goveia, S. Chatterjee, A. Batista, A. Muzikansky, C. C. Leow, L. Xu, T. T. Batchelor, D. G. Duda, D. Fukumura, and R. K. Jain
2016. Dual Targeting of Ang-2 and VEGF Receptors Prolongs Survival in Glioblastoma by Normalizing Tumor Vessels and Reprogramming Macrophages. *Proceedings of the National Academy of Sciences of the United States of America*, In press:1–6.

Pettet, G. J., C. P. Please, M. J. Tindall, and D. L. S. Mcelwain
2001. The Migration of Cells in Multicell Tumor Spheroids. *Bulletin of Mathematical Biology*, 63:231–257.

Pitt-Francis, J., P. Pathmanathan, M. O. Bernabeu, R. Bordas, J. Cooper, A. G. Fletcher, G. R. Mirams, P. Murray, J. M. Osborne, A. Walter, S. J. Chapman, A. Garny, I. M. M. van Leeuwen, P. K. Maini, B. Rodríguez, S. L. Waters, J. P. Whiteley, H. M. Byrne, and D. J. Gavaghan
2009. Chaste: A test-driven approach to software development for biological modelling. *Computer Physics Communications*, 180(12):2452–2471.

Pradel, L. P., A. Franke, and C. H. Ries
2018. Effects of IL-10 and Th2 cytokines on human macrophage phenotype and response to CSF1R inhibitor. *Journal of leukocyte biology*, 103:545–558.

Pradel, L. P., C.-H. Ooi, S. Romagnoli, M. A. Cannarile, C. H. Ries, H. Sade, and R. Dominik
2016. Macrophage susceptibility to emactuzumab (RG7155) treatment. *Molecular Cancer Therapeutics*, 15(12):3077–3086.

Pradel, L. P. and C. Ries
2015. ID: 92: Regulation of macrophage susceptibility to anti-CSF-1R antibody Emactuzumab treatment. *Cytokine*, 76(1):82.

Pyonteck, S. M., B. B. Gadea, H.-W. Wang, V. Gocheva, K. E. Hunter, L. H. Tang, and J. A. Joyce
2012. Deficiency of the macrophage growth factor CSF-1 disrupts pancreatic neuroendocrine tumor development. *Oncogene*, 31(11):1459–1467.

Qaiser, T., Y.-W. Tsang, D. Taniyama, N. Sakamoto, K. Nakane, D. Epstein, and N. Rajpoot
2018. Fast and Accurate Tumor Segmentation of Histology Images using Persistent Homology and Deep Convolutional Features. Pp. 1–39.

Qian, B., Y. Deng, J. H. Im, R. J. Muschel, Y. Zou, J. Li, R. A. Lang, and J. W. Pollard
2009. A distinct macrophage population mediates metastatic breast cancer cell extravasation, establishment and growth. *PLoS ONE*, 4(8).

Qian, B.-Z., J. Li, H. Zhang, T. Kitamura, J. Zhang, L. R. Campion, E. A. Kaiser, L. A. Snyder, and J. W. Pollard
2011. CCL2 recruits inflammatory monocytes to facilitate breast-tumour metastasis. *Nature*, 475(7355):222–5.

Qian, B. Z. and J. W. Pollard
2010. Macrophage Diversity Enhances Tumor Progression and Metastasis. *Cell*, 141(1):39–51.

Rademakers, S. E., J. Lok, A. J. V. D. Kogel, J. Bussink, and J. H. A. M. Kaanders
2011. Metabolic markers in relation to hypoxia ; staining patterns and colocalization of pimonidazole , HIF-1a, CAIX, LDH-5, GLUT-1, MCT1 and MCT4 Saskia.

Ren, X. and J. Malik
2003. Learning a Classification Model for Segmentation. *Proceedings of the Ninth IEEE International Conference on Computer Vision - Volume 2*, 1(c):10—-.

Ries, C. H., M. A. Cannarile, S. Hoves, J. Benz, K. Wartha, V. Runza, F. Rey-Giraud, L. P. Pradel, F. Feuerhake, I. Klaman, T. Jones, U. Jucknischke, S. Scheiblich, K. Kaluza, I. H. Gorr, A. Walz, K. Abiraj, P. A. Cassier, A. Sica, C. Gomez-Roca, K. E. de Visser, A. Italiano, C. Le Tourneau, J.-P. Delord, H. Levitsky, J.-Y. Blay, and D. Rüttinger
2014. Targeting tumor-associated macrophages with anti-CSF-1R antibody reveals a strategy for cancer therapy. *Cancer Cell*, 25(6):846–859.

Ripley, B. D.
1976. The Second-Order Analysis of Stationary Point Processes. *Journal of Applied Probability*, 49(2):303–318.

Robertson-Tessi, M., A. El-Kareh, and A. Goriely
2012. A mathematical model of tumor-immune interactions. *Journal of Theoretical Biology*, 294:56–73.

Romero Castro, E., G. Corredor, C. Lu, A. Madabhushi, X. Wang, and V. Velcheti
2018. A watershed and feature-based approach for automated detection of lymphocytes on lung cancer images. *Medical Imaging 2018: Digital Pathology*, (March):26.

Rovida, E. and P. D. Sbarba
2015. Colony-Stimulating Factor-1 Receptor in the Polarization of Macrophages: A Target for Turning Bad to Good Ones? *Journal of Clinical & Cellular Immunology*, 06(06).

Savitzky, A. and M. J. Golay
1964. Smoothing and Differentiation of Data by Simplified Least Squares Procedures. *Analytical Chemistry*, 36(8):1627–1639.

Schindelin, J., C. T. Rueden, M. C. Hiner, and K. W. Eliceiri
2015. The ImageJ ecosystem: An open platform for biomedical image analysis. *Molecular Reproduction and Development*, 82(7-8):518–529.

Scholl, S. M., C. Pallud, F. Beuvon, K. Hacene, E. R. Stanley, L. Rohrschneider, R. Tang, P. Pouillart, and R. Lidereau
1994. Anti-colony-stimulating factor-1 antibody staining in primary breast adenocarcinomas correlates with marked inflammatory cell infiltrates and prognosis. *Journal of the National Cancer Institute*, 86(2):120–126.

Seager, R. J., C. Hajal, F. Spill, R. D. Kamm, and M. H. Zaman
2017. Dynamic interplay between tumour, stroma and immune system can drive or prevent tumour progression. *Convergent Science Physical Oncology*, 3(3):034002.

Shi, J. and J. Malik
2000. Normalized Cuts and Image Segmentation. *Ieee Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905.

Sica, A., M. Erreni, P. Allavena, and C. Porta
2015. Macrophage polarization in pathology. *Cellular and Molecular Life Sciences*, 72(21):4111–4126.

