MANCHESTER
1824
The University of Manchester

Applications of Agent-Based Models in Cancer Research and Epidemiological Simulations

A thesis submitted to the University of Manchester for the degree of
Doctor of Philosophy
in the Faculty of Science and Engineering

2021

Dario Panada
School of Computer Science

Contents

Word Count: 51,211

List of figures

List of tables

List of publications

Panada, Dario, and Bijan Parsia. “PanaXea: A Framework for the Development and Parametrization of Agent-Based Models.” Proceedings of the 33rd European Modeling Simulation Symposium (EMSS 2021), September 2021

Panada, Dario, and Bijan Parsia. “The Trap of 2D in Artificial Models of Tumours. The Case for 3D in-Silico Simulations.” Proceedings of the 11th International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH 2021), July 2021.

Panada, Dario, and Bijan Parsia. “A Framework for Parameter Space Exploration in Multi-Scale Models” The 22nd International Workshop on Multi-Agent-Based Simulation (MABS2021), May 2021. (Oral Presentation)

Panada, Dario, et al. “In-Silico Modelling of Phenotypic Switching in Tumours: Investigating Potentials for Non-Invasive Therapies.” Proceedings of the 2021 IEEE the 9th International Conference on Bioinformatics and Computational Biology (ICBCB 2021), Feb. 2021. - Best Presentation Award

Abstract

Agent-based models have attracted the interest of cancer researchers thanks to their ability to simulate months worth of tumour growth in a matter of hours, significantly reduced costs compared to laboratory studies and ease with which multiple therapeutic approaches may be simulated and compared. In addition, agent-based models also provide granular insights into phenomena that may be hard to observe in live organisms, such as intra-cellular mechanics, gene expression rates and phenotypic properties of individual cells.

We developed an agent-based model simulating the growth of tumours in organisms. It aimed to investigate the potential of exploiting phenotypic switching in cancer cells for therapeutic purposes to develop non-invasive therapies. Our simulations suggest that most cells in the tumour body adopt the Warburg phenotype.

The model we developed included significantly more elements of cancer biology than others present in the literature. Critical aspects included angiogenesis, intra-cellular gene networks, hypoxia heterogeneity in the tumour mass, and solubles diffusion. Further, improvements on state of the art concerning model parametrisation and validation were also made, adding to the reliability and robustness of the model. Specifically, growth curves produced by our simulations were compared to those obtained from literature through an objective error function, a level of validation that is very rarely found in the literature.

In the context of tumour modelling, we also considered whether 2D and 3D simulations, two prevalent approaches, can be regarded as equivalent. We simulated the growth of a sample tumour in both dimensionalities, with results suggesting that 2D simulations do not capture vital spatial aspects and lead to erroneous results. While 3D simulations are often cited as significantly more computationally expensive, we also found that this is not the case. They can, in fact, be run on standard laptops or middle-to-small cloud instances.

Our bespoke framework, PanaXea, can rapidly develop simulations in various domains. For example, in our cases, the study of biological processes and epidemiology. The flexibility of our framework is evidenced, for example, by the ease with which agents may be used to represent humans in the case of the COVID model or biological cells in the case of the Warburg investigation.

The framework integrates with open source libraries to facilitate the visualisation and analysis of results and is made available in an open-source form. We present two smaller case studies of its use and a more comprehensive sample application related to epidemiological modelling in the context of COVID19 spread.

Declaration of Originality

A portion of my literature review includes original work initially submitted as part of my thesis for the degree of Master of Research (MRes) in Computer Science at The University of Manchester. The research area of my MRes program was the same as my PhD, and therefore several key papers in the area occur in both reviews of existing literature.

I hereby confirm that no other portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or institute of learning.

Acknowledgements

I would firstly like to thank my supervisor, Professor Parsia, for his invaluable guidance during my doctoral studies. His expertise and support allowed me to mature as a researcher and successfully develop my research. I would further like to thank him for his continuous encouragement and enthusiasm, even in the face of initial difficulties and rejection of work submitted for publication.

I would also like to acknowledge my first supervisor, Professor King, for his guidance during the first year of my PhD and his help in laying the foundations of my project and establishing the direction of my research.

Finally, I would like to extend my thanks to the academic and administrative staff in the Department of Computer Science at The University of Manchester. Despite challenges during my PhD years, including the COVID19 pandemic, they have always been supportive and understanding, always aiming to support students to the best of their abilities.

Chapter 1

Introduction

Agent-based simulations have historically attracted the attention of researchers in the modelling community for their potential to offer insights into complex phenomena. Applications range from the simulation of biological phenomena to urban planning, logistics and containment of infectious diseases, among others. Among the most appreciated strengths of agent-based models, we would like to mention their resolution: the possibility of collecting information about individual agents, the option of adding new features to existing models, and the flexibility with which multiple 'what-if' scenarios may be configured and observed.

Our research provides the following contributions to the research domain. Firstly, a comprehensive model of tumour development specifically designed to study phenotypic switching in cancer cells and potentials for non-invasive therapies. This is accompanied by theoretical protocols and software implementation concerning model parameterization and evaluation strategies. Secondly, a study comparing the simulation of tumour growth in 2D and 3D, detailing the dangers and pitfalls of the former. Then, a novel framework for the development of agent-based models, PanaXea. This aims to make the development of agent-based models accessible to researchers who do not have a background in computer science and easily integrate with major cloud vendors to allow the scaling of simulations. Finally, a further application of our framework, PanaXea, to the epidemiological modelling in the context of the COVID19 pandemic.

This thesis is primarily concerned with applying agent-based models to cancer research. Specifically, exploring phenotypic switching in cancer cells and the therapeutic implications of different proportions of the tumour mass acquiring the Warburg Phenotype.

The Warburg Phenotype [1, 2], named after Nobel laureate Otto Heinrich Warburg, describes a set of behavioural and physiological changes displayed by cancer cells. Specifically, the phenomenon whereby cancer cells convert glucose to lactate rather than undergoing aerobic respiration in mitochondria[1] even in the presence of oxygen.[2] While the causes of this are not fully understood, it has been suggested that byproducts of anaerobic respiration may provide molecules and other materials required to sustain the continuous proliferation of cancer cells. Further, the release of lactate and the consequent acidification of the environment weakens surrounding tissues and may aid tumour expansion.

Anaerobic respiration requires more glucose molecules to provide the same amount of energy[3] as its anaerobic counterpart, resulting in cancer cells requiring more glucose than their non-cancer counterparts. In turn, this has been suggested to increase the amount of glucose membrane transporters in cancer cells [3, 4]. Hence, we would

[1] Mitochondria - the 'powerhouse' of the cell. The part of the cell responsible for using available sugar to 'recharge' the energy stores of the cell.

[2] This refers to a change in metabolism where cells switch from an oxygen-based to a sugar/glucose based one.

[3] ATP - Adenosine Triphospate. The molecule in cells used to store store and carry readily-available energy. Ie: Energy which can be accessed without needing long or complex biochemical pathways.

like to investigate whether this difference between cancer and non-cancer cells could be exploited using glycoconjugate drugs (Ie: Drugs bound to glucose molecules) to target tumour cells selectively.

To this end, a comprehensive model was developed, including key elements of cancer biology such as angiogenesis[4], diffusion of solubles and intra-cellular gene networks. In addition, simulations of tumour growth were run hypothesizing different proportions of cancer cells acquiring the Warburg phenotype and resulting growth curves compared to those observed in laboratory studies. Results suggest that a majority of cancer cells acquire the Warburg Phenotype and that, therefore non-invasive therapies based on glycoconjugate may form a viable therapeutic approach.

In developing the aforementioned model, we also aimed to improve the state of the art concerning model parameterization and validation. Publications within our research domain often provide a single set of parameter values, if reported. Similarly, model validation is then commonly limited to 'eyeballing' of growth curves lacking any systematic comparison to results obtained from wetware experiments. We used established model parametrization techniques including grid search and random search to explore our model's sensitivity to different parameters and understand which regions in the space of parameter values would result in the growth curves produced by our model to align to those observed in the laboratory studies. We further developed a fitness function to objectively score our model's performance. Finally, we used decision tree unravelling to understand which parameters our model is most sensitive to. Having so stress-tested our model provides us with confidence that results were not merely due to a fortuitous combination of parameter values but rather that they are descriptive of the underlying biological phenomena.

During our process of literature summary, it came to our attention that models of tumour growth had been previously developed both in 2D and 3D environments. While the former may be advantageous as it is easier to visualize, less computationally expensive and perhaps easier to implement, it is also true that the latter more closely resembles actual tumour growth in humans. Therefore, to address whether differences between 2D and 3D models are significant, we reproduced a set of simulations in both environments, devising a protocol to translate parameter values, simulation setup and other properties of the model between the two setups. Our results suggest that 2D simulations do not adequately capture the complexity of tumour biology, therefore producing erroneous results. Specifically, in our case, we observed how regions of hypoxia[5] observed in 3D did not develop in the derived 2D simulation. Thus, the expression of hypoxia-inducible factors in cancer cells did not occur in the 2D simulation, leading to significant divergences in growth rates.

We have contributed a comprehensive model of tumour growth, enabling researchers

[4]The process by which the existing blood network develops and grows by elongating existing vessels and sprouting new ones.

[5]Conditions of low oxygen concentration.

to investigate significant problems in cancer research. We have also addressed the question of parametrization and evaluation in agent-based models, providing examples of how rigorous techniques concerning these matters may strengthen and increase confidence in the results. Finally, specific to the simulation domain of cancer biology, we have investigated differences between two common approaches: 2D and 3D modelling. Here, results suggest that 2D simulations may easily be unsuitable for research.

Our simulations were implemented using a framework we developed which is now open-source: PanaXea. In developing the framework, we adopted the following guiding principles. Firstly, it should be developed in a non-proprietary and widely adopted language. Our choice was for Python, one of the most popular languages with a smooth learning curve and a wealth of available libraries. Secondly, it should not be strongly opinionated and enforce minimal constraints on the developers. Thus, our framework enforces a minimal hierarchy allowing developers to design simulations according to their needs. Finally, it should easily integrate with major cloud vendors, allowing simulations to scale and bypassing the constraints of individual hardware. The framework was used for multiple undergraduate and master student projects at The University of Manchester, and a technical system description paper was accepted for publication at EMSS 2021.

PanaXea was further used in a project to investigate the epidemiological spread of COVID19. Specifically, it was used to explore the efficacy of progressively stronger containment measures and deconfinement strategies once the disease had been brought back under control. The framework simulated the spread of COVID both at an urban level, simulating a city the size of Manchester (UK), and at a national level, simulating the whole of the UK. Results suggested that weaker forms of lockdown are ineffective at containing the spread of the disease. Further, once lockdowns have successfully reduced infection rates, deconfinement should be carefully planned and staged to avoid successive waves being worse than the original. This project was accepted for presentation at the MABS 2021 workshop.

1.1 Thesis Organization

The research undertaken for this PhD encompasses several themes, which led to the publications listed earlier. Such work laid the foundations of this thesis. A summary of each chapter, major contributions and related publications is presented in Table 1.1.

Chapter	Content	Relevant Publications	Theme
Introduction	Introduction to cancer and cancer modelling, presentation of problems and relevance, summary of experiments and results, overview of thesis structure.		
Background	Provides an overview of both tumour biology and current modelling practices.		
Literature Review	Overview of existing literature in the domain of cancer modelling, strengths and limitations of current work.	ICBCB2021, MABS2021, SIMULTECH2021	
Current Challnges	Summary of major challenges and obstacles faced by researchers when modelling natural phenomena.		
Modelling Strategies	Description of PanaXea as a framework and justification of why it was developed. Point to how it is versatile and can be applied to different domains (epidemiological modelling, cancer...) In the context of cancer modelling, explain why 3D is better than 2D and common pitfalls that may occur as a result of using 2D.	EMSS2021, MABS2021, SIMULTECH2021	Modelling best-practices and frameworks
Warburg Investigation	Introduction to the Warburg effect, its significance for the development of non-invasive therapies, novelty of hybrid continuous/discrete modelling strategies and the interplay between the two, results and future work.	ICBCB2021	Cancer Modelling and Practical Investigative Applications
Decision Tree Unravelling	Use of decision tree unravelling to find regions of suitability in the parameter space of the Warburg model.	ICBCB2021	Modelling best-practices and frameworks
Epidemiological Modelling	Replication of COVID work, the adaptation of PanaXea to a network model for cost optimization, use to predict effects of lockdown and deconfinement.		Cancer Modelling and Practical Investigative Applications
Conclusion & Future Work	Summary of results, reflection on the feedback of independent peer-reviewers, self-assessment and outlining of aspects where the work could be expanded and improved upon.		

Table 1.1: Outline of thesis structure, including a summary of each chapter's content

Chapter 2

Background

2.1 Biological Background

2.1.1 Cancer

Cancer is a disease responsible for over 10 million deaths per year worldwide. Cancer is acquired by an individual when a lineage of cells acquires sufficient genetic mutations. The normal intracellular mechanisms that govern and regulate cell division fail or become otherwise inhibited, leading to uncontrolled proliferation. Sustaining proliferative signalling, evading growth suppressors, resisting cell death and enabling replicative immortality have been recognized as some of the defining hallmarks of cancer; see Hanahan et al. [5]

Cancer damages the host tissues in a variety of manners, including competing for access to nutrients, secreting excess waste and toxic by-products in the extracellular matrix, causing the acidification of the local environment and outright invading the physical space occupied by non-cancer cells. Eventually, if untreated, this compromises the functioning of organs and leads to the patient's death.

The most common types of cancer are bladder, breast, colon and rectal, kidney, and leukaemia, although many more exist and in principle, most tissues could give rise to forms of tumour. Cancer originates from a small sub-population of cells acquiring the required mutations to develop the potential for uncontrolled growth. In the initial stages, the tumour is known as a 'Hyperplasia', which indicates a group of cells dividing more rapidly than expected. Cells in this population will eventually begin mutating physically and becoming larger, with this stage of tumour development known as 'Dysplasia'. This eventually develops into in-situ cancer, which is confined to its site of origin, and finally to a malignant tumour that possesses the ability to invade neighbouring tissues and form metastases.

Metastases refer to secondary tumours formed as a result of cells detaching from the original cancer mass and entering the bloodstream. Eventually, these may bind to other organs and tissues and begin proliferating, giving rise to new tumour masses. As we will see in Section, 2.1.3, metastases limit the efficacy of anti-cancer therapies and are often associated with poor prognosis.

To date, the causes of cancer are not fully understood. Several genetic and behaviour risk factors have been identified, which are linked to an increased likelihood of developing cancer. Age is a primary risk factor, with older people more likely to develop cancer than younger ones. Substances that promote (or strongly correlate to) cancer development are known as carcinogens. These include the tobacco found in cigarettes and, indeed, smoking has been identified as one of the leading causes of cancer. Alcohol consumption and obesity also increase the risk of developing cancer, as does prolonged exposure to different forms of radiation. Finally, one's chance of developing cancer is also affected by the history of cancer in their family and genetic predispo-

sition. In most cases, the co-occurrence of multiple risk factors eventually facilitates the development of cancer in an individual.

Cancer Metabolism and the Warburg Effect

Cell metabolism refers to a set of chemical processes whereby cells replenish their energy reserves and produce building blocks needed for cell repair and growth. In humans, the first step is called glycolysis, which refers to the breakdown of glucose into two pyruvate molecules.

Under normal oxygen concentration conditions, cells then undergo a set of oxygen-reliant metabolic pathways commonly referred to as cellular respiration. In this scenario, pyruvate and oxygen are processed to release energy, with carbon dioxide and water molecules produced as by-products. The energy produced is used to 'recharge' adenosine diphosphate (ADP) molecules into adenosine triphosphate (ATP), which act as short-term energy storage units for ready use in cellular activities.

When the oxygen concentration is particularly low, for example, during periods of intense physical activity, cells may activate an alternative form of anaerobic metabolism which is not reliant on oxygen. Such an approach is less efficient in energy production, resulting in less energy being stored and made available for cellular processes per unit of glucose than its aerobic counterpart. Further, lactate is also produced as a by-product of this reaction. Lactate is an acidic waste product that collects in the inter-cellular space and must eventually be removed by the organism.

However, cancer cells have been observed to undergo anaerobic respiration even in the presence of sufficiently high oxygen concentrations to support cellular respiration. This phenomenon is commonly referred to as the Warburg Effect. The reasons behind this are, to date, not fully explained. Damages to the mitochondria (the cell's organelles responsible for cellular respiration) have been suggested as an explanation. However, other studies have reported these as functioning normally and reported aerobic and anaerobic activities co-occurrence. Alternative explanations indicate that cancer cells may have a higher tolerance to acidic environments. Therefore, releasing lactate in the inter-cellular space may facilitate the invasion process by selectively targeting and damaging healthy tissues. Another possible explanation is that cancer cells may benefit from the over-production of NADH molecules, a co-enzyme central to cell metabolism, to sustain continuous growth and proliferation.

Many of the alterations observed in the metabolic pathways of cancer cells may become targets for cancer therapy. In our Warburg Investigation (see Chapter 6), we will be focusing on the viability and potentials of selective targeting of cancer cells through glycoconjugate drugs. Therefore, the therapeutic approach motivating our investigation relies on exploiting the over-expression of glucose membrane transporters in cancer cells.

2.1.2 Vascularization and the Necrotic Core

Healthy tissues will generally feature a careful balance of vessels and cells, ensuring proper amounts of oxygen and nutrients can be delivered while waste products collected. However, such a balance is disrupted in tumour bodies, where uncontrolled proliferation leads to the existing network of blood vessels being insufficient to cater to the needs of the growing population.

This leads to the formation of gradients of hypoxia in the tumour mass. That is, to regions where oxygen concentration is insufficient to sustain cellular activity. This, in turn, results in the activation of Hypoxia-Inducible Factors (HIF), which mediate a set of processes aimed at helping cells survive in inhospitable environments. Among these is the secretion of Vascular Endothelial Growth Factors (VEGF), which stimulate the elongation of existing blood vessels and sprouting of new ones from existing. It is further noted that blood vessel structures observed in tumour masses often display an irregular topology.

Despite this, tumour growth rates often exceed those of the surrounding blood vessel networks. Furthermore, extreme hypoxia, mainly observed towards the centre of tumour spheroids, can lead to the formation of necrotic cores. That is, groups of tumour cells that are essentially starved of oxygen and nutrients or surrounded by waste products to the extent that their normal functioning is compromised to the point of death. As a result, these cells no longer contribute to tumour growth, which is observed predominantly on the surface of the tumour mass.

2.1.3 Current Therapies

Outright prevention of the development of cancer in individuals is to date impossible. Nevertheless, lifestyle choices and habits may significantly impact the probability of individuals developing cancer. Risk factors related to this include obesity, regular exposure to sunlight without adequate protection and consumption of carcinogens such as those contained in cigarettes and alcohol.

Tumours detected earlier are more likely to be successfully treated than those detected in later stages. Therefore, regular screenings to test for neoplastic tissues are advised, especially for ‘at-risk’ groups. For example, female adults are advised to test regularly for breast cancer, whereas male adults are advised to test for prostate cancer. Individuals may also look out for warning signs which could highlight the need for further testing. For example, in the case of melanoma (skin cancer), the appearance of new skin moles or the growth of existing ones may suggest the need for further testing.

In the instances where cancer has developed, early-stage interventions are more likely to succeed than later-stage ones. Specifically, tumours that have not entered metas-

tasis are more easily treatable than those that have. The reason for this is twofold. Firstly, a single-site tumour allows for targeted and localized intervention. Secondly, if successful, removal of the identified tumour should result in the absence of any other neoplastic tissue. Whereas, in the case of tumours that have entered metastasis, secondary tumours deriving from the original ones may give rise to relapses.

Several types of interventions may be prescribed to patients. At times, a combination of these may be deemed appropriate.

Surgery may be advised in case of early-stage tumours which have not yet entered metastasis. A surgeon will remove the neoplastic mass and some portion of surrounding tissue. Surgery is only appropriate for solid tissues (Eg: It cannot be used for leukaemia, a form of blood cancer) located in areas that may be accessed without compromising the patient's health. Further, certain brain tumours may not be suitable for surgery as they may be found in the regions that may not be accessed without compromising neural functions. There are no significant side effects to surgery compared to other kinds of therapy, although some time may be necessary for the resulting wound to heal.

Chemotherapy is one of the most commonly employed therapies. Chemotherapy refers to the subministration of substances that may reduce the ability of cancer cells to grow and proliferate and/or help the host organism fight the tumour. Chemotherapy treatments are usually provided in repeated cycles for maximal efficacy and to ensure tumours are eradicated. Chemotherapy may serve as a standalone therapy or complement other approaches. For example, surgical removal of a tumour may be followed by chemotherapy to kill leftover cancer cells. Chemotherapy also has the advantage of being able to target non-solid tumours and tumours located in parts of the body that may not be reached by surgery. However, treatment by chemotherapy may have many adverse side effects on the patient. These range from milder ones, such as tiredness and fatigue, to more severe ones, such as soreness and bleeding.

Radiation therapy, often referred to as radiotherapy, refers to high-energy radiation to kill cancer cells. This form of treatment may be localized or generalized. Patients often consider radiotherapy as an alternative to surgery and may be used in conjunction with other forms of treatment. Eg: Chemotherapy Some of the side effects of chemotherapy may be similar to radiotherapy, including the development of nausea and soreness. The type of radiation the patient will require depends on several factors, including the type of tumour, size, and location within the body. Radiotherapy may be subministered via internal or external beam radiation. In the latter case, a radiation-emitting machine is positioned on the patient or at a short distance. In the former case, a radiation source (e.g., a capsule containing radioactive material) is placed inside the body near the tumour. This is known as brachytherapy. Or, the patient may be injected with radioactive liquid, known as systemic therapy.

An alternative form of treatment is hormone therapy, otherwise referred to as en-

docrine treatment. Hormones may be used to block or reduce the growth of the tumour or alter the effects of other hormones within the body. Hormone therapy is often used together with different approaches. It typically serves one of three primary purposes: to reduce the volume of the tumour before other forms of treatment are enacted, reduce the risk of tumour relapses following therapy, and destroy cells that may have survived different forms of treatment of initial phases of relapse. Hormone therapy is unfortunately not exempted from side effects. Milder ones may include fatigue and tiredness but may range from digestion problems (constipation and diarrhoea), weaker bones, and compromised sexual functioning.

A person's immune system usually can identify and attack cancer cells. However, the efficacy of this mechanism may be limited. For example, some cellular subpopulations of the tumour may acquire mutations that make them less visible to the immune system. Or, they may express membrane proteins which reduce the immune system's ability to target them. Finally, they may also be able to change the way cells around them behave to act as a protective layer shielding cancer cells from the immune system.

Immunotherapy aims at improving the immune system's ability to respond and fight cancer. It may do so in several ways. Immune checkpoint inhibitors may override normal 'switches', which would prevent immune reactions. Under normal circumstances, these are important as they help modulate the immune system's response and prevent it from being over-aggressive. However, in the case of tumours, a more robust immune response may be desirable. Further, T-cells (a kind of white blood cell) may be extracted from the tumour environment, grown in a laboratory to reach larger quantities of these and injected back into the body. The idea behind this approach is to selectively develop those subpopulations of T-cells that are more specific to cancer cells, thus allowing for a more targeted response. Finally, monoclonal antibodies that recognize and bind to cancer cells may be developed in a laboratory and injected into the patient. This helps the immune system identify and target cancer cells.

Treatment efficacy varies considerably from patient to patient, depending on individual circumstances. A variety of factors may impact a prognosis, with the patient's age, stage of tumour growth and other medical conditions being among those most strongly correlated.

Further, in most cases, cancer treatment will result in some form of side effects. Whether these may be milder or stronger, they often negatively impact the patient's physical and mental health and may lead to debilitating conditions in the long run. Therefore, developing targeted therapies with reduced side effects is a primary research target in oncology and one we address as part of our research.

2.2 Cancer Modelling

Cancer research has traditionally been performed either in-vivo or in-vitro. The latter refers to the growth and study of cultures of cancer cells in an artificial environment. Eg: A petri dish. The former refers to studying tumours grown and developed in a live medium. Eg: Commonly, a mouse.

More recently, ex-vivo has been developed as a popular alternative to those previously indicated. This refers to experiments on tissue sections extracted from living organisms with minimal alterations to their conditions. Contrarily to in-vitro, where cells originate from existing cultures, in ex-vivo, cells are tissues removed from a different host or patient each time. And, differently for in-vivo, the experiments are performed outside of the host and in a laboratory environment. Advantages of ex-vivo include more closely mimicking what would be observed in a live being, without putting the host or patient at risk. This is especially useful, for example, in trials of experimental treatments which may have harmful side effects.

We will now briefly discuss the advantages and disadvantages of each approach and then introduce in-silico as a complementary strategy to obtain additional insights into its cancer biology. Regarding in-vivo, the first significant advantage is that of the tumour growing in an environment more closely resembling its natural one. This includes the development of a proper extracellular matrix, interaction with other cells, development of angiogenesis and the possibility to observe the process of metastasis. In addition, growth in a live versus artificial medium also presents significant differences.

Exposure to the extracellular matrix means cancer cells will be affected by the activity of neighbouring tissues. They will need to compete for access to nutrients and defend themselves from the immune system, among other aspects. The geometry of the tumour mass also differs between a Petri dish-like environment and a live host organism. Whereas in the former, cells will grow in a roughly regular, circular fashion, the geometry of the tumour mass in the latter may be more irregular and dictated by the morphology of existing tissues.

In-vivo does, however, also provide several drawbacks. In the first instance, growing tumours in-vivo is more complicated than in-vitro. Tumour cells must first be grown in a Petri dish-equivalent environment until the cancer body reaches a mass sufficient to be introduced in a host. After that, a suitable host needs to be selected (Eg: nude mice) and the tumour implanted by means of surgery.

The monitoring of tumours in-vivo is also more complex than in-vitro. Whereas in the latter case, marker solutions may aid medical imaging, in-vivo often requires that tumours be extracted from the host and measured. This often results in a limited number of measurements per experiment. Further, removal of the tumour from its host usually means the extraction process alters the cancer mass's properties. For

example, and related to our research case studies, extraction of the tumour will imply exposure to oxygen which will rapidly down-regulate expression rates of Hypoxia-Inducible Factors (HIF).

In-silico modelling is intended as a compliment, not a replacement, to in-vivo and in-vitro research. For example, in-silico tumour modelling simulates the growth and development of tumours in a computer. Common approaches employed in in-silico modelling will be further explored in Section 2.2.1. Still, we can anticipate that the overall objective is to produce simulations that accurately reflect what would be observed in empirical studies to obtain additional insights into cancer biology.

In-silico modelling offers several strengths which make it an attractive area of research to explore. Firstly, simulated experiments are significantly cheaper to run than their wetware counterparts. Smaller simulations may be run on normal laptops or computer equipment commonly available in offices. More computationally expensive experiments may require more powerful infrastructures, but even then, 'pay as you go' models of infrastructure as a service available from most cloud providers will still result in very cost-effective packages.

In-silico experiments are also faster to run than their wetware counterparts. This is because longer timeframes of tumour development, such as months or years of cellular growth, may be reproduced in the space of hours or at most days. This allows to rapidly complete preliminary exploration and tests, then deciding which hypothesis should be explored in a laboratory wetware environment and which others should instead be discarded. Combined with the considerations on cost mentioned earlier, this may allow for a significant reduction in experimental costs.

2.2.1 Overview of Main Approaches

There are two main approaches to tumour modelling: continuous and discrete. First, we will look more closely at the strengths and weaknesses and then consider how a third hybrid approach may be implemented.

Continuous Modelling

Continuous modelling consists of solving a set of continuous partial differential equations (PDEs) to establish the change in concentrations of various tissue types and solubles in time. Such a scenario may show the expansion of cancer tissue at the detriment of healthy tissues, with the expansion of the former being affected by glucose concentrations in the environment. Furthermore, glucose concentrations may themselves be affected by concentrations of endothelial cells (Ie: Cells will form the walls of blood vessels), acting as sources, and cell tissues, which act as sinks. Finally, the concentration of endothelial cells will depend on the activity of cancer cells which

promote the process of angiogenesis. Assuming C to represent the concentration of cancer cells, E that of endothelial cells, H that of healthy cells, and g that of glucose, we could set up a system described in Equations 2.1.

$$\frac{\partial H}{\partial t} = -\alpha C_t \tag{2.1a}$$

$$\frac{\partial C}{\partial t} = f(g)(\beta C_{t-1} + \nabla^2 C_t) \tag{2.1b}$$

$$\frac{\partial E}{\partial t} = \gamma E_{t-1} \omega C_t \tag{2.1c}$$

$$\frac{\partial g}{\partial t} = \nabla^2 g + \eta E_t - \mu C_t \tag{2.1d}$$

In equation 2.1a we establish that healthy tissues may only decay and that this rate is proportional to the concentration of cancer cells. Equation 2.1bshows that the change in concentration of cancer cells is proportional to the previous concentration of cancer cells at that position, allowing for migration from/to neighbouring positions. This is modulated by glucose concentration (g), where the function f would dictate that the concentration of cancer cells may grow or shrink depending on glucose availability. Equation 2.1c establishes that the growth rate of endothelial cells is proportional to the concentration of cancer cells and endothelial cells. This assumes an initial non-zero concentration of endothelial cells. And finally, equation 2.1d governs the diffusion of glucose. This is mediated by the source rate (η) and sink rate (μ) and the diffusion in/out of the position.

Continuous models may be less computationally expensive as there is no need to store each cell's state explicitly. When using this approach, only the agent concentration at each position is kept but not the state of each agent.

Firstly, their parametrization may not be straightforward. In the case of our example, the equation coefficients need to be established. In cases where the model is susceptible to any of them, even minor fluctuations in these may produce significantly different results. Parametrization is also a challenge present in discrete models.

Secondly, partial differential equations do not allow for incorporating elements of stochasticity by their very nature. In our simulations, we relied heavily on asserting that cancer cells would adopt a given phenotype (Eg: The Warburg Phenotype, see Section 2.1.1) with a certain probability. Having defined that, whether or not that phenotype was adopted was only established during the simulation's run. Such concepts may not be leveraged when using continuous models of this sort.

Continuous simulations further do not allow the modelling of the state of individual cells. They only model the cellular concentration at each position. These would have been unsuitable for some of our investigations, as we aimed to establish, for example, which proportion of the tumour mass adopted the Warburg phenotype. This required us to model the phenotype of individual cells explicitly.

Discrete Models

Discrete models represent cells (or groups) as entities, otherwise known as agents. Each of these has a state and behaviour that determine its evolution in time. For example, an agent representing a cancer cell may include properties in its state such as its progress in the cell life cycle, the expression rate of Hypoxia-Inducible Factors (HIF) and proliferation rate. But, on the other hand, its behaviour may determine that it should proliferate continuously and aggressively invade neighbouring spaces even when these are occupied by other cells.

Space may be represented either as a continuous or discrete environment, although the latter is more commonly found in the literature. In that case, agents would be assigned coordinates on a grid which may be either a 2D or 3D one. Note that one of our research themes explores the implications of opting for 2D over 3D. Depending on the model's configuration, a grid position may be occupied by one or more agents belonging to one or more classes. Positions may also differ in environmental properties. For example, in our simulations, different positions would have different glucose and oxygen concentrations.

Time is typically discretized into time-steps, or epochs, corresponding to real-world units. In many of our models, we decided to set each epoch as corresponding to a period of two hours, the length of the shortest phase of the cell life-cycle.

Discrete models have several points of strength, which make them especially well-suited to our research. Firstly, they offer the possibility of explicitly observing the phenotype of each cell. This is especially interesting in our investigations, where we were interested in reproducing and analyzing regions of phenotypic heterogeneity within the tumour. As an example, this proved to be a particularly useful feature in our investigation concerning expression rates of Hypoxia-Inducible Factors. Similarly, they also allow the inclusion of elements of stochasticity. That is aspects of an agent's behaviour that are non-deterministic, evaluated at model runtime and which may differ between different iterations of the same simulation. We leveraged this in our Warburg investigation, allowing for phenotypic heterogeneity across various cancer cell agents.

Discrete models further allow us to implement complex behaviours. If we consider the following statement, which was one of the prerequisites to our model: 'An agent's expression rate of Hypoxia-Inducible Factors mediates its proliferation rate, and is

itself affected by the concentration of oxygen at its current position.' we notice how this becomes challenging (if at all possible) to implement in a continuous model through a PDE. On the other hand, implementing a discrete model using a high-level language (better if following an object-oriented approach) is relatively straightforward.

Finally, and although perhaps subjective, discrete models may be easier to communicate. For example, the agents' behaviour may be described by visual means such as a flowchart diagram, thus making it accessible to an audience that does not possess a specialist background. On the other hand, PDEs may require a bit more expert knowledge in calculus.

Hybrid Models

While most of the models we found in literature adopted a discrete or continuous approach exclusively, we found a sizeable minority that included elements of both. These were particularly interesting as they could leverage and combine the strengths of the two methods. For example, our cancer models used to investigate the Warburg Effect and 2D vs 3D modelling use a hybrid approach.

In brief, hybrid models involve simulating specific aspects of cancer development by means of agents while delegating the simulation of others to a continuous model. In our case, cellular activity was simulated using agents, with each agent representing a group of cells. However, diffusion of solubles was delegated to a continuous model, which considered the concentration of cells at each position to establish solubles' source and sink rates.

This is appropriate for sub-problems where explicit modelling of each member is unfeasible. For example, in the context of diffusion, the modelling of individual molecules of glucose or even oxygen atoms is computationally infeasible (due to the sheer number of them) and unnecessary. That is because, contrarily to cells, we are not interested in the state of individual molecules or atoms, just in their concentrations at each position.

Further, modelling soluble diffusion using a continuous approach allows us to delegate the solving of the system of PDEs to an ad-hoc library, thereby relieving us from the engineeristic challenges posed by matters such as ensuring conservation of mass through iterations and special conditions at the environment's boundary.

When developing such a hybrid system, special care needs to be given to accurately transferring information between the continuous and discrete systems. For example, in the case of our models, the number of agents of each class (Eg: Cancer cells, endothelial cells) at each position is used to establish the source and sink rates of solubles, which is then used to parametrize an instance of the PDE solver. This is then

used to set soluble concentrations at each position for that epoch. It is important to note that this process needs to be repeated at the start of each epoch.

2.3 Reproducibility and Replicability in In-Silico Systems

Replicability is the possibility of other researchers obtaining the same results using the same artefact as the developers (Ie: Their exact same code), which is another benefit of in-silico models. Furthermore, distribution of the source code, accompanied by setup configuration and system images to run the experiments, allows any researcher to re-run and validate experiments. It also further provides additional studies to be derived from previous ones, with the confidence that no significant differences were introduced by accident.

Reproducibility, that is, the possibility of other researchers obtaining the same results as the developers using a different implementation which follows their specifications may also be facilitated by in-silico models. For example, by providing experimental setup and agent and process description using a formal, shared markup language or schema, for example: SBML [6].

In cases where developers do not share their setup or process descriptions using any proper tool but only provide a "natural language description" of their environment, in-silico models may still promote reproducibility. That is because in-silico models allow to more directly control the simulation's variables and reduce the likelihood of unintended side-effects being introduced.

For example, a wetware experiment attempting to control the temperature in a tank of aquatic creatures to study the effect on the metabolic rate of a species of fish would also affect the proliferation of the bacterial population. If the methodology to control this was not included in the paper, it would affect the study's reproducibility. However, an in-silico model may be able to address this by excluding bacteria altogether or making bacteria insensitive to changes in temperature. Of course, this would only make sense in cases where researchers did not want to consider the impact of bacteria, and these were intended purely as "noise".

Finally, in-silico models also benefit from the ease of monitoring and measuring. As mentioned earlier, with wetware strategies such as in-vivo tumour growth, there is a considerable risk of phenotypic properties of the cancer cells (Eg: Expression rates of Hypoxia-inducible factors) being altered during the process of removing the tumour mass from its host. This is not the case with in-silico models, where each agent's accurate measurements and environment position may be made without affecting the simulation itself.

Note however that the advantages illustrated in this section only occur in instances where the code is shared and made accessible. Or when at least sufficiently detailed

specifications are distributed to allow the solution to be accurately re-implemented. If such is not the case, then such in-silico solutions would have no significant advantage in terms of reproducibility and replicability compared to in-vivo and in-vitro.

Chapter 3

Literature Review

This chapter will present an overview of existing in-silico models of tumour growth. These have been grouped into thematic macro-categories, including models of tumour hypoxia (Section 3.1), therapeutic intervention (Section 3.2), tumour phase transition (Section 3.3), more advanced models of soluble diffusion (Section 3.4) and models of angiogenesis (Section 3.5). Articles reviewed were selected based on their relevance to our study, specifically when including elements of cancer biology we wanted to incorporate in our model.

Our review comprises mostly discrete, agent-based models as this is the category our research is closest to. However, continuous models were included when related to specific items we investigated, for example, diffusion of solubles. In addition, while our investigation did not include simulations of drug delivery to tissues, we did review existing work and, in our evaluation, discuss how this feature could be integrated into our model.

For each piece of work reported, we summarise the model's features. These aspects make it relevant to our investigation and, where applicable, outline some of its limitations. Finally, we conclude by positioning our work indicating which areas we improve on state of the art from literature.

3.1 Models of Tumour Hypoxia

Models of tumour hypoxia were mostly focused on spatio-temporal phenotypic distributions throughout the tumour. That is, how did regions of tumour hypoxia evolve in time. Simulations generally focused on specific items related to hypoxia, such as its impact on therapeutic intervention (Dasu et al. [7], Powathil et al. [8]), or its impact on cell metabolism (Venkatasubramanian et al. [9], Bertuzzi et al. [10]). There was no dominant approach between continuous or discrete, with literature presenting multiple examples of simulations implemented in either way. Or, as in the case of Powathil et al. [8], hybrid models incorporating elements of both.

Dasu et al. [7] developed a model incorporating both the effects of hypoxia[1] on tumour development and the effects of anti-angiogenic therapy. The authors highlighted the importance of not treating cancer hypoxia as a 'black-or-white' binary homogeneity, instead of allowing different tumour regions to express different hypoxic levels. They corroborated their studies by investigating over thirty initial blood vessel configurations and found that their approach offered results more closely matching medical literature.

Powathil et al. [8] developed a multi-scale model of tumour growth to obtain insights into the response to chemotherapy.[2] Their model incorporated both discrete and

[1] This refers to a condition where cells or tissues suffer from oxygen depletion. This negatively impacts the functioning of cells and may lead to irreversible damage and/or their death.

[2] Medical therapies aimed at reducing the growth of tumours and/or killing existing tumour cells by means introducing cytotoxic chemical compounds.

continuous elements, where the internal cell cycle of individual cells was modelled using a system of ordinary differential equations. Their model further featured oxygen heterogeneity and hypoxia as well as HIF[3] expression and drug transport.

Venkatasubramanian et al. [9] proposed a continuous model of a tumour spheroid that incorporated multiple aspects of cell metabolism, such as the glycolytic and oxidative patheways, alongside lactate secretion. A similar work was presented by Bertuzzi et al. [10], who demonstrated that for a sufficiently large tumour mass, a necrotic core[4] is formed. Further mathematical modelling of cell metabolism and the effect of nutrient availability on growth rates was proposed by Schuster et al. [11]. Basanta et al. [12] also modelled cell metabolism and found that invasive phenotypes of tumour cells are more likely to appear after a switch to glycolytic metabolisms.

Qutub et al. [13] developed a model that facilitated the prediction of HIF expression as a response to oxygen concentration. It is essential to highlight that this model exclusively addressed the relations between the two variables and was not concerned with the actual growth of tumour masses incorporating this phenomenon. Therefore, it differed from the research we propose. We will include such a relation into the model simulating cancer expansion, enabling us to compare our results to medical studies instead of limiting ourselves to a statistical tool for observing the correlation between different factors.

3.2 Models of Therapeutic Intervention

The models of therapeutic intervention we reviewed primarily took the form of continuous implementations. This contrasts with other feature domains, where a balance tends to be struck with continuous and discrete implementations where it is not the case that the discrete trend dominates. This peculiarity can probably be attributed to the nature of therapeutic intervention, which near-exclusively is concerned with drug delivery to tissues. This necessarily involves diffusion from blood vessels, and it therefore makes sense that mathematical models involving different kinds of differential equations should be preferred. A notable exception to this is found the work by Norton et al. [14], which models the response of the immune system. In such an approach, macrophages (white blood cells) are modelled as discrete entities.

Norton et al. [14] used an in-silico model to investigate the effect of macrophage infiltration on tumour progression.[5] They developed a discrete model where agents represented different classes of cells and included phenomena such as angiogenesis. The results were compared to in-vivo results, specifically mouse xenografts, and reported that while macrophage infiltration may somewhat decrease tumour growth,

[3]HIF: Hypoxia-Inducible Factors: A set of genes that are activated in the presence of low oxygen concentrations and mediate a series of phenotypic adaptations aimed at helping the cell survive.

[4]Region in the core of the tumour comprising mostly dead cells.

[5]Macrophage infiltration refers to a process where the body's immune system responds to the development of the tumour and attacks cancer cells.

the extent to which it does is limited and not substantial.

Komarova et al. [15, 16] used mathematical modelling to investigate the anti-tumorigenic properties of aspirin on colorectal cancer. They developed a 3D model that tested different dosages of aspirin for two weeks and reported that the treatment could be used to reduce cellular proliferation.

Coletti et al. [17] used continuous in-silico models to investigate the effects of immunotherapies on castration-resistant human prostate cancer. Their results showed that scheduling drug subministration is critical to achieving the desired results. Specifically, the 'high-tumour' and 'no-tumour' results both formed attractor states, with the result being determined by the timing of drug dispensation.

Scalerandi et al. [18] described a continuous model to study the effects of anti-angiogenic therapy, which is a therapy aimed at reducing vessel growth to limit the nutrients and oxygen delivered to the tumour. Partial differential equation models of nutrient, oxygen and VEGF[6] concentrations, as well as cancer cells and cell densities, were included. The authors highlighted the variety of drugs available for containing angiogenesis and focused on several classes for increasing anti-angiogenic potential. They concluded that tumours subjected to anti-angiogenic therapy grow more slowly and achieve a smaller mass overall than those grown in the absence of treatment. This is in agreement with the literature from medicine and biology.

