



UNIVERSITY COLLEGE LONDON

Developing Novel Non-Invasive Measures of Tumour Physiology.

Candidate Name:

Benjamin Oliver Richard HIPWELL

Supervisors:

Prof. Simon WALKER-SAMUEL

Dr. Douglas PENDSE

Dr. Thomas ROBERTS

A dissertation submitted in partial fulfilment of
the requirements for the degree of

Doctor of Philosophy

of

University College London.



August 2, 2020

Declaration

I, Benjamin Oliver Richard Hipwell, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

Cancer is the leading cause of death globally; ahead of heart disease, stroke and chronic obstructive pulmonary disease. There were 8.2 million attributed deaths in 2012 [1], and the number of new cases is expected to increase by 70% in the next two decades [2]. Development of novel new therapies has consequently become a multi-billion dollar industry. However, the development of new cancer therapies is limited by our ability to accurately quantify their effects. This thesis focuses on the development of novel non-invasive biomarkers for assessing changes in tumour microstructure.

In chapter 3, the capabilities of an in-house diffusion MRI technique known as Vascular, Extracellular and Restricted Diffusion for Cytometry in Tumours (VERDICT) are investigated. The technique was used to detect changes in tumour microstructure caused by a) tissue fixation and b) administration of temozolomide therapy.

In chapter 4, the development of a Monte Carlo tissue diffusion simulation framework is described. The simulation framework is then applied as a tool for validating diffusion MRI models, including VERDICT.

Chapter 5 presents an exploration of the potential applications of machine-learning based approaches within the field of diffusion MRI. In a preliminary study, a neural network is trained on synthetic diffusion MRI data, and applied to real-world *in-vivo* data to try and extract microstructural tissue features without the need for explicit model fitting.

The overall aim of this thesis was to assist in the development and validation of advanced diffusion MRI modelling techniques, and explore the future potential of synthetic data and machine-learning models in the extraction of new cancer biomarkers.

Impact Statement

The findings of this thesis provide further demonstration of the capabilities of advanced compartmental diffusion MRI models such as VERDICT. Further, the development of a sophisticated Monte Carlo simulation framework provides a powerful tool for validating complex diffusion models applied in tissues where the effect of perfusion in the microvascular network cannot be neglected. The exploratory research into using synthetic data to train machine learning algorithms demonstrates the huge potential for machine-learning based approaches within the field of diffusion MRI modelling.

The work presented in the first results chapter in this thesis demonstrates that compartmental models such as VERDICT are able to detect changes in the microstructure of tumour tissue caused by therapeutic response at an earlier time point than conventional methods.

In the second results chapter, for the first time, the development of a Monte Carlo simulation framework that incorporates the effects of tissue diffusion and microvascular blood flow is presented. The blood flow components have been simulated on vessel substrates extracted from real tumour networks. The simulation framework developed in this chapter has a wide range of applications. Foremost, it provides further insight into the behaviour of the complex mathematical models used to interrogate diffusion MRI data. Furthermore, it allows the relationship between the model parameters and the true microstructural tissue features to be investigated. This is crucial if such models are to be successfully translated into the clinical setting.

In the final results chapter, an additional application of the simulation framework is presented. Synthetic diffusion MRI data is used as the training dataset for a neural network model, that aims to ‘learn’ the direct relationship between the raw MRI signal and the microstructural tissue features, without the need for explicit model fitting. This work demonstrates the huge untapped potential of machine-learning models in the field of diffusion MRI, and provides a foundation for future research.

Acknowledgements

There are countless people without whom I am absolutely certain I would have never even come close to *starting* a PhD, let alone finishing it. To name just a few:

Firstly, I would like to express my sincerest gratitude to my primary supervisor Prof. Simon Walker-Samuel, for giving me the opportunity to work in CABI to begin with, for being an endless source of inspiration and new ideas, for his continued motivation and guidance and for coining the catchphrase ‘what would be *really* cool is...’. I would like to thank Dr. Thomas Roberts, for his invaluable friendship and support, for helping to steer me through the difficult early stages of my PhD and for teaching me everything I know about practical MRI physics, MATLAB and table tennis. I would also like to thank Dr. Daniel Stuckey and Dr. Tammy Kalber, for their advise (both scientific and general words of wisdom), for their kindness and moral support, and their generosity with their own time. I also wish to thank Prof. Mark Lythgoe, for his leadership within CABI and for cultivating a working environment in that is both friendly, laid back, and yet highly motivating. I would also like to thank Dr. Bernard Siow for his technical support, Valerie Taylor for her experimental know-how, Morium for her OPT wizzardry and assistance, and Dr. Arun Niranjana for his coding help.

I also wish to thank the staff in the UCL Medical Imaging CDT for accepting me onto the PhD program in the first instance. I would like to thank Prof. Sebastien Ourselin and Dr Gary Zhang for their excellent leadership of the CDT program.

I would like to thank my incredible friends inside and outside CABI. Dr. John Connell for the endless games of table tennis, and always being there for a laugh and a pint. Dr Laurence Jackson for the post-conference adventures in Vietnam. George Dwyer, for being a legendary flatmate and for all the ‘Wing Wednesdays’. Ollie Youle, for getting me through my Master’s degree, and all the ridiculous situations we’ve been through. All of the fantastic people in CABI, you know who you are, for making the day-to-day so much fun.

Finally, I wish to thank my family and loved ones to whom I owe so much. To my wonderful mum and sister, I wouldn’t have got this far without you. Your love and support has kept

me going through the really hard times. To my grandparents, for always believing in me and cheering me on. I can't thank you both enough for your unconditional love and support.