Sica, A., P. Larghi, A. Mancino, L. Rubino, C. Porta, M. G. Totaro, M. Rimoldi, S. K. Biswas, P. Allavena, and A. Mantovani
2008. Macrophage polarization in tumour progression. *Seminars in Cancer Biology*, 18(5):349–355.

Sica, A. and A. Mantovani
2012. Macrophage plasticity and polaization: in vivo veritas. *Journal of Clinical Investigation*, 122(3):787–795.

Sidibe, A., P. Ropraz, S. Jemelin, Y. Emre, M. Poittevin, M. Pocard, P. F. Bradfield, and B. A. Imhof
2018. Angiogenic factor-driven inflammation promotes extravasation of human proangiogenic monocytes to tumours. *Nature Communications*, 9(1).

Sirinukunwattana, K., S. E. Raza, Y. W. Tsang, D. R. Snead, I. A. Cree, and N. M. Rajpoot
2016. Locality Sensitive Deep Learning for Detection and Classification of Nuclei in Routine Colon Cancer Histology Images. *IEEE Transactions on Medical Imaging*, 35(5):1196–1206.

Slaoui, M. and L. Fiette
2011. Histopathology Procedures: From Tissue Sampling to Histopathological Evaluation. In *Drug Safety Evaluation: Methods and Protocols*, J. C. Gautier, ed., volume 691, chapter 4, Pp. 69–82. Humana Press.

Sompayrac, L.
1999. *How The Immune System Works*, 5th edition. Chichester: Wiley Blackwell.

Sousa, S., R. Brion, M. Lintunen, P. Kronqvist, J. Sandholm, J. Mönkkönen, P.-L. Kellokumpu-Lehtinen, S. Lauttia, O. Tynninen, H. Joensuu, D. Heymann, and J. a. Määttä
2015. Human breast cancer cells educate macrophages toward the M2 activation status. *Breast Cancer Research*, 17(1):101.

Spratt, J. A., D. von Fournier, J. S. Spratt, and E. E. Weber
1993. Decelerating growth and human breast cancer. *Cancer*, 71(6):2013–2019.

Stack, E. C., C. Wang, K. A. Roman, and C. C. Hoyt
2014. Multiplexed immunohistochemistry, imaging, and quantitation: A review, with an assessment of Tyramide signal amplification, multispectral imaging and multiplex analysis. *Methods*, 70(1):46–58.

Stanley, E. R., K. L. Berg, D. B. Einstein, P. S. W. Lee, F. J. Pixley, Y. Wang, and Y. G. Yeung
1997. Biology and action of colony-stimulating factor-1. *Molecular Reproduction and Development*, 46(1):4–10.

Starruß, J., W. De Back, L. Brusch, and A. Deutsch
2014. Morpheus: A user-friendly modeling environment for multiscale and multicellular systems biology. *Bioinformatics*, 30(9):1331–1332.

Stockmann, C., D. Schadendorf, R. Klose, and I. Helfrich
2014. The impact of the immune system on tumor: angiogenesis and vascular remodeling. *Frontiers in oncology*, 4(April):69.

Stout, R. D., C. Jiang, B. Matta, I. Tietzel, S. K. Watkins, and J. Suttles
2005. Macrophages Sequentially Change Their Functional Phenotype in Response to Changes in Microenvironmental Influences. *The Journal of Immunology*, 175(1):342–349.

Stout, R. D. and J. Suttles
2004. Functional plasticity of macrophages: reversible adaptation to changing microenvironments. *Journal of leukocyte biology*, 76(3):509–13.

Stoyan, D. and H. Stoyan
1994. *Fractals, Random Shapes and Point Fields: Methods of Geometrical Statistics.* Chichester: Wiley.

Stutz, D., A. Hermans, and B. Leibe
2018. Superpixels: An evaluation of the state-of-the-art. *Computer Vision and Image Understanding*, 166:1–27.

Swat, M. H., G. L. Thomas, J. M. Belmonte, A. Shirinifard, D. Hmeljak, and J. A. Glazier
2012. Multi-Scale Modeling of Tissues Using CompuCell3D. *Methods in Cell Biology*, 110:325–366.

Thompson, K. E. and H. M. Byrne
1999. Modelling the internalization of labelled cells in tumour spheroids. *Bulletin of Mathematical Biology*, 61(4):601–623.

Tin Kam Ho
1995. Random decision forests. *Proceedings of 3rd International Conference on Document Analysis and Recognition*, 1:278–282.

Turner, S. and J. A. Sherratt
2002. Intercellular adhesion and cancer invasion: A discrete simulation using the extended potts model. *Journal of Theoretical Biology*, 216(1):85–100.

Unwin, D. J.
1996. GIS, spatial analysis and spatial statistics. *Progress in Human Geography*, 20(4):540–551.

Upton, G. and I. Cook
1996. *Understanding statistics.* Oxford: Oxford University Press.

Usami, Y., K. Ishida, S. Sato, M. Kishino, M. Kiryu, Y. Ogawa, M. Okura, Y. Fukuda, and S. Toyosawa
2013. Intercellular adhesion molecule-1 (ICAM-1) expression correlates with oral cancer progression and induces macrophage/cancer cell adhesion. *International Journal of Cancer*, 133(3):568–578.

Van Liedekerke, P., M. M. Palm, N. Jagiella, and D. Drasdo
2015. *Simulating tissue mechanics with agent-based models: concepts, perspectives and some novel results*, volume 2. Springer International Publishing.

van Lieshout, M. N. M. and A. Baddeley
1996. A non-parametric measure of spatial interaction in point patterns. *Statistica Neerlandica*, 50(3):344–361.

Van Overmeire, E., B. Stijlemans, F. Heymann, J. Keirsse, Y. Morias, Y. Elkrim, L. Brys, C. Abels, Q. Lahmar, C. Ergen, L. Vereecke, F. Tacke, P. De Baetselier, J. A. Van Ginderachter, and D. Laoui
2015. M-CSF and GM-CSF receptor signaling differentially regulate monocyte maturation and macrophage polarization in the tumor microenvironment. *Cancer Research*, Pp. 35–43.

Varia, M. A., D. P. Calkins-Adams, L. H. Rinker, A. S. Kennedy, D. B. Novotny, W. C. Fowler, and J. A. Raleigh
1998. Pimonidazole: A novel hypoxia marker for complementary study of tumor hypoxia and cell proliferation in cervical carcinoma. *Gynecologic Oncology*, 71(2):270–277.

Vasiliadou, I. and I. Holen
2013. The role of macrophages in bone metastasis. *Journal of Bone Oncology*, 2(4):158–166.