3.3 Models of Tumour Phase Transition

Models of tumour phase transition investigate the transition of a tumour from relying on the existing blood vessel network (vascular growth) to promoting angiogenesis (avascular growth). These were implemented both as discrete models (Drasdo et al. [19]) as well as continuous models (Astaning et al. [20]). Both implementations were successful at providing insights into key tumour dynamics, such as growth rates and metabolic processes.

Drasdo et al. [19] developed a model of tumour growth to investigate the phenomena and conditions impacting cancer development during the transition from vascular to avascular growth.[7] Their model featured a high-resolution simulation, with individual cells being modelled. Their report, which complemented laboratory studies, suggested that nutrient depletion may determine the size of the necrotic core but not that of the overall tumour. Astaning et al. [20] developed a continuous model of the Warburg effect[8], describing the transition from an aerobic metabolism to a glycolytic one. The model presents a mathematical formulation of key metabolic processes, ex-

[6]Vascular Endothelial Growth Factor: A soluble secreted by cancer cells to stimulate the growth and elongation of existing blood vessels as well as sprouting of new ones to support the developing tumour mass.

[7]Angiogenesis growth refers to tumour development before angiogenesis can be observed, whereas vascular growth refers to tumour growth after angiogenesis has started.

[8]Term used to refer to a set of phenotyipic changes which occur in cancer cells and which are generally considered to promote tumour growth and expansion.

panding on previous work in this domain area and showing how pre-existing models may be expanded and combined.

3.4 More Advanced Models of Soluble Diffusion

The models reviewed in this section deep-dive on selected aspects of soluble diffusion. This recognizes that different substrates present distinct chemical and physical properties, and that a "one size fits all approach" to diffusion may not necessarily represent actual dynamics accurately. The models here reviewed are all continuous, and focus on modeling the diffusion of oxygen through tissues (Secomb et al [21, 22]) and through blood vessels (Goldman et al. [23]).

Secomb et al [21, 22] presented a model aimed exclusively at simulating oxygen diffusion in tissues. Starting with a pre-established vessel network, they considered factors such as oxygen diffusivity, tissue oxygen consumption rates and vessel surface area to model diffusion via a set of partial differential equations. Their study demonstrated how to obtain accurate oxygen gradients and provided reference values for model parametrisation.

Goldman et al. [23] presented one of the most comprehensive models of diffusion. Starting with a discussion on intra-vascular transportation and the carrying capacity of haemoglobin in red blood cells, they included differences in oxygen diffusion out of the vessels between larger arteries, arterioles and capillaries. The velocity of blood and partial de-oxygenation were also considered. However, the evolution of the blood vessel structure in response to chemotactic gradients was beyond this investigation's scope.

3.5 Models of Angiogenesis and Vascularization

Among the models of Angiogenesis and vascularization we reviewed, examples feature of both continuous (Cai et al. [24]) and discrete approaches (Welter et al. [25]). In the former case, different classes of agents are used to represent tip and trunk cells. The space is discretized as a Cartesian grid, with each position being potentially occupied by one or more agents. Chemotactic gradients of Tumour Angiogenic Factors (TAF) or Vascular Endothelial Growth Factors (VEGF) drive the instantiating of new agents which simulates the process of blood vessel elongation and sprouting.

In the latter case, on the other hand, partial differential equations are used to determine changes in endothelial cell concentrations in time. The concentration of these at time t is a function of the cells in the neighbouring positions - to prevent new blood vessels forming 'from nothing' - as well as the chemotactic gradient at that position.

Notable are also the works of Anderson et al. [26–29], which start with a continuous implementation and then from it derive a discrete equivalent one.

Anderson et al. [26–29] probably contributed the most papers to the sphere of vessel network evolution during angiogenesis. The previous article developed an advanced continuous model that included both VEGF and Fibronectin as contributing factors and further expanded that into a set of algorithms that could be used within discrete models. Their model offered one of the most realistic simulations, including observations of 'looping', increased growth rate when approaching the tumour and the so-called 'brush effect' in higher VEGF densities.

Cai et al. [24] developed a continuous model to explore the Spatio-temporal evolution of tumour masses, including angiogenesis and vessel collapse. They included specific features of angiogenesis, such as intravascular blood flow, interstitial fluid movement and blood rheology. Their results provided insights into the distributions of different types of cancer cells, such as proliferating, quiescent, or necrotic, during tumour growth.

Welter et al. [25] also developed a discrete model of tumour angiogenesis, including different classes of vessels, such as arteries, veins and capillaries. They also included multiple types of vessel dynamics, such as vessel dilation and cooption. Their model facilitated the identification of 'hot spots of increased vascular density and zones of pressure drops in the initial vasculature.

3.6 Positioning of our Work

While some of the papers described are closely related to our research, we believe we have identified an area that has not been explored yet and offers opportunities for novel insights. For example, while it is true that Dasu et al. [7] investigated hypoxia in growing tumours and Qutub et al. [13] formed a model to describe the relation between oxygen concentrations and HIF expression, the former did so without including HIF expression. In contrast, the latter offered a purely statistical model without applying it to a simulation of growing cancer.

We further reviewed 22 papers from our bibliography (see supplementary material). We evaluated them to establish whether they model the following biological phenomena, all included in our model: angiogenesis, diffusion of solubles, expression of HIF, and Warburg-mediated proliferation rate, oxygen uptake and VEGF secretion rate. We found that on average, only two of the six were included, with the most common being the diffusion of soluble (16/22 papers), followed by angiogenesis (10/22), HIF expression (9/22), Warburg-mediated proliferation rate (6/22), Warburg-mediated oxygen uptake (3/22) and, finally, Warburg-mediated VEGF secretion rate (2/22). No paper included all six aspects. We further found that only 5/22 papers explicitly described their model as 3D.

We want to acknowledge the work of Jiang et al. [30], which inspired our multi-scale approach to simulate cellular, intracellular and inter-cellular phenomena. To Stamatakos et al. [31], we credit the inspiration for implementing the cell life cycle. Gevertz et al. [32]'s research inspired our method of generating the initial vessel configurations and oxygen diffusion phenomena as well as cancer cells acting as VEGF sources and cancer endothelial as the sink. We have expanded upon their model by making endothelial cells act as sources of oxygen and cancer cells as sinks. To Anderson et al. [26], we owe our implementation of the angiogenesis algorithm, governing vessel elongation and sprouting. To Anderson et al. and to Spill et al. [33], we owe the differentiation between tip and trunk vessel cells, with the former able to elongate/sprout vessels and the former being static. Due to the work of Patel et al. [B2001], we acknowledge that because temporal scales and cell life are so different, we can consider oxygen and VEGF gradients to be in a steady state at the start of each time step and solve them as such. From Dasu et al. [7], we took the systems governing cell hypoxia and cell reactions to different oxygen concentrations. Qutub et al. [13] is credited with the curves and equations relating oxygen concentration to HIF expression in cancer cells, which is fundamental to our research. To Secomb et al. [22] [21], we attribute a continuous diffusion model very similar to the one we ultimately implemented. Finally, we thank Goldman et al. [23] and Hsu et al. [34] for their thorough and meticulous work on the more complex and detailed aspects of oxygen diffusion. Although our implementation is not identical to theirs, their papers still enabled us to decide which items should be included and helped us proceed with the confidence that our method could offer reliable results.

Chapter 4

Current Challenges

We will now present the main challenges inherent to cancer modelling related to engineeristic aspects (Eg: Implementation, computational requirements...) and biological elements. Eg: Modelling different aspects of tumour biology

4.1 Model Parametrization

The first issue we identified was that of model parametrization. That is to say, how do we assign parameter values? Biological models inherently operate in a large parameter space, owing to the vast amount of biological and physical phenomena that need to be accounted for. These span from the duration of the cell life-cycle to the parameters required to generate a function describing the relationship between oxygen concentration and expression rates of hypoxia-inducible factors to the diffusivity coefficients of different solubles. On top of these, some parameters do not define a natural biological or physical quantity but are instead related to the model's functioning. For example, deciding to what real-world time unit an epoch corresponds to.

Two main aspects need to be considered related to the challenge of model parametrization. Firstly, there is the question of parameter confidence. That is, how well-known the parameter value is and what assurance we have in that. As an example, it is relatively well-established that the duration of the cell life-cycle is on average twenty-four hours. That is to say, while individual cells may, for a variety of reasons, anticipate or exceed that threshold, we can expect that within a population of cells, that would be the time required for a cell to divide.

Parameters such as the diffusivity of glucose in tissue are more difficult to establish. The literature has measured and reported such values, but publications report significantly different values. A plausible explanation for this is that the value is extremely sensitive to various factors such as the morphology of the tissue being examined, its temperature, vascularization, etc. So, very similar studies may report distinct values.

Finally, for specific parameters, values may not be known. For example, it is reported that expression rates of Hypoxia-Inducible Factors increase exponentially as oxygen concentration falls below a given threshold. However, the exact coefficients are unknown, and developers must therefore decide how this relation is implemented.

In addition to parameter confidence, there is the matter of model sensitivity to specific parameters. That is to say, to what extent do relatively small changes in the values of single parameters affect the overall performance of the model, where a model is particularly sensitive to specific parameters, special care must be given to them.

In our models, we categorized parameters as green, yellow and amber depending on the confidence we had in their values (Ie: Whether they were reported in the literature, whether contrasting values were reported in different publications...) and how

sensitive the model was to them. For example, for parameters classified as amber, the parameter space was searched extensively using established strategies such as grid search and random search.

4.2 Diffusion Modeling

Modelling the diffusion process of solubles adds an additional challenge as it requires us to consider problems happening at a different scale from the cellular one. Indeed, while cellular processes (Eg: Cell growth) happen on a scale of hours, diffusion occurs at time scales smaller by orders of magnitudes to that.

There are also significant differences in the number of entities to consider when comparing cellular processes and diffusion. Whereas the number of cells in a tissue section may be approximated by looking at the ratio between the volume of said tissue and that of a cell, enumerating and modelling the individual oxygen atoms that diffuse is an intractable problem.

It, therefore, follows that explicit modelling of diffusion through discrete entities, as we do for cells, is not possible for solubles. We opted for a hybrid approach to approach this where a continuous model would be used to establish concentrations of solubles throughout the environment based on agent distributions.

That itself, however, adds further engineering challenges. For example, the need to ensure that the two sub-models (discrete and continuous) remain in sync, devising the best way to ensure that information is accurately transferred between them. That is a problem because the agent-based model discretizes space as a cartesian grid, whereas the continuous one used for diffusion, as the name suggests, makes use of a cartesian space.

Alongside spatial homogenization, there is also the matter of temporal alignment. As mentioned, processes of soluble diffusion happen at time scales by orders of magnitude inferior to cellular activities. But, due to the computational cost of solving the diffusion model, this shouldn't be solved more than once per epoch.

Finally, there is the challenge of designing the set of equations used to represent the process of diffusion. In our case, we used a simple diffusion equation with different classes of cells representing source and sink rates. Still, more sophisticated approaches which consider different diffusion coefficients (Eg: Across the blood vessel membrane vs in the interstitial space) could also be considered.

4.3 Model Validation

A key question to consider is the validity of a model. That is, how closely the tumour growth simulated in-silico replicates what would typically be observed in-vivo

or in-vitro. Unfortunately, a significant amount of publications does not address this issue or limit themselves to reporting that growth patterns (Eg: Exponential growth) are observed and consistent with what is expected.

Before discussing some examples from literature, it is worth illustrating how deciding which of two growth curves is closer. The expected may not be a trivial task. Consider the case of Figure 4.1. In Figure 4.1athe simulated tumour is consistently smaller than the expected volume. In Figure 4.1b, on the other hand, the tumour is closer to the anticipated volume for more of the simulation but eventually exceeds that value with the growth rate suggesting that difference would become more significant. This highlights that visual results may be misleading or difficult to interpret in the absence of an objective strategy (formula, algorithm) to address this issue.

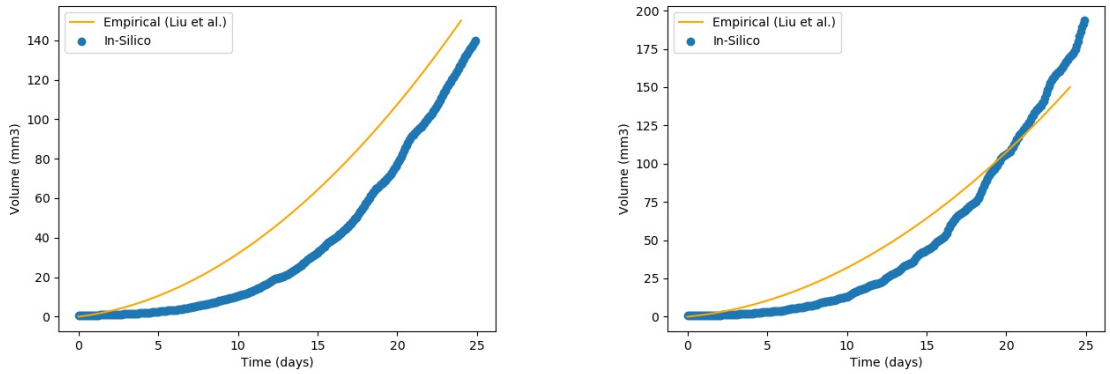

(a) (b)

Figure 4.1: Comparison of two cases where simulations yield curves with different behaviours, each ”good” or ”bad” in different ways. It may be hard to decide which of the two is best in these cases.

Another open question is determining whether a curve exhibits exponential growth patterns often cited as a hallmark of tumour growth. This is especially challenging when the measured time frame is short, as is the case in many experiments that focus on a few months of tumour development. Furthermore, especially in the early stages of growth, the curve may not yet have developed its ”elbow” and look linear for all intents and purposes. Hence, judging whether the growth is exponential or linear just by looking at the reported growth curve would be inappropriate. Instead, possible solutions could involve extending the simulation over a more extended period to understand the actual dynamics of the curve. Or fitting a function to the generated curve and observing whether this is linear or exponential. Although, this latter may not work for especially short simulations. Or, finally and perhaps more appropriately, not using ”linear vs exponential” as a metric for evaluation instead of comparing actual to expected volumes. Other indicators, such as the presence or absence of angiogenesis, may also be included.

To summarize these statements, we believe that an objective assessment of model performance should include some form of quantitative numerical analysis of the re-

sults—ideally, one which compares to actual laboratory measurements. As we will point out in several examples, this is not often observed. In the absence of that, results should at least be related to specific studies and evaluated qualitatively. E.g. A particular gene being over-expressed in a model, consistent with over-expression also observed in study X. In this instance, the magnitude of the over-expression is not addressed quantitatively, but results are at least compared to a specific study.

We would also like to point out that, in many instances, biological phenomena expected to be observed are mentioned. Eg: Increased cellular proliferation, angiogenesis... However, when the discussion and conclusion then say that they were observed in the model, little evidence supports this. Eg: What would the proliferation threshold have been to say that cells now divided at an 'enhanced' rate? Or, how much more development of the blood vessel network should there have been to state that angiogenesis was taking place?

A primary example of this is found in the work of Olsen et al. [35], where in the discussion the authors state, "As expected, cancer cells can grow much larger with angiogenesis, which results in normal cells essentially being removed from the system" and "Our tumour cells grow in the spheroid shape that is generally expected in cancer growth, whether or not angiogenesis is activated.". Again, therefore, not compare their results to existing empirical data from laboratory studies.

A further example is represented by Gerlee et al. [36], where tumour morphology is compared to a generally expected one but without any objective measure of similarity: "For normal oxygen concentrations we find tumours with a smooth circular morphology, while in low oxygen concentration we observe tumours with invasive fingered shape.". Similarly, tumour behaviours as a function of oxygen concentrations are once again compared subjectively rather than through a systematic comparison against laboratory data: " The change in finger width is more obvious for the two lower oxygen concentrations, where we can observe a clear shift towards smaller tumours with wider fingers as the matrix density is increased."

Similarly, Cai et al. [37] all reference expected behaviours and confirm their model aligns to them, but do not provide any form of numerical or objective comparison of this: "This phenomenon will disappear when cell migration ability is removed from the model. This pseudopalisade development result confirms the first mechanism of pseudopalisade formation proposed by Bratet al." or "Consequently, some dysfunctional vessels undergo collapse, which leads to local hypoxia. This result demonstrates the second mechanism of pseudopalisade formation.".

More recent work also displays the same weakness, for example, Shamsi et al. [38], which claims that: "One important issue illustrated in this work is the ability of glycolytic tumours to progress in poorly vascularized tissues without the need to rely on neovasculature." and "the computational model developed in this work elucidates the function of Warburg effect in the context of the tumour microenvironment and

determines what possible advantages it can confer to cancer cells", without however directly comparing results to existing pieces of laboratory work.

This is even more evident in pure mathematical proof of concept models (Eg: Bertuzzi et al. [10], Astanin et al. [20]), where often no attempt is made to relate the results back to clinical data. Or, as in Coletti et al. [17], where in their own words, the only analysis shows results to be " in qualitative agreement with experimental results", with again no standardized on the objective comparison.

More elaborate forms of analysis are employed, for example, by Kather et al. [39]. Specifically, some attempts to relate model outputs to laboratory observations are made. Results are compared to expected clinical outcomes by manually classifying histopathological quantification of stroma and lymphocytes labelled as "Stro low, Lym high" or "Stro high, Lym High". Then, comparing survival benefits between the groups as predicted by the model versus clinical data. This is a significantly more elaborate approach than those illustrated earlier and an example of the direction that research should ideally take when considering model validation. A possible weakness in this specific approach could be that. In contrast, the correlation between stroma and lymphocyte numbers and survival benefit, the causality between the two is based on a set of assumptions presented in the paper itself. To substantiate the validation quality, including additional objective measures could have been helpful. Eg: Tumour volume, vascularization...

To conclude, the relevance and importance of systematic and objective evaluation of models results can be summarized in the words of Professor Ioannidis [40]: "A research finding is less likely to be true [...] where there is greater flexibility in designs, definitions, outcomes, and analytical modes." Specifically, we would like to emphasize the need for shared strategies and mechanisms of model validation that can be used to ensure that a given model correctly replicates the dynamics and behaviours which would be expected in a laboratory. This ensures that the findings reported are grounded in the biology of the phenomenon being studied and not merely an artefact of the model.

4.4 Infrastructure

Our experiments required simulations to be run multiple times, with iterations scaling in the hundreds and even thousands. This implied running them on a single computer was impractical, even when this might have been an institutional research-purpose server. Instead, we found it convenient to leverage cloud infrastructure (in our specific case that offered by Amazon Web Services) to deploy hundreds of simulations in parallel.

Two main challenges in this area involved the provisioning of infrastructure and the collection and aggregation of results.

Infrastructure provisioning refers to starting a sufficient number of cloud servers to process the workload, setting up the environment (OS, development environments, configuration...), deploying the desired code version (often from some version control system), distributing experiments across the various servers and monitoring the execution of simulation while they are progressing.

Collection and aggregation of results refer to transferring summaries of completed simulations from the servers on which they were running to some form of the centralized repository in a format that yields to post-execution analysis. For example, each simulation's result may be stored as a CSV file (or some other form of structured content) and uploaded to a storage-purposed server.

To address these challenges, we designed our framework, PanaXea, to integrate with major cloud vendors (Eg: AWS) and leverage the resources they offered. For example, we used AWS EC2 instances to run our simulations, AWS CodePipeline to provision our instances and S3 buckets to store and aggregate results. To ensure that all simulations were running in the same environment and that differences in performance would not be attributable to different software versions, we ran our simulation within docker containers.

Chapter 5

Modelling Strategies

5.1 Introduction

Cancer research has been historically performed via in-vivo and in-vitro experiments. The former indicates experiments run on live animals or patients, whereas the latter refers to experiments performed in a laboratory environment such as cell cultures in a petri dish. While contributions from this research have been invaluable, in-vivo and in-vitro approaches do present several weaknesses. Some of these have been identified and discussed by Kam et al. [41], we now summarise their main points.

Regarding in-vivo experiments, the authors highlight the difficulty in controlling the experimental environment. Especially in complex organisms, several factors beyond the control of the analysts impact the development of the tumour mass. They make it difficult to ascertain which phenomenon or factor is the driving force behind an observed behaviour of cancer cells. Regarding in-vitro experiments, while these allow better control of the environment, it becomes hard to replicate the complexity of in-vivo environments. For example, it is challenging, if not impossible, to replicate the process of Angionesis in an in-vitro environment. Angiogenesis refers to the process whereby cancer cells stimulate the formation of new blood vessels and are considered central to cancer growth.

In-silico models represent a possible solution to address the limitation of in-vivo and in-vitro approaches. These are artificial models simulating the development of biological systems, offering researchers the possibility of incorporating processes that cannot be replicated in-vitro while at the same time allowing for a level of control of the experimental setup, which is not achievable in-vivo.

In addition to research and hypothesis testing, in-silico models could also develop personalised patient models. This would enable clinicians to predict whether a particular therapy would be successful on a given patient or whether it would present specific side effects. As a result, it would reduce the risk of patients undergoing ineffective or even damaging therapies reducing operational costs and increasing odds of success.

The strengths of in-silico models are also discussed further in-depth by Saeidnia et al. [42], Agur et al. [43] and Ekins et al. [44].

A challenge in the development of in-silico models is their parametrisation. Because of the multitude of parameters needed, and because these need to be informed by medical literature, multiple values are often drawn from different studies and publications or estimated based on various other studies. This is the case, for example, in McDougall et al. [45] where the probability of vessels sprouting as a function of Tumour Angiogenic Factors (TAF) concentration is estimated as a sigmoid relation. This is also seen in Anderson et al. [26], where due to the difficulty in obtaining exact measurements, it was assumed that epithelial cells could respond to any non-zero

TAF concentration. Furthermore, this is also the case in Kather et al. [39] where parameter values need to be estimated as they are not known from the literature.

Many existing frameworks share a limitation: a lack of support for parameter space exploration and no utility to deploy simulation instances to cloud services quickly. Eg: Amazon Web Services, Google Cloud, proprietary infrastructures. As part of our investigation, we developed a general-purpose framework for the development of hybrid agent-based models: PanaXea.

PanaXea is developed in Python, one of the most commonly adopted languages in the scientific community, and leverages widely used open-source libraries for data processing and analysis. This contrasts with other popular frameworks such as NetLogo [46] or GAMA [47], which use proprietary languages and limit the developers in the toolkits and libraries they may import. Popular frameworks exist in non-proprietary languages, including Repast [48] and MASON [49]. The main weakness of these, in our opinion, is that they use Java rather than Python. This is not a matter of personal preference: Python is increasingly becoming the language of choice for the scientific community, with experts from domains other than Computer Science often being fluent developers and a growing ecosystem of data analysis and visualisation tools. Eg: IPython Notebooks. It is also worth noting that Python has a gentler learning curve than other languages, with syntax often resembling natural English and implementation not requiring the developers to master more in-depth concepts such as those related to Object-Oriented Programming.

Some Python frameworks for agent-based Modelling do exist, among which we would like to acknowledge MESA [50]. Many of our design choices align with theirs, but critical improvements made on the implementation include us supporting three-dimensional environments. This is fundamental, for example, to accurately model biological phenomena as highlighted in multiple publications [51–53]. We further support Numerical Grids which, overlaid to model environments, allow to store numerical properties associated to each position. Eg: Nutrient concentration at a particular position. Finally, we further implement each epoch as a three part process; Thus allowing, within the same epoch, all agents to complete a set of steps before a further set of agent actions is initiated. The value of this is further explained in Section 5.2.1.

The Panaxea framework offers out-of-the-box tools to easily instantiate a model environment, define agents, populate environments with agents and synchronise the execution of all agents to a common schedule. Reproducibility is one of the critical features of the framework. Panaxea enforces an explicit definition of starting conditions such that the same experiment can be reproduced multiple times on different hosts. It also offers features to monitor the evolution of the model and easily export observations for post-execution analysis.

Further, PanaXea offers tools to automatically perform and evaluate parameter space

exploration (Eg: Grid Search, Random Search), reporting which parameter values lead to the best model performance. Some of these features are already present in other frameworks, such as OpenMole [54]. Especially compared to the latter, our minimalist toolkit for parameter space exploration is considerably smaller in terms of size (355MB OpenMole vs 7kb our toolkit). It is worth noting that OpenMole offers many additional functionalities. Still, we would not have been leveraging many of them while paying the cost of increased project size regarding our research objectives. Further, OpenMole is written in Java, whereas our framework is written in Python. Having developed our parameter space exploration toolkit in Python, we avoid requiring developers to set up two separate environments.

The framework further makes deployment to cloud infrastructures seamlessly. It is distributed with a Docker image set up to host the framework and tools to monitor simulations running on different cloud nodes and reconcile and aggregate results. While every effort was made for the framework to be developed in a clean, efficient manner, no assumption of users being proficient in Python is made. In the interest of the framework being approachable to as broad an audience as possible, knowledge of core concepts of programming and the basic syntax of Python should be sufficient to develop complex models.

The rest of this Section is structured as follows. First, we present an overview of our framework, the components and code extracts of its usage in Section 5.2. Code extracts, in particular, refer to the case study we provide in Section 5.4. Next, we implement a biological model of angiogenesis derived from existing literature and use our parameter space search toolkit to evaluate its sensitivity. Finally, we present our conclusions and final remarks in Section 5.5.5.

5.2 Methods

5.2.1 PanaXea - Our Framework

The model was implemented using PanaXea[1], an open-source framework we developed. This is written in Python and aimed at facilitating and promoting the development of agent-based models. The framework has a modular approach, exposing tools that may be adapted to suit the individual needs of developers. We enumerate the main components with a brief explanation of how they may be used. A class diagram detailing our implementation may be found in the Appendix; see Figure 2.

Steppables

Steppables refer to entities that are progressed as the simulation executes. That is entities whose inner logic is executed and whose state is updated once per epoch.

[1] https://github.com/DarioPanada/panaxea

There are two main classes of steppables: Agents (see Section 5.2.1) and Helpers (see Section 5.2.1).
Each steppable allows developers to define behaviours executed during the prologue, the main and the epilogue stages of an epoch. All steppables' prologue methods will be executed during an epoch before the main methods begin. Furthermore, all main methods will be executed before epilogue methods start. This allows, for example, for all agents to collect information about their environment without any other agent's action having changed environmental properties. Instead, agents would collect information during prologue and then execute actions in main or epilogue.

Agents

Agents are the main actors of a simulation and may represent any self-contained entity such as a person, a biological cell, etc. Agents have a position in an environment and one or more properties. Further, classes extending the agent base class also encapsulate the logic each member of such class would execute at each epoch and describe its behaviour. In our simulation, we implemented two classes of endothelial cell agents: Tip Cells and Trunk Cells. Following, we present an extract of how Tip Cells are implemented in PanaXea:

```
1   class TipCell(Agent):
2
3   def ___init___(self):
4       super(TipCell, self).___init___()
5
6   def step_main(self, model):
7       # Get current position
8       current_position = \
9           self.environment_positions[
10          "agent_environment"]
11
12      # p0 represents the probability of
13      # the cell remaining stationary
14      p0 = self.___calculate_p0(
15          model,
16          current_position
17      )
18      # p1-p4 represent probabilities
19      # of the cell moving to any
20      # of the 4 adjacent positions
21      p1, p2, p3, p4 = self.___calculate_p1to4(
22          model,
```

```
23              current_position
24          )
25
26          next_position = self.___get_next_position(
27              p0,
28              p1,
29              p2,
30              p3,
31              p4,
32              current_position
33          )
34
35          # If the cell has moved at
36          # least 180 positions along the x-axis
37          # we assume it has
38          # traversed the section of tissue
39          if next_position[0] > 180:
40              model.exit = True
41
42          # Add a trunk cell to the old position
43          t = TrunkCell()
44          t.add_agent_to_grid(
45              "agent_environment",
46              current_position,
47              model)
48
49          # Move itself to the new position
50          self.move_agent(
51              "agent_environment",
52              next_position,
53              model)
```

The example above illustrates the behaviour of TipCell agents. Namely, at each epoch, they decide whether they will move to an adjacent position or remain stationary. This decision is made in the *___get_next_position* method on line 26. A TrunkCell is instantiated at their previous position where they decide to move. Developers can leverage common model functionalities provided by the framework (Eg: add_agent_to_grid) and program their ad-hoc behaviour specific to their simulation.

Environments

This consists of classes to instantiate 2D and 3D Cartesian grids. Grids may be numerical, storing a value at each position or object grids, storing one or more agents at each position. For example, numerical grids were used to store TAF and Fibronectin concentrations in our implementation, whereas object grids were used to store endothelial cells. Environments also provide the functionality to move agents between positions and to find adjacent neighbours with the most or fewest agents in them. A code extract is provided below:

```
xsize = ysize = 200

NumericalGrid2D("taf_environment", xsize,
ysize, model)
NumericalGrid2D("fib_environment", xsize,
ysize, model)

ObjectGrid2D("agent_environment", xsize,
ysize, model)
```

This instantiates two 2D Numerical Grids and one 2D ObjectGrid. It also declares the environment size, names each environment to be referenced in the future and binds it to the model object, so it becomes part of the simulation.

Helpers

Helper classes interface with the simulation and are executed once per epoch as agents. However, they are not 'actors' in the simulation but rather perform an auxiliary function such as recording population size, changing environment properties, etc. In our example, we used helpers to update TAF and Fibronectin concentrations at each epoch. A Helper may be defined as:

```
class ConcentrationsHelper(Helper):
    """
    At the start of each epoch,
    updates TAF and Fibronectin
    concentrations
    """

    def ___init___(self):
        super(ConcentrationsHelper, self)
        .___init___()
```

```
11
12          def step_prologue(self, model):
13              self.___update_taf(model)
14              self.___update_fibronectin(model)
15
```

This is added to the schedule so they become part of the simulation as:

```
1  model.schedule.helpers.append(ConcentrationsHelper())
```

In this specific case, the Helper is responsible for updating soluble concentrations at each environment position before the start of each epoch. The Helper is not a member of the simulation per se, but its work is essential.

The Model

The model is the parent object of the simulation and holds all other elements such as environments, agents, the schedule, etc. Further, the model may also contain a set of simulation properties. For example, model properties include all values for parameters used in the equations in our simulation.

The Schedule

The schedule is a singleton class that holds all steppable: agents and helpers, part of the simulation. Each epoch consists in traversing the schedule and executing the logic of each steppable. By default, agents are stepped in the same order at each epoch unless the developer randomises this within their simulation. This is to avoid systematic bias owed to some agents consistently updating before/after others. The principle of schedule randomization at each epoch is applied throughout all our case studies as well as the Warburg Investigation, see Chapter 6. All steppables must be registered with the schedule to ensure that they are executed as part of the simulation.

5.2.2 Parameter Space Exploration

PanaXea emphasises parameter space exploration and sensitivity testing by enforcing a separation between a model and the set of parameter values used for a specific simulation. The model would define all aspects of the simulation, such as the strategy whereby environment properties are set, agent behaviours, etc. However, these would include variables whose value would be assigned at runtime as they are loaded from an external configuration file. For example, in our model, the equations generate probabilities for a tip cell to move to a position using a high number of variables.

The equations are defined in the model and are common to all simulations, but the actual value of each of the equations' variables would be different for each run.

PanaXea offers a script to generate experiments. Given a set of experiment parameters and value generation strategies, this produces a configuration file in the format of comma-separated values (CSV) where each column corresponds to a parameter and each row to a simulation. We support Grid Search and Random Search as parameter generation strategies [55]. However, by design, it is possible to add other algorithms with minimal changes to the existing codebase. The usage of an external configuration file also allows for experimental reproducibility.

For example, the following definition:

```
parameters = [
    {
        "name": "gamma",
        "search_strategy": "grid_search",
        "range": [0, 10, 2]
    },
    {
        "name": "alpha",
        "search_strategy": "random_search",
        "mean": 1,
        "stDev": 0.1
    },
    {
        "name": "epsilon",
        "search_strategy": "constant",
        "value": 5
    }

]
```

Will result in the experiment definitions being generated reported in Table 5.1.

experiment	gamma	alpha	epsilon
1	0	1.01	5
2	2	0.98	5
3	4	0.97	5
4	6	1.03	5
5	8	0.95	5
6	10	1.05	5

Table 5.1: Sample experiment definition generation with PanaXea. Parameter values for gamma are generated by grid search on range 0 to 10, parameter values for alpha are generated by random search on a distribution with mean μ=1 and standard deviation $\sigma = 0.1$ and parameter values for epsilon are kept constant at 5.

5.2.3 Analysis of Results

PanaXea decouples simulation state logging during execution from model serialisation. Tools are available to decide, at each epoch, which properties of the model the developers wish to record. This could be as simple as recording the entire state of the model (environments, agents and their properties...) or as a more fine-tuned approach. For example, the developers may wish only to record a subset of agent properties from certain agent classes. For example, a simulation of cellular biology may feature hundreds of gene expression rates across multiple types of cells. However, the developers may only be interested in the expression rates of specific genes for certain types of cells. Recording everything at each epoch would lead to significant memory usage for no benefit.
Simulation results may then be serialised in any desirable format. By default, PanaXea supports python 'pickles' (object serialisations). This is the most versatile and can represent any sort of information. Nevertheless, information could also be mapped to any format with minimal effort. Eg: Tabular: CSV, SQL, Semi-structured: XML, JSON...

Integration with other Systems

Simulation results can be saved by default as Python "pickles". That is, full-disk serialisations of Python objects which can then be re-hydrated. Ie: Loaded from memory into an active Python process. In addition, PanaXea supports loading existing PanaXea models created by other researchers and developers. However, integration with other tools is currently not provided out of the box and is a limitation. We, therefore, consider how the framework could be expanded such that models could be serialised into alternative, cross-compatible formats.

We consider Systems Biology Markup Language (SBML) in the first instance [56]. SBML is a format used to describe biological systems using XML, defining how different entities (Eg: Environments, agents, processes...) should be described.

Alongside technical specifications, XML Schema and RelaxNG Schemas validate documents. However, as pointed out on the official website, these can only verify that the syntax of the markup language is correct and not that the actual content is sensible. Eg: They can identify whether a chemical compound is correctly described, but not whether the compound is actually 'plausible' according to the laws of Physics.

We can envision a clear path to export PanaXea models to SBML, although some effort would be required. In the first instance, the model properties we would like to export must be identified. These would include our agent classes, their parameters and parameter values, environmental properties, and the agents' state. These would then need to be mapped to appropriate SBML entities. Finally, code would need to be written to extract the values we wish to persist into an appropriate data structure (Eg: A Python dictionary). Finally, an SBML or XML Python API (Eg: LibSML [6]) would be needed to create the SBML files.

Another exciting integration would be the Simulation Experiment Description Markup Language (SED-ML) [57]. This is also an XML-based description format, but alongside describing the biological phenomena and simulated entities, it also includes information about the model itself. This allows defining a simulation pipeline, including which software version to use, which simulations to run and what kind of analysis can be performed. The steps to integrate with this solution would include writing a SED-ML description of our simulation set. Further, some entry point methods in PanaXea would probably need to be added, which would read a SED-ML file and start the appropriate setup, execution and analysis pipeline.

Finally, another integration that would be worth considering is that with runBioSimulations [58]. This allows to aggregate multiple models and frameworks, combining these and delegating different tasks. For example, a framework designed to model the process of diffusion accurately could be coupled with another to obtain gene expression rates due to substrate concentration. The advantage of this would be that domain expertise previously encapsulated in individual models could be shared, reaching higher levels of biological completeness and overall more accurate models. runBioSimulations describes several interfaces that models need to implement to expose a standard API, so changes to the PanaXea codebase would need to be made to achieve this.

5.2.4 Extensibility

The framework is unprejudiced by design. Rather than enforcing modelling conventions, it exposes to developers a set of flexible tools that may be deployed and arranged as best suits the problem at hand. Emphasis is on such tools imposing as little overhead as possible, thus freeing computational resources for the actual simulations. Wherever possible, established libraries are encouraged, such as Numpy

for numerical calculations, Pandas for dataset management and Matplotlib for visualisation.

The framework is designed to integrate into pipelines which include separate stages for experiment generation, execution and analysis. Because of this, open and accepted standards such as CSV for experiment definition and python pickle files to store experiment outputs are encouraged. This latter point is significant, as it allows for results to be loaded into a developer's analysis software of choice, thus not constraining them to a preset showcase of tools distributed with our framework.

5.2.5 Cloud Deployment

Computational requirements of agent-based models often exceed the capabilities of everyday home or office computers, and the amount of time required for these to complete may make this undesirable regardless.

PanaXea has been designed with cloud deployment in mind. The framework is distributed as a python package installed via Pip (Python Package Index). A docker image is also made available, allowing for simple execution of any PanaXea project.

Finally, a toolkit for Amazon Web Services (AWS) is provided. This provides scripts to write experiment files to AWS Queues as messages. Each experiment may be individually downloaded and executed by EC2 instances (cloud computing instances) part of a designated fleet. Upon completion, the resulting pickle and any other generated output are automatically uploaded to an S3 Bucket (cloud storage) from where they can be downloaded.

The rationale for selecting AWS as a technology is that AWS Cloud Credits for Research support our research. However, the scripts could be easily adapted to interface with other major providers such as Google Cloud or Microsoft Azure, where a proprietary setup is preferred. We recommend an Apache Kafka instance to act as a queue service and a Kubernetes system to manage a cluster of Docker images where simulations would run. As a replacement to S3 Buckets, any storage server accepting incoming SCP requests should fulfil the task.

5.3 Elements of Our Model

We now present details of our model's implementation which are common both to the case studies in this chapter and our Warburg Investigation (Chapter 6).

5.3.1 Diffusion of Solubles

The solubles modelled include oxygen, glucose and VEGF. Endothelial cells act as sources of oxygen and sinks of VEGF and glucose, whereas cancer cells act as sinks of oxygen and glucose and sources of VEGF. Equation 5.4 modelled the diffusive transport of a soluble in the interstitial space. At the start of each epoch, the secretion and uptake rates are calculated based on the number of agents at each grid position and their properties. Following the solution of the continuous model, these are copied back to the discrete environment and made available to the agents.

$$\frac{\partial \kappa}{\partial t} = D_\kappa \nabla^2 \kappa + s_\kappa - u_\kappa \tag{5.1}$$

Where s_κ refers to the source rate and u_κ to the uptake rate of κ at a particular position. D_κ refers to the diffusion coefficient of κ in tissues. Positive concentration values are ensured by solving the Equation twice. If the first iteration yields positions with a negative concentration, cells are removed from the simulation, making the sink rate for such positions zero. Then the Equation is solved again. Oxygen is measured in terms of oxygen tissue tension (interstitial space) and oxygen tension in the blood (for source rates of endothelial cells), so the units are those of partial pressure (mmHg). The uptake rates are estimated as the reduction in partial pressure values at a specific position owing to the metabolic activity of cells. A similar approach has been found in Whitley et al. [59]. The glucose concentrations are calculated in μmol and VEGF as arbitrary units. As is further explained in Section 5.3.2, the latter case is because we are only interested in the direction of VEGF gradients and not in the actual concentration values.

Owing to the differences in the time scales of the molecular diffusion processes and the length of our epochs, we assume concentration gradients to be constant and therefore solve for the steady-state. FiPy [60] was used to solve the equations as mentioned earlier at each epoch using the finite volume approach. Our environment was set up using no-flux boundary conditions.

5.3.2 Angiogenesis

Angiogenesis refers to how existing blood vessels develop into new ones. In tumour environments, the current vessel network quickly becomes unable to provide the required nutrients to the rapidly growing tumour mass, which begins secreting vascular endothelial growth factors. The vessels in the nearby environment respond to this by elongating and sprouting new ones.

Our model differentiates between two classes of endothelial cells: tip cells and trunk cells. Trunk cells are static agents that act as placeholders for indicating an

oxygen source. Tip cells, on the other hand, exhibit a more dynamic behaviour. Endothelial cells will respond by elongating the upgradient given a minimum VEGF concentration, which is set as a model parameter. In elongating, the tip cell will move to the adjacent position with the highest VEGF concentration, and a trunk cell will be added to the previous position. A rest period of 12 epochs, i.e. 24 hours, is assumed between each tip cell's elongation.

5.3.3 Updating of Agent State

At each iteration, agents are updated in random order. This ensures that forms of systematic bias are not introduced in the simulation. For example, suppose an agent was constantly updated first. In that case, this could result in them having preferential access to resources (Eg: Space to proliferate, nutrients...) which would act at the detriment of agents which were consistently updated later in the epoch.

This is further consistent with the length of our epochs, corresponding to two hours of real-world time. It makes sense that cells should not always 'act' precisely in the same order (Eg: Cell A replicating first, followed by Cell B...) but rather that the order should be non-deterministic.

5.3.4 Spatial Discretisation

The proposed models simulate tumour growth in a 3D environment represented as a grid, where each grid element corresponds to a volume of $0.042mm^3$, with our environment having a size of 20x20x20 units. The model comprises discrete and continuous sub-models. The continuous sub-model is used to calculate soluble concentrations at each epoch. The soluble source and sink rates at each position are calculated based on the number of agents there and their properties. Each position will also have a concentration of solubles (glucose, oxygen and VEGF) updated at each epoch. The discrete sub-model comprises agents representing cells, specifically cancer and endothelial cells. We assume an average cell volume of $5,000\mu m^3$ [61, 62], and each cell agent represents 120 cells, which, due to sharing close spatial proximity, develop in perfect sync. This allows us to explore the heterogeneity of the tumour environment while reducing the computation costs and maintaining tractability. Up to 70 cancer cell agents can occupy each position.

5.3.5 Temporal Discretisation

Time is represented as epochs, with each epoch corresponding to a real-world duration of two hours. This corresponds to the shortest phase of the cell life cycle: mitosis. At the start of each epoch, the soluble concentration values for each position are updated. Following this, the agents are stepped in a random order, which results

in them updating their internal state and possibly their local position based on their previous state and their interaction with other agents and the environment. The high-level flow of the simulation is summarised in Figure 5.1.