To my grandfather, I wish you could have been here to see me cross the finish line. Knowing that you were proud of me gave me the fuel I needed for the final push. This is for you old chap.

Contents

Declaration	1
Abstract	2
List of Figures	13
List of Tables	23
List of abbreviations	24
1 Motivation and Thesis Overview	26
1.1 Motivation	26
1.2 Thesis Overview	30
2 Background	32
2.1 Cancer	32
2.1.1 History and Brief Overview	32
2.1.2 Tumour Microstructure	33
2.1.3 Quantitation of Tumour Microstructure	34

2.2	Fundamental Physics of MRI	35
2.2.1	Spin and Nuclear Magnetism	35
2.2.2	Relaxation and RF Excitation	41
2.2.3	Signal Reception and Spatial Encoding	45
2.2.4	Basic Pulse Sequences	48
2.2.5	Image Formation	53
2.3	Diffusion-Weighted MRI	55
2.3.1	Basic Overview	55
2.3.2	Acquisition	58
2.3.3	Analysis	65
2.3.3.1	Apparent Diffusion Coefficient	65
2.3.3.2	Diffusion Tensor Imaging	69
2.3.3.3	Intra-Voxel Incoherent Motion	71
2.3.3.4	Multi-Compartment Models	74
2.3.4	Validation and Microstructural Modelling	82
2.3.4.1	Histological Validation	82
2.3.4.2	Physical Phantoms	84
2.3.4.3	Numerical Phantoms	86
2.4	Summary	89
3	Imaging Microstructural Changes using VERDICT MRI	90
3.1	Motivation	90

3.2	Introduction	91
3.2.1	Tissue Fixation	92
3.2.1.1	Formaldehyde Fixation	93
3.2.2	Temozolomide Therapy	96
3.3	Study Outlines	96
3.3.1	Tissue Fixation Study	96
3.3.2	Temozolomide Study	97
3.4	Methods	98
3.4.1	Tissue Fixation Study	98
3.4.1.1	Animal Models	98
3.4.1.2	Tumour Development	99
3.4.1.3	Anaesthesia and Physiological Monitoring	99
3.4.1.4	Perfusion/Immersion Fixation	100
3.4.1.5	MRI Protocol	100
3.4.1.6	Model-Fitting	103
3.4.1.7	Statistical Analysis	105
3.4.2	Temozolomide Study	105
3.4.2.1	Animal Models	105
3.4.2.2	Tumour Development	106
3.4.2.3	Anaesthesia and Physiological Monitoring	106
3.4.2.4	Study Design	107
3.4.2.5	MRI Protocol	108

3.4.2.6	Data Analysis	109
3.4.2.7	Statistical Tests	110
3.5	Results	110
3.5.1	Tissue Fixation Study	110
3.5.1.1	Cell Radius	110
3.5.1.2	Intracellular Volume Fraction	111
3.5.1.3	VERDICT Diffusivity	112
3.5.1.4	ADC	113
3.5.2	Temozolomide Study	114
3.5.2.1	VERDICT 'BallSphereStick'	114
3.5.2.2	VERDICT 'BallSphereAstrosticks'	116
3.5.2.3	ADC	117
3.5.2.4	Volumetric Measurements	118
3.5.2.5	VERDICT Model Performance Comparison	119
3.6	Discussion	122
3.6.1	Tissue Fixation Study	122
3.6.2	Temozolomide Study	124
3.7	Chapter Summary	127
4	Development and Application of a Monte Carlo Model of Tissue Diffusion and Flow	129
4.1	Motivation	129
4.2	Introduction	130

4.3	Methods	132
4.3.1	Animal Models	132
4.3.2	Data Acquisition	133
4.3.3	Monte Carlo Simulation Framework	134
4.3.3.1	Substrate Generation	134
4.3.3.2	Simulation Dynamics	149
4.3.3.3	Diffusion Data Synthesis	154
4.3.4	Model Fitting	156
4.3.5	Verification	157
4.3.6	Application to Validation of Diffusion Models	158
4.3.6.1	VERDICT Validation	158
4.3.6.2	ADC and IVIM	160
4.3.6.3	Effect of Blood Flow	161
4.3.6.4	Noise Sensitivity	162
4.4	Results	163
4.4.1	Validation of Monte Carlo Framework	163
4.4.2	VERDICT Validation	166
4.4.2.1	Intracellular Walker (Volume) Fraction	166
4.4.2.2	Cell Radius	168
4.4.2.3	Vascular Walker (Volume) Fraction	170
4.4.2.4	ADC and IVIM models	171
4.4.3	Effect of Blood Flow Velocity	175

4.4.4	Noise Sensitivity	177
4.5	Discussion	180
4.5.1	Validation of Monte Carlo Framework	180
4.5.2	VERDICT Validation	182
4.5.3	Other Diffusion Models	185
4.5.4	Effect of Blood Flow Velocity	186
4.5.5	Noise Sensitivity	187
4.6	Chapter Summary	189
5	Application of Machine Learning to Real and Synthetic MR Data	193
5.1	Brief Overview of Machine Learning	193
5.2	Artificial Neural Networks	195
5.2.1	Brief Historical Background	195
5.2.2	Sigmoid Neurons	200
5.2.3	Neural Network Structure	202
5.2.4	How Neural Networks ‘Learn’	205
5.2.4.1	Cost Function	206
5.2.4.2	Gradient Descent	208
5.2.4.3	Backpropagation	214
5.2.5	Application to Medical Imaging	219
5.3	Study Overview	222
5.3.1	Methods	222