Vedaldi, A. and S. Soatto
2008. Quick shift and kernel methods for mode seeking. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 5305 LNCS(PART 4):705–718.

Veksler, O., Y. Boykov, and P. Mehrani
2010. Superpixels and supervoxels in an energy optimization framework. In *Computer Vision - ECCV 2010*, K. Daniilidis, P. Maragos, and N. Paragios, eds., Pp. 211–224. Springer.

Veta, M., J. P. W. Pluim, P. J. Van Diest, and M. A. Viergever
2014. Breast cancer histopathology image analysis: A review. *IEEE Transactions on Biomedical Engineering*, 61(5):1400–1411.

Ward, J. P. and J. R. King
1997. Mathematical Modelling of avascular-tumour growth. *IMA Journal of Mathematics Applied in Medicine and Biology*, 14:39–69.

Ward, J. P. and J. R. King
1999. Mathematical modelling of avascular-tumour growth. II: Modelling growth saturation. *IMA journal of mathematics applied in medicine and biology*, 16(2):171–211.

Webb, S. D., M. R. Owen, H. M. Byrne, C. Murdoch, and C. E. Lewis
2007. Macrophage-based anti-cancer therapy: Modelling different modes of tumour targeting. *Bulletin of Mathematical Biology*, 69(5):1747–1776.

Webster, J. D. and R. W. Dunstan
2014. Whole-Slide Imaging and Automated Image Analysis: Considerations and Opportunities in the Practice of Pathology. *Veterinary Pathology*, 51(1):211–223.

Wells, D. K., Y. Chuang, L. M. Knapp, D. Brockmann, W. L. Kath, and J. N. Leonard
2015. Spatial and Functional Heterogeneities Shape Collective Behavior of Tumor-Immune Networks. *PLoS Computational Biology*, 11(4):1–22.

Wilkie, K. P. and P. Hahnfeldt
2013. Mathematical models of immune-induced cancer dormancy and the emergence of immune evasion. *Interface Focus - Royal Society Publishing*, (June).

Williams, B. J., D. Bottoms, D. Clark, and D. Treanor
2018a. Future-proofing pathology part 2: Building a business case for digital pathology. *Journal of Clinical Pathology*, Pp. 1–8.

Williams, B. J., D. Bottoms, and D. Treanor
2017. Future-proofing pathology: The case for clinical adoption of digital pathology. *Journal of Clinical Pathology*, 70(12):1010–1018.

Williams, B. J., J. Lee, K. A. Oien, and D. Treanor
2018b. Digital pathology access and usage in the UK: Results from a national survey on behalf of the National Cancer Research Institute's CM-Path initiative. *Journal of Clinical Pathology*, 71(5):463–466.

Wyckoff, J., W. Wang, E. Y. Lin, Y. Wang, F. Pixley, E. R. Stanley, T. Graf, J. W. Pollard, J. Segall, and J. Condeelis
2004. A paracrine loop between tumor cells and macrophages is required for tumor cell migration in mammary tumors. Pp. 7022–7029.

Wyckoff, J. B., Y. Wang, E. Y. Lin, J. F. Li, S. Goswami, E. R. Stanley, J. E. Segall, J. W. Pollard, and J. Condeelis
2007. Direct visualization of macrophage-assisted tumor cell intravasation in mammary tumors. *Cancer Research*, 67(6):2649–2656.

Wynn, T. A., A. Chawla, and J. W. Pollard
2013. Macrophage biology in development, homeostasis and disease. *Nature*, 496(7446):445–55.

Wynn, T. A. and K. M. Vannella
2016. Macrophages in Tissue Repair, Regeneration, and Fibrosis. *Immunity*, 44(3):450–462.

Yuan, Y.
2014. Modelling the spatial heterogeneity and molecular correlates of lymphocytic infiltration in triple-negative breast cancer. *Journal of The Royal Society Interface*, 12(103).

Yuan, Z.-Y., R.-Z. Luo, R.-J. Peng, S.-S. Wang, and C. Xue
2014. High infiltration of tumor-associated macrophages in triple-negative breast cancer is associated with a higher risk of distant metastasis. *OncoTargets and therapy*, 7:1475–80.

Zhang, J., M. Marszałek, S. Lazebnik, and C. Schmid
2007. Local features and kernels for classification of texture and object categories: A comprehensive study. *International Journal of Computer Vision*, 73(2):213–238.

Zheng, Y., J. Bao, Q. Zhao, T. Zhou, and X. Sun
2018. A Spatio-temporal Model of Macrophage-mediated Drug Resistance in Glioma Immunotherapy. *Molecular Cancer Therapeutics*, P. molcanther.0634.2017.

Zhu, Q., S. Avidan, M.-C. Yeh, and K.-T. Cheng
2006. Fast Human Detection Using a Cascade of Histograms of Oriented Gradients. *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, Pp. 1491 – 1498.

Appendix A

Defining model boundaries via α-shapes

In this Section we define α-shapes, which provide a means of defining the boundary of a set of points by generalising the concept of a convex hull to permit concave boundaries (Edelsbrunner et al., 1983). Although there are several different ways of defining an α-shape, they can be most intuitively understood by considering spheres (or in 2D, discs) of radius α. If it is possible to place the edge of the sphere on a point such that the sphere contains no other points in the set, then that point is a member of the α-shape. This is shown in Figure A.1, where the α-shape is calculated for a 2D point set for two different values of α. For each point in the α-shape (coloured), a disk is shown which contains no other points from the set. If $\alpha = \infty$ then the disk has infinite radius of curvature and corresponds to the half-plane - in this case, the α-shape and the usual convex hull are identical.

As Figure A.1 demonstrates, as $\alpha \to 0$ more points will be included in the α-shape until the shape contains all points in the set. In order to capture our intuition of the "boundary" of the set, the value of α must be chosen carefully. We choose α to be the equilibrium cell radius, which ensures that the boundary of the spheroid represents the set of cells which a cell approaching the spheroid from outside could

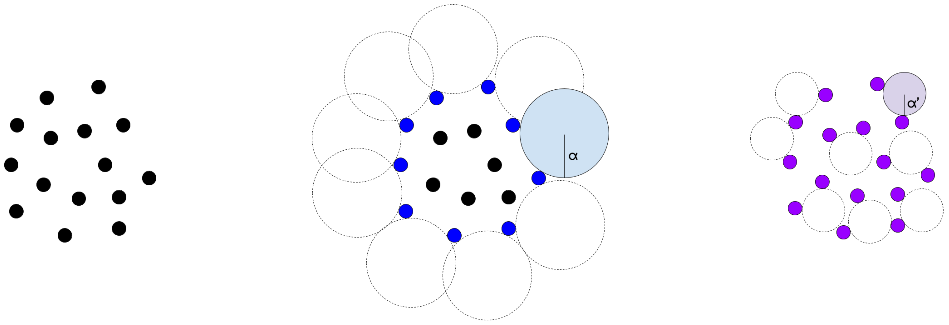

Figure A.1: Diagram visualising an interpretation of an α-shape. The α-shape, or α-hull, of a set of points can be conceptualised as those points which come into contact with the edge of ball of radius α which does not contain any other points. For the set of points on the left, we picture two α-shapes for two different values of α. In the central image, the α-hull is shown in blue and accurately captures what may be intuitively understood as the "edge" of the point set. The image on the right shows the α-hull when a smaller value of α, α', is used. In this case every point in the set is included in the α-shape, illustrating that the value of α must be selected with caution. In our model, we choose α to be the radius of a cell at equilibrium, meaning that the α-shape can be understood as referring to those cells which an external cell could readily "see".

touch without deformation.