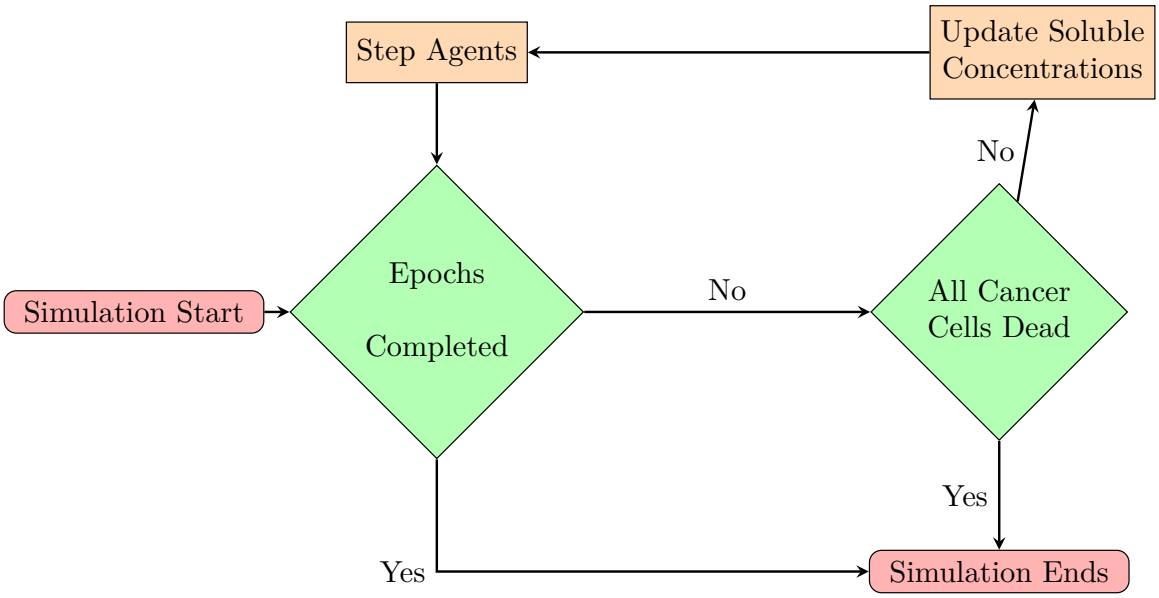

Figure 5.1: High-level overview of the simulation. An epoch consists of updating the soluble concentrations across the environment (oxygen, glucose, VEGF) and then stepping the agents (cancer cells, endothelial cells).

5.3.6 Dynamics and Properties of Cancer Cells

Cancer cell agents implement each stage of the cell life cycle, with the total duration of this being 12 epochs, corresponding to 24 hours. Upon successful completion, the cell will divide, producing two daughter cells at the current position. If the current position has reached carrying capacity, one of the daughter cell agents is transferred to the adjacent least populated position. If no such positions are available, the cell may not divide.

A probability of progressing into Synthesis defines the base likelihood of a cell progressing from Growth I into Synthesis. This is defined as a simulation parameter shared by all cancer cells. Upon completing Growth I, the cell will either transition into Synthesis during the following epoch or, if that is not the case, attempt to do so at each next epoch, always with the probability as mentioned earlier. Cancer cells act as sinks of oxygen and glucose and as sources of VEGF.

5.4 Case Study: Parameter Space Exploration

5.4.1 Introduction

We considered the discrete implementation of the model of angiogenesis proposed by Anderson et al. [26]. Such a model simulates a blood vessel network development

in a tissue environment in response to chemotactic stimuli promoted by a growing tumour mass. Figure 5.2 illustrates a sample evolution of the model as reported by the authors. To briefly summarise, the environment is discretised as a two-

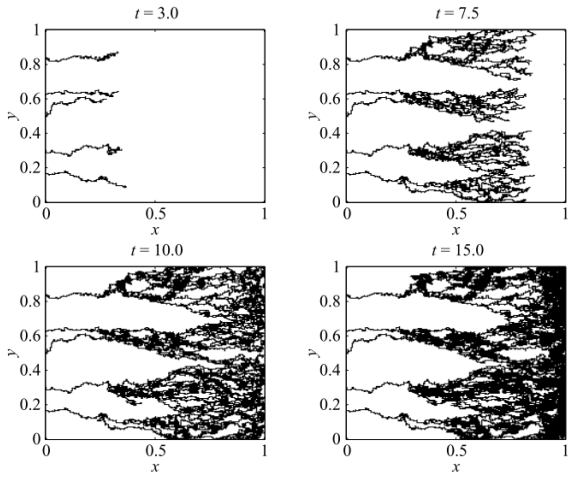

Figure 5.2: Development of existing blood vessels in time in response to chemotactic stimuli through elongation and sprouting, as reported by Anderson et al. [26]. Starting with 5 tip cells at x = 0, elongation is observed towards a tumour positioned at x = 1. As the vessels approach the tumour, sprouting is also observed.

dimensional Cartesian grid. Each position has a concentration of tumour angiogenic factor (TAF), which creates a chemotactic gradient, and Fibronectin, which creates a haptotactic gradient. The model is advanced by discrete steps called epochs. At each epoch, each tip endothelial cell's probability remains stationary or moving in one of four directions (up, down, left or right) is calculated.

The model is initially seeded with five tip cells, positioned along the y-axis at $x = 0$; see Figure 5.2. As an endothelial tip cell migrates to an adjacent position, an endothelial trunk cell replaces it at the original one, thus simulating the process of vessel elongation. Further, increasing TAF concentrations positively correlates with the probability of a cell sprouting. That is, of an existing blood vessel branching into two new ones. Finally, the interplay between TAF and Fibronectin concentrations determines probability values, with endothelial cells privileging positions with higher concentrations of both.

The model is well-suited to our research aims because of its detailed reporting of all required parameters as well as their values, alongside all equations used and a thorough explanation of its mechanics. This allowed us to re-implement it with confidence and explore the effect that changing parameter values have on the progression of the simulation.

The equations which govern endothelial movement are parametrised with 14 variables. Values for all of these are provided, but for many of them, neither their biological significance nor the process by which the value was obtained is discussed. In several instances, it was reported that trial and error was used to find values that allowed the simulation to behave as expected. However, specific strategies and ranges

tested were not disclosed.

5.4.2 Results

A simplified version of the model, in that it was seeded with a single endothelial tip cell at position x = 50 rather than multiple and did not account for sprouting, was used. Overall, the environment grid had a size of 200 x 200 units. The rationale was that this would simplify the interpretation of the results while remaining generalisable to more complex implementations. Initial TAF and Fibronectin concentrations were calculated using a logistic function along the x-axis of the environment. Hence, TAF values were highest at x = 200 and lowest at x = 0, and vice-versa for fibronectin concentrations.

The fitness of each simulation was calculated as inversely proportional to the number of epochs it took for the tip endothelial cell to traverse the environment and reach the tumour, see Equation 5.3. Each simulation was repeated ten times, with the fitness reported being the average of the runs. The model was allowed to run for a maximum of 500 epochs, after which the simulation would have been considered failed.

$$fitness = \frac{1}{n_{epochs}} \tag{5.2}$$

Figure 5.3: The function is used to establish the fitness of a given simulation in this case study. n_{epochs} refers to the number of epochs required by a tip cell to traverse the simulation space.

We studied the effect of changing values for parameters x_0, the chemotactic coefficient, and k, a coefficient present in all equations used to determine movement probabilities. We initially made use of grid search [63]. For k, which in the original publication was set to 0.75, we tested values from 0 to 10 in steps of 0.2. For x_0, which was initially set to 2,600, we tested values from 500 to 7,000 in steps of 500. The two parameters were explored independently. When x_0 was explored, the value of k remained constant with the original value used by the authors and vice-versa. While this case study demonstrates parameter space exploration through grid search, similar experiments using random search or other techniques would also have been supported by our framework, PanaXea. Results are shown in Figure 5.4.

Notably, it can be observed is that the model is not sensitive to x_0, and only sensitive to k when the value for this is below a certain threshold. x_0 is the chemotactic coefficient. Its assigned value in the model we are reproducing is 2,600; it represents the TAF gradient's impact on vessel elongation. The higher this value, the more endothelial cells should privilege regions of higher TAF concentration to those of higher VEGF concentration. In practice, this means that higher x_0 values should induce the vessels to elongate away from their origin and towards the tumour,

whereas lower values should promote the opposite. Therefore, we would expect a positive correlation with fitness, but this is not observed. This suggests that, in this model, the implementation of the chemotactic coefficient does not reflect its biological role.

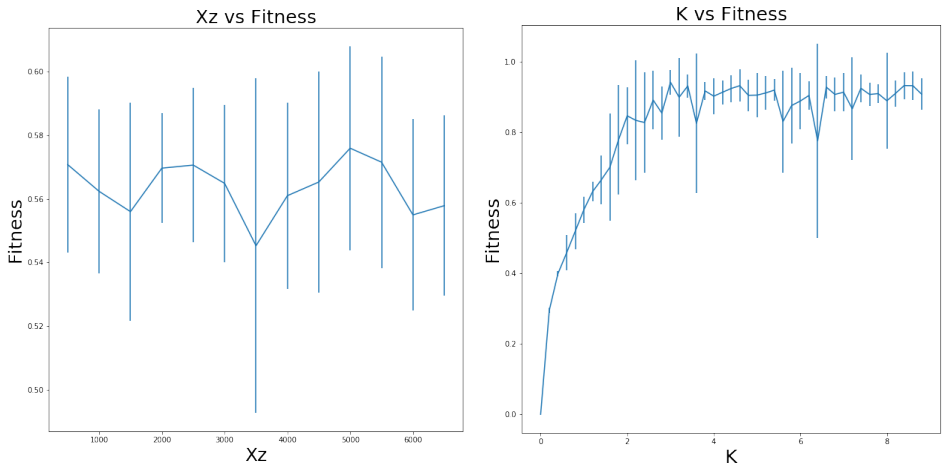

Figure 5.4: Grid search on parameters x_0 and k. x_0 was sampled from 500 to 7,000 in steps of 500, k was sampled from 0 to 10 in steps of 0.2. Simulations were repeated ten times per parameter value, with average values reported. Error bars show standard deviation. At each occurrence, simulations that did not complete within 500 epochs were excluded. For all values of $k > 0$, at least 80% of simulations were completed within 500 epochs.

The role of the parameter k is not fully explained. However, it is present as a coefficient in all terms of every Equation used to calculate movement direction probabilities. In the original model, it is given a value of 0.75. As can be seen, this leads to sub-optimal model performance, with the highest fitness occurring for $k > 1$. Accounting for stochasticity, increasing k beyond one does not lead to significant alterations in performance.

5.5 Case Study II: 2D vs 3D Modelling

5.5.1 Introduction

A decision that must be made early in the design process is whether the spatial discretisation will produce a two dimensional (2D) or a three dimensional (3D) environment. The case for 2D is that it closely resembles cancer growth in a Petri dish, a commonly adopted approach in laboratories, whereas 3D would allow the model to more closely mimic tumour growth in living tissues.

It has been suggested that advantages of 3D Modelling over 2D include better capturing of oncogene activation [53], protein expression and drug sensitivity [52, 64, 65] as well as more realistic biochemical and biomechanical environments [66] and a better translation of pathophysiological features of the tumour environment

[67]. Finally, 3D models have also been proposed to capture inter-cellular signalling pathways better [68].

An initial literature survey reveals many publications relying on 2D discretisation strategies to obtain insights into tumour biology. These include, among others, studies of blood vessel development in response to angiogenic stimuli [69], response to chemotherapy [51] and tumour oxygenation [70]. Although additional searches for publications in the field of tumour modelling revealed a large corpus of simulations using a 2D environment [24, 26, 71–73]. In cases where 3D was chosen an approach [35, 37, 74], it is not explained why this was decided for rather than opting for a 2D approach.

Previous related work [75] highlighted differences in simulation performance between 2D and 3D but did not assess the suitability or unsuitability of either.

Many studies using 2D models do not address the implications and potential limitations of modelling in 2D over 3D. For example, effects on the spatial distribution of agents, the evolution of the tumour mass and soluble diffusion are not addressed, nor is there any attempt, practical or theoretical, at mapping results back to a 3D environment. Finally, it is often unclear how decisions regarding the effective size of a grid position or the number of biological cells represented by an agent were taken. While these are model artefacts and properties of the simulation, changes in values could affect the model's output and support or invalidate a hypothesis.

With these items not addressed and the wealth of literature advocating 3D over 2D, and a substantial corpus of literature opting for both 2D and 3D approaches, we wish to investigate whether it is the case that 2D may offer a cheaper and more accurate alternative to 3D or if simulations should, on the other hand, be preferably or exclusively performed in three dimensions.

The rest of the Section is structured as follows. First, in Section 5.5.2 we present our model setup, parametrisation strategy and our approach to mapping values from a 3D to a 2D environment. Next, in our Section 5.5.3 we present our experimental findings for the various simulations in 2D and 3D, alongside measurements regarding tumour properties such as oxygenation levels, etc. Finally, in Section 5.5.4 we draw a conclusion regarding when 2D form an adequate approach to in-silico investigations or if 3D should be preferably or necessarily employed.

We implemented a simple agent-based model with continuous elements to include diffusion of solubles. Fundamental phenomena accounted for include cellular growth and proliferation, blood vessel development in response to secretion of vascular endothelial growth factors (VEGF) and Hypoxia and spatial heterogeneity with regards to oxygen concentration.

We believe that if despite the simplicity of our model, discrepancies in results are still observed between the 2D and 3D simulations, then these would also be present

in more complex models. Adding more biology would, therefore, only make these worse. In summary, the simplicity of our model allows us to explore the effects of 3D to 2D mapping of simulations while minimising additional work needed to implement more complex behaviours and streamline the process to map parameter values between dimensionalities.

Our initial setup in 3D was derived from starting conditions reported in the medical literature, see Appendix. A translation to 2D was then derived from this (See Table 9 in the Appendix), and simulations were run in 3D and 2D. Finally, resulting growth curves were compared to those reported in medical publications (See Figure 5.5) to determine whether the difference in error rates was significant. Where possible, parameter values were obtained from the literature. Where this was not possible, either because of specific values or unknown or because contrasting values were published, hyper-parameter tuning techniques such as grid search [76] and random search [55] were employed to explore a suitable search space.

5.5.2 Model overview

We now provide a high-level overview of our model, including key dynamics and properties of agents.

Cell life-cycle

The central concept behind cancer cells is that of uncontrolled proliferation. Given a sufficient oxygen concentration, cancer cells will keep growing and dividing until a physical constraint owing to their surroundings being saturated occurs. To account for various factors that may delay the division of individual cells, a probability value is set for cells to transition from Growth I (G1) into Synthesis. Upon completing G1, a cancer cell transitions into Synthesis with a given probability set as a simulation parameter. If the cell does not transition into Synthesis, it will attempt to do so at the following epoch with the same likelihood until it does. Once a cell progresses into Synthesis, it will complete the cell life-cycle and then divide into two daughter cells. The only exception will be if its surrounding environment is saturated, where it cannot further divide. That is because the surrounding environment being saturated indicates that there is physically no space left for the cell to divide.

A cancer cell may be Active, Quiescent or Dead. Active cells are growing and proliferating, quiescent cells are not progressing in the cell life-cycle and are secreting VEGF, and dead cells simply contribute to the tumour's volume but no longer have an active role. Two oxygen concentration thresholds are specified as simulation parameters: $O_{Hypoxia}$ and $O_{Critical}$, which determine the state of a cancer cell. The relation between oxygen concentration at a given position and the state of a cell is detailed in Equation 5.3. The transition from Active to Quiescent is reversible if

oxygen concentrations subsequently rise again above $O_{Hypoxia}$. Logically, a dead cell may, however, not transition back into other states. We assume $O_{Hypoxia} > O_{Critical} > 0$.

$$state = \begin{cases} Active, & \text{if } O \geq O_{Hypoxia}, \\ Quiescent, & \text{if } O \geq O_{Critical}, \\ Dead & \text{otherwise} \end{cases} \tag{5.3}$$

Diffusion of solubles

Our model incorporates oxygen diffusion from endothelial cells to cancer cells and VEGF from cancer cells to endothelial cells. Each soluble concentration is calculated at the start of each epoch, with values for each environment position updated. Given a soluble κ, Equation 5.4governs the diffusion process.

$$\frac{\partial \kappa}{\partial t} = D_\kappa \nabla^2 + s_\kappa - u_\kappa \tag{5.4}$$

Where D_κ is the diffusion coefficient, s_κ the source rate and u_κ the uptake (or sink) rate. The equations are solved using FiPy, [60] which implements the finite-volume method, and we set up our solution using no-flux boundary conditions.

Diffusion of solubles from the intracellular space to the extracellular matrix occurs by diffusion, which means the concentration outside cells may not be greater than that inside cells. As such, source rates of individual cells may be adjusted and reduced to avoid implausible scenarios such as soluble flow against a concentration gradient. The source rate effectively becomes zero where the inner and outer concentrations are equal.

Individual sink rates may also be adjusted to avoid negative substrate concentrations. However, a minimum uptake rate must be maintained owing, for example, to cells needing at least some level of oxygen. If a negative concentration is obtained, this indicates that the number of cells at a position exceeds the capacity of the blood vessel network and a number of cells equal to the amount required to restore positive concentrations is considered dead.

Angiogenesis

Angiogenesis refers to developing new blood vessels from existing ones in response to VEGF stimuli. Tumours secrete VEGF in conditions of Hypoxia to increase the oxygen and nutrient supply. In our model, we differentiate endothelial cells into Tip and Trunk cells.

Trunk cells are static agents. They act as sources of oxygen but do not proliferate or otherwise interact with the environment. Tip cells, on the other hand, exhibit

a more dynamic behaviour. Tip cells will migrate up-gradient to an adjacent grid position given a minimum VEGF concentration set as a model parameter. The position initially occupied will host a new Trunk cell. In-between elongations, Tip cells must complete one entire cell life-cycle. This is consistent with general cell-cycle dynamics, where following a division, cells must have time to grow and synthesise further biomass to sustain replication.

Experimental setup

We now illustrate our experimental setup, our derivation of 2D environments from 3D ones and parametrisation techniques.

The target curve

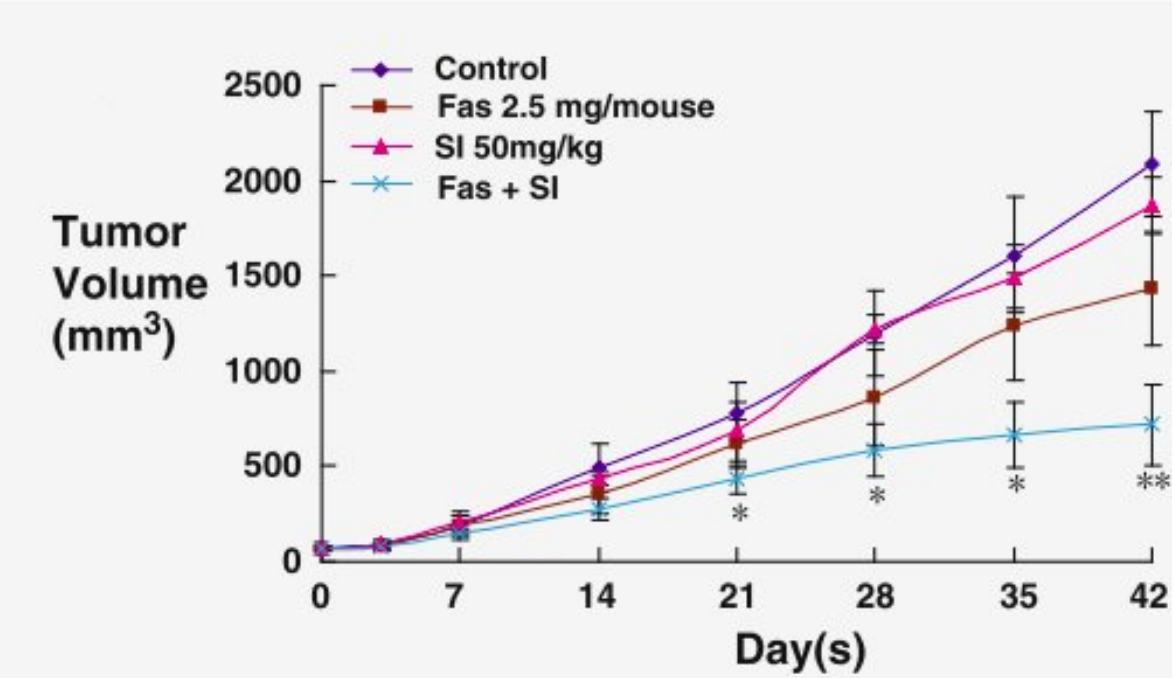

Figure 5.5: The target growth curves. We will be benchmarking our model against the 'control' curve over the first 25 days.

We will be benchmarking against the control curve reported by Chen et al. [77] and shown in Figure 5.5, comparing the tumour volume obtained in our simulation to that obtained in empirical studies. We will be considering the first two weeks of tumour growth to keep the problem computationally tractable and allow us to repeat large amounts of simulations on our infrastructure. In picking the publication we would use as a reference, we selected one that included reported readable growth curves with sufficient data points and smaller error bars.

As mentioned earlier, the environment is discretised as a Cartesian grid. Table 5.2summarises the effective sizes of grid positions, alongside other relevant spatial values such as the volume of cancer cells and initial tumour volume. Derived 2D values are also reported.

3D		2D	
Item	Value	Item	Value
Volume of Cancer Cell	2,000.00 μm^3	Area of Cancer Cell	158.74 μm^2
Volume of Grid Position	0.30 mm^3	Area of Grid Position	0.45 mm^2
Start Tumour Volume	1.20 mm^3	Start Tumour Area	1.13 mm^2

Table 5.2: Summary of spatial units for 2D and 3D simulations. 2D values were derived by assuming 3D structures were cuboids and then deriving the base area. Values are shown to two decimal places.

Model resolution

A decision that needs to be made upfront is the model's resolution. That is to say, how an agent will represent many cancer cells. An agent representing multiple cancer cells implies an assumption that all those cancer cells will share the same genotype and phenotypic state. I.e., They will be in perfect sync regarding the state in the cell life-cycle, uptake and secretion rates of solubles, expression rates of genes, etc. This may be a reasonable assumption when the number of cancer cells per agent is small (Ie: High resolution), as it is plausible that several cells in close proximity would share similar traits. On the other hand, the model's reliability may decrease as the number of cancer cells per agent increases, as we are then assuming that larger portions of the tumour are in perfect sync. This is an assumption that may be plausible. Assigning more cancer cells per agent compromises model reliability and computational cost. A high number of cancer cells per agent will mean overall fewer agents necessary but lowers the resolution (Ie: We capture a lower degree of cellular heterogeneity.), whereas a low number of cancer cells per agent provides a higher resolution but higher computational costs. Different resolutions also affect the maximum number of agents per grid position and the required initial number to achieve the desired start volume or area. We will be testing three resolutions: 300, 600 and 1,200 cancer cells per agent. These are summarised in Table 5.3.

Testing a resolution of fewer than 300 cancer cells per agent becomes computationally intractable. Decreasing the resolution to exceed 1,200 cancer cells per agent results in 2D simulations results in having a maximum number of agents per position of 1, which hinders either tumour growth or blood vessel development.

Cancer Cells per Agent		Max Agents per Grid Position	Initial Number of Agents
300	2D	9	24
	3D	500	2,000
600	2D	4	12
	3D	250	1,000
1,200	2D	2	6
	3D	125	500

Table 5.3: Summary of how different model resolutions affect the maximum number of agents held in a grid position and the number of agents required at the start to obtain the desired initial volume or area. As expected, increasing the number of cancer cells per agent decreases the carrying capacity of individual positions and the initial number of required agents.

Starting conditions

Cancer cells in an amount appropriate for a given resolution are seeded at the centre of the grid. Every position in the grid is seeded with one endothelial cell. Initial soluble concentrations are calculated and assigned to positions. This is done by calculating source and sink rates at each position based on the number of agents of each class present there and then used to solve for concentration values. The simulation is after that allowed to run its course.

5.5.3 Results

Each simulation set (e.g., 300 cancer cells per agent, 2D) comprises 200 simulations. These share the same parameters assigned as constant values but differ in values for those set by grid search or random search and those related to spatial quantities, such as the maximum number of agents per position. Each simulation produced a growth curve of the tumour compared to an expected growth curve derived from medical literature. This gave an indication of the error rate for such a simulation. A set's performance is calculated as the average error rate of all simulations in it. This allows, for each resolution, to compare the performance of 2D and 3D simulations. For 2D simulations, an extra step is needed to map areas back to volumes.

The error function

The error function compares the tumour growth curve produced by a simulation (in-silico curve) to the one observed in empirical laboratory studies (empirical curve). Ultimately, the error function produces a single value: The absolute mean error. This indicates how closely an in-silico growth curve aligns with the empirical growth curve.

Expected tumour volumes at different time points were obtained from the empirical growth curve (Figure 5.5, control curve).

A polynomial function was then fit these points using polynomial regression. The resulting polynomial function describes the expected volume (V_{exp}) at a day (d). This is reported as Equation 5.7, with coefficients reported to two decimal places. Note that our model's temporal resolution has one epoch corresponding to 2 hours. This was scaled when applying the error function, with a temporal resolution of 1 unit corresponding to 1 day.

$$V_{exp}(d) = 0.95d^2 + 18.10d - 65.68; 0 \leq d \leq 25 \tag{5.5}$$

Equation 5.6is a second-order polynomial, indicating that the growth rate of the tumour is exponential over the period investigated.

Given the actual volume at a day d reported in a simulation's growth curve, V_{act}^d, then the error ε_d is as reported in Equation 5.6the magnitude of the difference between the expected and actual value.

$$\varepsilon_d = |V_{act}^d - V_{exp}(d)| \tag{5.6}$$

The absolute mean error of a simulation (ε) is then calculated as the average of all errors calculate at each time-point:

$$\varepsilon = avg(\varepsilon_0, \varepsilon_1, ..., \varepsilon_{d-1}, \varepsilon_d); 0 \leq d \leq 25 \tag{5.7}$$

The absolute mean errors of each simulation in a simulation set were then averaged to obtain the aggregate average mean error of said simulation set.

Estimating volumes from areas

For 2D simulations, these returned tumour areas. The area was assumed to be the base of a 3D structure to compare these to expected volumes, and the corresponding volume was derived via a simple mathematical transformation.

The generic form of this basic transformation is provided in Equation 5.8, where V indicates volume and a area. It represents the derivation of a cube's volume from a square's area by obtaining the square's width and raising it to the cube.

$$V(a) = a^{\frac{3}{2}} \tag{5.8}$$

Resulting error rates

Figure shows the average error rates for 2D and 3D simulations across all resolutions5.6. For higher resolutions (Ie: Lower numbers of cancer cells represented

by a single agent), 3D clearly outperforms 2D. However, at lower resolutions (Ie: Higher numbers of cancer cells represented by a single agent) it is the case that the two approaches appear to be comparable or that 2D outperforms 3D. This is, in fact, misleading and will be addressed in our discussion.

We note that error rates in the case of 3D simulations are stable across all resolutions, with discrepancies of $\pm 150 mm^3$ being owed to the stochastic nature of the model. As for 2D, we note that the error seems to be linearly dependent and negatively correlated to the resolution. While it might appear 2D could be as good as 3D in some instances, this is not the case and is further elaborated in our discussion.

Finally, a Welch's t-test was performed comparing the populations of errors for 2D vs 3D at each resolution. Results were significant for across all resolutions for $p < 10^{-10}$. This t-test is appropriate as it does not assume equal population variance.

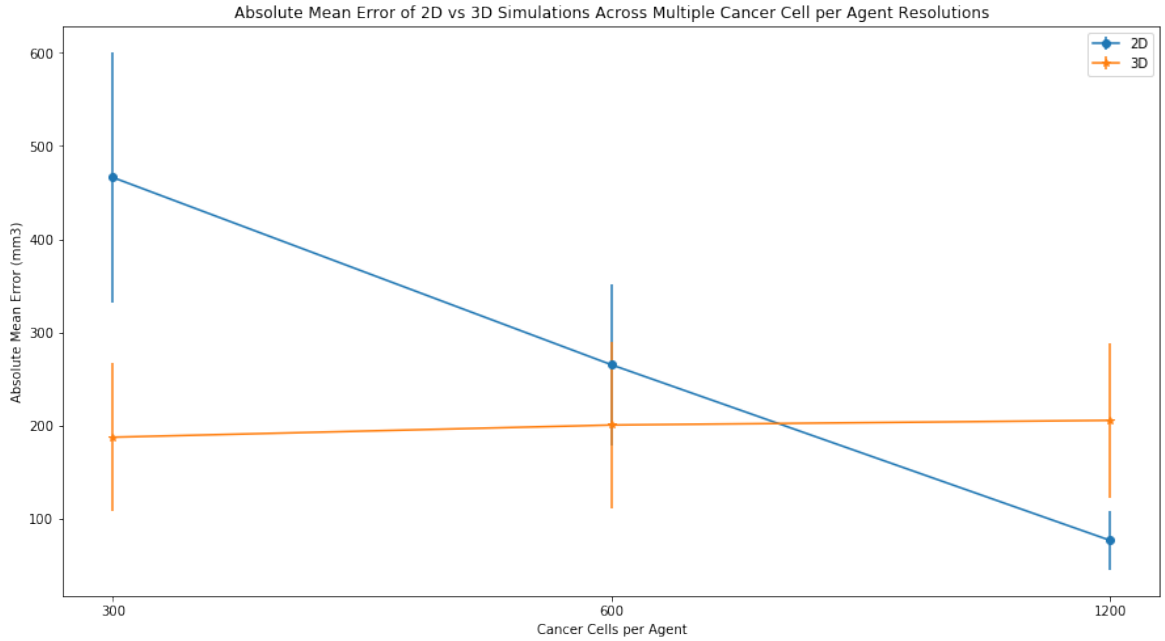

Figure 5.6: Summary of error rates for each resolution, 2D vs 3D. Error bars represent the standard deviation on the mean. A higher number of cancer cells per agent indicates a lower resolution. This means the simulation is modelling larger groups of cells together rather than allowing them to develop independently.

5.5.4 Discussion

Our results suggest a relatively small and stable error in 3D simulations, independently of the resolution. More interesting is the 2D case, which seems to outperform 3D at lower resolutions and whose error seems to increase linearly as the resolution is increased.

This might initially be interpreted as 2D outperforming 3D at all but the highest resolutions, suggesting that the former offers a viable if not preferable alternative to the latter. However, this is not the case. The apparent better performance of

2D is an artefact of the specific model configuration, which coincidentally allows it to minimise the absolute mean error. We will discuss how 2D simulations suffer a fundamental flaw due to dimensionality mapping from 3D. We will also emphasise the importance of validation strategies beyond growth curve comparison and consider elements such as soluble concentrations and phenotype ratios within the cancer cell populations. Without these additional verifications, our analysis would be incomplete and conclude that 2D forms an alternative to 3D, whereas this is not the case.

In the first instance, we note the direction of the mean error for 2D simulations. Errors are calculated as the difference between actual and expected volumes. So, a positive error means the model predicted a larger volume than expected and a negative one that the model predicted a smaller volume than expected. A summary is reported in Table 5.4. For resolutions of 300 and 600 cancer cells per agent, all 2D simulations consistently over-estimate tumour growth. However, for a resolution of 1,200 cancer cells per agent, a decisive majority of simulation over-estimates tumour growth. It, therefore, appears fair to state that our 2D models systematically over-estimate tumour growth.

Resolution (Cancer Cells per Agent)	Number of Positive Errors	Number of Negative Errors
300	196	0
600	194	0
1,200	187	12

Table 5.4: For each resolution, we report the number of 2D simulations which over-estimate tumour growth (positive error) against the number that under-estimate tumour growth (negative error).

Given that the magnitude of the error of 2D simulations increases as resolution increases and that 2D systematically over-estimate growth, we can say that as resolution increases, 2D simulations over-estimate growth by larger amounts.

To explain this, we need to consider several factors. First, as we increase resolution, we decrease the number of cells per agent, and therefore we increase the initial number of agents needed to keep the starting volume consistent. Given more initial agents, it makes sense that the final tumour volume should be greater. This is because, given the tumour volume at a given time is dependent on that at an earlier time ($V(t) = \kappa V(t-1)$), then a larger initial volume will produce a larger final volume. $V_x(t) > V_y(t) \longrightarrow V_x(t+n) > V_y(t+n)$. This given the assumption that all other factors are equal, which they were in the context of our simulation. Therefore the error should increase, but this does not explain why we observe this in 2D but not 3D.

The reason behind this lies in the effect that 3D to 2D mapping and resolution adjustments combined have on the maximum number of agents per position. For example, in 2D, a position may host between 2 and 9 agents. In 3D, between 125 and 500, see Tables 5.2 and 5.3. This means that in 3D, we may obtain a much higher cancer-to-endothelial ratio at each position than in 2D. CanceThis is because

cancers produce daughter cells, each capable of dividing, resulting in exponential growth. In contrast, a tip cell will only create another tip cell and trunk cell, the latter being unable to divide, leading to linear growth.

In both 3D and 2D, cancer cells will exceed endothelial cells, but, crucially, in 2D, space constraints mean the ratio will never increase to the point where oxygen becomes a limiting factor. This is owed to the source rate of oxygen being much higher than its sink rate and an endothelial cell supporting multiple cancer cells. Unless enough space is allowed at a position for a significant amount of cancer cell divisions, these will never exceed the capacity of the initial blood vessel network.

In summary, as observed in our case, in 2D simulations, oxygen never becomes a limiting factor of growth, meaning that a higher number of initial cancer cells, as is seeded at higher resolutions, results in a larger final tumour volume. However, in 3D, oxygen becomes a limiting factor and modulates and restricts tumour growth.

Our results raise important considerations for tumour modelling. Where simulations implement a 2D environment, appropriate scales and parameter values need to be chosen to avoid issues similar to those we described, which would make the models biologically implausible. Specifically, parameter values derived from 3D studies (such as empirical ones in laboratories) should not simply be projected over to 2D in-silico experiments without considering the change of dimensionality. However, publications rarely discuss the process of deriving their 2D environments from 3D ones.

Consideration should be given to remedial action to make 2D models viable, for example, by implementing strategies to increase sink rates or decrease source rates of oxygen. Alternatively and among other possible approaches, decoupling the maximum number of cells per position from the overlying 3D model from which a 2D one has supposedly been derived. While these approaches may produce valid results, the question then is what impact does the aforementioned set of remedial actions have on the reliability of results. For example, setting the maximum number of agents per cell to an arbitrary value may result in a model accepting unrealistic cell concentrations in tissues. Or, tweaking source and sink values may result in the model working with hyper- or under-oxygenated tissues.

More importantly, enforcing artificial constraints on the model subtracts from its emergent behaviours. The main value of these models lies in them revealing unexpected and insightful patterns from a set of simple, underlying rules. Developers should specify the behaviour of individual agents so that they mimic as closely as possible their biological counterparts. Developers should not try to direct the evolution of the model in a specific direction as they would be doing when imposing remedial strategies. To do so would be to constrain the simulation setting it up towards a particular conclusion or end-state, which would limit its usefulness towards gaining further understanding of biological phenomena.

It is worth adding some biological considerations to the limitations of the 2D models we discussed. Alongside unrestricted growth where oxygen is no longer a limiting factor, other properties of tumours will result as altered. In 2D, we will not observe regions of Hypoxia in the cancer mass nor the formation of a necrotic core. Secretion of VEGF will not occur, which in turn will lead to no angiogenesis or blood vessel development. Hypoxia-inducible factor (HIF) pathways would not be activated in simulations where these were implemented. HIF pathways mediate significant phenotypic alterations in cancer cells [78, 79], therefore not being observed in 2D. Consequently, it is clear how the impact of oxygen concentrations on 2D models significantly negatively affects accuracy and reliability.

We also emphasise the numerical scale of our simulations. That is the fact that we did not rely on results from a single simulation but, instead, averaged results from multiple runs. Confidence in our results is further reinforced by running 200 simulations per simulation set. A simulation set refers to, for example, the group of 2D simulations having 300 cancer cells per agent. This reduces the likelihood of results due to the model's stochastic nature, allowing us to infer that they are genuinely representative of the model's behaviour.

This contrasts with other publications surveyed, where the number of simulations run is not discussed. Given models often include elements of stochasticity, multiple iterations of these need to be evaluated to obtain a clearer picture of their underlying behaviour and evolution. Multiple runs also allow us to explore the search space for parameters whose value has not been established with certainty, assessing the model's sensitivity to changes in these and informing us about the impact on the simulation results. Search strategies for parametrisation, such as but not limited to grid and random search, should be explicitly mentioned and discussed. Distribution parameters should be reported, and the number of draws from each distribution and relevant information. This is necessary for experimental reproducibility.

We will now consider costs. We ran our simulations on Amazon AWS EC2 instances of class r4.large. These feature two vCPUs and 16GiB of RAM. These are between \$0.01 and \$0.001 per hour, depending on the specific plan selected. A single 2D simulation took on average one hour and a single 3D simulation on average four hours. We will also note that while we used a large fleet of instances to run a high number of simulations, individual simulations could be run in a reasonable time within an average office laptop.

Hence, the price of simulations in 2D or 3D is comparable. Even for the more expensive 3D simulations, a single simulation set comprising 200 simulations will still cost under \$1. Given the severe limitations of 2D we discussed and the negligible differences in computational and dollar cost, we think our investigation makes a robust case for 3D as the sole viable solution.

In summary, careless or approximate projections of 3D environments into 2D ones,

or the creation of 2D environments without consideration for their 3D counterparts, raises questions about the model's validity. This discussion is omitted, as is the case in many publications. This impacts the reliability of results as it is not clear what the model has achieved accuracy or resolution. Further, due to the negligible difference in costs between 2D and 3D, we believe there is a clear case for the latter to be preferred over the former.

5.5.5 Conclusion

We have developed an agent-based model of tumour growth with continuous elements to account for soluble diffusion and simulated tumour growth in 3D and derived 2D environments at multiple resolutions of cancer cells per agent. In addition, we have addressed the issue of parametrisation by wherever possible relying on values in the literature and where these were unknown using well-established search techniques such as grid search and random search.

Error rates for each simulation were calculated by comparing growth curves produced by our model to those reported in the literature. We compared error distributions and average error values for 2D and 3D at each resolution. Initial results suggested, purely based on error rates, that 3D simulations consistently produced low error rates and outperformed 2D at higher resolutions. On the other hand, 2D simulations outperformed 3D at lower resolutions, and their error depended on the resolution.

Further analysis revealed that the combined impact of cancer cells per agent resolution and 3D to 2D mapping significantly altered the topology of the 2D environment. Because of this, oxygen concentrations, a factor limiting growth in 3D, no longer affected tumour growth in 2D. As a result, tumour growth in 2D became a function of the initial volume predominantly and was no longer modulated by oxygen availability as is the case in 3D.

These results highlight the challenges of simulating a phenomenon that occurs in 3D, such as tumour growth, in 2D. Furthermore, neglecting to explicitly derive the 2D environment from a 3D one, as many studies do, raises questions about the model's validity and reliability of the results.

Our results suggest that 3D models should be preferred to 2D ones. However, in cases where 3D ones may not be implemented, a clear discussion should be provided regarding the derivation of the 2D model from a 3D one.

In conclusion, while it is true that a single study may appear as too limited evidence, we would like to emphasise how this highlights how "minor" discrepancies between 2D and 3D (Eg: Preventing the formation of Hypoxia) may significantly affect the model's performance. The problems we illustrate in this Section are likely to affect any model implementing more complex dynamics impacting their results and reliability.

Further, we believe it likely that other such discrepancies exist. To hypothesise: The topology of the vessel network developing during angiogenesis would be different between a 2D and 3D environment, again impacting the formation and/or distribution of regions of Hypoxia in the tumour.

While additional studies in this area may be beneficial, we believe that our investigation provides strong evidence supporting 3D models being reliable, robust and more correct than 2D ones.

Chapter 6

Warburg Investigation

6.1 Introduction

Cancer cells undergo an array of phenotypic alterations generally referred to as the Warburg effect. While the causes of such alterations are still being researched, in-silico models have previously suggested that they may benefit tumour expansion. Eg: Tissue acidification facilitating cancer growth, see Archetti [80] , Smallbone et al. [81] and Gerlee et al. [36]. These changes include switching from oxygen-based (aerobic) to glucose-based (anaerobic) metabolism even when sufficient oxygen is present. In cells where this occurs, the over-expression of membrane glucose transporters, especially GLUT1, has been observed. It has been suggested that glucose-conjugate drugs could be used to selectively target them, facilitating a therapy specific to tumour cells and reducing the incidence of side effects in patients.

A question easily derived from such a proposal is the proportion of cells adopting the Warburg phenotype and over-expressing GLUT1. The answer to this would likely be related to the efficacy of such a therapy. However, the literature reports contrasting findings, leaving the question open for investigation. Concerning our investigation, we propose to simulate tumour growth under various hypotheses of increasingly larger portions of the cancer cell population adopting the Warburg phenotype. Via systematic analyses of how closely the results of each simulation align with the results from laboratory studies, we will be able to suggest which hypothesis is most supported. Our model incorporates the fundamental aspects of cancer biology, including angiogenesis (tumour-stimulated development of new blood vessels), the expression of hypoxia-inducible factors (HIFs, a set of genes which mediate the tumour's response to low oxygen concentrations) and HIF-mediated characteristics, such as proliferation rate, VEGF secretion rate and metabolic rate, and the diffusion of solubles such as oxygen, glucose and VEGF. Contrary to similar studies, we simulate tumour growth in a 3D rather than 2D environment, further contributing to the model's accuracy and reliability, see Chapter 5.