5.3.1.1	In-Vivo Data	222
5.3.1.2	Synthetic Data	223
5.3.1.3	Artificial Neural Network Development	224
5.3.1.4	Applying the Trained Network to In-Vivo Data	227
5.4	Results	228
5.4.1	Intracellular Volume Fraction	228
5.4.2	Extracellular Volume Fraction	230
5.4.3	Vascular Volume Fraction	232
5.4.4	Cell Radius	234
5.4.5	Blood Flow Parameters	235
5.5	Discussion	236
5.6	Chapter Summary	238
6	Final Summary and Conclusions	241
	Bibliography	246
	Appendix A Perfusion-Fixation Protocol	280

List of Figures

2.1	A schematic showing nuclei (a) without the influence of an external magnetic field and (b) with an external magnetic field. The nuclei are here classically represented as individual magnetic dipoles that align parallel or anti-parallel to the field. The energy levels associated with the two configurations is shown in (c).	37
2.2	The distribution of spins within a macroscopic sample, with a spherical distribution of directions, all precess around the principle magnetic field direction \mathbf{B} at angular velocity ω .	39
2.3	Diagrams showing the orientations of the magnetic fields viewed from the rotating reference frame. (a) shows the general case where the longitudinal field B_0 and the B_1 field combine to create the total effective magnetic field B_{total} . (b) shows the case when the frequency of the rotating field ω is equal to the Larmor frequency ω_0 , where the B_1 becomes the total effective field. In both cases the net magnetisation vector M precesses around the total effective field vector. In (b) the magnetisation vector tilts by an angle $\alpha = 90^\circ$ into the xy plane.	43
2.4	Graph showing the characteristic relaxation and recovery of the z-component of the magnetisation following excitation.	44
2.5	A typical Free Induction Decay signal. By Nmr_fid_good_shim.svg: Gyro-Magicianderivative work: Imalipusram This file was derived from Nmr fid good shim.svg; CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=3031223	
2.6	A timing diagram showing a typical gradient echo pulse sequence. The diagram shows the order in which the various RF pulses and gradient waveforms occur during the acquisition. Shown below and labelled is the effect of the aforementioned pulses and gradients on the net magnetisation vector (red) and the individual spins (blue).	49

2.7	A timing diagram showing a typical spin echo pulse sequence. The diagram shows the order in which the various RF pulses and gradient waveforms occur during the acquisition. Shown below and labelled is the effect of the aforementioned pulses and gradients on the net magnetisation vector (red) and the individual spins (blue).	51
2.8	The construction of k-space from detected signal echos. Each digitally sampled echo with a particular phase encode value corresponds to a row of the k-space matrix.	53
2.9	Typical examples of a magnitude and phase image of a human brain. Images taken from http://mriquestions.com/making-an-sw-image.html on 22/05/2018	55
2.10	Illustration of the random thermal motion of a water molecule.	56
2.11	A schematic showing the basic pulse sequence structure of the pulsed-gradient spin-echo acquisition	59
2.12	A schematic showing the effect of the motion sensitising gradients and RF pulses of the PGSE sequence on the phase of a single spin in the xy plane. (a) shows the spin following the application of the 90° RF pulse, tipping the spin into the xy plane. (b) Following the first motion-sensitising gradient, the spin has gained a phase offset ϕ_1 . (c) Following the 180° RF pulse, the phase offset has been inverted. (d) Following the second motion-sensitising gradient, the spin gains a further phase offset ϕ_2 , which acts to cancel out the first phase offset due to the spin phase being flipped by the 180° pulse.	61
2.13	Figure showing (left) $b = 500$ and (right) $b = 1000$ images of the brain. Taken from https://mrImaster.com/characterise%20image%20dwi%20.html , accessed 30/05/2018	65
2.14	Figure showing the ADC map, calculated from the $b500$ and $b1000$ images shown in figure 2.13. The map demonstrates how the ADC value relates to the microstructural properties of the underlying tissue. In this case, the ventricles and Cerebrospinal Fluid (CSF) spaces within the brain appear bright as they have a higher ADC value. These regions are fluid-filled, meaning water molecules are free to diffuse. Conversely, the white and grey matter appears darker, as these regions are densely packed with cells and neurons meaning diffusion is restricted and the ADC value is lower. Image taken from https://mrImaster.com/characterise%20image%20dwi%20.html , accessed 30/05/2018	67

2.15	Schematic showing how a diffusion tensor can be visualised as an ellipsoid, with the principle axes direction/dimensions given by the eigenvalues of the diffusion tensor matrix.	70
2.16	Image showing the orientations of the capillaries. The image is a zoomed-in section of a segmented vessel network from a subcutaneous LS174T tumour model.	72
3.1	Structure of a typical protein	94
3.2	Methylene Bridge	95
3.3	A schematic showing the PGSE sequence used for the VERDICT acquisitions, with the key sequence parameters labelled.	101
3.4	A schematic of the timeline of the Temozolomide therapy study, showing the tumour inoculation day, imaging and therapy timepoints.. . . .	107
3.5	Box & Whisker plots showing the values of the cell radius parameter produced by the VERDICT models fitted to the <i>in-vivo</i> and <i>ex-vivo</i> data. LEFT: 'BallSphereStick' and 'BallSphere' models. RIGHT: 'BallSphereAstrosticks' and 'BallSphere'.	111
3.6	Box & Whisker plots showing the values of the intracellular volume fraction produced by the VERDICT models fitted to the <i>in-vivo</i> and <i>ex-vivo</i> data. LEFT: 'BallSphereStick' and 'BallSphere' models. RIGHT: 'BallSphereAstrosticks' and 'BallSphere'.	112
3.7	Box & Whisker plots showing the values of the diffusivity parameter produced by the VERDICT models fitted to the <i>in-vivo</i> and <i>ex-vivo</i> data. LEFT: 'BallSphereStick' and 'BallSphere' models. RIGHT: 'BallSphereAstrosticks' and 'BallSphere'.	113
3.8	Box & Whisker plots showing the values of the ADC fitted to the <i>in-vivo</i> and <i>ex-vivo</i> data.	114