Appendix B

Agent-Based Model Verification

Model verification is the process of ensuring that a particular software implementation captures an underlying model (Cooper et al., 2015). In this Appendix, we briefly demonstrate the correct implementation of parts of our code against theoretically predicted benchmarks.

B.1 Brownian Motion

In this Section we test the implementation of the Brownian motion force described in Section 4.3.4. For a diffusion constant D, a particle moving randomly under Brownian motion in d dimensions will on average be found a distance of $\sqrt{2dDt}$ away from its starting position after a time t. We test our implementation of random force in two dimensions against this metric for different values of D. Figure B.1 shows the average displacement of a single macrophage moving under random motion only in our simulations. Each point in Figure B.1 is the average displacement of 100 simulations. The solid lines are the benchmarks predicted by theory.

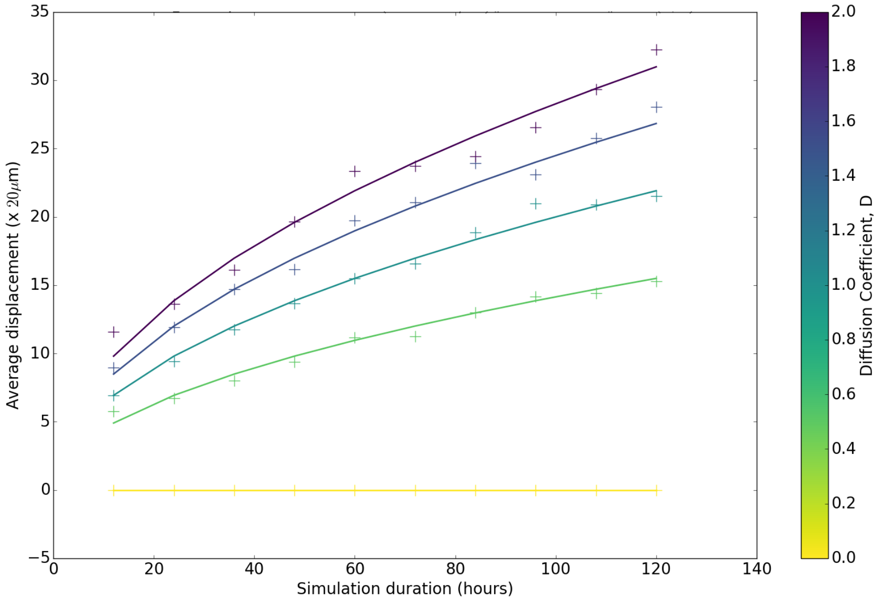

Figure B.1: Average displacement of a macrophage moving solely under Brownian motion, averaged over 100 simulations each for different values of D. The solid lines show the theoretically predicted displacement as D varies.

B.2 Chemotaxis

In this Section, we test the implementation of the chemotaxis model described in Section 5.2.3. To test our implementation, we place a macrophage on a known gradient of CSF-1, shown in Figure B.2. For a constant gradient such as that shown in Figure B.2, we expect the time the macrophage takes to cross the domain to scale linearly with the length of the domain. Figure B.3 shows the comparison between our model implementation and the predicted end time for different values of the chemotactic sensitivity coefficient, χ. The results indicate that the chemotactic force in our model is correctly implemented.

By solving the CSF-1 gradient numerically, it is possible that non-linear CSF-1 gradients will cause macrophages to deviate slightly from the predicted end time. This is because the CSF-1 gradient is solved on a fixed mesh whereas the equations of motion for the macrophages are ‘off-lattice’. If a macrophage is between two mesh points, then the concentration of CSF-1 at the macrophage is linearly interpolated

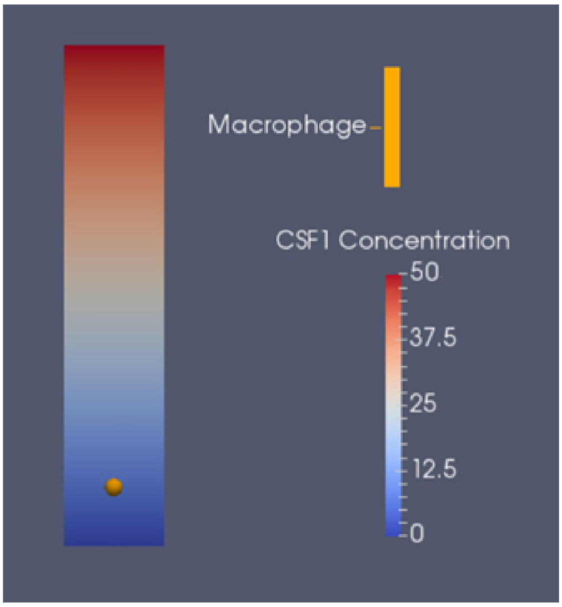

Figure B.2: Simulation set up for the runway test for chemotaxis implementation. A macrophage is placed at one end of a runway of fixed length, and a chemotactic gradient is prescribed across the runway.

using values from those mesh points. Where the CSF-1 gradient is non-linear, this will introduce a small numerical error. As this error is small in relation to the forces acting on the macrophage, including those due to chemotaxis, we neglect it in our calculations.