Our model expands on previous work in several ways. First, developing the model in 3D allows us to create an effective fitness function. As will be discussed, the fitness function works by comparing the volume of the virtual tumour to the expected one from the laboratory measurements throughout its growth, see Section 6.2.5. A 2D virtual environment would make it extremely difficult, if not impossible, to develop such a function. We should also consider the dynamics of tumour development. In the case of angiogenesis, blood vessels often form complex networks which develop spatially in three dimensions. The morphology of these affects the growth of the cancer mass significantly, which is captured accurately, which would not be possible in 2D.

Further advantages of 3D modelling over 2D include better capturing the oncogene activation [53], protein expression and drug sensitivity [52, 64, 65] as well as more realistic biochemical and biomechanical environments [66] and a better translation of

the pathophysiological features of the tumour environment [67]. Finally, 3D models also better capture inter-cellular signalling pathways [68].

Our results will be further supported by the decision to not rely on a single set of parameter values. Instead, we thoroughly explore the space of reasonable values. In drawing our conclusions, we can therefore be mindful of the sensitivity of our model to various parameters and whether specific thresholds that cause divergences in observed behaviours exist.

6.2 Our Model

We now present implementation details specific to the Warburg investigation. Note that this extends the general principles of our simulation, detailed in section 5.3.

6.2.1 Experimental Setup

As outlined in Chapter 1, it is a well-established hypothesis that cancer cells adopt what is known as the Warburg phenotype. However, the proportion of cells within a cancer population that adopt such a phenotype is still a matter of debate. We propose to address such a question by simulating tumour growth under the hypothesis of cells having varying probabilities of adopting such a phenotype. The Warburg Phenotype is observed in most tumour masses throughout their development as "a major driver of the cancer progression machinery, resistance to conventional therapies, and poor patient outcome" [82]. To align with this, we will quickly mention that in our simulation such a transition is irreversible. To account for the fact that some cells may not exhibit the Warburg Phenotype, this is not inherited by daughter cells upon division. That is, upon being added to the simulation following division each cell initially exhibits the non-Warburg Phenotype and, at each epoch, has a given probably of acquiring it. The Warburg Phenotype further mediates a variety of phenotypic alternations such as enhanced glucose uptake rates [83, 84], proliferation rates and expression of Hypoxia-Inducible Factors. It is also important to note that cells adopting the Warburg phenotype no longer depend on a the availability of a minimum oxygen concentration, whereas their non-Warburg counterparties do.

We tested the following five Probabilities of Warburg Switch (PWS): 5%, 13%, 20%, 40% and 80%. The rationale for selecting such values was to sample both the "high probability" and "low probability" spaces of PWS. Starting with PWS of 5% and 13%, and given that results were not significantly different, we progressively increased probabilities by greater amounts until eventually, we found that growth curves were significantly different between 80% and other values. Each of these was tested under the following six assumptions that non-Warburg cells require

the following minimum concentrations of oxygen for survival: 12mmHg, 15mmHg, 16mmHg, 17mmHg, 18mmHg and 20mmHg. These values were selected to sample the range of biologically plausible values for this parameter in a grid-search manner. A total of 30 PWS/minimum oxygen concentrations pairs were therefore considered. One hundred simulations were run with other parameter values drawn from a normal distribution for each.

A final word concerning the visualisation of results. We wish to investigate various PWS hypotheses and establish which allow our model to fit growth curves from empirical studies better, see Figure 6.2a. Therefore, our strategy for PWS sweeping has been to increase its value in 'steps' and create batches of simulations for each of these. Other parameters have been drawn from random distributions, as we discuss later. The claim we wish to make is that irrespective of other parameter values, there is a clear correlation between the choice of a PWS value and the model's performance. To illustrate this, we have visualised the relation between PWS, minimum oxygen concentration required by non-Warburg cells and absolute mean error. This allows us to illustrate how the correlation between fitness and PWS is maintained, although more or less strongly, across different value combinations of other parameters.

The high dimensionality of the search space means that if we proposed a similar visualisation for every other parameter, this would become intractable. However, it is worth remembering that while the aggregated visualisation does focus on a small subset of parameters, individual records which have been combined to produce it do capture the underlying heterogeneity of values across the entire parameter set.

6.2.2 The Cell Life-Cycle

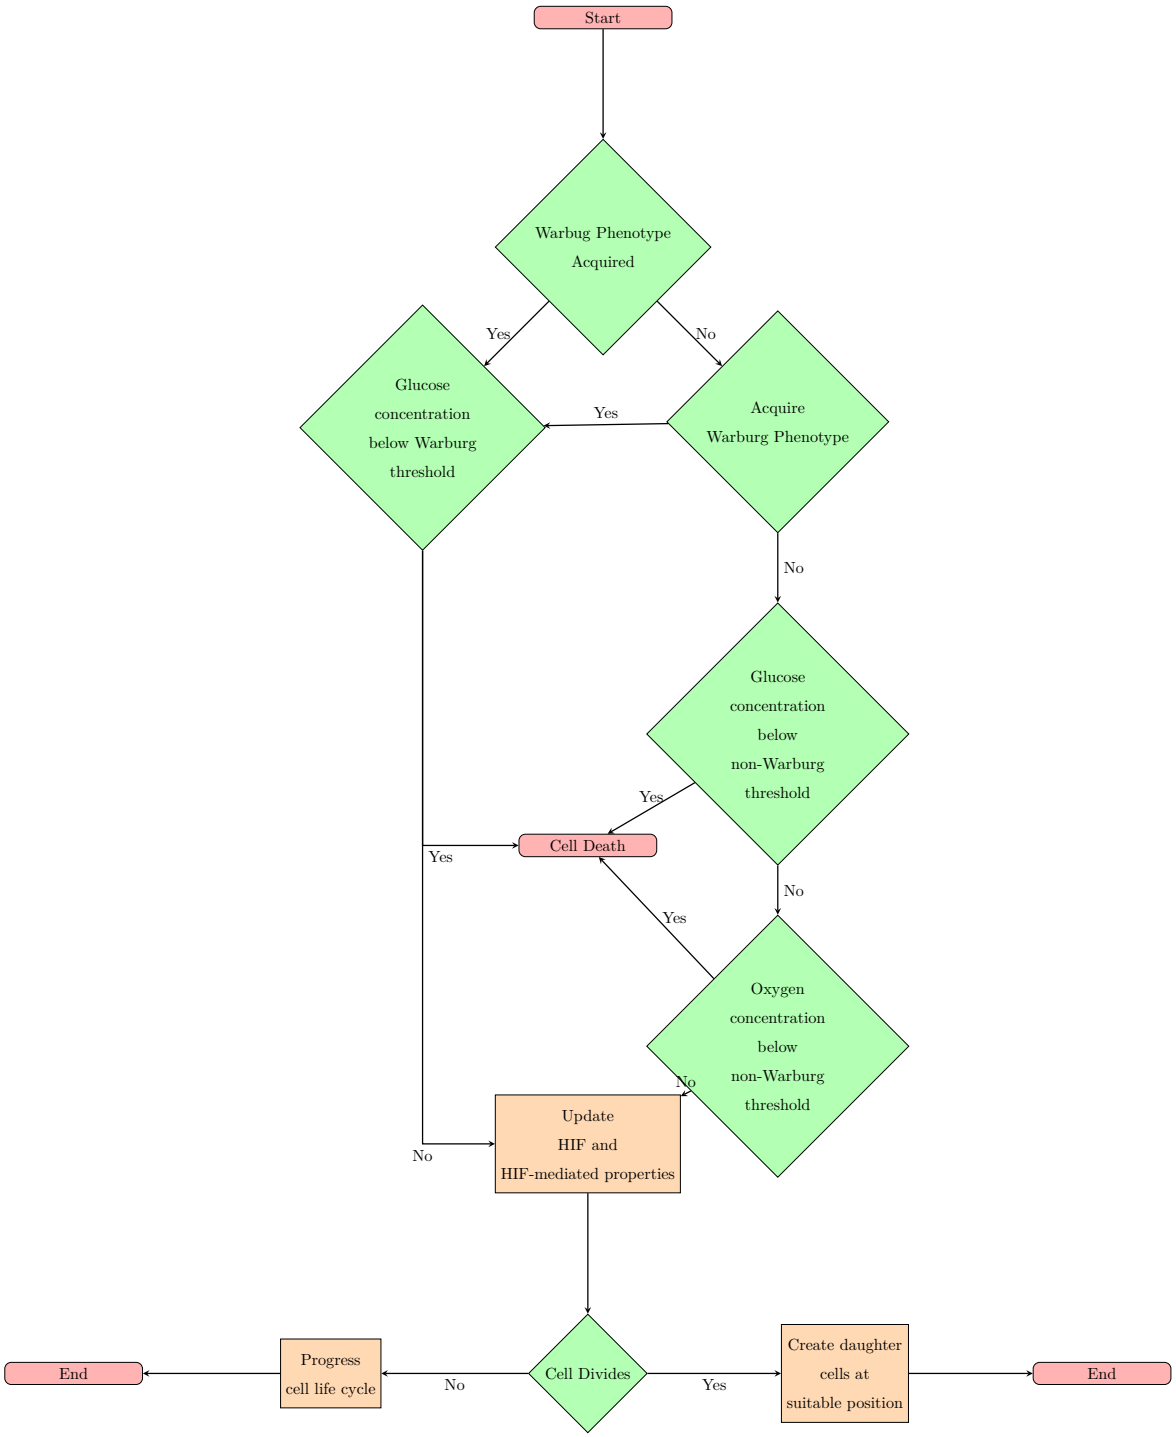

Figure 6.1: High-level summary of the modelling of cancer cells' epochs. Non-Warburg cells are modelled as dependent on oxygen and glucose, whereas Warburg cells are dependent on glucose only. The 'Progress cell life cycle' process includes deciding whether to progress from stage G1 into Synthesis. Creating a daughter cell at a suitable position involves identifying the adjacent position with the fewest agents. If all neighbouring positions are saturated, then division does not occur.

Figure 6.1 describes the behaviour of cancer cells specific to the Warburg Investigation. This is similar to what was described in our case studies but extends that to consider the probability of a cell acquiring the Warburg Phenotype.

HIFs are a set of transcription factors that respond to low oxygen concentrations, also known as hypoxia [85]. HIFs are known to mediate adaptations, which allow cells to survive in low-oxygen conditions, such as a transition from aerobic respiration to glycolysis. In cancer cells, HIF expression has been observed in conditions of normoxia (Ie: Normally expected oxygen concentrations) as well [78, 79, 86–88]. In our model, HIF expression mediates an increase in the probability of progressing into Synthesis, promotes VEGF secretion and reduces oxygen uptake rates.

We finally come to the matter of cancer cells adopting the Warburg phenotype, directly related to our research question. We model this as a boolean flag where cells either have or have not acquired such a phenotype. The probability of the cancer cells switching from a non-Warburg to a Warburg condition is set as a model parameter and shared by all cancer cells. Non-Warburg cells may, at each epoch, adopt the Warburg phenotype with such a probability. The transition is irreversible but not inherited by the daughter cells. The survival of Non-Warburg cells depends on minimum oxygen and glucose concentrations being available at their position. Warburg cells, on the other hand, only rely on glucose. However, the minimum glucose concentration required by Warburg cells is higher. Warburg cells will also begin expressing HIF at higher oxygen concentrations than their non-Warburg counterparts. In practice, this results in Warburg cells promoting angiogenesis, proliferating more rapidly and consuming less oxygen. All of this is consistent with tumour biology.

6.2.3 Starting Conditions

As mentioned previously, we assumed an environment represented by a grid of size 20 x 20 x 20 positions.

At the beginning of the simulation, every other position in the environment was seeded with a tip cell, forming a 'chessboard mosaic' initial blood vessel configuration. Ie: Given a position (x,y) with a tip cell, the positions at coordinates $(x+1,y),(x-1,y),(x,y+1),(x,y-1)$ would not have a tip cell. The positions at coordinates $(x+1,y+1),(x-1,y+1),(x+1,y-1),(x-1,y-1)$ would have a tip cell. This replicates what is observed in healthy tissues, where ratios of blood vessels to tissue are carefully maintained to ensure sufficient oxygenation is allowed. That, even to the point where capillaries are separated from each other even by only individual cells [89]. Every other position not occupied by a tip cell was seeded with a static 'tissue placeholder' that simply had the role of simulating the tissues' oxygen and glucose uptakes. The effect of such a placeholder value became negligible once the cancer cells took over a position.

A regular patch was selected from position (8,8,8) to position (11,11,11), boundary values included. This effectively created a patch of size four positions x 4 positions

x 4 positions, whose centre coincided with the environment grid's. Each position within such a patch was seeded with cancer cells in an amount equal to the 'Initial Number of Cancer Cell Agents' parameter. The cell life cycle stage of each cancer cell was randomised.

Determining the size of the initial tumour mass was especially challenging. Any volume observed in a laboratory would be an estimate that necessarily includes a margin of error comprising healthy tissue cells and dead cells. We decided to consider the central patch as a "target area" to seed cancer cell agents. Being unable to justify an exact number of cancer cell agents per position, we used random search to determine ranges of initial cancer population size, which would allow the model to perform. The size of the "target area" patch itself was estimated via a preliminary parameter sweep exploration. For example, we found that smaller patch sizes led to the initial tumour mass being too small, resulting in smaller volumes than expected.

Finally, the choice to randomise the initial state of each cancer cell agent is the only appropriate one. That is because of the extreme unlikelihood of all cells in a tumour being in the same phase of the cell life cycle.

6.2.4 Model Parameters

The cell life cycle is 24 hours, the minimum oxygen concentration for the survival of non-Warburg cells is 12 mmHg, the threshold for cells to enter hypoxia is 20 mmHg. In contrast, non-Warburg cells start expressing HIFs at an oxygen concentration of 75 mmHg. About the process of angiogenesis, a time interval of 24 hours is required between the division cycles of endothelial cells. Parameter values that could not be extracted from the literature were assigned through established parametrisation techniques such as grid search and random search, see Section 6.3.1.

6.2.5 Our Fitness Function

The development of a fitness function allows us to objectively evaluate the simulation results and obtain a measure of performance. Towards this aim, we compare the growth curves generated by our model to those reported by Liu et al. [90], shown in Figure 6.2a. Figure 6.2b shows how a sample curve generated by our model compares to a polynomial fit of the expected curve.

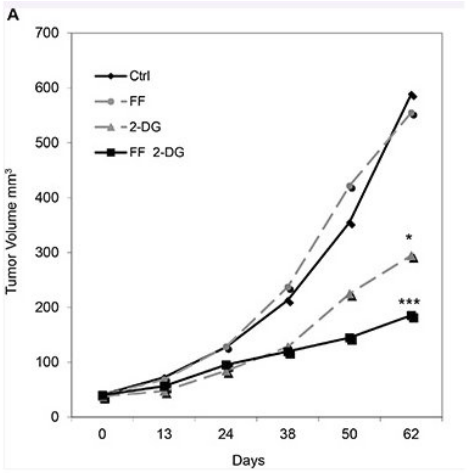

(a) Growth curves from empirical studies against which we will be comparing those generated by our model. These are as reported by Liu et al. [90], who present the methodology as: 'Experiments consisted in implanting human melanoma xenografts in nude mice. Mice (n = 8/group) were treated either FF only (100 mg/kg/day, oral gavage daily), 2-DG (41 ug/ml/hr, ALZET pumps subcutaneously implanted) only, or a combination of FF (100 mg/kg/day) and 2-DG (41 ug/ml/hr). Mice in the control group had saline oral gavage and ALZET pump delivery. Tumour volumes were measured weekly after treatment. Representative tumour volumes of one of two independent experiments are shown.

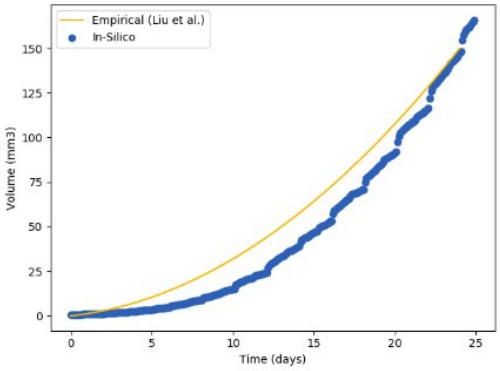

(b) Original growth curve reported by Liu et al. (Empirical) and the fit we produced to evaluate our simulations and a sample growth curve generated by the model (In-Silico).

Figure 6.2: Growth curves from medical literature and comparison of growth curve generated by one of our simulations to the polynomial fit of the expected curve. Our derived empirical curve has been smoothed to account for the unlikeliness that the growth rate of the tumour would be linear between points reported in Figure (a), with changes in this only occurring at those sampled points. Eg: Days 13, 24, 38, etc... Instead, we think that the straight lines representing growth between sampled points in Figure (a) are not representative of actual volumes in-between readings, with a smoothed line better approximating this.

Our initial work consisted in comparing growth curves under control treatment. As a first step, a second-degree polynomial $\beta : \mathbb{R} \longrightarrow \mathbb{R}$ was fit to the following data points of the control time-series: (0, 0), (13, 50) and (24, 150) to generate a benchmark function such that, for a given time-point t, $\beta(t)$ would return an estimate of the tumour volume observed in the laboratory setting. The reason for only using three data points was that the others were outside our domain of interest. That is, beyond the first month. An alternative approach, which we believe would have yielded similar overall conclusions to our experiment, would have been to use all points to generate the curve and then only use the part corresponding to the first month. We can then take a tumour volume value reported by our model at time t, v^t_{actual},

and obtain an absolute error, ε^t by comparing it to the expected tumour volume $v^t_{expected} = \beta(t)$, as detailed in Equation 6.1.

$$\varepsilon^t = abs(v^t_{actual} - v^t_{expected}) \quad (6.1)$$

If we then consider a time-series of tumour volumes reported by our model after simulating cancer growth of a period of time T, the average mean error of the simulation ε_{avg} could be calculated as the average of all $\varepsilon_i; 1 \leq i \leq T$. It then follows that a lower value of ε_{avg} would indicate the growth curve produced by the simulation more closely aligned to the expected one. In contrast, a higher value would indicate more discrepancy.

The use of a fitness function of this kind allows analysing large amounts of simulations rigorously and confidently determining which perform better. In previous studies, more subjective approaches would be adopted where authors would make a subjective judgment regarding the performance of a simulation. Ie: Eyeballing curves. This would make it challenging to assess multiple re-runs of a simulation or compare different studies. In cases where multiple people were involved in judging simulations, a question of standard alignment and bias would also be raised. We now have a systematic and consistent strategy to determine how well a simulation performs by developing a mathematical approach.

As we will see in the following sections, our fitness function can also be exploited to implement various parameter space exploration techniques, improving on the state of the art parametrisation techniques. Figure 6.3 shows samples of model simulations that provide good and bad fits of the expected curve.

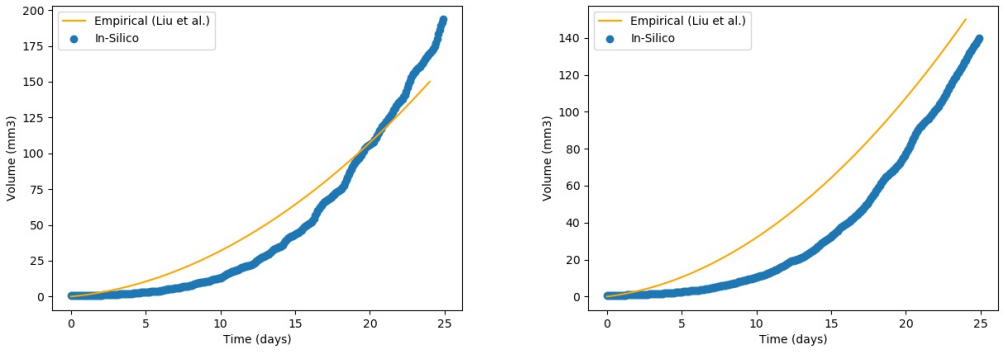

(a) Sample simulation resulted in a good fit. Absolute mean error of 12.40 mm^3.

(b) Sample simulation, which resulted in a poor fit. Absolute mean error of 20.27 mm^3.

Figure 6.3: Comparison of simulations that produced a good and poor fit of the expected growth curve for illustration purposes.

6.3 Results

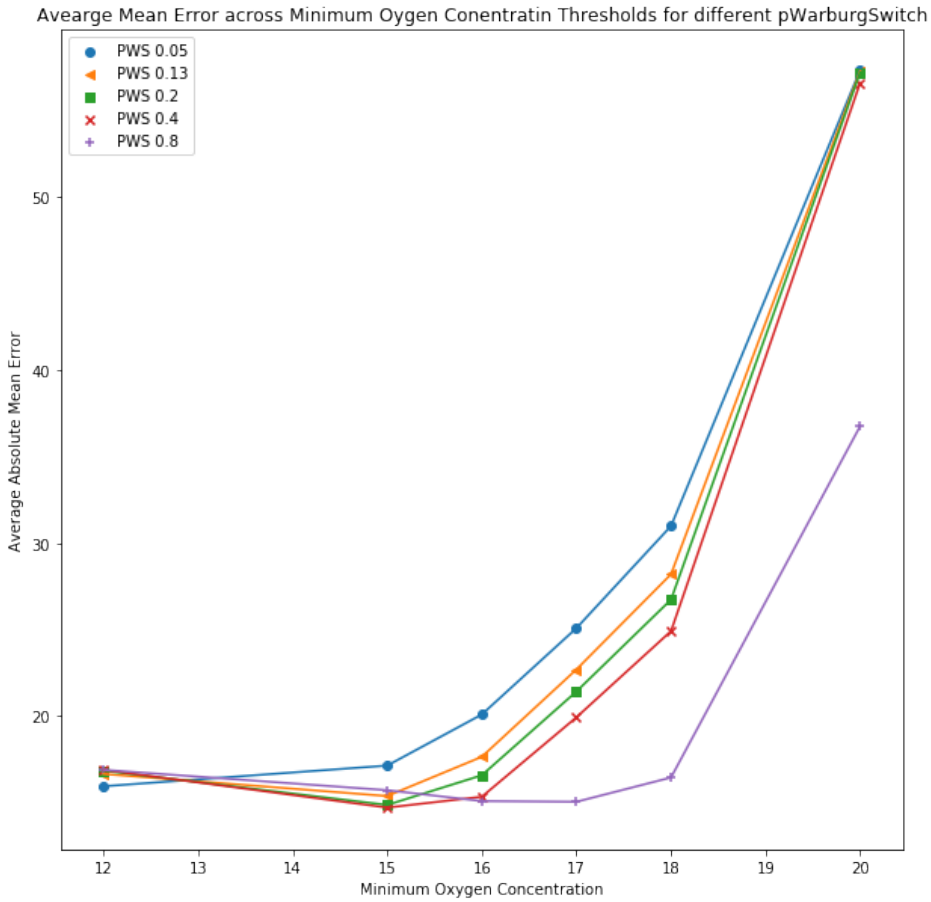

Figure 6.4: The average mean error for each PWS population across multiple minimum oxygen concentration thresholds. PWS refers to the probability of Warburg switch, the probability of cancer cells adopting the Warburg phenotype. As minimum oxygen concentration increases, pressure against populations with a lower proportion of Warburg Cells also increases. As a result, populations where the overall Warburg population is smaller will observe a slower growth due to non-Warburg cells dying. This, in turn, results in higher absolute mean error values.

Figure 6.4 suggests that higher probabilities of cancer cells adopting the Warburg Phenotype allow the model to produce growth curves that better fit the expected ones. This is particularly the case for higher values of minimum oxygen concentration. In particular, for higher oxygen requirements, the difference in error values between higher and lower PWS populations is comparable to the differences in volume between control and therapy curves in empirical studies (Figure 6.2a). This allows us to make an initial suggestion, later confirmed by statistical analysis (Table 6.1), that differences in errors are statistically significant.

The choice of the value range for Minimum Oxygen Concentration was based on this generally being reported as approximately between 16 and 18mmHg. However, we make allowances for a margin of error and consider a range of values close to this.

To assess whether such differences are significant, we compared absolute mean errors for every PWS group against the PWS=80% group for each level of minimum oxygen concentration. The resulting p-values are reported in Table 6.1.

	PWS 5%	PWS 9%	PWS 13%	PWS 17%	PWS 20%
Minimum Oxygen Concentration Threshold 0	0.01336	0.01110	0.01547	0.01306	0.07381
Minimum Oxygen Concentration Threshold 3	0.00945	0.01214	0.01714	0.00982	0.07868
Minimum Oxygen Concentration Threshold 6	0.00053	0.00224	0.00347	0.00273	0.06503
Minimum Oxygen Concentration Threshold 9	0.00016	0.00372	0.00539	0.00203	0.07551
Minimum Oxygen Concentration Threshold 12	0.00092	0.00006	0.00025	0.00113	0.05383
Minimum Oxygen Concentration Threshold 15	0.00031	0.00003	0.00017	0.00127	0.00202
Minimum Oxygen Concentration Threshold 18	0.02911	0.00011	0.00022	0.00459	0.00019

Table 6.1: Summary of p-values resulting from comparing populations of average mean errors from all PWS groups against PWS=80% for each minimum oxygen concentration threshold. This was done by using a two-tailed t-test for the means of two independent samples of scores as distributed in the SciPy Stats library: (https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.ttest_ind.html), where each group had the same number of members ($n = 100$).

Table 6.1suggests that for minimum oxygen concentration thresholds above 12, the average mean errors of all PWS groups are significantly different from those of the PWS=80% group for $p < 0.05$. This, together with our analysis of Figure 6.4 suggests that a higher probability of cells adopting the Warburg phenotype results in our model is more closely aligned with the expected growth curve.

Finally, we decided to investigate the direction of the error understand if this is consistent across all trials. That is if the model consistently over-estimates or under-estimates tumour volume at a given point in time across all simulations. This is shown in Figures 6.5a and 6.5b.

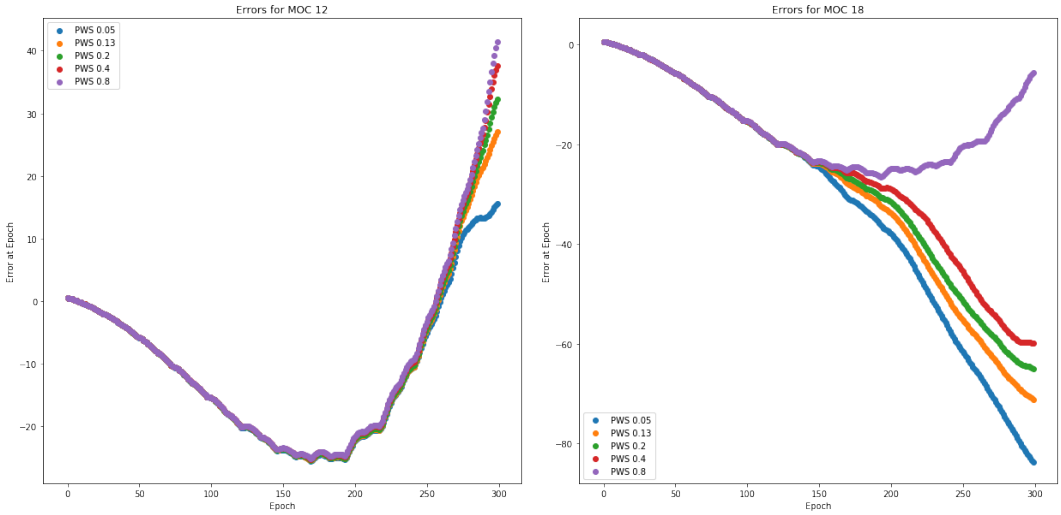

(a) We observe no significant difference over the initial part of the tumour growth for lower minimum oxygen concentration requirements. However, over the last 50 epochs, higher PWS values allow for larger tumours over the last period. This is probably due to the tumour entering avascular growth and tumour masses benefiting from a higher Warburg population which isn't sensitive to regions of oxygen deficiency.

(b) For higher minimum oxygen concentration requirements, we observe that all groups other than PWS=80% cannot grow at the required rates. The increasing negative magnitudes indicate that the tumour is lagging further behind the expected volume. On the other hand, the PWS=80% group can maintain a growth rate comparable to the expected one through a growing Warburg population.

Figure 6.5: Mean error at each epoch for each PWS group from the second round of experiments. PWS refers to the probability of Warburg switch, the probability of a cancer cell acquiring the Warburg phenotype at each epoch.

Results are consistent with the previous analysis presented, which suggest that, in the presence of a sufficiently high minimum oxygen concentration requirement, a higher probability of Warburg switch is required for the model to reproduce expected growth patterns. This is also shown in Figure 6.6, where we look at two sample growth curves taken from simulations with lower and higher probabilities of acquiring the Warburg phenotype and compare them to the expected growth curves.

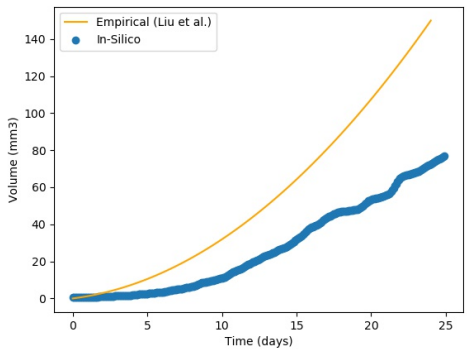

(a) Comparison of the growth curve generated by our model against the expected one for a simulation with minimum oxygen concentration levels for non-Warburg cells of 18mmHg and probability of acquiring the Warburg phenotype of 5%. Absolute mean error of 30.45 mm^3.

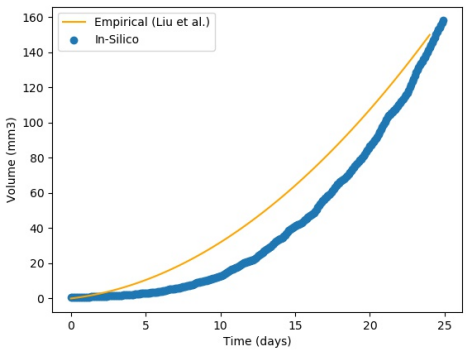

(b) Similar comparison, the minimum oxygen concentration of 18mmHg, but the probability of acquiring the Warburg phenotype is 80%. Absolute mean error is 14.85 mm^3.

Figure 6.6: Comparison of growth curves generated by our model against the expected ones. We compare lower and higher probabilities of acquiring the Warburg phenotype, with results suggesting that a higher probability leads to a more accurate simulation of tumour growth.

Figure 6.7 provides further insight comparing tumour properties for simulations where cells have a 5% probability of acquiring the Warburg phenotype and those where cells have an 80% probability. Figures 6.7c and 6.7d are especially of interest. These show changes in the proportion of tumours cells that have adopted the Warburg phenotype across time. For lower probabilities, we observe periodic sharp increases. These are due to the tumour mass growing to exceed the capacity of the existing vessel network. Hence, we observe a sharp decline in the number of non-Warburg cells which die because of lack of oxygen. This phenomenon is observed at regular intervals, and we believe it is one of the main reasons for differences in tumour volume.

We also considered the possibility that decreases in the size of the Warburg population could be due to insufficient glucose concentrations. Warburg cells do require higher amounts of glucose than non-Warburg ones. However, if that was the case, this behaviour should be observed independently of the probability of cells acquiring the Warburg phenotype. It should be more prominent for simulations with a higher PWS, as more cells would be Warburg cells. The fact this is only observed at lower probabilities, on the other hand, supports our original explanation.

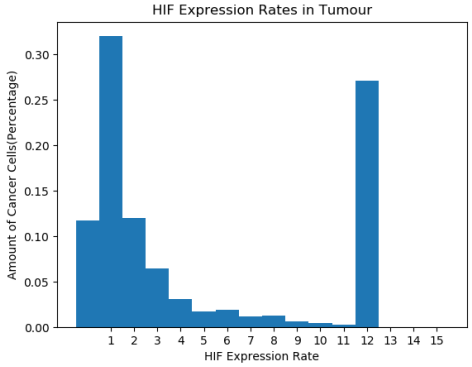

(a) Distribution of HIF expression rates at simulation end for a simulation with minimum oxygen concentration for non-Warburg cells of 18mmHg and probability of acquiring the Warburg phenotype of 5%. In most cases, this suggests cells have sufficient oxygen for survival. Ie: Without acquiring the Warburg phenotype. Conversely, the absence of significant amounts of cells with higher HIF expression means that cells rapidly die because of lack of oxygen, with a peak around the expression rate of 12 is due to the small Warburg population.

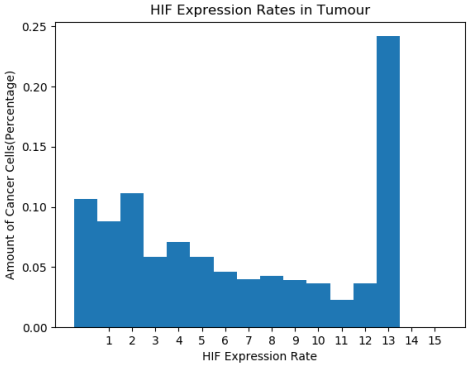

(b) Distribution of HIF expression rates at simulation end for a simulation with minimum oxygen concentration for non-Warburg cells of 18mmHg and probability of acquiring the Warburg phenotype of 80%. The presence of a higher proportion of Warburg cells allows for the survival of a larger part of the tumour mass in conditions of oxygen that would not normally support cellular life. Hence, allowing for a more heterogeneous spread of HIF expression rates.

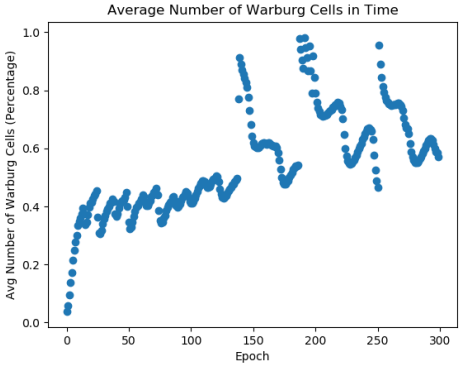

(c) Warburg population size for a simulation with minimum oxygen concentration for non-Warburg cells of 18mmHg and probability of acquiring the Warburg phenotype of 5%. Past approximately epoch 100, we observe the existing blood vessel structure unable to support the tumour mass. Hence, non-Warburg cells will deplete, resulting in a sharp increase in the proportion of Warburg cells. This will lower again as non-Warburg cells proliferate once more until exceeding the vessel's capacity and rapidly depleting again. This oscillating behaviour is repeated across periods of approximately 50 epochs.

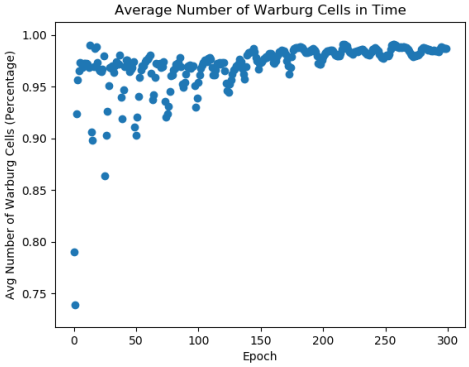

(d) Warburg population size for a simulation with minimum oxygen concentration for non-Warburg cells of 18mmHg and probability of acquiring the Warburg phenotype of 80%. A high probability of cancer cells acquiring the Warburg phenotype leads to these rapidly taking over the tumour mass. In later stages of growth, when blood vessels are unable to support the tumour mass, we observe a near-total dominance of Warburg cells.

Figure 6.7: Further comparisons of tumour properties for simulations with a probability of acquiring the Warburg phenotype of 80% (left) and 5% (right). We present a summary of the distribution of HIF expression rates at the simulation end and the proportion of Warburg to non-Warburg cells across time.

6.3.1 Parameter Space Search

Our model has 21 parameters which are categorised as described in Table 6.2.

Parameter Category	Description	Number of Parameters
Amber	Preliminary investigations suggest that the model is highly sensitive to this parameter, and either there are no known values for it in literature or literature reports conflicting values.	4
Yellow	Either preliminary investigations suggest the model is highly sensitive to this parameter or there are no known values for it in literature, or conflicting values are reported. But, it is not the case both conditions are met.	5
Green	Either preliminary investigations report a low sensitivity to this parameter, and its value is fairly certain from literature, or the value is thoroughly established in the literature.	12

Table 6.2: Summary of definitions for parameter categorisation as Amber, Yellow or Green and number of parameters that falls into each category.

While values for green parameters may be kept constant across simulations, it becomes hard to justify the choice of using a single value for yellow or amber ones. We, therefore, took the opportunity while working towards our main research question to explore the applicability of grid search [76, 91] and random search [55], used to tune hyperparameters of traditional machine learning models, to our model. We first explored the applicability of grid search. In the first instance, we selected the parameters we wanted to tune. Grid search involves selecting a set of values to test for each parameter under consideration and then running the model with each possible value combination. The number of parameter value combinations that have to be tested is the Cartesian product of all parameter value sets. Ie: If we have parameter sets p_1, p_2 and p_3 where $|p_1| = 2$, $|p_2| = 3$ and $|p_3| = 4$ then we would have 24, $(2 \times 3 \times 4)$, parameter combinations to test. Grid search suffers from the 'curse of dimensionality where, as the number of parameters increases, the computational cost eventually becomes intractable. We, therefore, restrict our application of grid search to the four amber parameters detailed in Table 6.3.

Parameter	Description	Rationale	Units
Minimum VEGF Concentration	Minimum concentration of Vascular Endothelial Growth Factors at a position for Endothelial cells to engage in angiogenesis.	No literature value, high correlation (-0.79).	Arbitrary Units
Initial Number of Cancer Cells	Initial number of cancer cells at each position in the initial seeding patch.	Difficult to extrapolate from the graphs reported in the paper we use to obtain empirical data, high correlation(0.98)	Scalar (Number indicating th
Maximum Glucose Uptake Rate	Glucose uptake rate of cells that have adopted the Warburg phenotype.	High correlation (-0.85), value unknown from literature.	pMol/time unit
Minimum Glucose Concentration Warburg	Minimum glucose concentration required for the survival of cells that have adopted the Warburg Phenotype.	No literature value, high correlation (-0.74)	pMol

Table 6.3: Summary of amber parameters and rationale for why they have been classified as such. Correlation values refer to Pearson coefficients between parameters and final tumour volume obtained through initial explorations involving single parameter sweeps. mmHg/time unit refers to the change in partial pressure at the position over the time unit considered for diffusion solving owed to the action of the cell acting as a source or sink of the soluble.

A summary of parameter values used in our grid search is provided in Table 6.4.

Parameter	Values for Grid Search
Minimum VEGF Concentration	1, 3,5 ,7
Initial Number of Cancer Cells	1, 3, 7, 10, 13, 17, 20
Maximum Glucose Uptake Rate	10, 15, 20, 25, 30
Minimum Glucose Concentration Warburg	10, 15, 20, 25, 30

Table 6.4: Summary of values for amber parameters explored in Grid Search. Value ranges were obtained by observing the model's behaviour in an initial set of preliminary investigations involving single parameter sweeps. Generally speaking, values smaller than the lower end of each range or greater than the higher end of each range caused the model to produce obviously unsuitable results. A total of 700 simulations were run, corresponding to each permutation of parameter values.

- A model was classified as Non-performing if either:
  - At simulation end, the simulated tumour volume was less than 90% of the

expected one;

- At the simulation end, the simulated tumour volume was at least 13% larger than expected.

- Performing, otherwise.

The distribution of results at simulation end is summarised in Figure 6.8. A substantial majority of simulations is non-performing, indicating the low payout against a high computational cost of grid search.

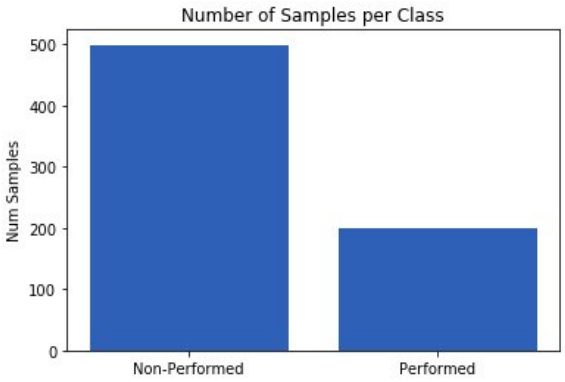

Figure 6.8: Summary of distributions between non-performing and performing simulations following grid-search on amber parameters. A majority of simulations was found to be non-performing.

We then consider parameter space exploration using random search. As a quick summary, random search involves generating values from parameters by drawing them from a distribution, where one distribution is defined for each parameter. One hundred simulations were run, details of the distributions are provided in Table 6.5.

Parameter	Mean	Standard Deviation
Minimum VEGF Concentration	5	1
Initial Number of Cancer Cells	14	2
Maximum Glucose Uptake Rate	20	3
Minimum Glucose Concentration Warburg	15	3

Table 6.5: Summary of distribution parameters for random search on amber parameters. Normal distributions were used. At each drawing, if a value extracted was less than 0, the sample was discarded and re-drawn. In the case of Minimum VEGF Concentration, an upper bound of 10 was also in place. Values picked for means are approximately the median values of the corresponding ranges for grid search.

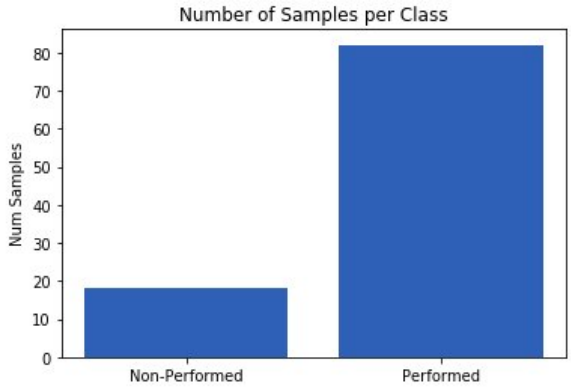

Figure 6.9: Summary of distributions between non-performing and performing simulations following random-search on amber parameters. A substantial majority of simulations was found to be performing. Therefore, the same principle for classification between performing and non-performing was applied.