3.9	The parameter values produced by the VERDICT 'BallSphereStick' model. (a) Intracellular Volume Fraction (b) Extracellular Volume Fraction (c) Vascular Diffusivity (d) Cell Radius. On each plot, the values produced from fitting the model to the control and therapy group data are shown, with error bars corresponding to the standard deviations. Asterisks are used to denote time-points where a significant difference was found between the control and therapy group.	115
3.10	The parameter values produced by the VERDICT 'BallSphereAstrosticks' model. (a) Intracellular Volume Fraction (b) Extracellular Volume Fraction (c) Vascular Diffusivity (d) Cell Radius. On each plot, the values produced from fitting the model to the control and therapy group data are shown, with error bars corresponding to the standard deviations. Asterisks are used to denote time-points where a significant difference was found between the control and therapy group.	116
3.11	The parameter values produced by the ADC model. Asterisks are used to denote time-points where a significant difference was found between the control and therapy group.	117
3.12	Side-by-side plots showing exemplary VERDICT data from one of the therapy cohort animals. The symbols represent the measured data, and the lines represent the fits from the two VERDICT models.	119
3.13	Therapy group: The final value of the maximum likelihood objective function achieved by the VERDICT 'BallSphereStick' (Solid line) and 'Ball-SphereAstrosticks' (Dashed line) models for each imaging time point post therapy, averaged across all study subjects, with standard deviation across the subjects also shown.	120
3.14	Control group: The final value of the maximum likelihood objective function achieved by the VERDICT 'BallSphereStick' (Solid line) and 'Ball-SphereAstrosticks' (Dashed line) models for each imaging time point post therapy for the control dataset, averaged across all study subjects, with standard deviation across the subjects also shown.	121
4.1	Left: H& E stained histology of a giant cell tumour (bone). Right: H & E stained histology of a neurofibroma tumour. By Nephron - Own work, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=14831566	135
4.2	Flow diagram showing the algorithm used to pack spherical cells into the tissue substrate.	137

4.3	An illustration of the process of checking proposed cells for overlap.	139
4.4	An example output from the sphere packing algorithm described above. The parameters used to generate the substrate shown were mean radius $10\mu m$, standard deviation $5\mu m$. A large standard deviation was only used in order to give a clearer illustration of the variation in cell radius that can be achieved using this algorithm.	140
4.5	A random vessel network generated using random seed points joined together using a MST algorithm. The figure shows a network produced using 1000 vertices ($vesseldensity = [1000, 1000, 1000]$, with root vertex at $(0, 0, 0)$).	142
4.6	The resulting simulated vessel network produced from the vessel structure of figure 4.5. The vessel radius in this case was specified as $5\mu m$ throughout the network.	143
4.7	Multiple vessel networks generated within a single simulation substrate, using the methods described in part 1.	144
4.8	Images showing the various processing stages involved in extracting the vessel structures for use in the simulation framework. (a) A surgically resected LS174T human colon adenocarcinoma tumour. The tumour was grown subcutaneously on the right flank of an immunocompromised mouse. (b) The resected tumour following optical clearing using the benzyl alcohol, benzyl benzoate (BABB) protocol [3]. (c) OPT image of the cleared tumour, resolution $5 - 10\mu m$. (d) Segmentation produced by Gaussian-subtraction, Frangi filtering, thresholding and skeletonisation. (e) Schematic representing the conversion of the segmented data into a 3D spatial graph format. (f) Resulting 3D graphical representation of the full tumour vascular network.	145
4.9	Segmented OPT data from an entire subcutaneous tumour, skeletonised and converted into graph format. The vessel network has been overlaid on structural MR data of the subcutaneous tumour on which the OPT imaging was performed.	146
4.10	Illustration of the REANIMATE modelling pipeline. The pipeline includes the complete process for estimating the steady-state blood flow values for complete vessel networks extracted from real tumours [4]	147

4.11	Left: A full simulated tumour vessel network, produced using the REANIMATE framework. Middle: 4 examples of voxel-sized portions cropped from the complete tumour network, two from the middle of the tumour and two from the edge. Right: A complete simulation substrate comprising spherical cells packed around the vessel network from one of the example voxels.	149
4.12	Validation of the Monte Carlo simulation framework performed by generating synthetic diffusion-weighted signal from a substrate comprising parallel cylinders (radius $8\mu m$), then comparing the synthetic diffusion-weighted signal with signal generated using a corresponding analytical model. Also shown is synthetic signal generated using the Monte Carlo simulation tool included in the Camino diffusion toolkit, using a parallel cylinders substrate (radius $8\mu m$).	164
4.13	Validation of the Monte Carlo simulation framework performed by generating synthetic diffusion-weighted signal from a substrate comprising packed spheres (radius $8\mu m$), then comparing the synthetic diffusion-weighted signal with signal generated using a corresponding analytical model. Also shown is synthetic signal generated using the Monte Carlo simulation tool included in the Camino diffusion toolkit, using a spherical mesh substrate (radius $8\mu m$).	165
4.14	Parameter correlation plots for the intracellular walker fraction plotted against fitted model parameters produced using the VERDICT ‘BallSphereStick’ model (top row) and ‘BallSphereAstrosticks’ model (bottom row). The plots show the ground truth values (x-axis) taken from the simulation, plotted against the fitted parameter values produced by the VERDICT models. From left to right: Intracellular walker fraction vs intracellular volume fraction, intracellular walker fraction vs cell radius, intracellular walker fraction vs vascular volume fraction. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The solid black line shows the 1-1 correlation line. The dashed black line is the line of best fit fitted to the average of the four voxels.	166