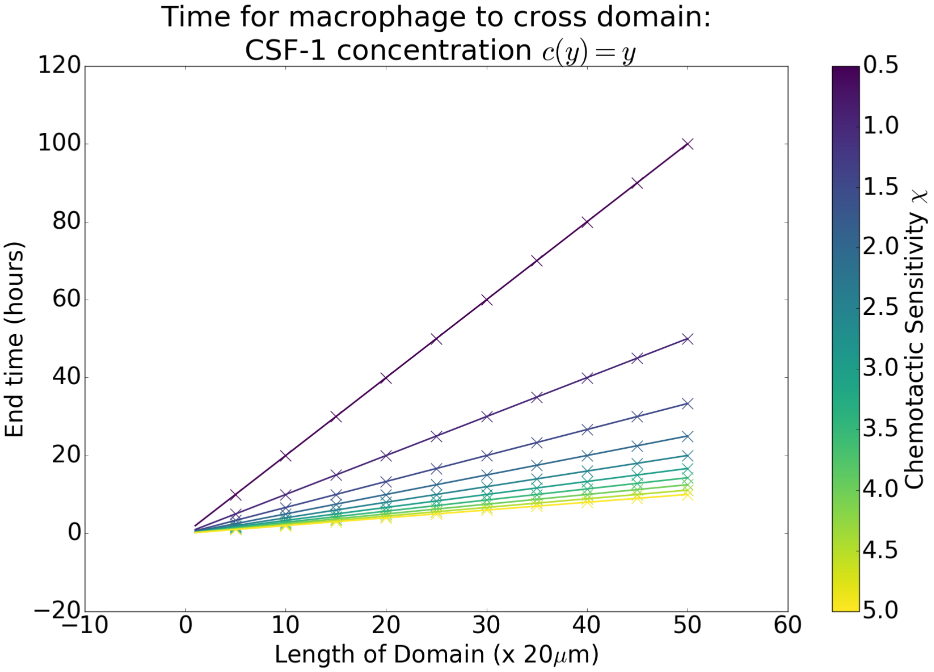

Figure B.3: Time taken for macrophages to cross the runway in the chemotaxis implementation test. Crosses show the time taken for macrophages to cross the runway for different chemotactic sensitivity coefficients and for different lengths of the runway. Solid lines show the predicted time taken for the macrophage to cross the runway.

Appendix C

Default parameters for agent-based modelling

In the table that follows, we state the parameter values that were used for the modelling in Chapters 4 and 5. Where a range is given, the values indicate the upper and lower bounds used in parameter sweeps.

The values used were chosen to generate physically realistic results, and do not represent experimentally measured values.

Symbol	**Parameter**	**Value**	**Units**
dt	Timestep	1/120	Hours
R_{int}	Radius of interaction	1.5	Cell diameter
ω_{∞}	Oxygen boundary value*	1.0	Dimensionless
ω_{q}	Quiescence threshold*	0.3 - 0.7	Dimensionless
ω_{h}	Hypoxia threshold*	0.1 - 0.7	Dimensionless
τ_{ρ}	Critical hypoxic duration	8 - 16	hours
τ_{apop}	Apoptosis duration	48	hours
ν	Damping coefficient	1	μg hours^{-1}
μ	Spring constant	45.0	μg (Cell diameter)$^{-1}$ hours^{-2}

Symbol	**Parameter**	**Value**	**Units**
$s_{i,j}$	Resting spring length	1.0	Cell diameter
D	Random motility coefficient	0.01	(Cell diameter)^2 hours^{-1}
η	Average cell cycle length	8 - 32	hours
τ_η	Cell cycle duration	0.75η- 1.25η	hours
κ	Oxygen consumption rate	0.03	hours^{-1}
D_ω	Oxygen diffusion coefficient	1	(Cell diameter)^2 second^{-1}
β	Surface tension coefficient	5	μg hours^{-2}
D_c	CSF-1 diffusion coefficient	1	(Cell diameter)^2 second^{-1}
κ_c	CSF-1 production rate	0.03	hours^{-1}
L	Size of bounding domain edge	200	Cell diameter
c_∞	CSF-1 boundary value*	0.0	Dimensionless

Table C.1: Summary of default parameter values used to generate simulations of tumour spheroids.

*Scaled with regard to the oxygen concentration on the spheroid boundary.

Appendix D

Variance/mean ratio curves of the RITA dataset

D.1 Examples of VMR curves for the RITA dataset

Figure D.1 shows examples of the VMR curves which were used as a feature for clustering examples from the RITA dataset in Chapter 3.

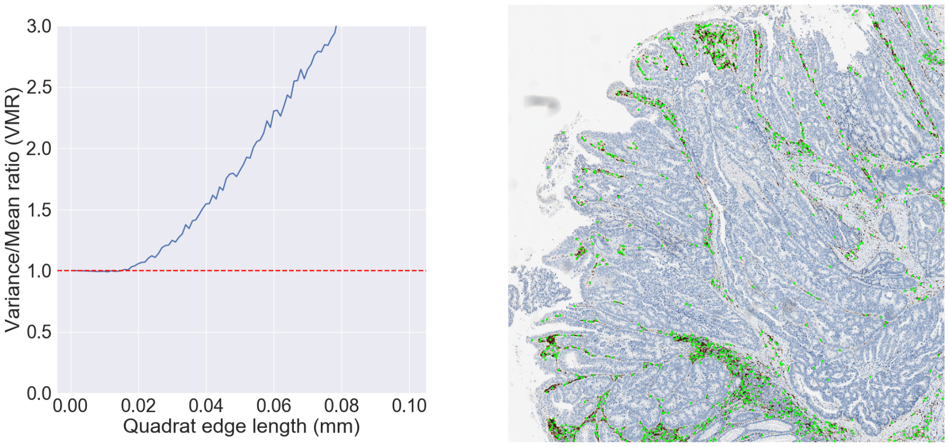

(a) Variance/mean ratio curve showing clustering of macrophages on length scales greater than 0.02 mm

(b) Image used to generate the VMR curve in Figure D.1a.

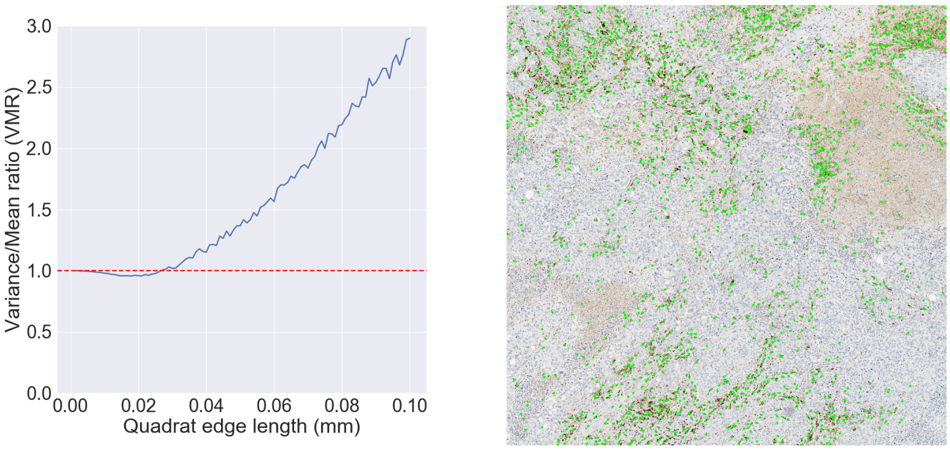

(c) Variance/mean ratio curve showing clustering of macrophages on length scales greater than 0.03 mm

(d) Image used to generate the VMR curve in Figure D.1c.