Comparing Figures 6.8 and 6.9 we can draw several conclusions. First, random search produces a significantly higher ratio of performing to non-performing simulations. In addition, random search is also able to obtain several performing simulations compared to grid search with a significantly lower number of trials. In essence, our results suggest that random search significantly outperforms grid-search on our model.

Random search is also known for handling higher dimensionality problems than grid search. This means we consider more parameters than we would with grid search. We, therefore, decided to perform a random search on the union of amber and yellow parameters. A summary of yellow parameters is provided in Table 6.6.

Parameter	Description	Rationale	Units
Minimum Glucose Non-Warburg	Minimum glucose concentration required for the survival of non-Warburg cells.	Conflicting values in literature but a low correlation with the model's performance. (-0.12)	mmHg
Oxygen Diffusivity	Diffusivity of oxygen in tissues.	A large range of values reported in the literature and highly correlated to the model's performance. (0.70)	$\frac{cm^2}{s}$
Minimum Probability of Synthesis	Minimum probability of a cell progressing from G1 into Synthesis. This increases with the expression of Hypoxia-Inducible Factors.	This is a model artefact, so it cannot be based on literature, plus a high correlation with model performance. (0.60)	Probability
Glucose Diffusivity	Diffusivity of Glucose in tissues.	A value for this is provided in literature but highly correlated with the model's performance. (0.77)	$\frac{cm^2}{s}$
Probability of Warburg Switch	Probability of a non-Warburg cell adopting the Warburg phenotype.	No literature value is found but fairly low correlation with the model's performance. (0.4)	Probability

Table 6.6: Summary of yellow parameters and rationale for being classified as such. Correlation values refer to Pearson coefficients between parameters and final tumour volume obtained through initial explorations involving single parameter sweeps. mmHg/time unit refers to the change in partial pressure at the position over the time unit considered for diffusion solving owed to the action of the cell acting as a source or sink of the soluble.

A summary of values used to generate distributions from which to draw values for yellow parameters is provided in Table 6.7.

Parameter	Mean	Standard Deviation
Minimum Glucose Non-Warburg	25	2
Oxygen Diffusivity	0.3	0.1
Minimum Probability of Synthesis	0.5	0.2
Glucose Diffusivity	0.001	0.0005
Probability of Warburg Switch	0.5	0.2

Table 6.7: Summary of distribution parameters for random search on yellow parameters. Normal distributions were used. At each drawing, if a value extracted was less than 0, the sample was discarded and re-drawn. In the cases where values refer to a probability, if the value is drawn was not between 0 and 1 (included), this was discarded and re-drawn. Decisions on values for each mean/standard deviation were made considering the results of preliminary investigations involving single parameter sweeps.

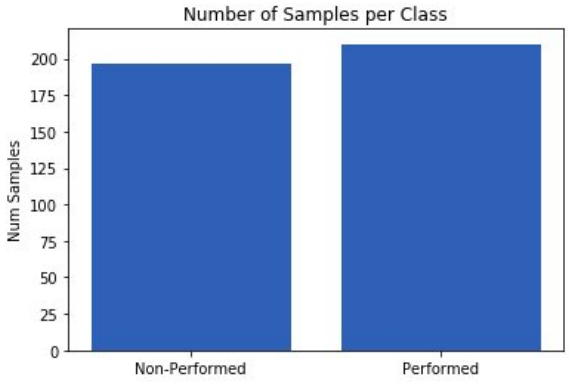

Figure 6.10: Summary of distributions between non-performing and performing simulations following random-search on the union of amber and yellow parameters. A total of 400 simulations were run. A slight majority of simulations were found to be performing.

Figure 6.10summarises the split between performing and non-performing simulations when a random search is performed over the yellow and amber parameters union. A majority of simulations are still performing, but this is now a slight majority rather than the substantial majority, which was observed for amber parameters only, see Figure 6.9.

We can speculate over several possible reasons for this. Firstly, it could be that either Figure 6.9, random search on amber parameters only, or Figure 6.10, the union of amber and yellow, were either the result of a lucky draw in the former case or an unlucky draw in the latter case. Experiments would need to be repeated with more simulations to ascertain whether the results hold. Yet, random search is known to converge rapidly, so this hypothesis appears unlikely. Alternatively, it could be that for one of the yellow parameters, there exists a smaller range of values leading to a performing simulation that can be obtained from our distributions. If that were the case, which appears more likely, that would explain the re-balancing of classes observed. Again, re-running experiments drawing values from different distributions would be required to explore this hypothesis.

Regardless, we have been able to show the suitability of random search to explore a highly-dimensional search-space in non-standard domain-specific models. As a result, we used these parameters to achieve our results. Specifically, without performing parameter space exploration, we would have been limited to ”hand-picked” values, which would have neither informed us about the sensitivity of our model to them nor given us the same level of confidence in the robustness of our results.

6.4 Conclusion

The proportion of the tumour mass and location of regions that adopt the Warburg phenotype is a question of significant scientific interest and central to the possibility of developing drugs with fewer side effects. Therefore, we explored the hypothesis of cancer cells having varying probabilities of adopting such a phenotype and compared

the resulting growth curves to the expected ones from empirical studies. In our model, we included the key aspects of the Warburg effect, such as the enhanced uptake rate of glucose, proliferation rate and expression of HIFs.

Across all assumptions of the minimum oxygen concentration required by non-Warburg cells for survival, a higher PWS leads to the model performing better. In all cases, the difference in absolute mean errors between lower (PWS=5%) and higher (PWS=80%) PWS (Figure 6.5b) simulations are statistically significant. In most cases, it is also greater than the difference observed between the empirical control and treatment curves (Figure 6.2a). That last statement is to say that if we were to say that the differences in our in-silico generated curves were not significant, then neither would the differences between the control and treatment curves in the study we take as a reference.

Our results, therefore, suggest that a sizeable majority of the tumour population must adopt the Warburg phenotype for cancer to grow successfully. As a result, it would be likely that a majority of the cancer cells over-expressed glucose transporters and would therefore make suitable targets for glucose-conjugate drugs.

From a biological perspective, we can suggest the following explanations for the results observed. First, as tumours grow, the existing vessel network eventually cannot provide sufficient oxygen and other nutrients to the cancer mass. This marks the transition from vascular growth to avascular growth and coincides with the beginning of angiogenesis and other HIF-mediated phenomena. The question then was whether such 'remedial action' could be enacted via only a minority of cancer cells adopting the Warburg phenotype or if a more substantial number was needed. Our simulations suggest that nearly the entire cancer cell population must contribute and adopt the Warburg phenotype to ensure its survival and growth during the avascular stage.

To summarise the biological interpretation of our results, we can say that the avascular development of the tumour mass has a precondition that the cancer cells should adopt the Warburg phenotype. It is not enough for a minority of the cancer cell population to adopt this phenotype, stimulate sufficient angiogenic development for supporting non-Warburg growth and lead proliferative activity. In contrast, the remaining cells maintain a more quiescent state.

We developed a model to simulate early tumour growth and the inception of the avascular phase. The model could be further extended to simulate further growth and development.

In the first instance, we consider the heterogeneity of the tumour mass. While previous works have modelled cancer bodies as homogeneous regarding their properties, we allow each cancer cell agent to differ in HIF expression rates, nutrient uptake rate, and other factors. This is a more realistic description of the underlying

biology. The tumour micro-environment is highly diverse, which assumes all cancer cells share the same state, an unrealistic one. In particular, our model allows different regions of the tumour to suffer varying levels of hypoxia and therefore interact with the environment in various manners. This is a crucial inclusion and necessary for the results to be considered realistic. Allowing the cells to exhibit different probabilities of acquiring the Warburg phenotype, and therefore different glucose uptake rates is also fundamental to the adequate exploration of our research question. We ultimately want to investigate drugs based on glucose-conjugates, so we must have a non-homogeneous model that simulates different cells with different glucose uptake rates. Otherwise, it would be impossible to address our research question as the drug would affect the entire tumour mass equally by default.

Chapter 7

Decision Tree Unravelling

7.1 Introduction

Upon completing our Warburg investigation, we found ourselves having repeated thousands of iterations of our simulation, each with different sets of parameter values for yellow and amber parameters. This resulted from parameter space exploration and attempting to identify those combinations of parameter values that most closely allowed our model's growth curve to align to that derived from literature. Each of these simulations can be represented as a tuple $(\vec{P}, f)$. Where $\vec{P}$ is a vector of parameter values for a given simulation and f is the fitness of such a simulation. Namely, perform or non-perform. $f \in \{Perform, Non-Perform\}$ (See Section 6.3.1 for a detailed explanation of how these classifications are assigned.)

We, therefore, attempted to make use of this generated data to answer the question of which areas in the parameter space region for yellow and amber parameters would lead to simulations that could be labelled as performed. To this end, we decided to train a decision tree classifier over our data to predict, given a hypothetical vector of parameter values $\vec{P_h}$, its corresponding fitness. We decided to train a decision tree rather than another model as this can be easily visualised and interpretable. Compared to other types of models, which act as "black boxes", where it is not possible to understand why or how they reached a given outcome, with decision trees, it is easy to see how different combinations of feature values eventually lead to a leaf node.

Briefly put, a decision tree classifier is a standard and established machine learning model. It belongs to the subset of machine learning techniques known as 'supervised learning', meaning that a set of labelled training data is required to set up the model before it can be used to make predictions. As with other classifiers, once trained, it aims to predict which of a set of classes the sample belongs to given a set of features.

To understand decision trees, we need to introduce the concept of entropy. In essence, entropy refers to the heterogeneity in class labels within a set of samples. For example, a set s_1 with class labels $\{A, A, B, B\}$ would have an entropy of 100%, whereas a set s_2 with class labels $\{A, A, A, A\}$ would have an entropy value of 0%.

Decision trees build a tree data structure where each transition from a node at a higher level (Ie: Closer to the root node) to a node at a lower level (Ie: Closer to leaf nodes) leads to a lower level of entropy. To achieve this, at each non-leaf node, they consider how the data could be split into two groups such that the maximum reduction in entropy could be achieved. For example, if at a node N splitting data on feature M being greater than x has been established to lead to the greatest reduction in entropy, then N would acquire two child nodes N_r and N_l. Data with $M > x$ would be pushed to N_r, whereas data with $M \leq x$ would be pushed to N_l. Such a process is then recursively repeated until one of the following conditions is met:

1. All samples at a given node share the same class;
2. Optionally, more than x% of samples at a given node share the same class;
3. Optionally, there are fewer than n samples at a given node;
4. Optionally, the tree has reached a maximum depth d.

The setting of hyper-parameters x (rule 2), n (rule 3) and d (rule 4), if applicable, need to be established by the user and may lead to significant differences in performance.

After training, a new sample can be classified by starting at the tree's root and traversing it based on the sample's feature values until a leaf node is reached. Then, the sample is assigned the classification most common among the samples at that leaf node.

7.2 Initial investigation on amber Parameters

We initially considered the following amber parameters: Minimum VEGF Concentration, Initial Number of Cancer Cells, Maximum Glucose Uptake Rate and Minimum Glucose required for Warburg Cells.

For our initial investigation, our training set comprised simulations generated via the grid search setup illustrated in Tables 7.1 and 7.2.

Parameter	Values
Min Glucose Warburg	0.5, 1, 1.5, 2,3
Max Glucose Uptake Rate	10,15,20,25,30
Min VEGF Concentration	1,2,5,7,8
Initial Num Cancer Cells	5,10,13,15,20,25

Table 7.1: First grid search ranges for initial investigation on amber parameters.

Parameter	Values
Min Glucose Warburg	0.7,1.2
Max Glucose Uptake Rate	8,12,17
Min VEGF Concentration	7,13,17
Initial Num Cancer Cells	3,5,7

Table 7.2: Second grid search ranges for initial investigation on amber parameters.

Samples were classified as nonperforming, where either at epoch 300 we had fewer than 270,000 cancer cell agents or at epoch 300 we had over 340,000 cancer cell

agents, or the model exceeded 300,000 cancer cell agents before epoch 270. A sample is classified as performed if none of those conditions was met.

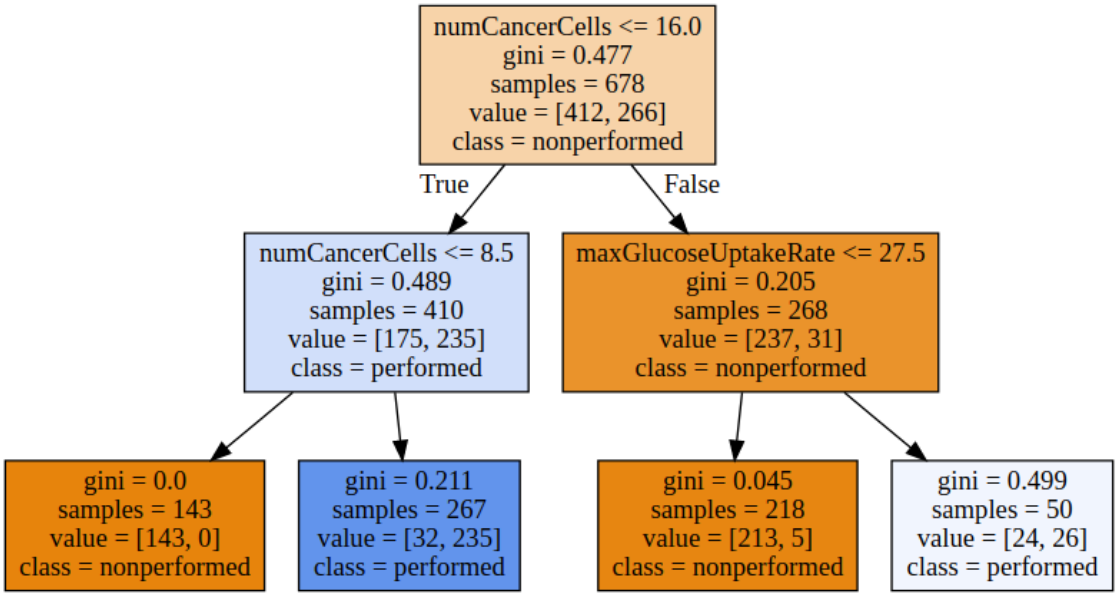

Figure 7.1: Decision tree trained to predict whether a given set of simulation parameter values would have resulted in the model producing a growth curve aligning with what was observed in a laboratory environment (performing) or not (nonperforming). The aim was to describe the parameter space explored by our simulations using grid search and random search to understand which parameter combinations would lead the model to perform or not perform. Transitions leading to the left of a node indicate the split condition is True, whereas those to the right show the split condition is false. A maximum depth of 3 was used as a stop condition. For example, a simulation with fewer than 8.5 (Ie: 9) initial cancer cells agents would have led to the model not performing. In this instance, the model determined that the initial number of cancer cells (numCancerCells) is the most important feature and that too large an amount of these (16) will lead to the model not performing unless the uptake rate of Warburg cells (maxGlucoseUptakeRate) is greater than 27.6.

For example, the tree displayed in Figure 7.1, while purposely limited in depth to 3, already provides several valuable pieces of information. Firstly, the training algorithm established that the initial number of cancer cell agents very closely correlates to the final performance of the model. But also, the maximum glucose uptake rate of Warburg cells is another feature that closely correlates to model performance. We can observe this because among all possible features, in this case, each amber parameter being a candidate feature, the decision tree's training algorithm selected these to split the data on.

Even with a minimal maximum depth (Ie: 3), the algorithm can establish lower and upper boundaries for the parameter numCancerCells and an upper boundary for maxGlucoseUptakeRate.

A few notes about this specific tree: a total of 804 simulations, derived from our earlier study (see Chapter 6), were used as a training set. Of these, 480 had labels non-performed and the remaining performed. To obtain an estimate of the accuracy of our predictor model, we made use of cross-validation (k=10), which reported an

average accuracy of 89.5%.

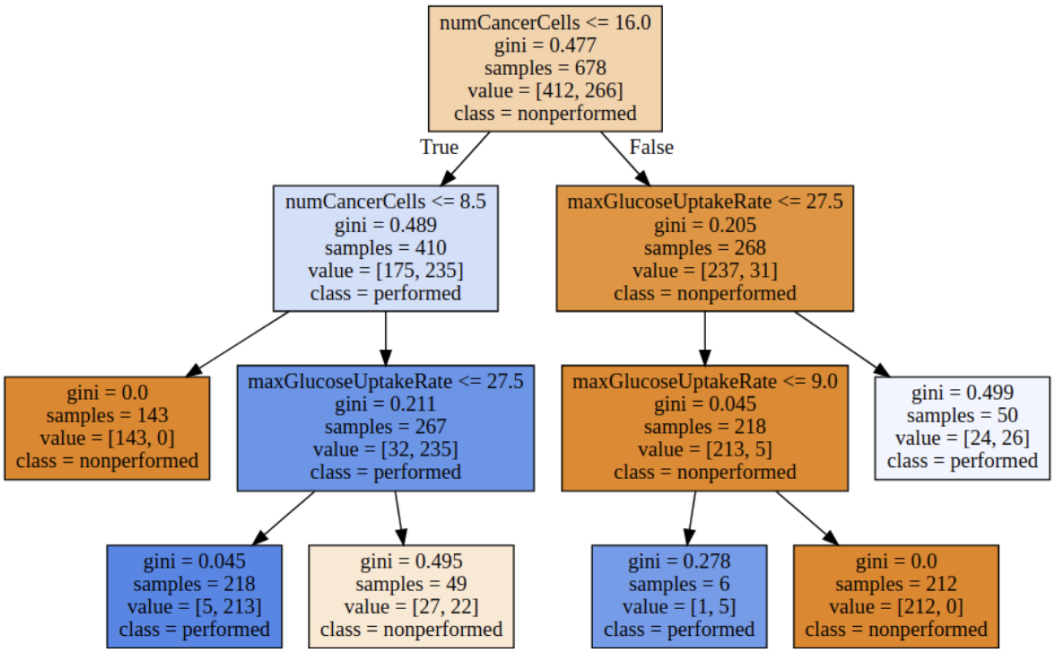

Figure 7.2: Decision tree classifier growth with a max depth of 4. The structure of the first two layers is comparable to our earlier tree grown with a maximum depth of 3, but additional insights are obtained of boundaries for parameter values.

Figure 7.2 shows that allowing the decision tree from Figure 7.1grows one additional layer. Using cross-fold validation with $k = 10$ we obtained an average accuracy of 90%. We can see that while the structure of nodes closer to the root node remains unchanged, additional insights are obtained regarding boundaries for parameter values. Yet, such a depth is still not sufficient for all parameters to be included.

So far, we were able to obtain the following insights on parameter boundaries for amber parameters:

- With an initial number of cancer cell agents values inferior to 9. The model always classifies nonperforming;
- Where the initial number of cancer cells is between 9 and 16, the deciding factor is maximum glucose uptake rate which, if exceeding 27.5, yields the model to be nonperforming;
- Where the initial number of cancer cells exceeds 16, the model will most likely not perform. Edge cases are observed with maximum glucose uptake rate values inferior to 9 where the model performs. Albeit for a sufficiently low number of samples, it can probably be ignored. This can probably be safely ignored for a maximum glucose uptake rate greater than 27.5 as it concerns a small number of samples and has a 50:50 split.

We then repeated a further iteration of grid search on amber parameters, using the range reported in Table 7.3.

Parameter	Values
Min Glucose Warburg	3, 10, 15, 20, 25
Max Glucose Uptake Rate	10, 15, 20, 25, 30
Min VEGF Concentration	1, 3, 5, 7
Initial Num Cancer Cells	1, 3, 7, 10, 13, 17, 20

Table 7.3: Third grid search ranges for initial investigation on amber parameters.

This resulted in 698 simulations (2 did not complete due to instances timing out), of which 480 were classified as non-perform and the remaining as perform. The resulting trees with max-depth three and max-depth five are shown in Figures 7.3 and 7.4 respectively. This reported an accuracy of 88% using cross-validation with $k = 5$.

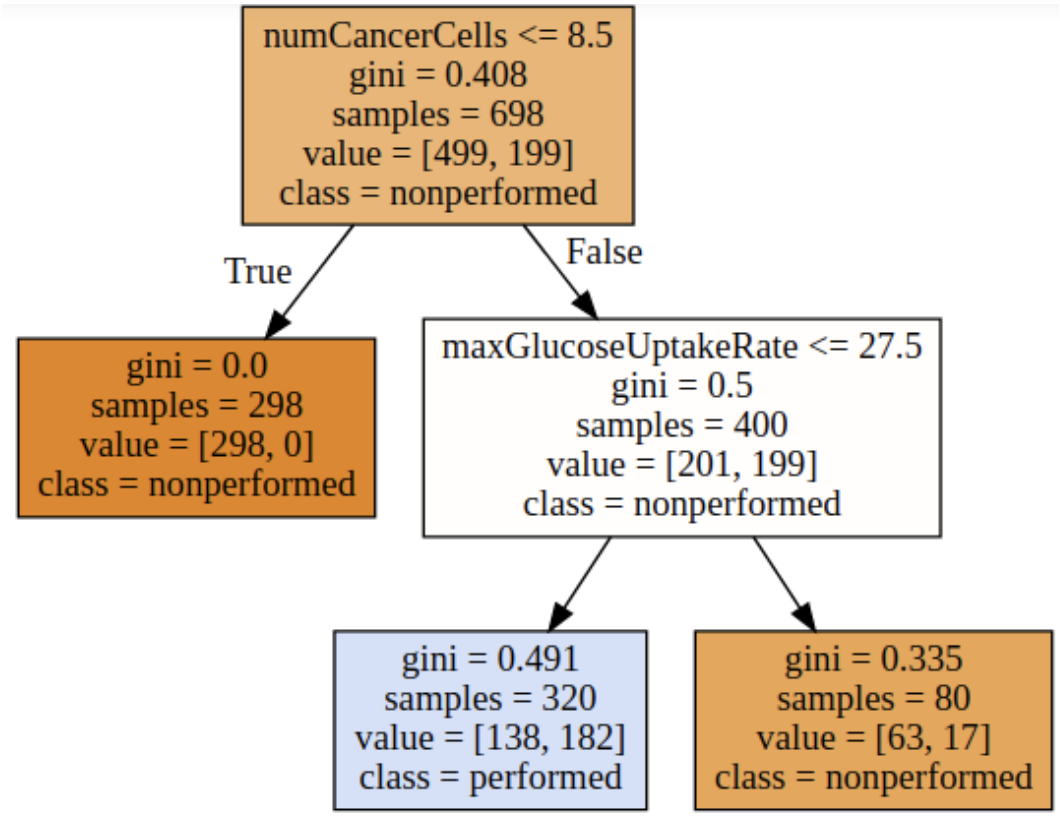

Figure 7.3: Decision tree with max-depth of 3 developed from the second set of grid search iterations. See Table 7.3. This reported an accuracy of 77% using cross-validation with $k = 5$.

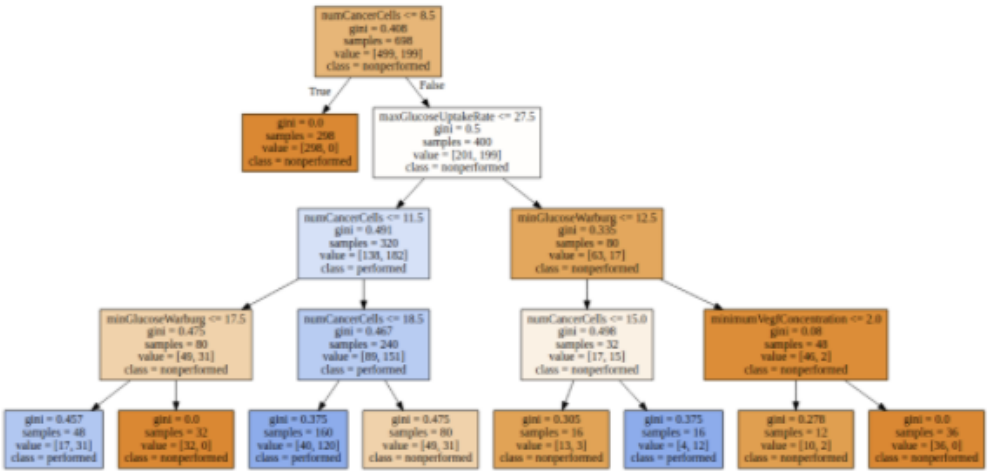

Figure 7.4: Decision tree with max-depth of 5 developed from the second set of grid search iterations. See Table 7.3. This reported an accuracy of 88% using cross-validation with $k = 5$.

The tree with a max depth of 5 shown in Figure 7.4 is of particular interest. For the first time, it also includes nodes that split data on additional parameters, including minimum VEGF concentration to stimulate angiogenesis and minimum glucose concentration required for the survival of Warburg cells. Yet, the depth of such a decision tree yields poorly to manual analysis.

7.3 Random Search on Amber Parameters

We then explored the search space of amber parameters through random search. Finally, we ran 100 simulations with values drawn from normal distributions parametrised illustrated in Table 7.4.

Parameter	Mean	Standard Deviation
minGlucoseWarburg	15	3
maxGlucoseUptakeRate	20	3
numCancerCells	14	2
minVegfConcentration	5	1

Table 7.4: Setup of distributions used for random search against amber parameters. A total of 100 simulations were generated. If, in the process, a simulation with any parameter having a value less than 0 were generated, such a sample was discarded and re-drawn.

The first observation we can make in this context is that, as is often reported in the literature, random search was more effective than grid search at finding suitable value combinations in the parameter space search. Indeed, out of 100 simulations, 82% were classified as performing. This is a significant improvement on grid search, where most simulations were classified as non-perform.

We trained decision trees to maximum depths of 3 and 5, with results shown in Figure 7.5 and 7.6 respectively.

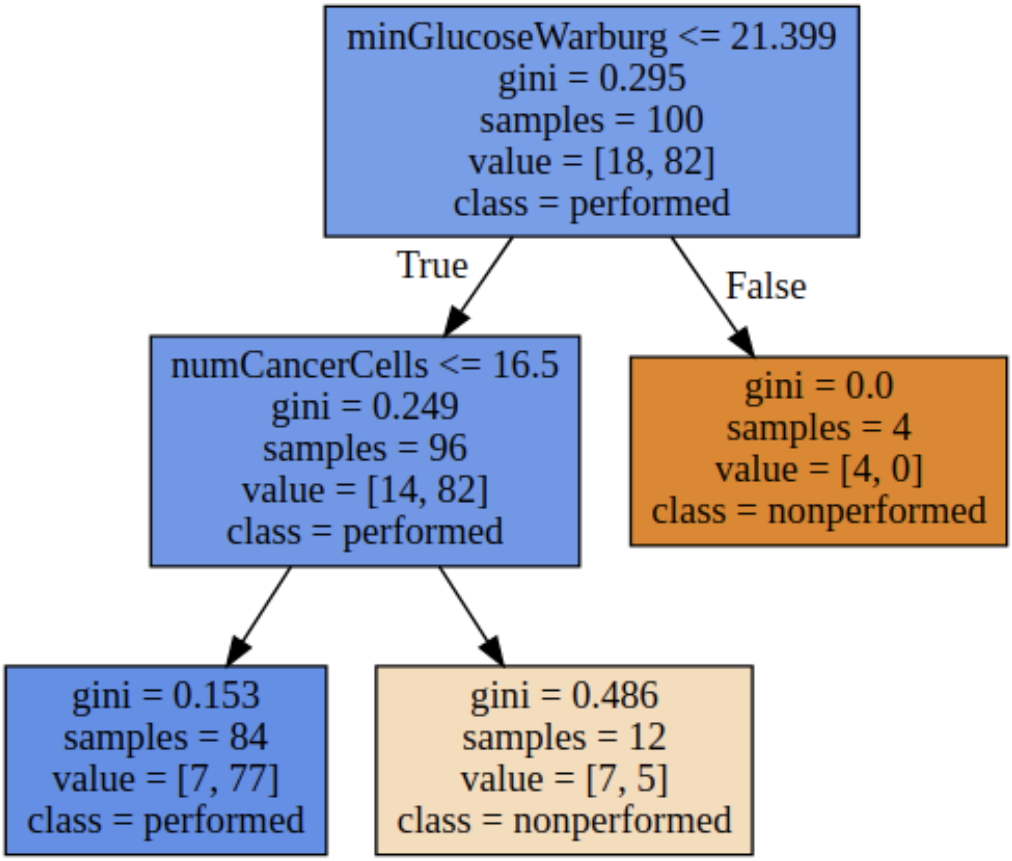

Figure 7.5: Decision tree trained on simulations generated via random search on amber parameters. The tree was grown with a maximum depth of 3, and cross-validation with k=5 reported an average accuracy of 84.97%.

From the tree grown to a maximum depth of 3 (Figure 7.5), we can note that:

The model does not perform when:

- minGlucoseWarburg is less than 21.30 (although only 4% of the samples, so potentially could be disregarded);
- more importantly, in cases where the number of cancer cell agents is greater than 16.

And that the model does perform when:

- numCancerCells is less than 16.

These are similar parameter ranges to those identified with grid search, but significantly fewer samples were needed to obtain them with random search.

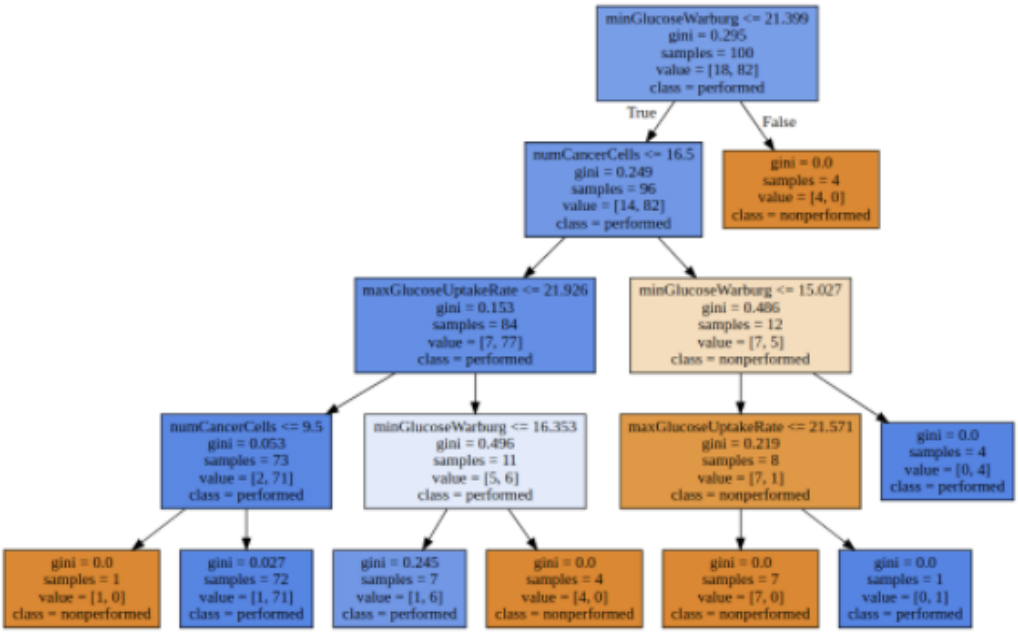

Figure 7.6: Decision tree trained on simulations generated via random search on amber parameters. The tree was grown with a maximum depth of 5, and cross-validation with k=5 reported an average accuracy of 90%.

The decision tree trained with a max depth of 5 (Figure 7.6) provides further insight.

In particular, branching where numCancerCells is less than 16, the model will perform where maxGlucoseUptakeRate is less than or equal to 21. However, if it is greater, the model can still perform where minGlucoseWarburg is below the threshold of 16.353. We also observe the minGlucoseWarburg has to be below a certain threshold (21.399). Otherwise, the model does not perform.

Where numCancerCells is greater than 16, the model will still perform where minGlucose Warburg is greater than 16 OR maxGlucoseUptakeRate greater than 21.5.

Indeed, results suggest that where numCancerCells is greater than 16, the model only performs where particularly high glucose requirements hinder tumour growth.

Model	Grid Search Data	Random Search Data
Max-Depth 3	77.79%	84.97%
Max-Depth 5	84.38%	90.01%

Table 7.5: Summary of performance (cross-validation) on decision trees grown to max-depths of 3 and 5 on data obtained from grid search and random search on amber parameters. Random search consistently outperforms grid search with considerably fewer samples.

Table 7.5 summarises the performance of the various decision tree classifiers. In principle, we can tell both grid search, and random search led to classifiers being able to correctly learn the data when trained with a maximum tree depth of 3 and

5. As expected, more tree layers led to better performance. Interestingly, classifiers trained on data generated via random search consistently outperformed their grid search counterparts and did so with considerably fewer samples to train on.

Parameter	Grid Search	Random Search
minGlucoseWarburg	17.5 (maxGlucose-UptakeRate $\leq$ 27.5), 12.5 other-wise	21.39 at root
maxGlucoseUptakeRate	27.5	21.5 - 21.9 (At different nodes but same level)
minVegfConcentration	2	Not observed
numCancerCells	9-18	9-17

Table 7.6: Summary of value boundaries obtained for amber parameters by decision tree classifiers trained using grid search and random search data.

Table 7.6 reports value boundaries identified for amber parameters by decision tree classifiers trained on data generated by grid search and random search.

Values are comparable for the initial number of cancer cell agents but differ over the remaining parameters. Further, the parameter describing minimum VEGF concentration for angiogenesis to initiate was not observed in trees trained over random search data. The latter is possibly explainable by such parameters being poorly correlated to model performance. This theory is supported because the value occurs farther away from the root node, even in trees trained on grid search.

7.4 Random Search on Amber and Yellow Parameters

Contrarily to grid search, the computational complexity of random search does not increase as the number of parameters increases. That is to say, while the computation cost of grid search grows exponentially as new features are added, that of random search increases linearly. This is known as grid search's "curse of dimensionality". After the successful nature of our random search investigation on amber parameters (Section 7.3), we decided to extend this to the complete set of amber and yellow parameters. Once again, values for these were drawn from normal distributions whose properties are reported in Table 7.7.

Parameter	Mean	Standard Deviation
minGlucoseWarburg	15	3
maxGlucoseUptakeRate	20	3
numCancerCells	14	2
minVegfConcentration	5	1
minGlucoseNonWarburg	25	2
oxygenDiffusivity	0.3	0.1
minPSynthesis	0.2	0.1
glucoseDiffusivity	0.001	0.0005
pWarburgSwitch	0.5	0.2

Table 7.7: Summary of values used to generate normal distributions from which values for amber and yellow parameters were drawn.

Four hundred simulations were generated, where 190 were classified as non-perform and the remaining as performing. A decision tree classifier with a maximum depth of 3 and another with a maximum depth of 5 was trained and shown in Figures 7.7 and 7.8 respectively.

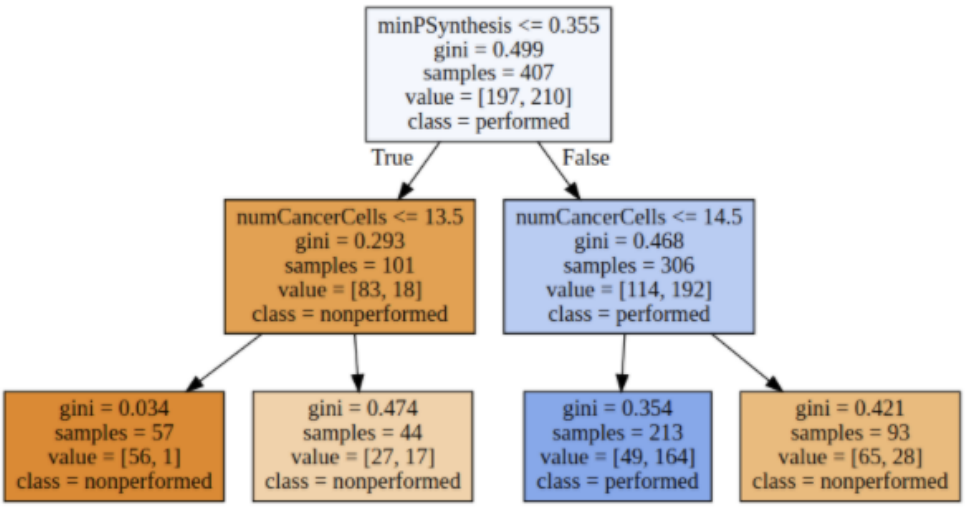

Figure 7.7: Decision tree with a maximum depth of 3 trained on simulations using random search to generate yellow and orange parameters values. With cross-validation (k=5), the tree reported an average accuracy of 71.98%.

Figure 7.8: Decision tree with a maximum depth of 5 trained on simulations using random search to generate values for yellow and orange parameters. With cross-validation (k=5), the tree reported an average accuracy of 84.03%.

While grown to a limited depth only, these trees provide insights into yellow parameters. These were not included in our previous exploratory rounds. First, this exploratory round revealed the minimum probability of a cell progressing into synthesis as one of the most important parameters. This is evidenced by both trees selecting this as a feature for the root node. It further provides a threshold value, roughly indicating that where this is not at least 35% the simulation will most likely be categorised as nonperforming. Secondly, among the amber and yellow parameters, it identifies the initial number of cancer cell agents as the second most important feature providing an upper boundary for this value.

7.5 Full Investigation

Finally, we allowed a decision tree classifier to grow to unbounded depth over the set of simulations generated from random search over yellow and amber parameters. We identified the node in the decision tree which led to classification performing and was the most descriptive. By descriptive, we mean that overall it provided the most information about the most features. A summary of boundaries obtained is provided in Table 7.8.

Parameter	Original Classification	Obtained Boundaries
Initial Number of Cancer Cells	Amber	≤ 13
Glucose Diffusivity	Yellow	≥ 0.000278
Minimum Glucose Non-Warburg	Yellow	≤ 18.4
Maximum Glucose Uptake Rate	Amber	≤ 23.2
Minimum Probability of Synthesis	Yellow	$0.247 \leq MPS \leq 0.359$

Table 7.8: Summary of parameter boundaries obtained by examining the nodes of a decision tree classifier trained over the results of simulations generated by random search. We provided at least one boundary (upper or lower) for two amber and three yellow parameters. In cases of parameters whose value is a real number, results are reported to three significant figures.

Thanks to this process, we recategorised parameters, of which two were originally amber and three yellow, as green. As a result, we were able to identify previously unknown value thresholds that, when assigned to such parameters, would lead simulations to be categorised performed.

In this final part of our investigation, we have summarised the technique to address uncertainty in parameter values. We have demonstrated that traditional decision tree classifiers can learn which parameter combinations allow our model to perform as expected. By examining the resulting tree structure, we may infer boundaries for parameter values. In further experiments, we can use this information to reduce the subset of the parameters classified as amber or yellow. If we repeated the investigation over a broader set of parameters, we believe we would eventually classify the vast majority, if not all, of the parameters as green.

Chapter 8

Epidemiological Modelling

8.1 Introduction

The COVID 19 pandemic disrupted our regular research work and provided us with an opportunity to devolve our knowledge and skills in agent-based modelling to epidemiological studies. In addition, the rapid advent of the pandemic introduced drastic confinement strategies, including 'stay at home' orders and the closing of social and entertainment venues. From the scientific community members to policymakers to the general public, everybody shared several questions: When would the pandemic be over or under control, and to what extent were lockdown measures useful and necessary?

In a way, the situation was unprecedented. To find the previous case of a widespread epidemic in Europe, we have to look back to the 1918 influenza pandemic known as the Spanish flu. Nevertheless, we cannot use that as a measure of comparison. A century ago, healthcare and medical facilities could not be compared to those of the present day, and research capabilities were nowhere near as developed as today. As an example, DNA sequencing techniques were not introduced until the 1970s. Further, the geopolitical scenario which influenced public policy could not have been more different. Whereas today, at least in Western Europe, we are privileged by the absence of local active military conflicts, World War One has actively fought a century ago.

We investigated modelling as a tool to predict the evolution of the pandemic and assess which mitigation strategies would be most effective. Broadly, we were interested in two main research questions. Firstly, the lightest lockdown we could introduce prevented widespread COVID outbreaks among the population and avoided the UK NHS being overwhelmed. Secondly, once the pandemic had been successfully contained, how should deconfinement and the relaxing of lockdown policies be carried out?

The work was initially inspired by the technical reports of Ferguson et al. [92] which advised on UK government policy. However, we differed in our implementation in that they made use of a network model whereas we made use of an agent-based model. While the latter turned out to be more computationally expensive, it allowed us greater flexibility in devising agents' behaviour, for example, by allowing an agent's network of interactions to change from day to day instead of being static.

The challenge of rapid context switching from the modelling of biological tissues to human behaviour and epidemiology is significant. Nevertheless, we successfully implemented a model to simulate COVID spread in Manchester (UK). Our results suggest that in the absence of vaccines and immune people, as was the case in the Spring of 2020, anything other than a heavy lockdown is insufficient to contain the spread of the disease. Policies should include the closing of entertainment and social venues. Most businesses should remain closed (or otherwise have staff work from

home). Households with members testing positive should self-isolate. Crucially, our model also suggested that schools may probably remain open without this significantly affecting the course of the pandemic.

Regarding post-lockdown deconfinement, our results suggested that unregulated rapid reopenings will make a second wave worse than the first. These forecasts were obtained in the early Summer of 2020 as many European countries were preparing to reopen. Sadly, but perhaps understandably, after several months of harsh lockdowns, the situation quickly returned to something very closely resembling normality. This became apparent in early Autumn of 2020 and throughout the following winter. Many European countries, including Italy, France and the United Kingdom, experience second waves with daily cases and death rates superior to the original epidemic of Spring 2020.

The gradual reopening, which we had auspicated in our report, did not occur. This would have allowed progressive development of immunity without overwhelming the healthcare system and might have prevented further sudden lockdowns as enacted in the United Kingdom around Christmas 2020.

The prospect at the time of writing, in the Spring of 2021, is a somewhat more serene one. The vaccination campaigns have picked up the pace, and many European countries are heading towards the 60% of population vaccinated threshold, which should allow for herd immunity. While the situation in other countries is still concerning, the general hope is that more and more people will be immunised as vaccine supply chains strengthen.

From a policy-making perspective, the hope is that the COVID19 experience has taught us how to promptly and effectively respond should there be a further epidemic. In addition, and within our remit, we have learned several important lessons in epidemiological modelling and model scalability, which we will share throughout the rest of this chapter.