4.15	Parameter correlation plots for the substrate cell radius plotted against fitted model parameters produced using the VERDICT ‘BallSphereStick’ model (top row) and ‘BallSphereAstrosticks’ model (bottom row). The plots show the ground truth values (x-axis) taken from the simulation, plotted against the fitted parameter values produced by the VERDICT models. From left to right: Substrate cell radius vs fitted cell radius, substrate cell radius vs intracellular volume fraction, substrate cell radius vs vascular volume fraction. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The solid black line shows the 1-1 correlation line. The dashed black line is the line of best fit fitted to the average of the four voxels.	168
4.16	Parameter correlation plots for the substrate vascular walker fraction plotted against fitted model parameters produced using the VERDICT ‘BallSphereStick’ model (top row) and ‘BallSphereAstrosticks’ model (bottom row). The plots show the ground truth values (x-axis) taken from the simulation, plotted against the fitted parameter values produced by the VERDICT models. From left to right: Vascular walker fraction vs fitted vascular volume fraction, vascular walker fraction vs intracellular volume fraction, vascular walker fraction vs fitted cell radius. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The solid black line shows the 1-1 correlation line. The dashed black line is the line of best fit fitted to the average of the four voxels.	170
4.17	Parameter correlation plots for the ADC model. From left to right: Cell Radius vs ADC, Intracellular Walker Fraction vs ADC, Vascular Walker Fraction vs ADC. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The dashed black line is the line of best fit fitted to the average of the four voxels.	172
4.18	Parameter correlation plots for the IVIM model. Top row from left to right: substrate cell radius vs D, substrate cell radius vs D*, substrate cell radius vs pseudo-diffusion volume. Middle row from left to right: intracellular walker fraction vs D, intracellular walker fraction vs D*, intracellular walker fraction vs pseudo-diffusion volume. Bottom row from left to right: vascular walker fraction vs D, vascular walker fraction vs D*, vascular walker fraction vs pseudo-diffusion volume. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The dashed black line is the line of best fit fitted to the average of the four voxels.	173

4.19	Parameter correlation plots for the substrate vascular walker fraction plotted against the fitted vascular volume fraction parameter produced using the VERDICT ‘BallSphereStick’ model (left) and ‘BallSphereAstrosticks’ model (right) for substrates with scaled blood flow velocities. The substrates were all generated from voxel 3, from the VERDICT validation experiments previously described. The colours represent the blood flow velocity scaling factor (Blue = 0, Red = 2)	176
4.20	SNR sensitivity results produced by the VERDICT ‘BallSphereStick’ (BSS) model. Top left: A bar graph showing the signal SNR value plotted against the mean absolute error of the fitted cell radius parameter. Remaining plots: Show the fits of the VERDICT BSS model to the synthetic data. Each plot shows gradient strength plotted against signal amplitude for the synthetic VERDICT data corrupted with various levels of noise. The points (stars, squares, triangles) show the synthetic data and the lines show the analytically calculated signal calculated from the parameters of the BSS model fitted to the data.	177
4.21	SNR sensitivity results produced by the VERDICT ‘BallSphereAstrosticks’ (BSA) model. Top left: A bar graph showing the signal SNR value plotted against the mean absolute error of the fitted cell radius parameter. Remaining plots: Show the fits of the VERDICT BSA model to the synthetic data. Each plot shows gradient strength plotted against signal amplitude for the synthetic VERDICT data corrupted with various levels of noise. The points (stars, squares, triangles) show the synthetic data and the lines show the analytically calculated signal calculated from the parameters of the BSA model fitted to the data.	179
5.1	A graphical representation of the most simple embodiment of a neural network unit. Shown are Inputs $x_1, x_2 \dots x_n$, and output function f [5].	196
5.2	A graphical representation of a neural network node, showing the integration function g and activation function f [5].	197
5.3	A schematic showing a network comprising multiple layers of perceptrons.	198
5.4	A graph showing a comparison between a typical threshold (dotted line) and corresponding sigmoidal (solid line) activation function.	201
5.5	An illustration showing the fully-connected representation of a typical feed-forward neural network, comprising an input layer with 10 nodes, two 9-node hidden layers and a 5-node output layer.	203

5.6	An illustration showing the steps of the gradient descent algorithm on a simplified function with a single parameter. Starting at a random parameter value, the algorithm determines the gradient of the cost function at that point, and then updates the parameter according to equation 5.9, then repeats the process until the minimum is found.	211
5.7	An illustration of the gradient descent algorithm being applied to a non-convex function. When the algorithm initialises at the left-hand point, it converges to a local minimum. When the algorithm initialises on the right-hand point, it finds the global minimum.	212
5.8	A simple network with one neuron in each layer, with the layer number indicated by L . Each neuron has an associated activation, a , and an associated bias, b . The neurons in consecutive layers are connected by weights w , which belong to the layer to which they connect.	215
5.9	Plot of number of network hidden layers vs final test dataset error. Each network was trained on 2941 instances of synthetically generated VERDICT MRI data, and tested on 735 instances of synthetic data. Each network comprised an input layer of 220 neurons, an output layer of 4 neurons, and a number of hidden layers each with 112 neurons. Each network was trained for 50000 epochs.	225
5.10	A schematic showing the neural network architecture that was trained on synthetic diffusion MRI data. The network has 5 layers (not including the input layer), with dimensions 108, 108, 108, 108, 4. The input layer comprised 220 nodes.	226
5.11	A side-by-side comparison showing the intracellular volume fraction parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data.	228
5.12	Histogram showing the f_{icvf} value distributions for both the fitted and learned results.	229
5.13	A side-by-side comparison showing the extracellular volume fraction parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data.	230
5.14	Histogram showing the f_{ees} value distributions for both the fitted and learned results.	231