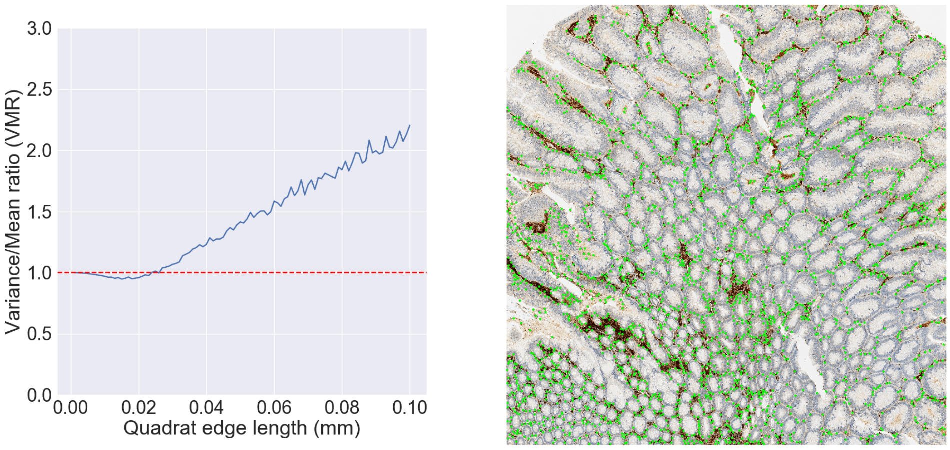

(e) Variance/mean ratio curve showing clustering of macrophages on length scales greater than 0.03 mm

(f) Image used to generate the VMR curve in Figure D.1e.

Figure D.1: Characteristic behaviours of the variance/mean ratio curve.

Appendix E

Agent-based simulations of *in vivo* geometries

This Appendix contains snapshots from agent-based models simulating *in vivo* geometries, described in Section 5.5. We present snapshots at a range of intervals from simulations conducted in four different tumour geometries.

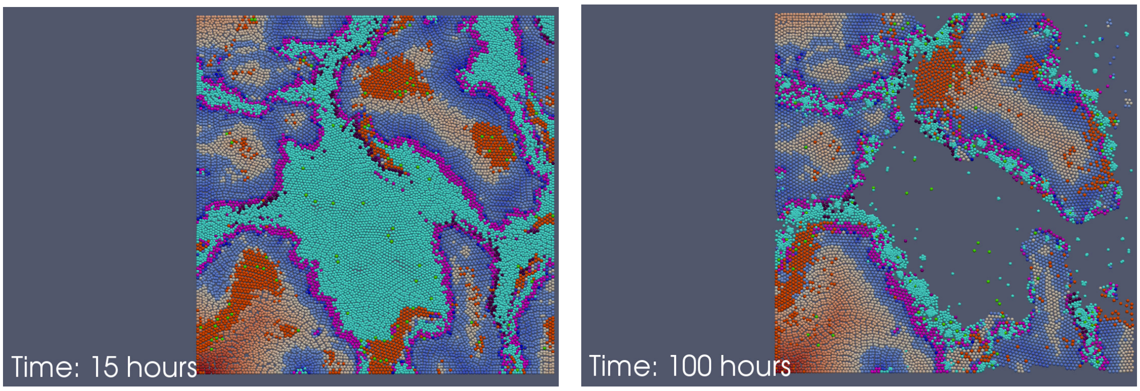

Figure E.1: Snapshots of our agent-based model initalised using oxygen environments from histological images