8.2 Previous Work

We start with previous work in epidemiological modelling (Section 8.2.1), followed by related work explicitly dealing with COVID. See Section 8.2.2. Our review will be concerned with discrete and agent-based models only. Primarily, this is because the vast majority of literature in the field uses these approaches. And also, because that is the approach we decided to implement for our investigation and so that more relevant.

8.2.1 Epidemiological Modelling

Cliff et al. [93] presented a model to simulate the spread of influenza epidemics in Australia. The size of their model was considerably larger than other studies we reviewed, comprising 19.8 million agents and one of the few cases where the number of agents matched the actual population size on a 1:1 ratio. They used an in-house modelling framework (AceMod) and simulated five scenarios averaging results across ten trials or a total of 50 simulations. The environment is discretised by breaking down Australia into local areas, which are further subdivided into geographical locations, and the model is parameterised using data from the 2006 Australia Census.

Another model of influenza is proposed by Ong et al. [94], who use a proprietary framework to investigate the spread of the disease within a ward where the environment is discretised as a Cartesian grid. The model was parameterised by conducting a field survey to obtain the activity patterns of individuals in a ward of Tan Tock Seng Hospital in Singapore. This model is particularly interesting as, rather than focusing on larger areas such as cities or countries, it aims at replicating the heterogeneity of agent behaviours in small spaces.

A model to optimise the age-specific vaccine distribution is presented by Özaltın et al. [95]. They considered ten configurations with three different vaccine coverage levels, three different response delay times and six different basic reproductive ratios, including approximately 560,000 individuals. They used a probabilistic method to establish whether an agent would become infected based on factors such as age, whether they were vaccinated, etc. and parametrised their model from available census data.

Parker et al. [96] presented an agent-based model to investigate the disease dynamics of T. b. rhodesiense human African trypanosomiasis (rHAT). This helped understand the mechanics behind the transmission of the disease and was essential to develop appropriate mitigation strategies. The model comprised 16,024 human agents and 5,250 tsetse agents. Cattle agents were also included, although details about these were not presented. The simulation was repeated 100 times with results averaged. The model was developed using a custom in-house framework written in Python 2.7. The authors used standard deviation to measure the stability of simulation results. The environment was discretised in regions of $1km^2$, census data was used to establish population parameters, and agent dynamics including mating and feeding were included.

Several pieces of work investigate H1N1 influenza. Wang et al. [97] developed a Repast Symphony model that used extended SEIR techniques to model the spread of flu in an environment comprising 1,000 agents across five scenarios, each simulated five times with results averaged. The model included parameters such

as the duration of the latent phase of the disease, fatality rate, etc. and reported growth curves were validated against those reported by the Yunnan Center for Disease Prevention and Control (Yunnan CDC). Ponnambalam et al [98] developed a model based on US census data from 2010 which made use of a network model to describe the different states an agent could be in and how they could transition between them. Their simulation is replicated 100 times. Yet another model of H1N1 is proposed by Peng et al. [99], to simulate how many people will get infected under various strategies, help decision-makers optimise systems to prevent epidemics. They simulate four containment strategies and include up to 1,000 agents. The model is implemented using the SEDIR approach and discretises the environment with grid-like coordinates. The model is developed using Agent Analyst, plus some Repast integration.

A particularly comprehensive model of H1N1 spread is presented by Khalil et al. [100]. They use existing census data to model the spread of the disease in Egypt. They model five different scenarios and use a custom extension of the SIR model, which they validate against other similar SIR models. Finally, Mei et al. [101] modelled the spread of H1N1 in urban areas. The model included 6,000,000 agents. It simulated three different scenarios, each repeated 30 times and averaged results. The model was parameterised using census data from the Changsha region in China. Interestingly, in contrast with other work we reviewed, the authors also discuss the infection process. They report the probability of an agent becoming infected depends on multiple factors, including how many infectious persons co-occupy the room, how long each contact lasts, the type of activity, and the infectiousness category of the contagious person.

A publication focused on the behaviour of individuals during a pandemic was presented by Aleman et al. [102]. They modelled the diffusion of SARS in a community of 5 million individuals parametrised with information from the Greater Toronto Area Census. In terms of the modelling of pandemic influenza, we would like to mention the work of Arsalan et al. [103], who investigated different measures to manage pandemic influenza to decide on the best intervention strategies. Their model was developed using proprietary Java code and was part of a minority that performed sensitivity testing on their parameters. The model was parameterised with data from the US census. And finally, Imoto et al. [104] explored the effect of immunity through vaccination against an emerging influenza pandemic in the Tokyo metropolitan area. The model featured over 1 million agents and explored multiple scenarios of progressively larger subsections of the population being vaccinated. They used a SEIR methodology and were also part of the minority, which used parameter sensitivity testing.

Perhaps of particular interest to our investigation is the work of Shojaati et al. [105], who investigated how ‘good practices’ (Hygiene, washing hands, etc.) affect the spread of MRSA in care homes. This is especially interesting when thinking back to

the high rates of mortality observed in care homes in the Spring of 2020. The model explored 5 'what-if' scenarios, each examining more rigorous compliance with the best hygiene and social distancing practices. The model used a SEIR approach and was developed using AnyLogic.

Several pieces of work reviewed the spread of HIV/AIDS. This is particularly interesting due to considerably different transmission mechanisms not based on the aerial transmission of viral particles. Paratus et al. [106] modelled the spread of HIV in the city of Maryland. This was achieved by implementing a network of sexual encounters, through which agents positive to the disease could infect other agents. Escudero et al. [107] presented a model to simulate the diffusion of HIV among people who inject drugs. Their model featured a population with 100,000 agents, and they simulated model scenarios, such as those involving therapy and assistance, vs those which didn't. The model was parameterised with data obtained by observational studies in New York City. Graw et al. [108] presented another study concerning HIV diffusion among people who inject drugs parametrised with data published by the infectology centre of Latvia collected between 1987 and 2010. They use a network of relations between different agents to model the spread of the disease, including approximately 3,000 agents simulating various scenarios with increasing likelihood of transmission. Finally, Wang et al. [109] modelled the spread of HIV using data from the Yunnan province of China. Their simulation featured a more advanced discretisation of the environment than other HIV models reviewed, including the use of GIS to establish heterogeneity in population density and establish more densely populated and rural areas.

8.2.2 Models of COVID

Kerr et al. [110] developed a pure-Python model of COVID spread. In their simulation, agents initially represented individuals, with this then being 'scaled' such that an agent described multiple people as the size of the simulation increased. Simulation parameters were assigned values derived from empirical data in census databases. Each agent could be healthy or infected with various degrees of severity. While infected, an agent became a possible vector of infection. The infection period could end with the agent recovering or dying. Agents were assigned a social network derived from randomly selected subsets of agents who frequented their same locations (Homes, schools/workplaces, or social and community settings.) The infection mechanism was implemented such that each encounter between an infected and non-infected agent resulted in the former being infected by the latter with a given probability derived from empirical data. Social distancing and lockdown measures were implemented using 'network clipping', reducing the number of interactions between agents. The authors' limitation lies in the degree of realism captured. To optimise for performance and allow the model to scale, CovaSim cannot include heterogeneity in the quality interactions between agents and simplifies these

phenomena greatly. Further, another limitation shared by many pieces of work (including others) lies in the difficulty of obtaining accurate data about the spread and infection rates of the epidemic.

Silva et al. [111] proposed a Python-based implementation of a SEIR model of COVID spread. Their model is focused on exploring the impact and efficacy of various forms of lockdown and containment, ultimately arguing that weaker lockdown conditions are insufficient to stop the spread of the disease. Agents represented individuals who, similarly to other models, interact with other agents in their homes and workplaces. The probability of infection resulting from the encounter between an infected and a healthy agent was empirically estimated. Parameter values were drawn from literature, although we would like to point out that several key ones (Eg: Probability of contagion) were sourced from technical reports which were not themselves peer-reviewed or independently verified. Eg: The reports by Ferguson et al. [112]

Gomez et al. [113] propose a similar model concerning Bogotá in Colombia. In their SEIR model, agents interact with other agents in their network, with infection occurring with a given probability upon their encounter between an infectious and healthy agent. This study acts more as a proof of concept framework design. The authors remark that parameter values were not set to real ones, and no reference is provided for these: " INFEKTA may be able to provide more accurate results if its parameters are set to real ones". The simulation explores different types of lockdown strategies and the impact on disease spread results, being consistent with general expectations that harsher lockdowns are more effective at containing the spread of COVID.

Kano et al. [114] propose a Python-based cellular automaton implementation of a SEIR model of COVID spread, focused on determining the impact of lockdowns on various forms of business activities. The study highlights the development of an economic gap between owners of essential and on-essential businesses, with the former becoming more prosperous due to stable revenue and fewer expenses. In contrast, the former will most likely struggle financially. Further, the study suggests that as people run out of money, they are less likely to comply with confinement measures, leading to higher infection rates. This study is interesting as it considers the impacts of lockdown beyond those most commonly considered, including infection or death rates and hospitalisation. However, parametrisation could be improved (All parameter values are estimated).

An interesting study is proposed by Almagor et al. [115]. This explores the effectiveness of tracing apps at reducing infection rates and considers possible downsides such as significant increases in demands for testing. The authors' results suggest that increased testing results in lowered infection rates. The simulation is implemented as a SIR model, with agents representing individuals and contacting

other agents in their households, family, workplaces, schools, friendships, and social circles. Social networks are dynamic, for example, with the likelihood of an encounter with a family member being higher than that of an encounter with a classmate, itself being more elevated than an encounter with one of the neighbours. This improves on other discussed examples, where social networks were fixed and agents always interacted with the same other agents. The model also considers heterogeneity in agent compliance with quarantine and social distancing, again improving on other examples where the perhaps unrealistic expectation that there would always be full compliance was assumed. Among the limitations of their implementation, the authors include their assumption that tests are 100% accurate and that notifications to users always work flawlessly. Further, adoption rates and compliance with guidelines are consistent across all demographic groups.

8.3 Methods

8.3.1 Overview

We simulated the spread of COVID in a city the size of Manchester. A synthetic population was generated, where each agent corresponded to a person. Time was advanced in discrete time steps (or epochs), corresponding to one hour. Each agent had a schedule that governed where they would be at each epoch. At home, at work...) Each agent's schedule differed to allow for heterogeneity in the population. See Section 8.3.4. Agents could become infected by exposure to another infected person. Once infected, they would go through an incubation period after which they either became asymptomatic or developed mild or acute symptoms. Agents could either die due to contracting the disease or become immune. We assumed that only symptomatic agents could die due to contracting the disease.

We ran the model under different hypotheses of containment strategies. For example, we were enacting progressively stronger lockdowns and observing the result this had on the spread of the epidemic. Milder lockdowns simply required households with infected members to self-isolate. Progressively, harsher lockdowns included closing social venues more businesses until eventually, a near-complete lockdown saw social venues, schools, and most workplaces closed.

8.3.2 Model Parameters

Parameter	Value
Probability of Becoming Asymptomatic Following Incubation	30%
Probability of Developing Mild Symptoms Following Incubation	50%
Probability of Developing Acute Symptoms Following Incubation	20%
Incubation Time (Mean)	120 Hours
Incubation Time (Standard Deviation)	72 Hours
Infection Length (Mean)	168 Hours
Infection Length (Standard Deviation)	48 Hours
Probability of Becoming Infected Upon Exposure (Mean)	80%
Probability of Becoming Infected Upon Exposure (Standard Deviation)	1%
Number of Households	226,640
People per Household (Mean)	2
People per Household (Standard Deviation)	1
Number of Schools	166
Number of Social Centers	3,000
Number of Workplaces	2,000
Age (Mean)	33
Age (Standard Deviation)	10
Age Above Which Somebody is Considered an Adult	21
Number of People Initially Infected	100
Maximum Number of Interactions per Hour	20
Hospital Size (Mean)	750 Agents
Hospital Size (Standard Deviation)	100 Agents
Reduction in Probability of Death due to Hospital Admission (Mean)	0.5%
Reduction in Probability of Death due to Hospital Admission (Standard Deviation)	0.01%
Reduction in Infection Length due to Hospital Admission (Mean)	48 Hours
Reduction in Infection Length due to Hospital Admission (Standard Deviation)	2 Hours
Number of Hospitals	8
Number of People Tested Daily	1,000

Table 8.1: Summary of model parameters for our COVID19 simulation. Where mean and standard deviation values are provided, values were drawn from a normal distribution parametrised as such. Parameters were extracted from the works of Ferguson et al. [116] [92]

Probability of death due to COVID is dependent on an agent's age. Age-dependent probabilities are summarised in Table 8.2.

Age (Years)	Probability of Death
<9	0.00%
10-39	0.40%
40-49	1.30%
50-59	3.60%
60-69	8.00%
$\geq$ 70	14.80%

Table 8.2: Probability of death due to COVID19 infection for different age groups, given that the infected agent was symptomatic. Elderly groups are significantly more likely to die than younger people. Whether an agent dies is established at the end of their infection period, with death occurring with the probability assigned to their age group.

8.3.3 High-Level Architecture

As illustrated earlier, the model compromised a synthetic population of agents. Each agent represented one person. The starting point to generate the population was the approximate number and size of households in Manchester. This was obtained from government census data. Next, each house was assigned a number of agents consistent with average household size, ensuring that there would be at least one adult. After this initial population seeding, each agent would be randomly assigned a social centre and a workplace or school based on their age.

Time was discretised in epochs, corresponding to one hour of real-world time. During an epoch, an agent would be at one of the following locations: Their home, school, or workplace, depending on their age or social centre. There, they would meet up to a fixed number of people who were also at that location and, with a given probability, infect them or be infected by them. Each agent was assigned a schedule to determine which locations they would be at each hour of the day. Public transport and commuting between locations were not implemented in this first implementation.

Upon instantiation, the model would continue until the entire population was dead, immune, or a maximum number of epochs had elapsed. At the start of an epoch, each agent would be transferred to their designated location. At each location, agents would be split into 'interaction bubbles'. This is based on the idea that each person only interacts with up to a certain amount of people per hour. Then, new infections per interaction bubble would be calculated based on the number of people who already had contracted COVID in that group.

8.3.4 Agents

Each agent represents a person. Each agent instance stores its demographic properties, medical properties and timetable. Whereas some parameter values are shared across all agents, others are personalised per agent. Namely, those which in Table 8.1 are assigned fixed values are shared. Whereas those for which distribution mean and standard deviation are provided, differ between agents. For example, in the case of the parameter "Probability of Becoming Infected Upon Exposure" for each agent, the value is drawn from a normal distribution with a mean of 80% and a standard deviation of 1%. Note, that is ceiled at 100% and floored at 0%.

Demographic Properties

An agent's demographic properties report those they report at an official census. These are their age, household, school or workplace and assigned social centre.

Medical Properties

These include properties that are of particular interest with regard to our study. These are the probability of becoming asymptomatic or developing mild or acute symptoms following covid incubation, the incubation time, the infection length and the probability of becoming infected upon exposure.

The Timetable

The Master Timetable

Each agent is assigned a timetable that determines where they will be at each hour of the day. Agents may be at one of the following:

- At home;
- At work (adults) or school (kids);
- At their designated social centre;
- At a hospital.

For each agent, timetables are derived from a 'master timetable' which dictates the following principles:

- Agents are at home during the first part of the day, midnight to 8 am included;
- Agents are at school/work between 9 am and 6 pm included;

- Agents are at home during the rest of the day.

Custom Agent Timetables

Hour	Probability	Hour	Probability
00:00	0%	12:00	10%
01:00	0%	13:00	10%
02:00	0%	14:00	10%
03:00	0%	15:00	10%
04:00	0%	16:00	20%
05:00	33%	17:00	66%
06:00	66%	18:00	66%
07:00	66%	19:00	33%
08:00	33%	20:00	33%
09:00	10%	21:00	33%
10:00	10%	22:00	10%
11:00	10%	23:00	0%

Table 8.3: Probability of an agent's timetable deviating from the master timetable at each hour. For example, the master timetable would have an agent at home at 5 am. However, upon instantiating an agent, they have a probability of 33% to deviate from that and either be at work or school (depending on their age) or at a social centre. Where an agent's timetable deviates, the location that replaces the one specified in the master timetable is randomly picked with equal probability among those candidates. The probabilities reflect our estimated likelihood of somebody deviating from the expected location at a given time. For example, it is more likely for there to be heterogeneity in people's locations at 6 am than at 5 am.

An agent's timetable may differ from the master timetable at each hour with probabilities reported in Table 8.3. For example, the master timetable says an agent would be home at 5 am. But, with a probability of 33%, a specific agent's timetable may have them instead of being somewhere else. For example, they might go to the gym early in the morning.

An agent's custom timetable is generated upon instantiation and remains constant throughout the simulation's execution. In this process, the location indicated by the master timetable is replaced. The new location with which it will be replaced is drawn randomly and equally from all those available. For example, an adult whose 5 am location is to change will be assigned to their social centre or workplace with equal probability.

Timetable Alterations due to COVID Symptoms

Agents may change their timetable when developing mild or acute COVID symptoms. These changes are reversible and only persist throughout the symptoms. Agents revert to their original schedule once they recover. Upon contracting symptoms, agents alter their timetable as follows:

- Children with mild or acute symptoms remain at home throughout the entire day;
- Adults with mild symptoms remain at home when according to their timetable they would otherwise be at a social centre;
- Adults with acute symptoms stay at home throughout the entire day.
- Any agent who is admitted to a hospital will remain there throughout the infection period;

Timetable Alterations due to Containment Strategies

Containment strategies (see Section 8.3.6) may alter an agent's schedule. For example, a containment strategy may prescribe venues of a certain type (such as schools) to be closed. As a result, when agents would generally be at that location, they will instead be at home.

Agent Updating during Epochs

An agent's behaviour is affected by whether they contracted the disease and any containment policy that may currently be in place.

As an agent is stepped during an epoch, they will come into contact with a set of other agents who share their household, workplace and school/workplace. Anytime a non-infectious agent meets an infectious one, the latter will infect the former with a given probability. Unless the former already recovered, in which case they are immune.

If an agent becomes infected, the incubation period will begin as of the following epoch. Similarly, containment policies or other effects derived from contracting the disease will also start in the next epoch.

We made this decision considering two factors: Firstly, to remove any bias in the system because some agents become infectious before others simply because of their position in the schedule. And secondly, to account for the physiological reality that viral particles would need some time to 'settle' and infect host cells before the person can be considered infectious.

8.3.5 The Infection Process

Perhaps unsurprisingly, our model assumes that an agent becomes infected through exposure to an infectious person. A non-infected person becomes infected with a given probability upon such an event. See Table 8.1.

At the start of each epoch, agents at each location are split into 'interaction bubbles'. That is, groups of up to n agents where n is the maximum number of hourly

interactions an agent can have. See Table 8.1. For example, if a given workplace has 72 people in it and n=20, this will result in 3 interaction bubbles of size 20 plus a further interaction bubble of size 12. On the other hand, a household with four people would produce a single interaction bubble of size 4. Interaction bubbles are randomly re-generated for each location at the start of each epoch to account for the possibility of agents entering/leaving a particular place according to their schedule.

After interaction bubbles are generated, each agent in a bubble is evaluated against every other person in the bubble. If a non-infected person meets a contagious one, they become infected with a given probability. Conversely, if a non-infected person doesn't become infected due to meeting another infectious agent, then meeting another contagious agent in the same epoch may have them acquire the disease due to this second encounter.

Upon being infected, an agent starts incubating the disease. During this time, they are asymptomatic and not contagious. After that, they develop the disease and are therefore infectious. An agent who developed the disease may be asymptomatic or display mild or acute symptoms.

Finally, at the end of the infection, an agent may recover and become immune (They may not be infected again) or die and therefore be removed from the simulation.

Any agent who develops acute symptoms is eligible for admission to a hospital. An agent will be admitted to any hospital with free space and remain there until they recover or die. Suppose an agent with acute symptoms cannot be admitted to a hospital because of lack of space (No hospital has spare capacity). In that case, they will keep applying until they are admitted as space frees up, recover or die.

Admission to a hospital involves a marginal decrease in infection length and probability of dying due to COVID. The actual amounts by which these are decreased are specific to each agent and drawn from a normal distribution parametrised with mean and standard deviation reported in Table 8.1.

8.3.6 Containment Strategies

Containment strategies represent policies implemented to limit the spread of COVID. They simulate real-world interventions such as closing specific venues, enforcement of social distancing or full lockdowns.

Our model allows us to specify different categories of containment strategies that will be enacted when the number of people infected with COVID exceeds a given percentage. For example, to close all schools when the rate of people having COVID exceeds 5%.

Any combination of policies may be specified for a simulation. A summary of supported policies is provided in Table 8.4

Policy	Effect	Parameter
Social Distancing	Reduces the maximum number of contacts a person may have during an epoch to n.	n, the new maximum number of contacts
Close Venue Type	Closes all venues of type t	t, the type of venue. One of: school, social center, workplace
Household Self Isolates	All household members will quarantine when a household member develops COVID symptoms with gravity g.	g, the type of symptoms. If set to 'mild', families self-isolate when anybody develops mild symptoms. If set to 'acute', families only self-isolate if a member develops acute symptoms.
Workplace Partial Lockdown	x% of workplaces are closed. Agents are either furloughed or instructed to work from home.	x, the percentage of workplaces that should be closed

Table 8.4: Summary of containment policies that may be implemented in a simulation. All policies take an additional parameter: the percentage of the population that should have contracted COVID upon which the policy is implemented. For example, a simulation could be set up to close schools when 5% of the population contracts COVID and workplaces when 20% of the population contracts COVID.

A sample containment strategy policy set for a simulation could be:

social-center_closed_0.01|self_isolate_mild_0.01|school_closed_0.01|
reduce_social_contact_to_10_0.01|0.5_workplace_partial_lockdown_0.05|
0.8_workplace_partial_lockdown_0.1

Which translates to:

- With 1% of the population having COVID: Social centres are close, Families with a member having mild symptoms self-isolate, schools are closed, The number of hourly contacts per agent is reduced to a maximum of 10;
- With 5% of the population having COVID: All the previous measures apply, and 50% of workplaces are closed;
- With 10% of the population having COVID: All the measures as mentioned earlier apply, and 80% of workplaces are closed.

Policies are cumulative. When policies applicable for 5% or more of the population have become infected, those policies that started applying those enacted at the 1% threshold remain in place. That is unless a more restrictive policy of the same

category now applies. Eg: If the number of maximum social contacts decreased from 15 to 10, the new policy would apply.

Policies are also incremental. When at the 10% threshold, we close 80% of workplaces. This simply adds a further 30% of workplaces to the original 50%.

Policies are de-activated as soon as the number of COVID-positive cases falls below the applicable threshold.

Population Sampling

In practice, it is impossible for any government to continuously, thoroughly monitor the spread of the epidemic. Therefore, to better align our simulation to real-world monitoring, we introduced sampling strategies.

At the start of each day, a random number of agents n (see Table 8.1) were selected and tested. The positivity ratio of the population, which determined the triggering of containment strategies, was established based on the sampled sub-population.

8.3.7 Simulation Setup

Agents and Venues

At the start of the simulation, agents are assigned to households. The number of agents assigned to each household is drawn from a normal distribution parametrised as per values reported in Table 8.1. Each household must have at least one adult living in it as a constraint. Each adult living in a household is randomly assigned a workplace and social centre. Each child in a household is also randomly assigned a social centre, but all children from the same household will be sharing the same school.

After that, n people are randomly selected and infected with COVID to simulate an initial seed of a contagious population. See Table 8.1 for parameter values.

Containment Strategies

We simulated six scenarios, each comprising different sets of containment strategies. These are summarised in Table 8.5 Containment strategies are activated when 0.0005% of the population has contracted COVID.

Scenario	Containment Strategies
None	None
Minimal	Families with members testing positive are instructed to self-isolate
Self Isolation	Close social centres, families with an infection self-isolate, close schools, some level of social distancing (hourly contacts reduced from 20 to 15)
Soft Lockdown	Close social centres, families with an infection self isolate, close schools, social distancing (hourly contacts reduced from 20 to 10), non-essential activities closed (50% closed, in practice, this means 50% of workplaces were randomly selected and closed, with agents who would normally be there staying at home instead)
Full Lockdown	Close social centres, families with an infection self isolate, close schools, have social distancing (hourly contacts reduced from 20 to 5), non-critical activities closed (80% closure)
Insane Lockdown	Close social centres, families with an infection self-isolate, close schools, have social distancing (hourly contacts reduced from 20 to 5), all workplaces closed

Table 8.5: Summary of simulated scenarios, each scenario comprises different sets of containment strategies. Containment strategies were activated when the proportion of the population infected with COVID exceeded 0.0005%.

Deconfinement Strategies

The second part of our investigation was concerned with deconfinement strategies. We resumed a simulation that had run for 42 days under full lockdown conditions and changed containment policies in turn to each of those detailed in Table 8.6. Again, the same thresholds of infection rates for policy activation were used.

Name	Description
Near Full Lockdown	Close social centres, families with an infection self isolate, close schools, heave social distancing (hourly contacts reduced to 5), non-critical activities closed (closed activities reduced from 80% to 70%)
Contained No Schools	50% of workplaces open with strong social distancing (10 contacts per hour). Schools are closed. Social Centers are closed.
Contained Schools	50% of workplaces open with strong social distancing (10 contacts per hour). Schools are open. Social Centers are closed.
Near Normal	Household with members who contracted COVID self-isolate, mild social distancing (hourly contacts reduced to 15)

Table 8.6: Summary of deconfinement policies. To explore these, we used a model that ran for 42 days under full lockdown and changed containment policies to those detailed above.

Simulation Execution

Simulations for each containment strategy were repeated five times, with epidemiological curves subsequently averaged. The union of all simulations concerning a specific containment strategy is hereby referred to as a simulation set.

For each simulation in a simulation set, parameters are drawn from a distribution (e.g., People's age) were re-drawn each time. This meant that each simulation in a set was executed with similar but distinct populations. We think this is appropriate as it adds stochasticity to the model, confirms its robustness, and avoids results due to the peculiarity of a specific population of agents but rather be more generalised to a set of populations with certain features.

8.4 Results

8.4.1 Impact of Containment Strategies

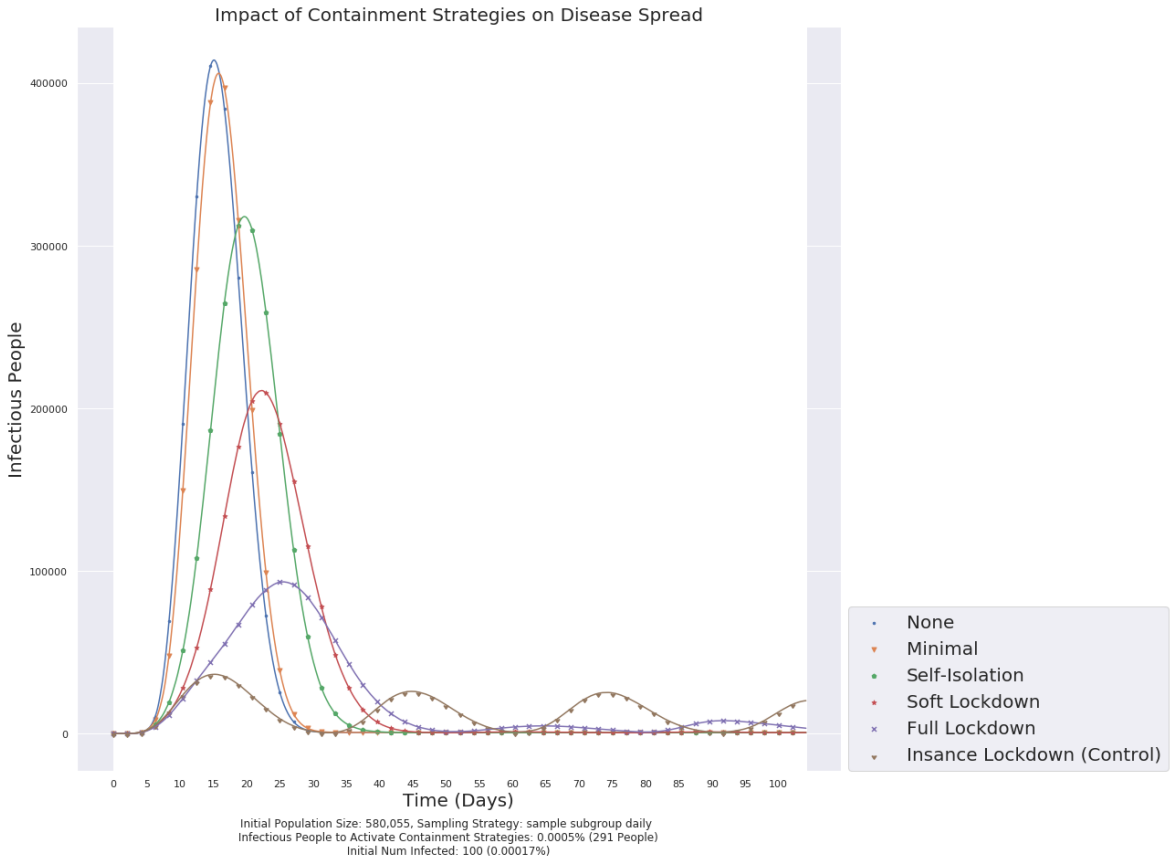

Figure 8.1: Expected impact of containment strategies of the spread of COVID19 across just over three months. Harsher lockdown measures, including the closure of schools, social centres, and most workplaces, more than halve the number of cases.

The impact of containment strategies is reported in Figure 8.1. No containment strategies or just asking families with positive individuals to self-isolate results in a peak of 68% of agents being contemporarily infected and, ultimately, the entire population having contracted COVID throughout the epidemic. It also results in a very sharp rise in cases, and therefore in the overwhelming of the healthcare system.

The first genuinely impactful set of measures is observed under the self-isolation scenario. This includes closing schools and social centres alongside the enforcement of social distancing by reducing hourly contacts from 20 to 15. Under these circumstances, the number of infectious people at the epidemiological peak is reduced by 25%, and the epidemic's progression is delayed by ten days.

Soft lockdown and full lockdown, which also close increasing amounts of non-essential businesses and further reduce hourly contacts, respectively reduce the number of infectious people at peak by 50% and 75%.

Insane lockdown represents an unrealistic scenario where the totality of social

centres, schools and workplaces are closed. This is unfeasible in practice, as essential services, business and retail necessarily must remain open. However, it provides a lower benchmark to compare against. Under this scenario, cases rise is purely due to infections spreading within households. Repeated oscillations are due to confinement measures periodically being lifted and re-implemented. As a result, the number of infected people continuously oscillates around and never exceeds the threshold of 0.0005% significantly.

Perhaps not unexpectedly, more robust lockdown measures contain the spread of COVID better than milder ones. However, the significance of these results is that they suggest that mild containment measures, such as telling families to self-isolate exclusively, do not significantly impact the progression of the disease. This is important in the light of the reluctance of European governments in the Spring of 2020 to act promptly and decisively. On the other hand, our results suggest that harsh measures, including the closing of non-essential businesses, schools and workplaces, must be implemented as soon as the curve of infections rises again.

8.4.2 Deconfiment

Returning to normality after lockdown is a sensitive and, to an extent, emotive process. Policymakers will be under great pressure from the public to lift restrictions as the number of cases and hospitalisations falls. But, especially after months of lockdown, people will understandably be eager to return to their everyday lives as soon as possible.

A risk associated with this, especially in the partial absence of vaccines or effective cures, is for the epidemic to start growing exponentially again, leading to the need for restrictions to be re-imposed a few weeks later.

This was observed, with some delay, in the Autumn of 2020. We speculate that the short delay between the first and second waves in Europe can be attributed to several factors. Firstly, the full lockdowns imposed in many countries in the Spring of 2020 successfully drove COVID numbers to significantly low levels. As a result, the growth of new cases throughout the Summer remained relatively small, at least at first. Secondly, the warmer season may also have impacted the spread of COVID. In Summer, people tend to spend more time outdoors with lower transmission risk and keeping windows open and better ventilating close spaces.

Nevertheless, by late Summer and throughout the extensive vaccine rollout in the Spring of 2021, European countries experienced widespread second and third waves of COVID, which triggered various forms of lockdowns and restrictions. In this second part of the investigation, we will examine the impact of different deconfinement policies and suggest whether other courses of action would have led to alternative outcomes.

We assume that deconfinement strategies would be enacted after 42 days of 'Full Lockdown' (see Table 8.5). After that, we would simulate policy relaxation per strategies described in Table 8.6. The threshold for policy activation was left unchanged. Results are summarised in Figure 8.2.

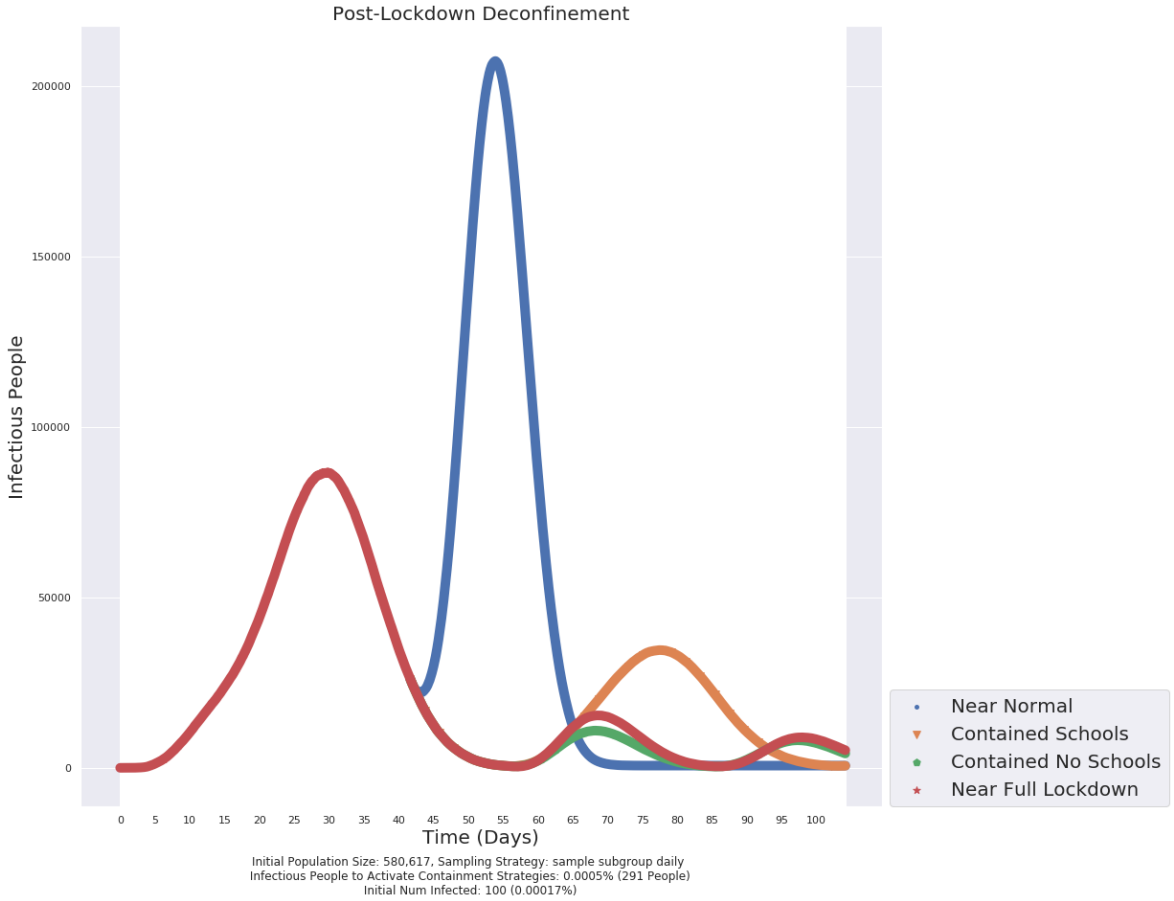

Figure 8.2: Summary of deconfinement policies. The first 42 days of simulation implement the 'Full Lockdown' scenario. After that, different policy relaxation strategies are implemented. We maintained a 0.0005% of population infected threshold for activation. Full lockdown, on average, came into effect at the end of the first week.

Results from our deconfinement investigations clearly show that a rapid return to normality (or near normalcy) where most of the population has not been vaccinated will rapidly develop into a second wave, even worse than the original one. That can be explained by the fact that most people would not yet have contracted the disease and therefore would not have acquired immunity. Upon the population being allowed to mix, as usual, it would not take long for the disease to propagate rapidly and uncontrollably.

As shown in Figure 8.3, the initial intervention with policies of full lockdown can drive down the number of infectious people and contain the spread of the epidemic. However, less than half of the population would at this point have contracted COVID and be therefore immune. A sudden reopening in the absence of herd immunity would drive the outbreak of a second wave.

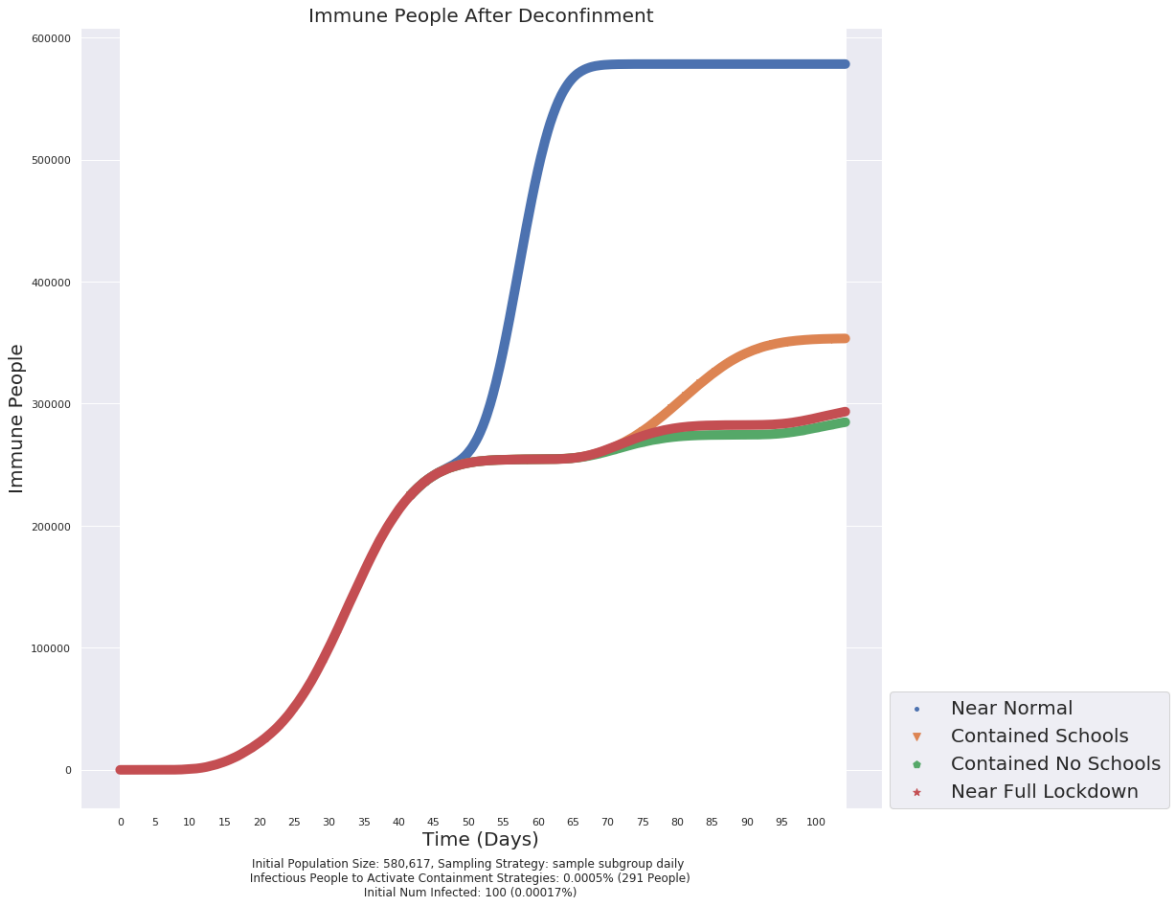

Figure 8.3: The number of immune people in the population. Day 42 is the beginning of deconfinement, coinciding with a significant reduction in infectious people due to the previous full lockdown. However, less than half of the population has contracted the disease and is immune. This is considerably below the herd immunity threshold commonly assumed to be at 60%.

Nevertheless, while results (Figure 8.2) suggest a full reopening is unfeasible, they also show that greater flexibility is possible even in the first deconfinement. Specifically: The difference between 'Near Full Lockdown' and 'Containment No Schools' is negligible. This suggests that a transition may be made from the total closing of all non-essential businesses to a more relaxed scenario where some workplaces can reopen. In addition, a slightly more tolerant maximum number of hourly contacts is also possible, increasing from five to ten.

Interesting to point out, the 'Contained No Schools' scenario develops into an alternation of periods of restrictions and no restrictions. The oscillating patterns observed between days 42 and the simulation's end suggest that these measures are sufficient to drive the number of infected people down below the activation threshold. At that point, with measures withdrawn, the epidemiological curve will start growing again until measures are once again triggered. Both peaks under 'Contained No Schools' are significantly lower than the original peak observed during the first part of the simulation, suggesting that this set of deconfinement policies would not result in public health services being overwhelmed.

Figure 8.3 also shows how under 'Contained No Schools', the proportion of the

population that contracted the disease increased by approximately 8% during the deconfinement phase. So, during approximately two months with, on average across our multiple iterations, half of the population now has contracted the disease. We can therefore speculate that, at the current rate, this would result in 60% of the population having contracted COVID and thus being immune within approximately another three months. Potentially, even with further relaxations allowing more venues and workplaces to open.

On the other, while perhaps ideally balancing public health and economic constraints, the scenario we have just discussed thoroughly disregards social elements. It assumes social venues would remain closed, significantly limiting the possibilities for interaction outside one's main family or workplace bubble. It further assumes schools should be kept closed, with detriment to pupils whose learning is hindered and parents of younger kids who must ensure somebody is available at home to care for children.