5.15	A side-by-side comparison showing the vascular volume fraction parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data. . .	232
5.16	Histogram showing the f_{vasc} value distributions for both the fitted and learned results.	233
5.17	A side-by-side comparison showing the cell radius parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data.	234
5.18	A side-by-side comparison showing the blood-flow related parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data. The left-hand plot shows the vascular diffusion coefficient D_v , and the right-hand plot is the mean blood flow velocity output from the neural network (cms^{-1}).235	235

List of Tables

2.1	A table showing the formulation of the various compartmental models discussed in the previous section. The mathematical form, fitted parameters and most commonly-represented compartments are shown. S is the diffusion signal, b is the b-value, \hat{G} is the gradient direction, d is the diffusivity, I is the identity tensor, γ is the nuclear gyromagnetic ratio, g is the gradient strength, δ is the gradient duration, l is the coil length, \mathbf{n} is the principle fiber direction, \mathbf{G} is the gradient vector, L is the cylinder radius, R_s is the sphere radius, θ, ϕ are the tensor orientations, $d_{ }$ is the diffusivity parallel to the principle fiber direction, d_{\perp} is the diffusivity perpendicular to the principle fiber direction [6].	81
3.1	The gradient parameters used for the VERDICT acquisitions in this study, along with the corresponding b-value for each combination.	102
3.2	The constraints applied to the parameters during the model fitting process.	104
3.3	Mean and standard deviations of the various fitted parameters at each study time point for the therapy and control groups. For the VERDICT models, the parameter values from the ‘BallSphereStick’ (BSS) and ‘BallSphereAstrosticks’ (BSA) models are shown.	118
3.4	Mean and standard deviations of the volumetric tumour measurements calculated from manual segmentation of the structural (T2-weighted) images of the therapy and control groups, for all time points.	118
4.1	Table showing the summary vascular parameters of each voxel cropped from the network, as well as the overall parameters of the network as a whole. . .	159
4.2	Table showing the ranges over which the substrate parameters were varied.	160

List of abbreviations

ADC	Apparent Diffusion Coefficient
AI	Artificial Intelligence
ANN	Artificial Neural Network
BABB	Benzyl Alcohol, Benzyl Benzoate
BGD	Batch Gradient Descent
CHARMED	Composite Hindered And Restricted ModEl of Diffusion
CR	Complete Response
DNN	Deep Neural Network
DTI	Diffusion Tensor Imaging
DW-MRI	Diffusion-Weighted Magnetic Resonance Imaging
EPI	Echo Planar Imaging
FID	Free Induction Decay
FSEMS	Fast Spin Echo Multi-Slice
GEMS	Gradient Echo Multi-Slice
HREM	High Resolution Episcopic Microscopy
IVC	Individually Ventilated Cages
IVIM	Intra-Voxel Incoherent Motion
mBGD	mini-Batch Gradient Descent
ML	Machine Learning
MRA	Magnetic Resonance Angiography
MRI	Magnetic Resonance Imaging
MSE	Mean Squared Error
MST	Minimum Spanning Tree
NBF	Neutral Buffered Formalin

NMR	Nuclear Magnetic Resonance
NN	Neural Network
NODDI	Neurite Orientation Dispersion and Density Imaging
OCT	Optical Coherence Tomography
PBS	Phosphate Buffered Saline
PCL	Poly(η -CaproLactone)
PFA	Paraformaldehyde
PGSE	Pulsed Gradient Spin-Echo
PR	Partial Response
RECIST	Response Evaluation Criteria in Solid Tumours
RF	Radio Frequency
RNN	Recurrent Neural Network
ROI	Region Of Interest
SGD	Stochastic Gradient Descent
SNR	Signal to Noise Ratio
SOM	Self Organising Maps
SVM	Support Vector Machine
UHF	Ultra High Field
VERDICT	..	Vascular Extracellular and Restricted Diffusion for Cytometry in Tumours

Motivation and Thesis Overview

1.1 Motivation

The development of new cancer therapies is limited by our ability to accurately quantify their effects. The Response Evaluation Criteria in Solid Tumours (RECIST), first published in 2000 [7] and updated in 2009 [8], are the most commonly used method for assessing whether or not an administered therapy has been effective in treating a patients' cancer. The process involves assessing the overall tumour burden at baseline via measurement of the longest diameter of up to five target lesions. This will usually involve imaging the affected organ(s) with a modality such as MRI, CT or ultrasound. The sum of the longest diameters is calculated and used as the reference value for characterising the degree of tumour response. A decrease of the sum of longest diameters by more than 30% is classed as Partial Response (PR), and complete disappearance of all target lesions is classed as a Complete Response (CR).

In many cases the bulk tumour volume is a sensitive predictor of disease progression or response to therapy. However, an increasing number of newly-developed 'novel' cancer treatments are not guaranteed to cause a significant change in the bulk volume of the tumour lesions. This is due to the fact that many of these new therapeutics are cytostatic,

rather than cytotoxic, in nature. Cytostatic therapies work predominantly by inhibiting the uncontrolled growth of cancer cells, and can prevent the tumour from spreading or undergoing metastasis. As a result such therapies do not directly cause the widespread cell death associated with cytotoxic drugs. Other new treatments may target specific metabolic signalling pathways within the cells. It is therefore essential to develop new non-invasive imaging biomarkers that are sensitive to the changes in tumour pathophysiology caused by new treatment methods. Microstructural tissue features such as cellularity, cell radius and microvascular properties have been hypothesised as being useful short-term and long-term markers for therapy response [9, 10].