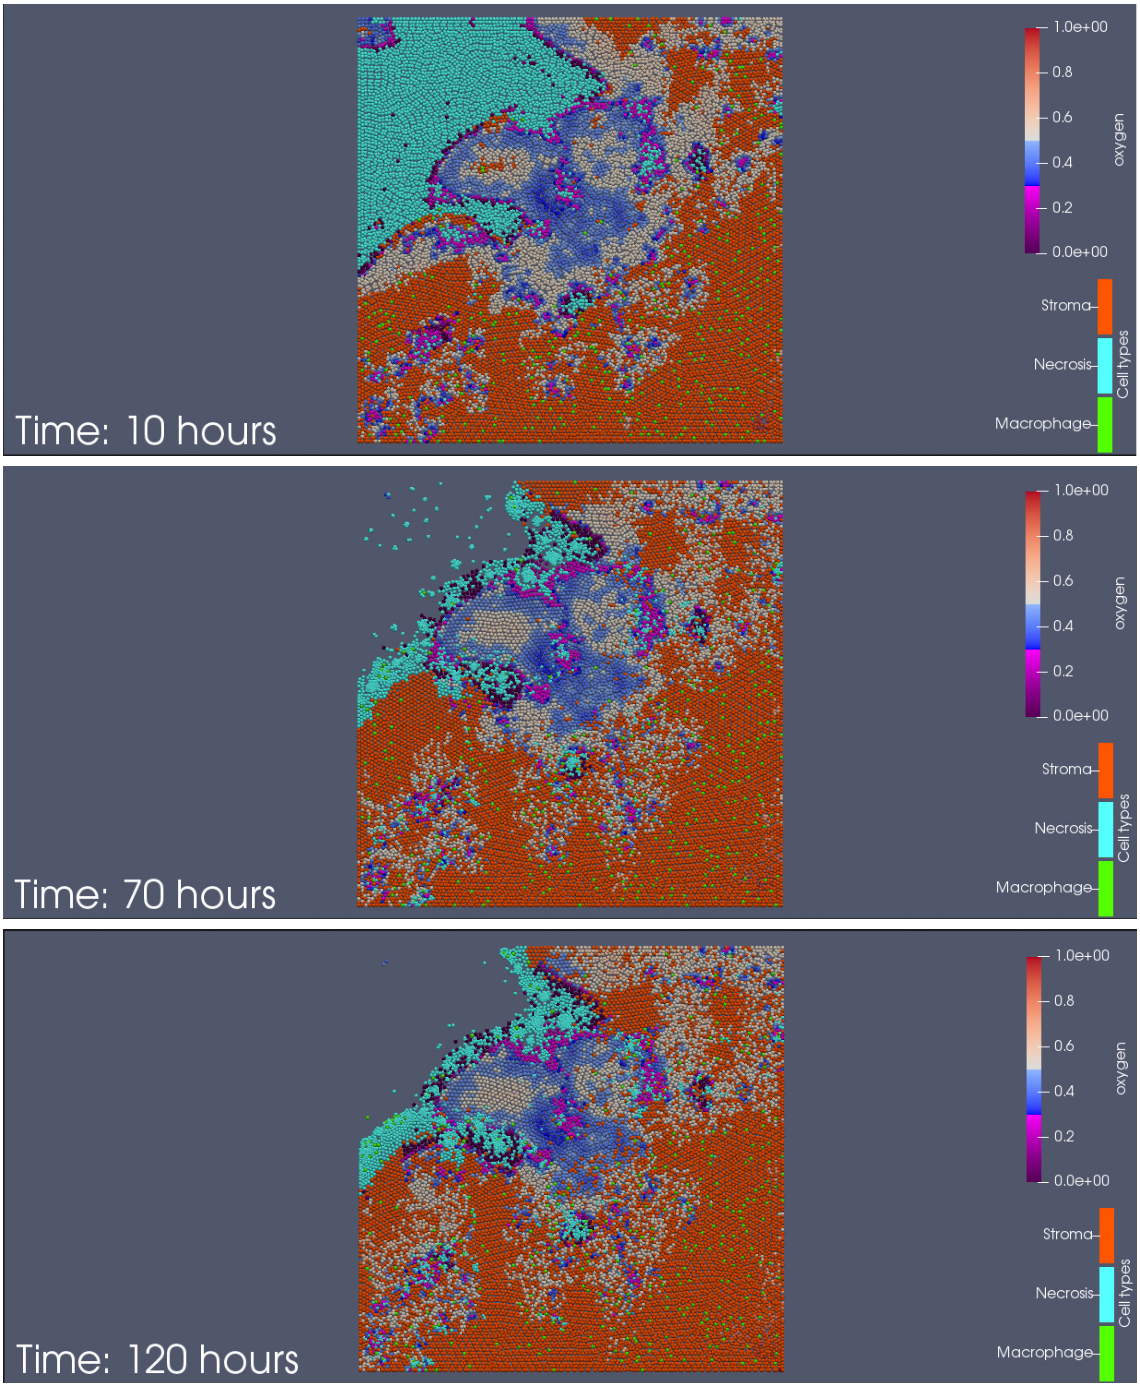

Figure E.2: Snapshots of our agent-based model initalised using oxygen environments from histological images