We did explore a scenario equivalent to 'Contained No Schools' except for open schools. Unsurprisingly, this is referred to as 'Contained Schools'. Under this assumption, all other limits would remain in place, but all schools would be allowed to stay open. This results in the post-deconfinement COVID peak to include twice as many infected agents, but still without overwhelming the healthcare service and many fewer infected people than there would be had restrictions been completely lifted. This scenario further has the added benefit of addressing the needs of kids and families, minimising disruption to their education and allowing at least some level of social interaction.

8.4.3 Comparison with Observed Data

As identified by other studies studies, see Kerr et al. [110], comparison with empirical data reporting infection rates and/or death numbers is especially difficult. This is due to several factors related to the quality of data available. In the first wave of the epidemic in March 2020, the availability of tests was limited. This was often reserved for patients admitted into hospitals and medical staff or other 'key workers'. Because of this, many asymptomatic patients or people with mild symptoms who recovered at home went undetected and were not reported in the official numbers. Further, delays were encountered in reporting numbers, again making a difficult comparison.

Because of this, a curve comparison similar to the one performed in our investigations of cancer would be inappropriate. If we look at the numbers reported for the first wave by the UK Government for Greater Manchester for infection rates and death, [117] we note that the cumulative total of cases to the end of May 2020 was under 3,000 people. Whereas, our results (Figure 8.1) suggest that even in a full lockdown scenario, cases would be of 90,000.

Instead, we will be performing a qualitative analysis comparing the dynamics of the epidemic. We will first consider the time it took for the epidemic to peak, followed by an exam of the amount of time necessary for infection rates to fall back to an amount comparable with what was observed in Manchester.

Considering the Lockdown scenario in our model (Figure 8.1), the peak of infections was observed on day 26. That is when there were the most infectious people. Regarding actual observed data, we will consider April 6th 2020, as this was when most cases were recorded. This is an estimate that should be accepted with a margin of error. If anything, it could be the day when most people were admitted to hospitals and tested positive, whereas the actual peak could have been at a time close to that. We will consider March 3rd as the start of the epidemic as that is the earliest date for which we have data available, giving us a figure of 34 days.

This gives a difference of 8 days between our predicted and observed infection peaks. We believe this is acceptable for several reasons. Firstly, a margin of discrepancy between simulated and actual results is expected. Secondly, in this case, the observed peak is later than expected. This is consistent with the reporting methodology employed during the first wave, where cases from earlier days were aggregated into the figures of later days as they became available. This could have contributed to shifting the peak forward in the actual data. Finally, we also need to consider variations in the test targets used to report actual data. That means that the tested population sample on some days may have been more representative of the infection rates in the whole population than that on other days. And that therefore, individual values may be less meaningful than day averages over a period of time. While the absolute number of positive cases in actual data is highest on April 6th, the 3-day average is comparable for the preceding 6 days. This could indicate that the actual peak may have occurred in a day in that week, reducing the gap between actual and expected values.

Comparing dates for the end of the first wave is perhaps an even more complex process. The reason behind this is that, while our model had the same rate of testing throughout the simulation (Ie: The entire population was tested each day), test rates in the United Kingdom increased throughout the epidemic. We note that in reported data, 7-day averages of infected people fell from 45 people per day (and note, this is an average of averages) in April 2020 at epidemic peak to 23 people per day in late May when restrictions were lifted. So, on average, we can estimate the first wave lasted approximately three months.

This is in striking contrast to what was predicted by our model, where the first wave ended approximately 50 days (just under two months) after the epidemic began. A possible explanation can be found in Figure 8.3, where our model suggested that by day 45, nearly half of the population would have contracted COVID and therefore been immune. This makes a double implication: The first is about the

actual number of people who contracted the disease, and the second is about the correlation between being contracted and being immune.

Regarding the first implication, our model assumed that people continued behaving and mixing as normal until lockdown measures were implemented. However, this might not have been the case as people began voluntarily shielding or reducing contacts. At the same time, many companies and institutions started asking or allowing staff to work remotely where possible, even before formal measures were implemented. This could have impacted the number of people exposed to the virus, reducing those who acquired immunity and therefore prolonging the first wave.

Regarding the second implication, assuming that everybody who became infected automatically developed full immunity may also be an inaccurate approach. If we allow that even a minority of infected people could contract the disease again, this might have contributed to prolonging the first wave.

An alternative explanation could lie in levels of compliance with lockdown measures. Our model assumed that once lockdown measures were put in place, the entire population would comply with them. However, we know that was not the case and that a whole set of violations ranging from discrete visits to relatives to media-acclaimed "COVID parties" took place at all levels of society. This could have facilitated the spread of the virus, prolonging the duration of the first wave of the pandemic.

And a final possible explanation may lie in the dynamics of the infected people to tests ratio. While the latter was decreasing, the former was increasing, meaning that as there were fewer infected people, it became statistically more likely for them to be captured by a test. Hence, looking at the overall amounts of positive cases, the impression remains that the epidemic was stalling, whereas it was just testing capacity which had increased. But, again, this would not be captured by our model as we monitored the entire population at each epoch rather than sampling a subset of this as occurred in reality.

Finally, we briefly look at our prediction that a sudden return to normal or near-normal conditions after the first lockdown would have eventually led to a second wave worse than the first. This was, unfortunately, the case in November 2020 through to spring 2021, where the United Kingdom (and Manchester) experienced more cases and higher mortality than in the spring of 2020. However, it is unlikely that a more gradual reopening at the end of the first wave followed by significant partial closures in the Summer of 2020 would have allowed for a different scenario. This, because of a lower circulation of the virus in the summer seasons, which would not have permitted or at least significantly reduced the acquisition of herd immunity, thus again presenting a similar scenario in fall.

8.5 Conclusion and Further Work

Predicting the spread of epidemics is a task that presents many challenges. On the one hand, there is the matter of human behaviour. While it may be convenient to assume people will always follow recommendations and policies, we know this is not always the case. Reflecting upon this, the mind probably wanders to sensationalised 'COVID parties' and outright violations for no legitimate purpose. While that is part of the story, there is much more to it. For example, essential workers such as factory employees did not necessarily have access to individual protection devices (Eg: Face masks), nor was it always possible to maintain social distancing. Workers in the so-called 'gig economy' comprise a large portion of undocumented migrants who may not have been eligible for public healthcare, including testing and screening, and may therefore have unwillingly become super-spreaders visiting multiple households per day. And, even the most conscientious person might have accidentally bumped into somebody else turning a corner while shopping for groceries. No two cases are alike, and modelling the heterogeneity of human interaction would make for a separate study in its own right.

There is then the question of the risk of transmission associated with different events. For example, two people having a meal outdoors may lower transmission risk than the same scenario but indoors. A classroom with an open window is better ventilated than one with closed windows, so supposedly safer. And what about the scenario of a busy, open road where people wear face masks but not all in the appropriate way? For example, some are not covering their nose, some temporarily taking it off to sip water. In the absence of precise data, many governments opted for a blanket ban on all social activities. However, understanding the actual risks of transmission is necessary to justify specific restrictions and may allow lighter and as-effective lockdowns in the future.

Nevertheless, while considering the above points may improve the quality of predictions and accuracy of the model, it is unlikely that would significantly affect the overall results to the points that new ones would go as far as contradicting those we reported. We, therefore, summarise the main contributions of investigation as follows:

Firstly, with regards to initial containment strategies. Early, decisive action is fundamental in containing the spread of COVID. Lighter forms of lockdown, especially those only limited to emphasising social distancing and closing a small subsection of venues, are insufficient. Instead, more robust lockdowns, including closing schools and social venues alongside a majority of workplaces, should be enacted.

Secondly, with regards to post-lockdown deconfinement. Contrarily to what was enacted in many European countries, a rapid return to normal is not advised. In

the absence of a vaccinated population and where confinement policies are not re-introduced, this will lead to a second wave comparable to or even worse than the original. Instead, a phased reopening with progressively more sectors and venues should have been introduced. This will allow the community to develop herd immunity without overwhelming the healthcare system gradually.

In conclusion, we successfully implemented a model of COVID spread in a city with demographics comparable to Manchester. Our model simulated people's daily lives, including participation in activities at their household, workplace or school and social centres. Infection was achieved by means of exposure to another infectious person.

We simulated various containment policies, ranging from light-touch lockdowns only recommending enhanced distancing and self-isolation to heavier forms that included closing social centres, schools, and nearly all households. Our results suggested that heavier lockdown measures were required to control the spread of the epidemic.

We further simulated deconfinement policies that being returned to normal societal norms should be phased. We reported that an unphased return, with a rapid return to normal or near-normal, would result in a second wave of the epidemic, even worst than the original. Instead, a phased approach with gradual reopenings should be preferred.

Our work provides evidence of the strengths of agent-based models in the domain of epidemiology, showing how they can be used to obtain information about the efficacy of various public policies in the context of a growing pandemic. Among its strengths is the ability to model individual agents, each with their specific schedule and behaviour and therefore obtain detailed insights concerning the number of infected people in time and the value of various containment policies.

Finally, we identified several ways to improve the project and different research leads that could be used to develop it further.

Chapter 9

Conclusion & Future Work

9.1 Research Summary

Agent-based models are a powerful yet perhaps underappreciated approach to simulating and understanding natural and social phenomena. Their strength lies in the possibility of simulating features of individual agents and therefore include heterogeneity in their behaviour. This has proven especially useful in our Warburg Investigation, where the efficacy of glucose-conjugate drugs is linked to the proportion of the tumour body that acquires the Warburg phenotype and expresses enhanced glucose uptake rates. Furthermore, in the case of our simulation of COVID diffusion, heterogeneity in the schedule of individual agents allowed a more accurate societal representation.

9.2 Summary of Contributions

9.2.1 Contribution 1 - Model Parametrization

The challenge of model parametrization was the first we addressed. Indeed, a simulation may only be as good as its parameters implying that special care must be given to this process. The strategy by which we addressed this issue is two-fold. In the first instance, a preliminary investigation indicated the correlation between each parameter and the model's performance. Results of this, combined with the presence or absence of evidence of the parameter's value from literature, allowed us to categorize each parameter value as 'green', 'yellow' or 'amber'. After that, grid and random search techniques were used to explore the parameter space of those parameters with classification yellow and amber.

Subsequently, treating individual parameters as features of experiment sets and their respective performance as a target value allowed us to treat our simulations as a machine learning task and therefore train decision tree classifiers to attempt and predict a simulation's performance from its parameters. We aimed to train a decision tree to accurately predict whether a simulation would 'perform' or 'not perform' based on a combination of parameter values. That is, whether the resulting tumour volume predicted by the simulation would be within a certain volume threshold of the expected result.

After that, visualizing and 'unravelling' the decision tree and following paths to 'performing' leaf nodes allowed us to establish upper and lower value thresholds for parameters for which these were previously unknown. This allowed us to circumscribe so-called regions of interest in the parameter space. That is, a combination of parameter values previously unknown would lead the model to perform and accurately predict the tumour's volume.

9.2.2 Contribution 2: 2D vs 3D Models

Always within the realm of model parametrization, we considered the projection of 3D tumour models into 2D ones. While commentary in the literature suggests that the former are preferable to the latter, many studies still rely on 2D models. We investigate the effect of approximate translation of parameters measured in a 3D environment to 2D models, finding that the results do not accurately reflect the realities of tumour biology. Our simulations found that our 2D model would not allow cancer cells to experience oxygen starvation, therefore neglecting significant effects of the Warburg phenotype and altering the nature of the tumour micro-environment.

9.2.3 Contribution 3: Investigating the Warburg Phenotype

Following these preliminary studies, we focused on our main research question: Investigating the proportion of cancer cells in a tumour mass that acquire the Warburg phenotype. To this end, we developed a comprehensive model of tumour biology. This included cancer cell agents, intra-cell gene networks, oxygen-dependent expression rates of Hypoxia-Inducible Factors (HIF), HIF-mediated oxygen and glucose uptake rates and VEGF (Vascular Endothelial Growth Factors) secretion rates as well as HIF-mediated proliferation rates. The model further included diffusion of solubles including oxygen, glucose and VEGF.

We considered different scenarios, each assuming different probabilities of cancer cells acquiring the Warburg phenotype and minimum oxygen concentrations required for the survival of non-Warburg cells. Each scenario was simulated multiple times, with growth curves averaged to produce one average growth curve per scenario. Average growth curves were compared to an expected growth curve obtained by fitting a polynomial function to laboratory data obtained from the literature. Error rates at each time-point representing the difference between the expected volume and that produced by the simulation were averaged to obtain an absolute mean error for each scenario. Finally, absolute mean errors were compared, allowing us to suggest that exclusively higher probabilities of Warburg switch (80%) resulted in low error rates across all assumptions of minimum oxygen concentrations required for the survival of non-Warburg cells. In practice, this means that for our model to reproduce growth curves observed in empirical studies, the near-totality of cancer cells must acquire the Warburg phenotype.

Most cancer cells acquiring the Warburg phenotype is relevant to the development of new anti-cancer drugs as it suggests a potential of effectiveness for those leveraging the principle of glyconjugation. That is, binding cytotoxic compounds to sugar molecules. Where cancer cells over-express glucose membrane transporters, they would be absorbing more of these compounds than their non-cancer counterpart,

effectively providing an avenue to targeted therapies. Our contribution here lies in the fact that literature had often reported contrasting evidence regarding the question of whether a majority of cancer cells adopts the Warburg phenotype. Through our simulations, we suggested that observed tumour volumes could not be achieved unless a majority of cancer cells did acquire the Warburg phenotype.

9.2.4 Contribution 4: Our Framework, PanaXea

Our simulations were built with PanaXea, our custom framework for agent-based modelling. PanaXea is built in Python, therefore not requiring users to learn a new programming or scripting language and reducing entry barriers, especially for scientists without a background in Computer Science. Further, the use of Python allows us to leverage a wealth of open source libraries and modules for, among other purposes, data analysis and visualization.

PanaXea is designed as a lite toolkit for the development of agent-based models. It exposes a set of core utilities, including interfaces for steppables (agents and helpers), 2D and 3D objects and numerical environments, the model's schedule and the model object. This aims to facilitate the onboarding of new users, rapidly getting them up to speed and allowing them to start developing their models.

Finally, PanaXea also integrates with major cloud services providers such as Amazon Web Services (AWS). To this end, our project repositories also distribute docker images to facilitate the deployment of the project alongside examples of how different AWS technologies (Eg: Simple Queue Systems, S3 Object Storage, EC2 Compute-Optimized Instances,...) can be leveraged to design a cloud architecture aimed at running hundreds of simulations in parallel and automatically collect and analyze results.

PanaXea emphasizes the importance of model parameterization and sensitivity testing. This is especially important given that many models from the literature do not justify the strategies whereby they assign parameter values. To this end, out-of-the-box tools for parameter space exploration techniques such as grid search and random search are provided.

We presented this set of functionalities through a case study where we reproduced a simplified version of an existing model of Angiogenesis previously published in the literature. Such a model uses a high number of parameters, although their role or relevance in the context of biological mechanisms is not always explained. We assessed the model's sensitivity to two of them: X_0 and κ. We found no correlation between the former and the model's performance, whereas the latter correlates with performance over the value range 0 to 1 and no longer shows any correlation after that. This part of the study emphasized PanaXea's strengths in model parameterization and the necessity to adequately explore a model's parameter space

to avoid uncorrelated parameters that add complexity without providing any net benefit.

9.2.5 Contribution 5: A Model to Predict the Spread of COVID

The COVID 19 pandemic of March 2020 raised additional challenges to our research work. But, it also provided us with an opportunity to explore other modelling domains. We used our aforementioned framework, PanaXea, to model the diffusion of the disease in a city of geography and demographics comparable to Manchester (UK). We leveraged the strengths of agent-based modelling by personalizing the schedule and state of each agent representing a person.

We simulated the effect of various types of confinement policies, simulating more restrictive forms of lockdown progressively to understand what kind of actions would be effective or ineffective at contrasting the spread of COVID. Our results suggested that early, decisive action is required to mitigate the epidemiological peak successfully. We then further looked at deconfinement strategies. That is how to gradually ease restrictions after the disease has been brought back under control. Here, our results suggest that a sudden return to normality with full re-openings should be avoided in the absence of a vaccinated population. This would lead to a second epidemiological wave even worse than the original.

In summary, the different aspects of our research highlight the strengths of agent-based models in multiple domains (tissue modelling, epidemiological modelling), emphasizing how the ability to simulate the state and behaviour of individual actors (be these cells, people...) provides a level of insight which could hardly be achieved by means of alternative modelling approaches. On the other hand, we also remark several potential weaknesses that could limit the predictive power of ABMs or result in erroneous results if not correctly addressed. Among these, the need for careful calibration in the translation from 3D to 2D environments and, of course, the special care that needs to be given to model calibration and sensitivity testing.

9.3 Future Work

9.3.1 The PanaXea Framework

There are several aspects of the PanaXea framework which could be improved. Firstly, the matter of performance should be addressed. At present, the framework progresses the simulation by means of a set of nested Python ‘for’ loops. For example, by repeatedly iterating through all agents in the schedule and calling their step methods.

This approach contributes to the simplicity of the framework. It makes the codebase

accessible even to users who may be new to it and may not necessarily have a formal background in Computer Science. However, the downside of this approach takes the form of a performance cost which slows down the execution of simulations. Alternative approaches should be explored, which both allow maintaining the existing simplicity of the code while at the same time optimizing its execution. These could include moving away from using classes to represent agents to other data structures (Eg: Named tuples, dictionaries, matrixes...), or using code optimizers and preprocessors (Eg: CPython) to produce a build with lower computational costs and better performance. It may also be worth moving to programming languages better known for their performance and speed (C++, Rust...). Finally, exploring parallelization of operations within a single simulation by virtue of multi-threading or parallel runs on multiple nodes or cloud instances may also be beneficial.

Performance is also affected by our strategy to model agents. Each agent is an instance of a Python class in the existing implementation that extends the framework's base Agent class. This is done in the interest of simplicity and intuitiveness to make the framework as approachable as possible to users without a formal technical background. However, this also means we pay the cost of additional overheads due to the instantiation of Python objects, limiting the framework's scalability in terms of the maximum number of agents included in a simulation. Therefore, alternative approaches, such as the use of Python's 'named tuples', should be explored to enhance performance.

9.3.2 2D vs 3D Investigation

We are grateful to the reviewers of the SIMULTECH21 conference for providing valuable feedback on this part of our research work during the review and acceptance phases of our submission.

In the first instance, we would expand the scope of our work. Indeed, our investigation concerned just a single tumour model observed in 2D and 3D. We would recommend expanding this to include multiple types of tumours, or at least multiple starting conditions in terms of initial volume and levels of vascularization.

Secondly, we would extend the model's evaluation to include additional elements beyond tumour volume. Indeed, this point could also apply to our Warburg Investigation (see Section 9.3.3). For example, oxygenation levels across the tumour microenvironment would be especially appropriate as many of our investigations were interested in establishing the prevalence of hypoxia throughout the tumour. As a correlated measurement, comparing average expression rates of Hypoxia-Inducible Factors in time between 2D and 3D could also provide insights into the 3D to 2D translation quality. This is because HIF expression rates are correlated to oxygen concentration and could act as a valuable proxy to monitor to obtain information about the evolution of spatial-temporal concentrations of the latter.

Finally, alongside exploring and identifying cases where 3D-to-2D translations do not work, it would also be interesting to establish which approaches do, on the other, lead to the correct functioning of the model even in 2D. For example, thinking back to our specific scenario, how could we ensure that cancer cells in a 2D model experience the same levels of oxygen starvation as their 3D counterparts? And, under which conditions would we be satisfied that the heterogeneity in regions of hypoxia produced in the 2D simulation was appropriate and representative of what could be observed in 3D?

In conclusion, an ambitious aim to expand this work further could be the creation of a protocol (or at least some 'best practices) to translate 3D models to 2D ones. This should include the transformation of volume units, the initial distribution of agents and any adaptation of parameter values. Especially, those related to diffusion.

9.3.3 The Warburg Investigation

Previous works have treated the source and sink rates of solubles (oxygen, VEGF, glucose...) as constants. Still, allowing such values to vary across the environment will enable us to obtain more realistic concentrations at each position. Given that hypoxia had such a crucial role in our investigation, calculating accurate concentration values was primarily important. In addition, our model can be easily expanded to account for factors such as the permeability of the vessel membrane and the effect of free oxygen and haemoglobin-bound oxygen.

Many of the previous works developed in-silico models of tumour growth in a 2D environment, citing the high computational costs of 3D ones. While the cost difference is significant, efficiently leveraging cloud resources has allowed us to overcome this obstacle and develop ours in 3D. Furthermore, this allowed us better to capture the Spatio-temporal evolution of the tumour mass and provide a more realistic description of its morphology and that of the blood vessel structure, as previously discussed.

Furthermore, our model accommodates intracellular gene networks and gene regulation complexity. Rather than implementing proliferation rates, uptake rates, and others as single values, continuous functions elaborate on the output of similar functions at the previous levels. For example, the proliferation rate is expressed as a probability proportional to the HIF expression rate, dependent on oxygen levels. With minimal effort, the complexity of this can be increased to multiple levels of gene regulation or even, as is often observed, multiple genes regulating the expression of others.

9.3.4 Exploring Glyconjugation

In-silico drug delivery as a strategy for clinical investigation has been the object of numerous publications [118–121].

A pharmacokinetic model describing the release, diffusion, and absorption of drugs in the interstitial space could simulate the delivery of glycoconjugate drugs. Differential equations would be used to describe the diffusion process. In contrast, source and sink rates of the soluble under study would be properties of its diffusivity across vessel membranes and cell uptake rates.

The environment would be broken down into multiple compartments, which, similarly to Zhan et al. [121] could represent the normal tissues, circulatory system, tumour intravascular space and tumour extracellular space. Each would have its diffusion dynamics and act as a source/sink from/to adjacent compartments. Specifically, we would expect the intravascular tumour space to have a higher uptake rate of the said drug than the normal tissues owing to higher expression rates of glucose membrane transporters.

The drug's effect on cancer cells would be evaluated through a pharmacodynamics model, which would describe the impact of the drug on the tumour. In addition, such a model could be used to explore the effect of different concentrations submnistered over various periods to determine which treatment, if any, is most effective.

As pointed out in previous studies [121], comparing tumour volumes across different in-silico trials may not be possible due to the short time window simulated. Therefore, one of the proposed empirical correlations, such as cell survival rate, may be more appropriate.

Of particular interest would be combining a pharmacokinetic model with our model of tumour angiogenesis, therefore exploring how the effects of the therapy change as these are implemented earlier or later in the development of the tumour. Also of interest would be observing the extent of the tumour sufficiently vascularized to deliver the drug successfully and the possible dynamics observed in hypoxic regions, which would become less of a target for therapy.

Finally, the development of such a model would help assess the impact of such therapy on healthy tissues. While it is true that the expectation is that cancer cells would be most affected, it is also reasonable to believe that an impact on healthy cells would be observed. Minimizing this is crucial for the therapy to be "targeted" and with as few side effects as possible.

9.3.5 COVID Modelling

We want to discuss how our work could be further expanded and improved.

The first and perhaps most apparent is to extend the scope from a single city to a region or even the entire country. This would be useful to obtain a more holistic view of the pandemic, including potential differences between rural and urban areas and how population density affects the epidemiological curve. However, this also presents several challenges, which we will now discuss.

Firstly, there is the matter of computational complexity. PanaXea, our framework, creates each agent as an instance of a Python object. This simplifies the development of the model but will also cause memory overheads which impact performance as the number of agents increases requiring more powerful servers. Expanding to include more agents and locations also means epochs will take longer to complete, as each of these will need to be evaluated. Vertical scaling, which uses more powerful servers, may work up to a point thanks to cloud resources that can offer hundreds of GB of RAM but may not be sufficient for particularly large models. Finally, horizontal scaling, so to parallelize evaluations, may be considered but would require re-engineering the simulation code and potentially re-engineering the underlying framework code itself. However, if achieved, this would allow the running of larger and more complex simulations.

The simulation's setup would also need to be enhanced in several ways. Thus far, agents have been randomly assigned schools, workplaces and social centres. That is based on the assumption that, in simulating a single city, all locations should be commutable to daily. However, that would not be the case when simulating larger territories. It is, for example, unfeasible that somebody living in Southampton would commute to Aberdeen daily. So instead, assigned places should be allocated based on the distance from their household or habitual place of residence.

Nonetheless, people travel medium and long distances as much as several times per week. That would need to be accounted for, perhaps with a hybrid schedule system where people were assigned different households, workplaces, etc. and switched between them periodically. Ie: If infected, carrying the disease with them.

The infection system itself would need further investigating. We treat all encounters with an infected person as equally likely to transmit COVID. While that may be appropriate to understand the macro-nature of the pandemic, the actual nature of viral transmission is probably more complicated. For example, the distance between individuals, ventilation, temperature, among other factors, all probably influence the outcome of an encounter. This, too, would need to be accounted for.

As mentioned in the introduction, there is the question of compliance with restrictive measures. We assumed that there would be full compliance with government policy,

whereas empirical evidence shows that is not the case. There are even cases of sanctioned mixing even in the harshest lockdown scenarios. For example, at least one household member must go and shop for groceries at least once a week. Even for the most resilient households that can fully operate from home, it is improbable and probably only applicable to a minority that they would not be able to leave their premises for months.

Agent scheduling, too, could be improved. We assumed that agents would exclusively be assigned a single household, workplace or school and social centre. While this might be appropriate for the first two, we know that in reality, people visit various other places daily and that their schedule might, and often would, change from day to day. This could perhaps be easily rectified by assigning multiple social centres to an agent or possibly changing their schedule from day to day.

Model parametrization could also be improved. At the very least, sensitivity testing would be needed to establish which parameters the model is most sensitive to and the way in which results are affected. Once this information was obtained, a methodology similar to that of our Warburg investigation (see Chapter 6) comprising Grid Search and Random Search would be useful in parameterizing the model and obtaining correct parameter values. Was the model was adequately calibrated on past data from 2020, 2021 and 2022 then it could become a predictor tool for future waves of COVID-19 or other air-born viral diseases.

And finally, there is the matter of confinement strategies. Perhaps the most obvious to introduce would be to account for the use of face masks. The reason for not including them in this investigation was that in the Spring of 2020, widespread shortages meant they were not available to the general public, and their use was therefore not mandated by government policy. However, the scenario changed during the following autumn in winter, and if developing the project further, it would make sense to include them. Similarly, with vaccine rollout accelerating now, it would be interesting to investigate how deconfinement policies should be shaped to account for this. For example, progressively easing restrictions for everybody as more people are vaccinated or perhaps only for those who have received their first dose?

Bibliography

[1] Peter Vaupel, Heinz Schmidberger, and Arnulf Mayer. “The Warburg effect: essential part of metabolic reprogramming and central contributor to cancer progression”. In: International Journal of Radiation Biology 95.7 (July 2019), pp. 912–919. ISSN: 13623095. DOI: 10.1080/09553002.2019.1589653/FORMAT/EPUB. URL: https://www.tandfonline.com/action/journalInformation?journalCode=irab20.

[2] Rob A. Cairns. “Drivers of the Warburg phenotype”. In: Cancer journal (Sudbury, Mass.) 21.2 (Mar. 2015), pp. 56–61. ISSN: 1540-336X. DOI: 10.1097/PPO.0000000000000106. URL: https://pubmed.ncbi.nlm.nih.gov/25815844/.

[3] Leszek Szablewski. “Expression of glucose transporters in cancers”. In: Biochimica et Biophysica Acta (BBA) - Reviews on Cancer 1835.2 (Apr. 2013), pp. 164–169. ISSN: 0304-419X. DOI: 10.1016/J.BBCAN.2012.12.004.

[4] Carly C. Barron et al. “Facilitative glucose transporters: Implications for cancer detection, prognosis and treatment”. In: Metabolism 65.2 (Feb. 2016), pp. 124–139. ISSN: 0026-0495. DOI: 10.1016/J.METABOL.2015.10.007.

[5] D Hanahan and R A Weinberg. “The hallmarks of cancer”. In: Cell 100.1 (Jan. 2000), pp. 57–70. URL: https://www.ncbi.nlm.nih.gov/pubmed/10647931.

[6] Benjamin J. Bornstein et al. “LibSBML: an API Library for SBML”. In: Bioinformatics 24.6 (Mar. 2008), pp. 880–881. ISSN: 1367-4803. DOI: 10.1093/BIOINFORMATICS/BTN051. URL: https://academic.oup.com/bioinformatics/article/24/6/880/194657.

[7] Alexandru Daşu, Iuliana Toma-Daşu, and Mikael Karlsson. “The effects of hypoxia on the theoretical modelling of tumour control probability.” In: Acta oncologica (Stockholm, Sweden) 44.6 (2005), pp. 563–571. ISSN: 0284-186X. DOI: 10.1080/02841860500244435.

[8] Gibin G. Powathil et al. “Modelling the effects of cell-cycle heterogeneity on the response of a solid tumour to chemotherapy: Biological insights from a hybrid multiscale cellular automaton model”. In: Journal of Theoretical

Biology 308 (Sept. 2012), pp. 1–19. ISSN: 00225193. DOI: 10.1016/j.jtbi.2012.05.015.

[9] Raja Venkatasubramanian, Michael A. Henson, and Neil S. Forbes. "Incorporating energy metabolism into a growth model of multicellular tumor spheroids". In: Journal of Theoretical Biology 242.2 (Sept. 2006), pp. 440–453. ISSN: 00225193. DOI: 10.1016/j.jtbi.2006.03.011.

[10] A. Bertuzzi et al. "ATP Production and Necrosis Formation in a Tumour Spheroid Model". In: Mathematical Modelling of Natural Phenomena 2.3 (June 2007), pp. 30–46. ISSN: 0973-5348. DOI: 10.1051/mmnp:2007002. URL: http://www.mmnp-journal.org/10.1051/mmnp:2007002.

[11] Stefan Schuster et al. "Metabolic pathways analysis". In: (2015). DOI: 10.1042/BST20150153.

[12] D. Basanta et al. "Evolutionary game theory elucidates the role of glycolysis in glioma progression and invasion". In: Cell Proliferation 41.6 (Dec. 2008), pp. 980–987. ISSN: 0960-7722. DOI: 10.1111/j.1365-2184.2008.00563.x. URL: https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2184.2008.00563.x.

[13] Amina a Qutub and Aleksander S Popel. "A computational model of intracellular oxygen sensing by hypoxia-inducible factor HIF1 alpha." In: Journal of cell science 119.Pt 16 (2006), pp. 3467–3480. ISSN: 0021-9533. DOI: 10.1242/jcs.03087. arXiv: NIHMS150003.

[14] Kerri Ann Norton, Kideok Jin, and Aleksander S. Popel. "Modeling triple-negative breast cancer heterogeneity: Effects of stromal macrophages, fibroblasts and tumor vasculature". In: Journal of Theoretical Biology 452 (Sept. 2018), pp. 56–68. ISSN: 10958541. DOI: 10.1016/j.jtbi.2018.05.003. URL: https://doi.org/10.1016/j.jtbi.2018.05.003.

[15] Natalia L. Komarova et al. "Aspirin and the chemoprevention of cancers: A mathematical and evolutionary dynamics perspective". In: Wiley Interdisciplinary Reviews: Systems Biology and Medicine 12.5 (Sept. 2020), e1487. ISSN: 19395094. DOI: 10.1002/wsbm.1487. URL: http://doi.wiley.com/10.1002/wsbm.1487.

[16] Tadanobu Shimura et al. "A comprehensive in vivo and mathematic modeling-based kinetic characterization for aspirin-induced chemoprevention in colorectal cancer". In: Carcinogenesis 41.6 (July 2020), pp. 751–760. ISSN: 0143-3334. DOI: 10.1093/carcin/bgz195. URL: https://academic.oup.com/carcin/article/41/6/751/5696845.

[17] Roberta Coletti, Andrea Pugliese, and Luca Marchetti. “Modeling the effect of immunotherapies on human castration-resistant prostate cancer”. In: Journal of Theoretical Biology 509 (Jan. 2021), p. 110500. ISSN: 10958541. DOI: 10.1016/j.jtbi.2020.110500.

[18] M. Scalerandi and B. Sansone. “Inhibition of Vascularization in Tumor Growth”. In: Physical Review Letters 89.21 (2002), p. 218101. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.89.218101. URL: http://link.aps.org/doi/10.1103/PhysRevLett.89.218101.

[19] Dirk Drasdo and Stefan Höhme. “A single-cell-based model of tumor growth in vitro: Monolayers and spheroids”. In: Physical Biology 2.3 (Sept. 2005), pp. 133–147. ISSN: 14783975. DOI: 10.1088/1478-3975/2/3/001. URL: https://iopscience.iop.org/article/10.1088/1478-3975/2/3/001%20https://iopscience.iop.org/article/10.1088/1478-3975/2/3/001/meta.

[20] Sergey Astanin and Luigi Preziosi. “Mathematical modelling of the Warburg effect in tumour cords”. In: Journal of Theoretical Biology 258.4 (2009), pp. 578–590. ISSN: 00225193. DOI: 10.1016/j.jtbi.2009.01.034.

[21] T W Secomb et al. “Theoretical simulation of oxygen transport to brain by networks of microvessels: effects of oxygen supply and demand on tissue hypoxia.” In: Microcirculation (New York, N.Y. : 1994) 7.4 (2000), pp. 237–247. ISSN: 1073-9688. DOI: 10.1080/mic.7.4.237.247. URL: http://www.ncbi.nlm.nih.gov/pubmed/10963629$%5Cbackslash$nhttp://onlinelibrary.wiley.com/resolve/openurl?genre=article%7B%5C&%7Dsid=nlm:pubmed%7B%5C&%7Dissn=1073-9688%7B%5C&%7Ddate=2000%7B%5C&%7Dvolume=7%7B%5C&%7Dissue=4%7B%5C&%7Dspage=237$%5Cbackslash$nhttp://informahealthcare.com/doi/abs/10.1080/mic.7.4.237.247$%5Cbackslash$nAll%20Papers/S/Secomb.

[22] T W Secomb and R Hsu. “Simulation of O2 transport in skeletal muscle: diffusive exchange between arterioles and capillaries.” In: The American journal of physiology 267.3 Pt 2 (1994), H1214–H1221. ISSN: 03636135. URL: http://www.ncbi.nlm.nih.gov/pubmed/8092288.

[23] Daniel Goldman. “Theoretical Models of Microvascular Oxygen Transport to Tissue”. In: Building 15.8 (2011), pp. 795–811. DOI: 10.1080/10739680801938289.Theoretical.

[24] Yan Cai et al. “Coupled modelling of tumour angiogenesis, tumour growth and blood perfusion”. In: Journal of Theoretical Biology 279.1 (2011), pp. 90–101.

ISSN: 00225193. DOI: 10.1016/j.jtbi.2011.02.017. URL: http://dx.doi.org/10.1016/j.jtbi.2011.02.017.

[25] M. Welter, K. Bartha, and H. Rieger. "Vascular remodelling of an arterio-venous blood vessel network during solid tumour growth". In: Journal of Theoretical Biology 259.3 (Aug. 2009), pp. 405–422. ISSN: 00225193. DOI: 10.1016/j.jtbi.2009.04.005.

[26] a R Anderson and M a Chaplain. "Continuous and discrete mathematical models of tumor-induced angiogenesis." In: Bulletin of mathematical biology 60.5 (1998), pp. 857–899. ISSN: 00928240. DOI: 10.1006/bulm.1998.0042.

[27] Paul Macklin et al. "Multiscale modelling and nonlinear simulation of vascular tumour growth". In: Journal of Mathematical Biology 58.4-5 (2009), pp. 765–798. ISSN: 03036812. DOI: 10.1007/s00285-008-0216-9.

[28] Steven R. McDougall, Alexander R A Anderson, and Mark A J Chaplain. "Mathematical modelling of dynamic adaptive tumour-induced angiogenesis: Clinical implications and therapeutic targeting strategies". In: Journal of Theoretical Biology 241.3 (2006), pp. 564–589. ISSN: 00225193. DOI: 10.1016/j.jtbi.2005.12.022.

[29] A R A Anderson S.R. McDougall and J A Sherratt M.A.J. Chaplain. "Mathematical Modelling of Flow Through Vascular Networks: Implications for Tumor-induced Angiogenesis and Chemoterapy Strategies". In: Bulletin of Mathematical Biology 64 (2002), pp. 673–702. ISSN: 00928240. DOI: 10.1006/bulm.2002.0293.

[30] Yi Jiang et al. "A multiscale model for avascular tumor growth." In: Biophysical journal 89.6 (2005), pp. 3884–3894. ISSN: 00063495. DOI: 10.1529/biophysj.105.060640. URL: http://dx.doi.org/10.1529/biophysj.105.060640.

[31] G. S. Stamatakos et al. "An advanced discrete state-discrete event multiscale simulation model of the response of a solid tumor to chemotherapy: Mimicking a clinical study". In: Journal of Theoretical Biology 266.1 (2010), pp. 124–139. ISSN: 00225193. DOI: 10.1016/j.jtbi.2010.05.019. URL: http://dx.doi.org/10.1016/j.jtbi.2010.05.019.

[32] Jana L. Gevertz and Salvatore Torquato. "Modeling the effects of vasculature evolution on early brain tumor growth". In: Journal of Theoretical Biology 243.4 (2006), pp. 517–531. ISSN: 00225193. DOI: 10.1016/j.jtbi.2006.07.002.

[33] F. Spill et al. “Mesoscopic and continuum modelling of angiogenesis”. In: Journal of mathematical biology 70.3 (2015), pp. 485–532. ISSN: 14321416. DOI: 10.1007/s00285-014-0771-1. arXiv: arXiv:1401.5701v1.

[34] Richard Hsu and Timothy W. Secomb. “A Green’s function method for analysis of oxygen delivery to tissue by microvascular networks”. In: Mathematical Biosciences 96.1 (1989), pp. 61–78. ISSN: 00255564. DOI: 10.1016/0025-5564(89)90083-7.

[35] Megan M. Olsen and Hava T. Siegelmann. “Multiscale agent-based model of tumor angiogenesis”. In: Procedia Computer Science 18 (2013), pp. 1016–1025. ISSN: 18770509. DOI: 10.1016/j.procs.2013.05.267. URL: http://dx.doi.org/10.1016/j.procs.2013.05.267.

[36] P. Gerlee and A. R.A. Anderson. “A hybrid cellular automaton model of clonal evolution in cancer: The emergence of the glycolytic phenotype”. In: Journal of Theoretical Biology (2008). ISSN: 00225193. DOI: 10.1016/j.jtbi.2007.10.038.

[37] Yan Cai et al. “Mathematical Modelling of a brain tumour initiation and early development: A coupled model of glioblastoma growth, pre-existing vessel co-option, angiogenesis and blood perfusion”. In: PLoS ONE 11.3 (2016), pp. 1–28. ISSN: 19326203. DOI: 10.1371/journal.pone.0150296.

[38] Milad Shamsi et al. “Mathematical Modeling of the Function of Warburg Effect in Tumor Microenvironment”. In: January (2018), pp. 1–13. DOI: 10.1038/s41598-018-27303-6.

[39] Jakob Nikolas Kather et al. “In silico modeling of immunotherapy and stroma-targeting therapies in human colorectal cancer”. In: Cancer Research 77.22 (2017), pp. 6442–6452. ISSN: 15387445. DOI: 10.1158/0008-5472.CAN-17-2006.

[40] John P.A. Ioannidis. “Why Most Published Research Findings Are False”. In: PLOS Medicine 2.8 (July 2005), e124. ISSN: 1549-1676. DOI: 10.1371/JOURNAL.PMED.0020124. URL: https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.0020124.

[41] Yoonseok Kam, Katarzyna A. Rejniak, and Alexander R.A. Anderson. “Cellular modeling of cancer invasion: Integration of in silico and in vitro approaches”. In: Journal of Cellular Physiology 227.2 (2012), pp. 431–438. ISSN: 00219541. DOI: 10.1002/jcp.22766. arXiv: NIHMS150003.

[42] Mohammad Abdollahi Soodabeh Saeidnia Azadeh Manayi. “The Pros and Cons of the In-silico Pharmaco- toxicology in Drug Discovery and Development”. In: August (2013). DOI: 10.3923/ijp.2013.176.181.

[43] Zvia Agur et al. “Employing dynamical computational models for personalizing cancer immunotherapy”. In: EXPERT OPINION ON BIOLOGICAL THERAPY 2598.August (2016). DOI: 10.1080/14712598.2016.1223622.

[44] S Ekins, J Mestres, and B Testa. “In silico pharmacology for drug discovery : applications to targets and beyond”. In: December 2006 (2007), pp. 21–37. DOI: 10.1038/sj.bjp.0707306.

[45] Steven R. McDougall, Alexander R.A. Anderson, and Mark A.J. Chaplain. “Mathematical modelling of dynamic adaptive tumour-induced angiogenesis: Clinical implications and therapeutic targeting strategies”. In: Journal of Theoretical Biology 241.3 (Aug. 2006), pp. 564–589. ISSN: 00225193. DOI: 10.1016/j.jtbi.2005.12.022. URL: http://www.ncbi.nlm.nih.gov/pubmed/16487543%20https://linkinghub.elsevier.com/retrieve/pii/S0022519305005564.

[46] Uri Wilensky. NetLogo. http://ccl.northwestern.edu/netlogo/. Northwestern University, Evanston, IL: Center for Connected Learning and Computer-Based Modeling, 1999. URL: http://ccl.northwestern.edu/netlogo/.

[47] Patrick Taillandier et al. “Building, composing and experimenting complex spatial models with the GAMA platform”. en. In: GeoInformatica 23.2 (Apr. 2019), pp. 299–322. ISSN: 1384-6175, 1573-7624. DOI: 10.1007/s10707-018-00339-6. URL: http://link.springer.com/10.1007/s10707-018-00339-6 (visited on 04/06/2021).

[48] Michael J North et al. “Complex adaptive systems modeling with Repast Simphony”. In: Complex Adaptive Systems Modeling 1.1 (2013). DOI: 10.1186/2194-3206-1-3.