Currently, in order to measure the aforementioned features (cell radius, cellularity/cell density etc.) it is necessary to perform a biopsy. This involves taking a sample of the tumour tissue using a needle or via surgery [11], which is a highly invasive procedure that also suffers from a very limited sampling area and poor repeatability. Histology is the most commonly-used technique for measuring the microstructural properties of resected tissue. This involves 'fixing' the samples in substances such as ethanol or formaldehyde and applying a histochemical or immunohistochemical stain [12, 13]. Studies have shown that the fixative substances can themselves alter the microstructural features of tissue [14, 15]. There is therefore a clear motivation for exploring alternative techniques for accessing histologic features of tumour tissue, that can overcome the current limitations and issues presented by the current clinical standard of tumour biopsy.

Diffusion-weighted magnetic resonance imaging (DW-MRI) has been suggested as a possible alternative to tumour biopsy and histology. Although conventional MRI is limited to a spatial resolution of a few hundred micrometers at best, diffusion-weighted imaging is sensitive to the geometry of tissue at a microscopic scale. This is due to the fact that DW-MRI is sensitive to the random *motion* of water within the tissue, which itself is influenced by the microstructural geometry. In the clinic, diffusion-weighted MRI has been

used previously to distinguish malignant growths from benign as well as assess tumour therapy response through calculation of the apparent diffusion coefficient (ADC) [16]. The ADC value of tissue is a broad representation of the freedom with which water can diffuse through the tissue, and can be related to inherent tissue properties such as cell density, interstitial fluid pressure and cell radius [17–19]. The ADC is therefore a powerful biomarker for studying cancer. However, the fact that it is influenced by so many microstructural tissue properties is a limitation as well as an advantage as it is impossible to know which of the contributing factors are responsible for any given change in the ADC value.

One of the methods proposed for separating out the contributions from each microstructural parameter is multi-compartment diffusion modelling [6]. The aim of multi-compartment modelling is to represent the total diffusion signal as a combination of signals from discrete compartments, with associated mathematical diffusion models, that are chosen to represent the underlying structure of the tissue. The complexity of multi-compartment models ranges from relatively simple bi-exponential fits to more complex in-depth biophysical models [9, 20, 21]. Recently developed compartmental models such as VERDICT (Vascular, Extracellular and Restricted Diffusion for Cytometry in Tumours) have demonstrated the ability to access voxel-specific histologic features such as cell radius and intra/extracellular volume fraction non-invasively. These techniques therefore have the potential to greatly increase our ability to measure the microstructural changes that occur in tumours naturally over time and in response to administered therapies. However, with the increasing complexity of the mathematical models involved the need for thorough validation becomes especially important.

The validation of these novel quantitative imaging techniques is an ongoing challenge. It is essential to establish the correspondence between tissue parameters calculated by fitting mathematical models to the diffusion signal and the corresponding microstructural properties of the tissue. Typically, validation can be done using histological analysis, biological

phantoms and *in-silico* modelling, or a combination of these [22–26]. Histological analysis is often considered the gold-standard technique for validating quantitative imaging techniques as it allows for a direct measurement of the tissue microstructure. As previously mentioned however, the processing steps involved may cause distortion of the tissue on a cellular level. It is also difficult to validate any parameters associated with tissue perfusion, although specific tissue stains that target microvascular structures can provide a useful surrogate measure of *in-vivo* perfusion via the measurement of microvessel density [27]. Biological phantoms provide the opportunity to generate DW-MRI validation datasets based upon the real motion of water molecules within a representative model (a ‘phantom’) of the appropriate tissue. The phantom is scanned with an MRI scanner to produce diffusion-weighted imaging data, which can then be used to validate the quantitative analysis technique against the known dimensions of the phantom. Computer simulations can be used to generate synthetic DW-MR data from a virtual tissue model with known dimensions while offering the user full control over all parameters. A simulation-based approach to validation allows highly complex microstructural environments to be examined, but are difficult to implement and rely on the user’s knowledge of how different simulation parameters may influence the diffusion signal [28, 29].

Techniques such as compartment diffusion models have been shown to provide useful insights into the microstructural properties in a variety of tissues. However, they are not without their limitations. Compartmental diffusion models use a simplified representation of the underlying tissue geometry such as packed spheres representing cells, or cylinders representing bundles of axons. Simple geometries must be used so that mathematical diffusion models can be developed that can closely approximate the expected diffusion-weighted MR signal. Such limitations can be overcome by moving away from an explicit mathematical modelling approach and instead attempting to infer a direct relationship between the raw data and tissue microstructure using a machine learning based approach.

Machine Learning (ML) is a branch of computer science that involves the use of algorithms that 'learn' how to perform a specified task without being explicitly programmed by the user. Machine-learning algorithms range in complexity from simple gradient-descent linear regression to deep neural networks and are used in a vast range of applications. Previously, machine learning has been applied to structural MRI data in order to classify and grade brain tumours [30]. It has also been applied to DW-MRI data in order to measure white matter permeability in the brain [31], classify temporal lobe epilepsy [32] and predict major depressive disorder [33]. Applying similar techniques to both synthetic and real DW-MR data from tumours may assist further development of novel microstructural biomarkers for cancer diagnosis and disease progression.