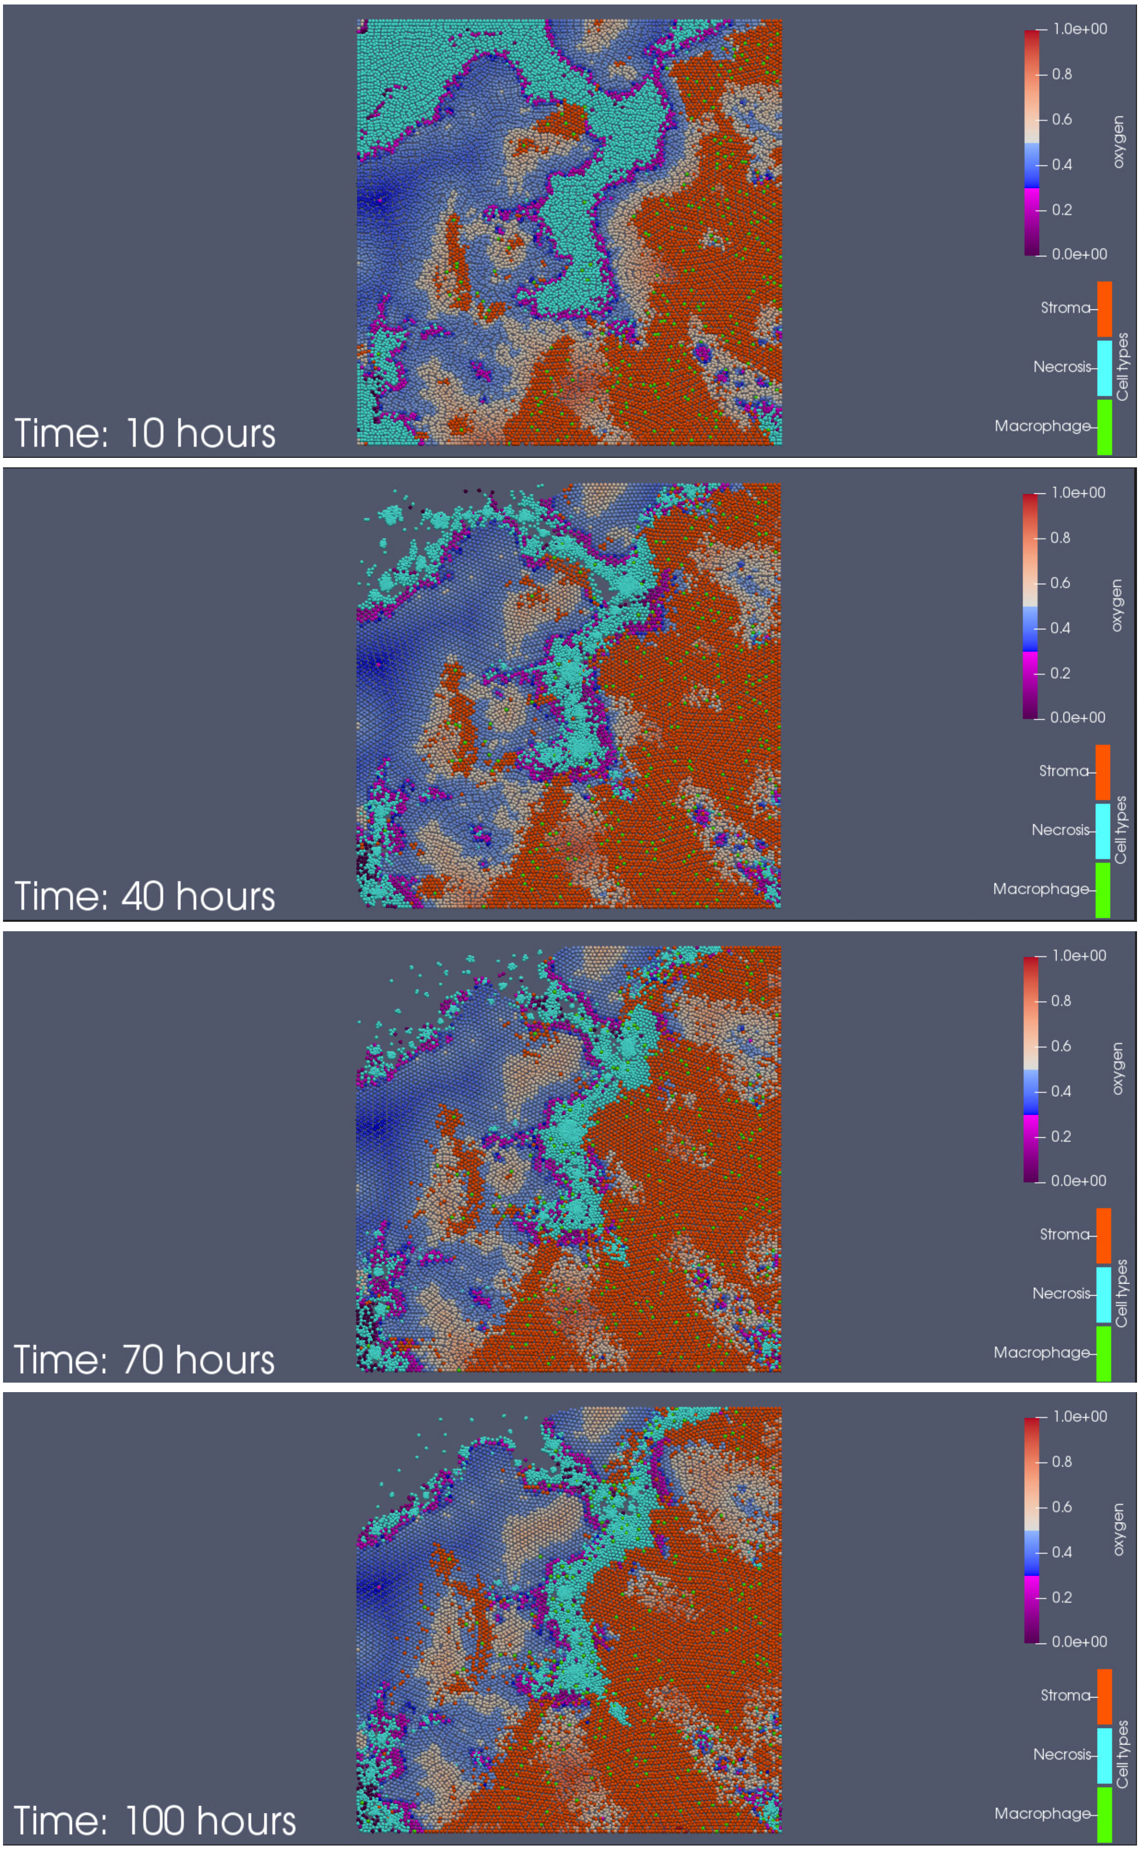

Figure E.3: Snapshots of our agent-based model initalised using oxygen environments from histological images

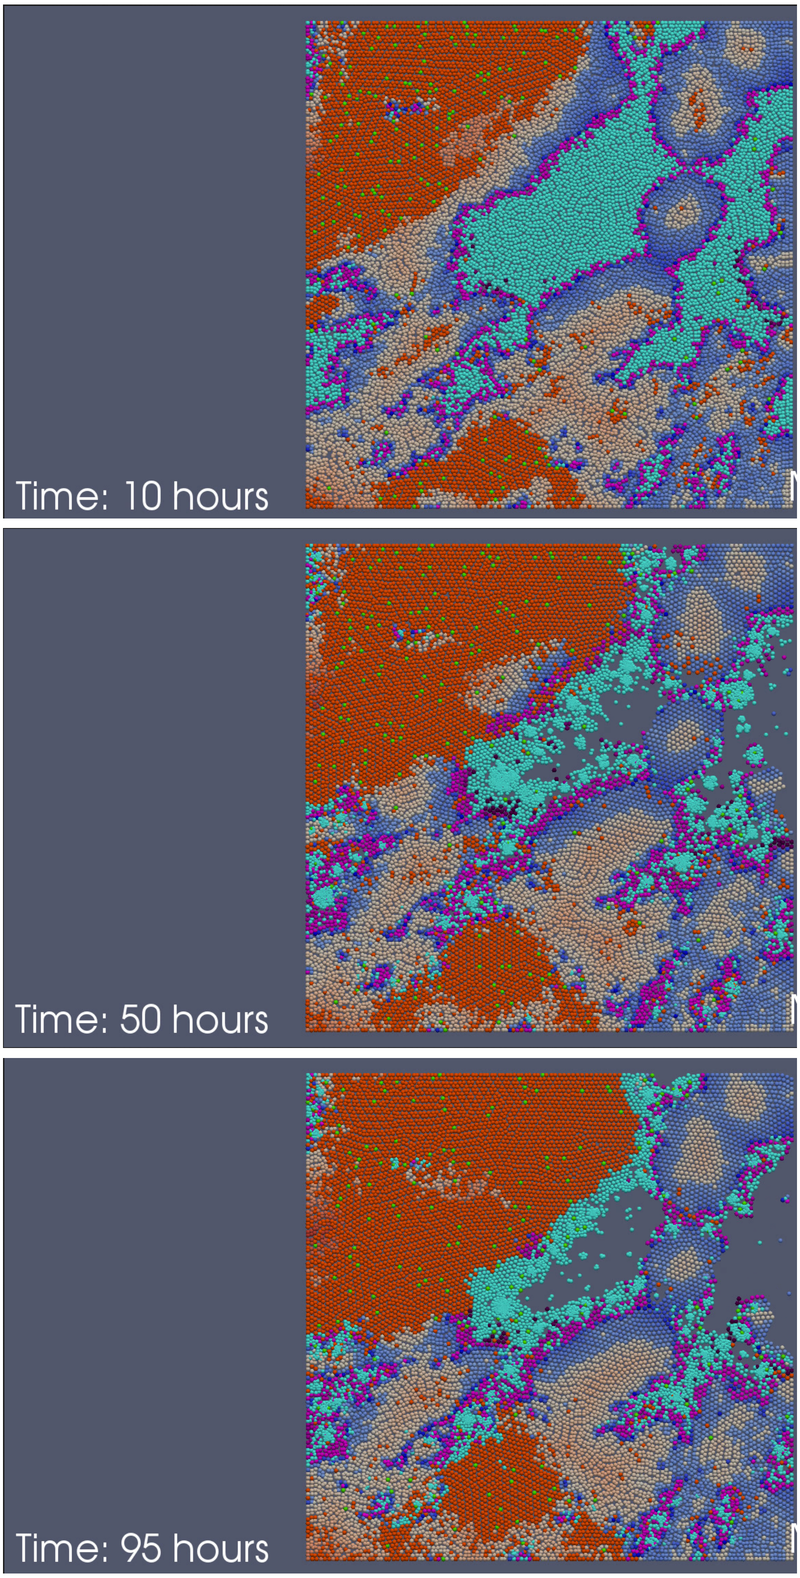

Figure E.4: Snapshots of our agent-based model initalised using oxygen environments from histological images