[49] Sean Luke et al. “MASON: A Multiagent Simulation Environment”. In: Simulation 81.7 (2005), pp. 517–527. DOI: 10.1177/0037549705058073.

[50] David Masad and Jacqueline Kazil. “Mesa: An Agent-Based Modeling Framework”. In: Jan. 2015, pp. 51–58. DOI: 10.25080/Majora-7b98e3ed-009.

[51] J. Sinek et al. “Two-dimensional chemotherapy simulations demonstrate fundamental transport and tumor response limitations involving nanoparticles”. In: Biomedical Microdevices 6.4 (2004), pp. 297–309. ISSN: 13872176. DOI: 10.1023/B:BMMD.0000048562.29657.64.

[52] Donglai Lv et al. Three-dimensional cell culture: A powerful tool in tumor research and drug discovery. Dec. 2017. DOI: 10.3892/ol.2017.7134.

[53] M. Pickl and C. H. Ries. “Comparison of 3D and 2D tumor models reveals enhanced HER2 activation in 3D associated with an increased response to trastuzumab”. In: Oncogene 28.3 (Jan. 2009), pp. 461–468. ISSN: 14765594. DOI: 10.1038/onc.2008.394.

[54] Sebastien Rey-Coyrehourcq Romain Reuillon Mathieu Leclaire. “OpenMOLE, a workflow engine specifically tailored for the distributed exploration of simulation models”. In: Future Generation Computer Systems 29.8 (2013), pp. 1981–1990. URL: http://www.openmole.org/files/FGCS2013.pdf.

[55] James Bergstra JAMESBERGSTRA and Umontrealca Yoshua Bengio YOSHUABENGIO. “Random Search for Hyper-Parameter Optimization”. In: Journal of Machine Learning Research 13 (2012), pp. 281–305. ISSN: 1532-4435. DOI: 10.1162/153244303322533223. arXiv: 1504.05070.

[56] Sarah M Keating et al. “SBML Level 3: an extensible format for the exchange and reuse of biological models”. In: Molecular Systems Biology 16.8 (Aug. 2020), e9110. ISSN: 1744-4292. DOI: 10.15252/MSB.20199110. URL: https://onlinelibrary.wiley.com/doi/full/10.15252/msb.20199110%20https://onlinelibrary.wiley.com/doi/abs/10.15252/msb.20199110%20https://www.embopress.org/doi/abs/10.15252/msb.20199110.

[57] Dagmar Waltemath et al. “Reproducible computational biology experiments with SED-ML - The Simulation Experiment Description Markup Language”. In: BMC Systems Biology 5.1 (Dec. 2011), pp. 1–10. ISSN: 17520509. DOI: 10.1186/1752-0509-5-198/FIGURES/5. URL: https://bmcsystbiol.biomedcentral.com/articles/10.1186/1752-0509-5-198.

[58] Bilal Shaikh et al. “runBioSimulations: an extensible web application that simulates a wide range of computational modeling frameworks, algorithms, and formats”. In: bioRxiv (Mar. 2021), p. 2021.03.05.433787. DOI: 10.1101/2021.03.05.433787. URL: https://www.biorxiv.org/content/10.1101/2021.03.05.433787v1%20https://www.biorxiv.org/content/10.1101/2021.03.05.433787v1.abstract.

[59] Jonathan P Whiteley et al. “Mathematical Biology Mathematical modelling of oxygen transport to tissue”. In: J. Math. Biol 44 (2002), pp. 503–522. DOI: 10.1007/s002850200135.

[60] Jonathan E Guyer, Daniel Wheeler, and James A Warren. “FiPy: Partial Differential Equations with Python”. In: Computing in Science Engineering 11.3 (1988), pp. 6–15. ISSN: 15219615. DOI: 10.1109/MCSE.2009.52. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4814978.

[61] Larry S. Cohen and George P. Studzinski. “Correlation between cell enlargement and nucleic acid and protein content of hela cells in unbalanced growth produced by inhibitors of DNA synthesis”. In: Journal of Cellular Physiology 69.3 (1967), pp. 331–339. ISSN: 10974652. DOI: 10.1002/jcp.1040690309.

[62] Theodore T Puck, Philip I Marcus, and Steven J : Cieciura. CLONAL G R O W T H OF M A M M A L I A N CELLS IN VITRO GROWTH CHARACTERISTICS OF COLONIES FROM SINGLE HELA CELLS WITH AND WITHOUT A FEEDER LAYER*. Tech. rep.

[63] James Bergstra et al. “Algorithms for Hyper-Parameter Optimization”. In: Advances in Neural Information Processing Systems (NIPS) (2011), pp. 2546–2554. ISSN: 10495258. DOI: 2012arXiv1206.2944S.

[64] Styliani Melissaridou et al. “The effect of 2D and 3D cell cultures on treatment response, EMT profile and stem cell features in head and neck cancer 11 Medical and Health Sciences 1112 Oncology and Carcinogenesis”. In: Cancer Cell International 19.1 (Jan. 2019). ISSN: 14752867. DOI: 10.1186/s12935-019-0733-1.

[65] Yoshinori Imamura et al. “Comparison of 2D- and 3D-culture models as drug-testing platforms in breast cancer”. In: Oncology Reports 33.4 (Apr. 2015), pp. 1837–1843. ISSN: 17912431. DOI: 10.3892/or.2015.3767.

[66] Kayla Duval et al. Modeling physiological events in 2D vs. 3D cell culture. June 2017. DOI: 10.1152/physiol.00036.2016.

[67] Jessica Hoarau-Véchot et al. Halfway between 2D and animal models: Are 3D cultures the ideal tool to study cancer-microenvironment interactions? Jan. 2018. DOI: 10.3390/ijms19010181.

[68] Angelika Riedl et al. “Comparison of cancer cells in 2D vs 3D culture reveals differences in AKT-mTOR-S6K signaling and drug responses”. In: Journal of Cell Science 130.1 (2017), pp. 203–218. ISSN: 14779137. DOI: 10.1242/jcs.188102.

[69] Feilim Mac Gabhann, James W Ji, and Aleksander S Popel. “Computational Model of Vascular Endothelial Growth Factor Spatial Distribution in Muscle and Pro-Angiogenic Cell Therapy”. In: (). DOI: 10.1371/journal.pcbi.0020127. URL: www.ploscompbiol.org.

[70] Anne C. Skeldon et al. “Modelling and detecting tumour oxygenation levels”. In: PLoS ONE 7.6 (2012), pp. 1–29. ISSN: 19326203. DOI: 10.1371/journal.pone.0038597. arXiv: arXiv:1011.1669v3.

[71] Le Zhang et al. Developing a multiscale, multi-resolution agent-based brain tumor model by graphics processing units. Tech. rep. 2011. DOI: 10.1186/1742-4682-8-46. URL: http://www.tbiomed.com/content/8/1/46.

[72] Stephen Turner and Jonathan A. Sherratt. “Intercellular adhesion and cancer invasion: A discrete simulation using the extended potts model”. In: Journal of Theoretical Biology 216.1 (2002), pp. 85–100. ISSN: 00225193. DOI: 10.1006/jtbi.2001.2522.

[73] Xiaoqiang Sun et al. Multi-scale agent-based brain cancer modeling and prediction of TKI treatment response: Incorporating EGFR signaling pathway and angiogenesis. Tech. rep. 2012. DOI: 10.1186/1471-2105-13-218. URL: http://www.biomedcentral.com/1471-2105/13/218.

[74] Abbas Shirinifard et al. “3D multi-cell simulation of tumor growth and angiogenesis”. In: PLoS ONE 4.10 (2009). ISSN: 19326203. DOI: 10.1371/journal.pone.0007190.

[75] A St et al. “MATHEMATICAL AND Mathematical Modelling of Flow in 2D and 3D Vascular Networks: Applications to Anti-Angiogenic and Chemotherapeutic Drug Strategies”. In: COMPUTER MODELLING Mathematical and Computer Modelling 41 (2005), pp. 1137–1156. DOI: 10.1016/j.mcm.2005.05.008. URL: www.sciencedirect.comwww.elsevier.com/locate/mcm.

[76] Luis Miguel Rios et al. “Derivative-free optimization: a review of algorithms and comparison of software implementations”. In: J Glob Optim 56 (2013), pp. 1247–1293. DOI: 10.1007/s10898-012-9951-y. URL: https://link.springer.com/content/pdf/10.1007%7B%5C%%7D2Fs10898-012-9951-y.pdf.

[77] Yi Chen et al. “Combined Src and ER blockade impairs human breast cancer proliferation in vitro and in vivo”. In: (). DOI: 10.1007/s10549-010-1024-7.

[78] Beatrice Philip et al. “HiF expression and the role of hypoxic microenvironments within primary tumours as protective sites driving cancer stem cell renewal and metastatic progression”. In: Carcinogenesis 34.8 (2013), pp. 1699–1707. DOI: 10.1093/carcin/bgt209. URL: https://academic.oup.com/carcin/article-abstract/34/8/1699/2463397.

[79] Gregg L. Semenza. “HIF-1 mediates the Warburg effect in clear cell renal carcinoma”. In: Journal of Bioenergetics and Biomembranes 39.3 (2007), pp. 231–234. ISSN: 0145479X. DOI: 10.1007/s10863-007-9081-2.

[80] Marco Archetti. “Evolutionary dynamics of the Warburg effect: Glycolysis as a collective action problem among cancer cells”. In: Journal of Theoretical Biology 341 (Jan. 2014), pp. 1–8. ISSN: 00225193. DOI: 10.1016/j.jtbi.2013.09.017. URL: https://pubmed.ncbi.nlm.nih.gov/24075895/.

[81] Kieran Smallbone et al. “Metabolic changes during carcinogenesis: Potential impact on invasiveness”. In: Journal of Theoretical Biology 244.4 (Feb. 2007), pp. 703–713. ISSN: 00225193. DOI: 10.1016/j.jtbi.2006.09.010. URL: https://pubmed.ncbi.nlm.nih.gov/17055536/.

[82] Peter Vaupel et al. “Revisiting the Warburg effect: historical dogma versus current understanding”. In: The Journal of Physiology 599.6 (Mar. 2021), pp. 1745–1757. ISSN: 1469-7793. DOI: 10.1113/JP278810. URL: https://onlinelibrary.wiley.com/doi/full/10.1113/JP278810%20https://onlinelibrary.wiley.com/doi/abs/10.1113/JP278810%20https://physoc.onlinelibrary.wiley.com/doi/10.1113/JP278810.

[83] Rakesh K. Tekade and Xiankai Sun. “The Warburg effect and glucose-derived cancer theranostics”. In: Drug Discovery Today 22.11 (Nov. 2017), pp. 1637–1653. ISSN: 1359-6446. DOI: 10.1016/J.DRUDIS.2017.08.003.

[84] Adewale Fadaka et al. “Biology of glucose metabolization in cancer cells”. In: Journal of Oncological Sciences 3.2 (July 2017), pp. 45–51. ISSN: 2452-3364. DOI: 10.1016/J.JONS.2017.06.002.

[85] B. H. Jiang et al. “Hypoxia-inducible factor 1 levels vary exponentially over a physiologically relevant range of O2 tension”. In: American Journal of Physiology-Cell Physiology 271.4 (Oct. 1996), pp. C1172–C1180. ISSN: 0363-6143. DOI: 10.1152/ajpcell.1996.271.4.C1172. URL: http://www.ncbi.nlm.nih.gov/pubmed/8897823%20http://www.physiology.org/doi/10.1152/ajpcell.1996.271.4.C1172.

[86] Rupert Courtnay et al. “Cancer metabolism and the Warburg effect: the role of HIF-1 and PI3K”. In: Molecular biology reports 42.4 (2015), pp. 841–851. ISSN: 15734978. DOI: 10.1007/s11033-015-3858-x.

[87] Patrick H. Maxwell, Christopher W. Pugh, and Peter J. Ratcliffe. “Activation of the HIF pathway in cancer”. In: Current Opinion in Genetics and Development 11.3 (2001), pp. 293–299. ISSN: 0959437X. DOI: 10.1016/S0959-437X(00)00193-3.

[88] Huasheng Lu, Robert A. Forbes, and Ajay Verma. “Hypoxia-inducible factor 1 activation by aerobic glycolysis implicates the Warburg effect in carcinogenesis”. In: Journal of Biological Chemistry 277.26 (2002), pp. 23111–23115. ISSN: 00219258. DOI: 10.1074/jbc.M202487200.

[89] Robert Weinberg. “The Biology of Cancer, Second Edition (Google eBook)”. In: (2013), p. 960. URL: https://books.google.com/books/about/The_Biology_of_Cancer.html?id=MzMmAgAAQBAJ.

[90] Huaping Liu et al. “Combining 2-deoxy-D-glucose with fenofibrate leads to tumor cell death mediated by simultaneous induction of energy and ER stress”. In: Oncotarget 7.24 (June 2016), pp. 36461–36473. ISSN: 1949-2553. DOI: 10.18632/oncotarget.9263. URL: http://www.oncotarget.com/fulltext/9263.

[91] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian Optimization of Machine Learning Algorithms. Tech. rep. 2012. URL: http://papers.nips.cc/paper/4522-practical-bayesian-optimization-of-machine-learning-algorithms.pdf.

[92] Neil M Ferguson et al. “Impact of non-pharmaceutical interventions (NPIs) to reduce COVID- 19 mortality and healthcare demand”. In: March (2020).

[93] Oliver M. Cliff et al. “Investigating spatiotemporal dynamics and synchrony of influenza epidemics in Australia: An agent-based modelling approach”. In: Simulation Modelling Practice and Theory 87 (2018), pp. 412–431. ISSN: 1569-190X. DOI: https://doi.org/10.1016/j.simpat.2018.07.005. URL: https://www.sciencedirect.com/science/article/pii/S1569190X18300972.

[94] Boon Som Ong et al. “An Individual-Based Model of Influenza in Nosocomial Environments”. In: Computational Science – ICCS 2008. Ed. by Marian Bubak et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 590–599. ISBN: 978-3-540-69384-0.

[95] Osman Y. Özaltın, Özden O. Dalgıç, and Fatih S. Erenay. "Optimal distribution of the influenza vaccine". In: Proceedings of the Winter Simulation Conference 2014. 2014, pp. 1411–1420. DOI: 10.1109/WSC.2014.7019995.

[96] Simon Alderton et al. "A Multi-Host Agent-Based Model for a Zoonotic, Vector-Borne Disease. A Case Study on Trypanosomiasis in Eastern Province, Zambia". In: (2016). DOI: 10.1371/journal.pntd.0005252.

[97] Jiasheng Wang et al. "Use of GIS and agent-based modeling to simulate the spread of influenza". In: (2010), pp. 1–6. DOI: 10.1109/GEOINFORMATICS.2010.5567658.

[98] Loganathan Ponnambalam et al. "A multi-agent model for adaptive vaccination during infectious disease outbreaks". In: 2016 International Conference on Computing Technologies and Intelligent Data Engineering (ICCTIDE'16). 2016, pp. 1–5. DOI: 10.1109/ICCTIDE.2016.7725354.

[99] Shuangyun Peng et al. "A simulation study of H1N1 space-time epidemic based on agent-based modeling". In: 2010 18th International Conference on Geoinformatics. 2010, pp. 1–4. DOI: 10.1109/GEOINFORMATICS.2010.5567699.

[100] Khaled M. Khalil et al. "An agent-based modeling for pandemic influenza in Egypt". In: 2010 The 7th International Conference on Informatics and Systems (INFOS). 2010, pp. 1–7.

[101] Shan Mei et al. "Simulating city-level airborne infectious diseases". In: Computers, Environment and Urban Systems 51 (2015), pp. 97–105. ISSN: 0198-9715. DOI: https://doi.org/10.1016/j.compenvurbsys.2014.12.002. URL: https://www.sciencedirect.com/science/article/pii/S0198971514001367.

[102] Dionne M. Aleman, Theodorus G. Wibisono, and Brian Schwartz. "Accounting for individual behaviors in a pandemic disease spread model". In: Proceedings of the 2009 Winter Simulation Conference (WSC). 2009, pp. 1977–1985. DOI: 10.1109/WSC.2009.5429727.

[103] Arsalan Paleshi et al. "Simulation of mitigation strategies for a pandemic influenza". In: Proceedings of the 2011 Winter Simulation Conference (WSC). 2011, pp. 1340–1348. DOI: 10.1109/WSC.2011.6147855.

[104] M M Saito et al. "Enhancement of Collective Immunity in Tokyo Metropolitan Area by Selective Vaccination against an Emerging Influenza Pandemic". In:

PLoS ONE 8.9 (2013), p. 72866. DOI: 10.1371/journal.pone.0072866. URL: www.plosone.org.

[105] Narjes Shojaati et al. “MRSA Transmission in a Personal Care Home Facility: A Spatially Explicit Agent Based Modeling Approach”. In: 2017 IEEE International Conference on Healthcare Informatics (ICHI). 2017, pp. 368–373. DOI: 10.1109/ICHI.2017.36.

[106] Parastu Kasaie et al. “Projecting the impact of implementing pre-exposure prophylaxis for HIV among men who have sex with men in Baltimore city”. In: 2015 Winter Simulation Conference (WSC). 2015, pp. 3140–3141. DOI: 10.1109/WSC.2015.7408439.

[107] Daniel J. Escudero et al. “Acute HIV infection transmission among people who inject drugs in a mature epidemic setting”. In: AIDS 30.16 (Oct. 2016), pp. 2537–2544. ISSN: 14735571. DOI: 10.1097/QAD.0000000000001218. URL: https://pubmed.ncbi.nlm.nih.gov/27490641/.

[108] Frederik Graw, Thomas Leitner, and Ruy M. Ribeiro. “Agent-based and phylogenetic analyses reveal how HIV-1 moves between risk groups: Injecting drug users sustain the heterosexual epidemic in Latvia”. In: Epidemics 4.2 (June 2012), pp. 104–116. ISSN: 17554365. DOI: 10.1016/j.epidem.2012.04.002. URL: https://pubmed.ncbi.nlm.nih.gov/22664069/.

[109] Qingxiong Cui et al. “Exploring HIV/AIDS epidemic complex network of IDU using ABM and GIS”. In: 2009 Chinese Control and Decision Conference. 2009, pp. 1090–1095. DOI: 10.1109/CCDC.2009.5191828.

[110] Cliff C. Kerr et al. “Covasim: An agent-based model of COVID-19 dynamics and interventions”. In: PLOS Computational Biology 17.7 (July 2021), e1009149. ISSN: 1553-7358. DOI: 10.1371/JOURNAL.PCBI.1009149. URL: https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1009149.

[111] Petrônio C.L. Silva et al. “COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions”. In: Chaos, Solitons & Fractals 139 (Oct. 2020), p. 110088. ISSN: 0960-0779. DOI: 10.1016/J.CHAOS.2020.110088.

[112] Neil M Ferguson et al. “Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand”. In: 20 (Mar. 2020). DOI: 10.25561/77482. URL: http://spiral.imperial.ac.uk/handle/10044/1/77482.

[113] Jonatan Gomez et al. “INFEKTA—An agent-based model for transmission of infectious diseases: The COVID-19 case in Bogotá, Colombia”. In: PLOS ONE 16.2 (Feb. 2021), e0245787. ISSN: 1932-6203. DOI: 10.1371/JOURNAL.PONE.0245787. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0245787.

[114] Takeshi Kano et al. “An agent-based model of the interrelation between the COVID-19 outbreak and economic activities”. In: Proceedings of the Royal Society A 477.2245 (Jan. 2021). ISSN: 14712946. DOI: 10.1098/RSPA.2020.0604. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rspa.2020.0604.

[115] Jonatan Almagor and Stefano Picascia. “Exploring the effectiveness of a COVID-19 contact tracing app using an agent-based model”. In: Scientific Reports 2020 10:1 10.1 (Dec. 2020), pp. 1–11. ISSN: 2045-2322. DOI: 10.1038/s41598-020-79000-y. URL: https://www.nature.com/articles/s41598-020-79000-y.

[116] Neil M. Ferguson et al. “Strategies for containing an emerging influenza pandemic in Southeast Asia”. In: Nature 437.7056 (2005), pp. 209–214. ISSN: 00280836. DOI: 10.1038/nature04017.

[117] Deaths in Manchester | Coronavirus in the UK. URL: https://coronavirus.data.gov.uk/details/deaths?areaType=ltla&areaName=Manchester.

[118] Myrianthi Hadjicharalambous, Peter A. Wijeratne, and Vasileios Vavourakis. “From tumour perfusion to drug delivery and clinical translation of in silico cancer models”. In: Methods 185 (Jan. 2021), pp. 82–93. ISSN: 1046-2023. DOI: 10.1016/J.YMETH.2020.02.010.

[119] Rakesh Gupta and Beena Rai. “In-silico design of nanoparticles for transdermal drug delivery application”. In: Nanoscale 10.10 (Mar. 2018), pp. 4940–4951. ISSN: 20403372. DOI: 10.1039/C7NR07898F. URL: https://pubs.rsc.org/en/content/articlehtml/2018/nr/c7nr07898f%20https://pubs.rsc.org/en/content/articlelanding/2018/nr/c7nr07898f.

[120] Gavin Fullstone et al. “Modelling the transport of nanoparticles under blood flow using an agent-based approach”. In: Scientific Reports 5 (2015), pp. 1–13. ISSN: 20452322. DOI: 10.1038/srep10649. URL: http://dx.doi.org/10.1038/srep10649.

[121] Wenbo Zhan, Moath Alamer, and Xiao Yun Xu. “Computational modelling of drug delivery to solid tumour: Understanding the interplay between chemotherapeutics and biological system for optimised delivery systems”. In: Advanced Drug Delivery Reviews 132 (July 2018), pp. 81–103. ISSN: 0169409X. DOI: 10.1016/j.addr.2018.07.013. URL: https://linkinghub.elsevier.com/retrieve/pii/S0169409X18301790.

[122] W F Mueller-Klieser and R M Sutherlandt. OXYGEN TENSIONS IN MULTICELL SPHEROIDS OF TWO CELL LINES. Tech. rep. 1982, p. 256. URL: https://www.nature.com/articles/bjc198241.pdf.

[123] Samuel Bernard and Hanspeter Herzel. “Why do cells cycle with a 24 hour period?” In: Genome informatics. International Conference on Genome Informatics 17.1 (2006), pp. 72–79. ISSN: 09199454. DOI: 10.11234/gi1990.17.72.

[124] Mark E. Daly et al. “Acute effects on insulin sensitivity and diurnal metabolic profiles of a high-sucrose compared with a high-starch diet”. In: American Journal of Clinical Nutrition 67.6 (1998), pp. 1186–1196. ISSN: 00029165. DOI: 10.1093/ajcn/67.6.1186.

[125] F. Kallinowski et al. “Glucose Uptake, Lactate Release, Ketone Body Turnover, Metabolic Micromilieu, and pH Distributions in Human Breast Cancer Xenografts in Nude Rats”. In: Cancer Research 48 (1988), pp. 7264–7272. ISSN: 15387445.

[126] S. R. McKeown. “Defining normoxia, physoxia and hypoxia in tumours - Implications for treatment response”. In: British Journal of Radiology 87.1035 (2014). ISSN: 00071285. DOI: 10.1259/bjr.20130676. URL: www.coylab.com.

[127] Sandeep Sharma and Deepa Rawat. Partial Pressure Of Oxygen (PO2). 2018. URL: http://www.ncbi.nlm.nih.gov/pubmed/29630271.

Appendices

Parameter	Value
κ	0.75
D	0.00035
α	0.6
x	0.38
ε_1	0.45
ε_2	0.45
p	0.34
β	0.05
Γ	0.1
h	0.005
n	0.1
Ψ	0.75

Table 1: Summary of model parameters kept constant and not subject to grid search.

Parameter	Description	Assignment Type	Details	Unit
VEGF Diffusivity in Tissue	Rate at which VEGF diffuses in tissues	random search	mean: 1, stDev: 0.01	positions/second
VEGF Secretion Rate	Rate at which VEGF is secreted by cancer cells	grid search	min: 1, max 9, step size: 2	a.u.
Length of Cell Life-Cycle	Number of epochs between cellular divisions	constant	12	epochs
Endothelial Division Delay	Time taken between division of endothelial cells	constant	12	epochs
Epochs	Number of epochs the simulation will run for	constant	300	2 hours/epoch
Hypoxic Threshold	Threshold below which cells start secreting VEGF	constant	20	mmHg
Maximum VEGF Concentration per Position	Maximum amount of VEGF that can be contained at a position	constant	10	a.u.
Minimum Oxygen Concentration for Cancer Cell Survival	Oxygen concentration below which cancer cells die	constant	10	mmHg
Minimum VEGF Concentration for Angiogenesis	Minimum VEGF required for endothelial cells to respond and beginnig the process of angiogenesis	grid search	min: 1, max 9, step size: 2	a.u.
Oxygen Diffusivity in Tissue	Rate at which oxygen diffuses in tissues	random search	mean: 1, stDev: 0.01	positions/second
Oxygen Secretion Rate	Rate at which oxygen is secreted by blood vessels	constant	190	mmHg
Oxygen Uptake Rate	Rate at which oxygen is absorbed by cancer cells	constant	20	mmHg
Probability of Cell Progressing into Synthesis	Probability of a cell progressing from G1 into Synthesis	grid search	min: 0.2, max 0.8, step size: 0.2	probability

Table 2: Summary of additional model parameters for our 2D v 3D investigation. Only those which are resolution-independent are provided here.

Warburg Investigation

Parameter	Description
Base Oxygen Emission Rate	Describes the rate at which oxygen flows out of blood vessels, also oxygen source value for diffusion.
Base Oxygen Metabolic Rate	Base sink value of oxygen for diffusion purposes changes as HIF expression increases.
Endothelial Division Delay	Number of Epochs an Endothelial Cell needs to exist for before it can divide.
Glucose Diffusivity	Glucose Diffusivity.
Glucose Secretion Rate	Source value of glucose for diffusion purposes.
Glucose Uptake Rate Non-Warburg	Glucose uptake rate (sink value for diffusion purposes) of non-Warburg cells.
Hypoxic Threshold	Oxygen Concentration at which HIF expression starts increasing.
Maximum Glucose Uptake Rate	Glucose uptake rate (sink value for diffusion purposes) of Warburg cells.
Maximumimum HIF Expression Rate	Maximum HIF Expression Rate of cancer cells.
Minimum Glucose Non Warburg	Minimum glucose concentration required by non-Warburg cells.
Minimum Glucose Warburg	Minimum glucose concentration required by Warburg cells.
Minimumimum HIF Expression Rate	Minimum HIF Expression Rate of cancer cells.
Minimumimum Oxygen Concentration	Minimum oxygen concentration needed for the survival of cancer cells.
Minimumimum Probability of Progression into Synthesis	Base probability of a cell progressing from G1 into Synthesis. (Can increase as HIF expression increases)
Minimumimum VEGF Concentration	Minimum VEGF concentration needed for endothelial cells to respond and start angiogenesis.
Num Cancer Cells	Initial number of cancer cells.
Oxygen Diffusivity	Oxygen diffusivity.
Probability of Warburg Switch	Probability of a cancer cell acquiring the Warburg phenotype.
Total Cell Life-Cycle Length	Total length of cell life-cycle (G1+S+G2+M).
Ultra-Hypoxic Threshold	Oxygen Concentration at which HIF expression is highest.
Warburg Hypoxic Threshold	Hypoxic Threshold of Warburg Cells.

Table 3: Summary of model parameters and their descriptions for the Warburg Investigation.

Parameter	Correlation	Correlation Type	Absolute Correlation	Correlation Label
Base Oxygen Emission Rate	0.42	Positive	0.42	Medium
Base Oxygen Metabolic Rate	-0.31	Negative	0.31	Medium
Endothelial Division Delay	-0.23	Negative	0.23	Low
Glucose Diffusivity	0.77	Positive	0.77	High
Glucose Secretion Rate	0.87	Positive	0.87	High
Hypoxic Threshold	-0.09	Negative	0.09	Low
Maximum Glucose Uptake Rate	-0.85	Negative	0.85	High
Minimum Glucose Non Warburg	-0.92	Negative	0.92	High
Minimum Glucose Warburg	-0.74	Negative	0.74	High
Minimum VEGF Concentration	-0.79	Negative	0.79	High
Minimumimum HIF Expression Rate	0.38	Positive	0.38	Medium
Minimumimum Oxygen Concentration	-0.12	Negative	0.12	Low
Minimumimum Probability of Progression into Synthesis	0.63	Positive	0.63	High
Num Cancer Cells	0.98	Positive	0.98	High
Oxygen Diffusivity	0.70	Positive	0.70	High
Probability of Warburg Switch	0.39	Positive	0.39	Medium
Ultra-Hypoxic Threshold	0.34	Positive	0.34	Medium

Table 4: Summary of Pearson Correlations between each parameter and model performance from initial exploration. Single parameters were swept while others were kept fixed. The following parameters were omitted: Cell Life-Cycle Length (Known), Maximum and Minimum HIF Expression Rates (Model Artifact). Low correlation is established as less than 0.2, high as >0.6, medium as in-between values.

Parameter	Confidence
Base Oxygen Emission Rate	Green
Base Oxygen Metabolic Rate	Green
Endothelial Division Delay	Green
Glucose Diffusivity	Yellow
Glucose Secretion Rate	Green
Hypoxic Threshold	Green
Max Glucose Uptake Rate	Amber
Maximum HIF Expression Rate	Green
Min Glucose Non Warburg	Yellow
Min Glucose Warburg	Amber
Min P Synthesis	Yellow
Minimum HIF Expression Rate	Green
Minimum Oxygen Concentration	Yellow
Minimum VEGF Concentration	Amber
Num Cancer Cells	Amber
Oxygen Diffusivity	Yellow
Probability of Warburg Switch	Yellow
Total Cell Life-Cycle Length	Green
Ultra-Hypoxic Threshold	Green

Table 5: Summary of original parameter classification as green, yellow and amber based on the presence of values from the literature, the confidence of those parameters and initial correlation with model performance.

Parameter	Confidence	Changed
Base Oxygen Emission Rate	Green	No
Base Oxygen Metabolic Rate	Green	No
Endothelial Division Delay	Green	No
Glucose Diffusivity	Green	No
Glucose Secretion Rate	Green	No
Hypoxic Threshold	Green	No
Max Glucose Uptake Rate	Green	No
Maximum HIF Expression Rate	Green	No
Min Glucose Non Warburg	Green	No
Min Glucose Warburg	Green	No
Min P Synthesis	Amber	No
Minimum HIF Expression Rate	Green	Yes (Original Amber)
Minimum Oxygen Concentration	Green	No
Minimum VEGF Concentration	Amber	No
Num Cancer Cells	Yellow	No
Oxygen Diffusivity	Yellow	No
Probability of Warburg Switch	Green	Yes (Original Yellow)
Total Cell Life-Cycle Length	Green	Yes (Original Yellow)
Ultra-Hypoxic Threshold	Yellow	No

Table 6: Summary of new confidence values after further investigation via decision tree unravelling and grid search.

Parameter	Value	Units
Base Oxygen Emission Rate	40	mmHg/time unit
Base Oxygen Metabolic Rate	25	mmHg/time unit
Endothelial Division Delay	12	epochs
Glucose Diffusivity	Determine by search-space exploration.	centimeters squared/second
Glucose Secretion Rate	190	pMol/time unit
Glucose Uptake Rate Non-Warburg	10	pMol/time unit
Hypoxic Threshold	20	mmHg
Maximum Glucose Uptake Rate	Determine by search-space exploration.	pMol/time unit. Refers to time unit used in solving diffusion.
Maximumimum HIF Expression Rate	16	Arbitrary Units
Minimum Glucose Non Warburg	Determine by search-space exploration.	pMol
Minimum Glucose Warburg	Determine by search-space exploration.	pMol
Minimumimum HIF Expression Rate	1	Arbitrary Units
Minimumimum Oxygen Concentration	12	mmHg
Minimumimum Probability of Progression into Synthesis	Determine by search-space exploration.	Probability value
Minimumimum VEGF Concentration	Determine by search-space exploration.	Arbitrary Units
Num Cancer Cells	Determine by search-space exploration.	Scalar value
Oxygen Diffusivity	Determine by search-space exploration.	centimeters squared/second
Probability of Warburg Switch	Determine by search-space exploration.	Probability value
Total Cell Life-Cycle Length	12	epochs
Ultra-Hypoxic Threshold	3	mmHg
Warburg Hypoxic Threshold	75	mmHg

Table 7: Summary of parameter values. Where values were assigned by grid search or random search this was also indicated.

Parameter	Description	Value	Units
Base Oxygen Emission Rate	Describes the rate at which oxygen flows out of blood vessles, also oxygen source value for the purpose of diffusion.	40[70]	mmHg/tim
Base Oxygen Metabolic Rate	Base sink value of oxygen for diffusion purposes, changes as HIF expression increases.	25 [122]	mmHg/tim
Endothelial Division Delay	Number of Epochs an Endothelial Cell needs to exist for before it can divide.	12 [123]	epochs
Glucose Secretion Rate	Source value of glucose for diffusion purposes.	190 [124]	pMol/time
Glucose Uptake Rate Non-Warburg	Glucose uptake rate (sink value for diffusion purposes) of non-Warburg cells.	10 [125]	pMol/time
Hypoxic Threshold	Oxygen Concentration at which HIF expression starts increasing.	20 [126]	mmHg
Maximum HIF Expression Rate	Maximum HIF Expression Rate of cancer cells.	16[45]	Arbitrary
Minimum HIF Expression Rate	Minimum HIF Expression Rate of cancer cells.	1[45]	Arbitrary
Total Cell Life-Cycle Length	Total length of cell life-cycle (G1+S+G2+M).	12 [70]	epochs
Ultra-Hypoxic Threshold	Oxygen Concentration at which HIF expression is highest.	3 [88]	mmHg
Warburg Hypoxic Threshold	Hypoxic Threshold of Warburg Cells.	75 [87]	mmHg
Minimum Oxygen Concentration	Minimum Oxygen Required by Non-Warburg Cells	12 [127]	mmHg

Table 8: Summary of parameters classified as green. mmHg/time unit refers to the change in partial pressure at the position over the time unit considered for diffusion solving owed to the action of the cell acting as a source or sink of the soluble. Resources cited refer to the original papers consulted to obtain baseline values for parameters. In instances, these have been further adjusted via preliminary exploration by means of parameter sweeping and considerations inherent to our model's structure and spatial discretisation.

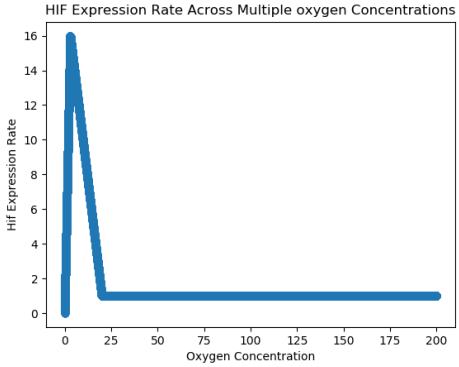

(a) HIF expression rates (a.u.) as a function of oxygen concentrations (mmHg) for non-Warburg cells. Expression begins as cells experience hypoxia and rapidly increase up to peaking under extreme hypoxia. As oxygen concentrations fall below those required for the cell's survival, they rapidly fall again.

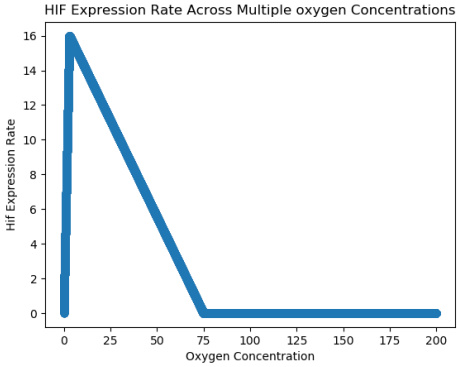

(b) HIF expression rates (a.u.) as a function of oxygen concentrations (mmHg) for Warburg cells. Compared to non-Warburg cells, expression of HIF begins even in the presence of higher oxygen concentrations.

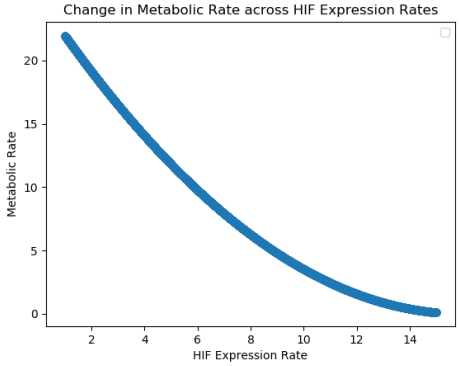

(c) Changes in metabolic rate (mmHg) as a function of HIF expression rates. Oxygen uptake decreases from the base level to 0 as HIF expression increases towards maximum levels.

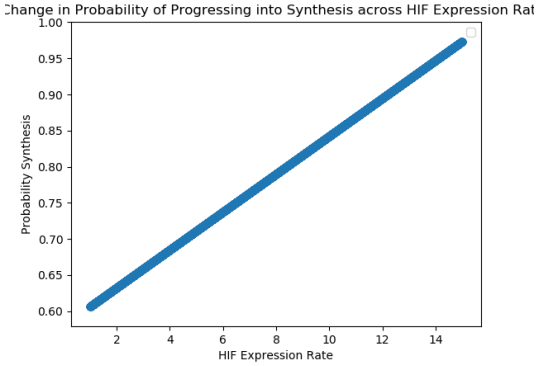

(d) Changes in probability of a cell progressing into synthesis as a function of HIF expression rates. As this is a parameter that is investigated in parameter tuning a sample representation is given for a base probability of 60%.

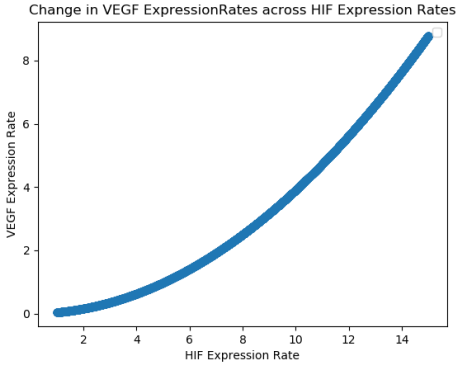

(e) Changes in VEGF secretion rate (a.u.) as a function of HIF expression rates. VEGF secretion increases as HIF expression increases to stimulate further blood vessel growth and development as a response to hypoxia.

Figure 1: Summary of HIF expression rates as a response to different oxygen concentrations for Warburg and non-Warburg cells as well as the relations between HIF and HIF-mediated phenomena.

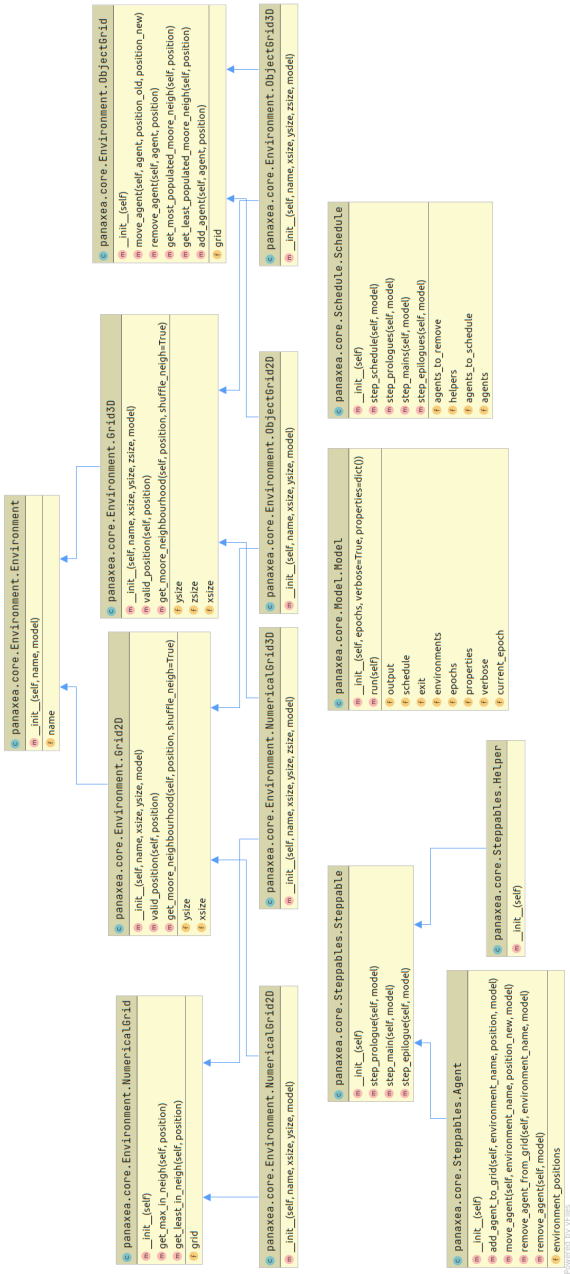

Figure 2: Class Diagram for PanaXea. Note environments may be either 2D or 3D, and in each variant may hold objects (Ie: Agents) or numerical values. Steppables include both agents and helpers.

3D		
Item	Value	Unit
Volume of Cancer Cell	2,000.00	um^3
Cancer Cells per Agent	300.00	cancer cells/agent
Volume of Agent	600,000.00	um^3 /agent
Volume of Grid Position	300,000,000.00	um^3
Max Agents per Grid Position	500.00	agents
2D		
Item	Value	Unit
Area of Cancer Cell	158.74	um^2
Cancer Cells per Agent	300.00	cancer cells/agent
Area of an Agent	47,622.03	um^2 /agent
Area of Grid Position	448,140.47	um^2
Max Agents per Grid Position	9.00	agents

Table 9: Summary of value translations from 3D to 2D. Mapping of volume to the area is done by deriving the area of the base from a cube. An agent's area is the product of the volume/area of an agent by the cancer cells per agent. The volume of a grid position is arbitrary but consistent across all simulations. Max agents per grid position is the ratio between the volume/area of the grid position and the volume/area of an agent.