1.2 Thesis Overview

Chapter 2 presents the theoretical background and context relevant to the project overall including the fundamental physics of MRI, and an overview of diffusion-weighted MRI and cancer. Chapter 3 presents the background, methods and results from two practical applications of VERDICT MRI for measuring tissue microstructure. Firstly, in detecting a change in the estimated cell radius of tumour tissue before and after formaldehyde fixation. Secondly, in detecting the response to temozolomide therapy in a preclinical murine model of glioblastoma. Chapter 4 describes the development of an *in-silico* Monte Carlo model of water diffusion and blood flow in tumour tissue. The chapter describes the theory and methods involved in the simulation, some examples of typical simulations performed, and the methods involved in producing synthetic diffusion MRI data from the results. The second half of chapter 4 presents the application of the synthetic diffusion MRI data as a tool for validating the VERDICT MRI framework. VERDICT data was generated from simulations on substrates with varied microstructural dimensions and parameters.

Different VERDICT models were fitted to the synthetic data and the resulting tissue metrics compared against the known values from the simulations. An investigation into the SNR-sensitivity of the VERDICT models was also performed by corrupting the synthetic data with increasing levels of Rician noise. Chapter 6 explores the possibility of utilising increasingly-popular machine learning based techniques to extract microstructural features from raw diffusion MRI data. The chapter begins with a review of the background and development of machine learning to date and gives details of the theory behind some of the most popular algorithms. The main experimental results sections of the chapter then follow. Firstly, training a neural network on synthetically generated DW-MRI data to investigate whether it is possible to learn microstructural features from the raw signal without the need for explicit mathematical modelling. Secondly, the possibility of testing a neural network trained on synthetic data on *in-vivo* data is investigated. Finally chapter 7 presents the overall conclusions and outlines the areas where future work can build upon this work.

Background

This chapter contains the background and theory underlying cancer, MRI, and diffusion MRI in particular. The background of machine learning is presented later on in chapter [6](#). This introduction is necessary to provide context to the research presented within this thesis and as a precursor to the ideas discussed in later chapters.

2.1 Cancer

2.1.1 History and Brief Overview

The word ‘cancer’ is the Latin word for the Greek ‘karkinos’, meaning ‘crab’. It was first used to describe the disease itself by Hippocrates around 400BC. However, most of our understanding of cancer was developed following the discovery of the DNA double helix in 1953. The term cancer refers to a group of diseases that are characterised by the uncontrolled proliferation and growth of cells. It can affect any of the tissues in the body, though the most common form is cancer originating in epithelial cells, known as carcinoma. Cancer begins as a normal functioning cell that undergoes some form of DNA damage, which causes the growth and proliferation of the cell to spiral out of control. As the cancer develops, the proliferating cells become structurally and functionally different

to the cells from which the disease originated, a process known as *differentiation*.

The severity of cancer is determined by how well differentiated it is, which is assessed via histological analysis of the tissue. Assessing the grade of the disease is a key part of the patient care pipeline, as it is the main factor that influences the choice of treatment. Analysis of the tumour microstructure, therefore, is of the utmost importance for determining the current state of the disease. Changes in the microstructural properties of the tissue can be indicative as to whether or not a tumour has responded to administered therapy, and so being able to access these properties non-invasively is the pervasive goal of microstructural imaging in cancer.

2.1.2 Tumour Microstructure

As previously mentioned, the microstructural environment of tumours is a topic of great importance within the field of oncology as it plays a central part in both the diagnosis, progression and treatment of the disease. Commonly, the microstructure of tumour tissue is described as three discrete compartments: the extracellular, intracellular and vascular space. Each compartment has very distinct characteristics.

The intracellular compartment contains various cells such as immune cells, macrophages and the cancer cells themselves. The cells themselves are each surrounded by a semi-permeable membrane, leading to the phenomenon of restricted diffusion discussed later on [34]. The shape and size of cancer cells tends to vary to a greater extent than in healthy tissue, especially in poorly-differentiated cancers [35, 36]. A range of cell radii of between 14-60 microns has been observed in breast cancer cells taken from patients [37].

The extracellular space plays an important role in the transport of molecules throughout the tumour via the presence of interstitial fluid. Tumours are known to exhibit a large

extracellular space, and greater interstitial fluid convection than in healthy tissue. These factors combined with greater interstitial fluid pressure can lead to poor drug retention [38, 39]. The extracellular matrix (ECM) consists of a collagen ‘scaffold’ that gives the tissue rigidity, but may also act to hinder the diffusion of water molecules within the extracellular space [40].

The vascular space of tumours is important as it provides the tumour with a source of oxygenated blood and nutrients in order to maintain viability. As previously mentioned, the vast majority of cancers are carcinomas originating from epithelial cells. The tissue produced by these cells is not vascularised, preventing tumour lesions from growing larger than 2mm diameter without creating new vessels. Tumours stimulate the growth of new blood vessels via a process known as angiogenesis [41]. Angiogenesis in tumours is dysregulated compared to healthy tissue and leads to the formation of highly chaotic and leaky vessel network that is characteristic of tumour tissue.

2.1.3 Quantitation of Tumour Microstructure

Currently, in order to measure the properties mentioned above it is necessary to perform a biopsy. This involves taking a sample of the tumour tissue using a needle or via surgery [11], this is a highly invasive procedure that also suffers from a very limited sampling area and poor repeatability. Histological analysis is most often used to measure the microstructural properties of the resected tissue, this usually involves fixing the sample in substances such as ethanol or formaldehyde and applying a histochemical or immunohistochemical stain [12, 13]. Fixative substances such as ethanol and formaldehyde may themselves damage the microstructure, thereby confounding the results. These effects have mostly been studied via sectioning of fixed samples or in-vitro [14, 15]. There is therefore a clear motivation for finding an alternative method for measuring the microstructural properties

of tumour tissue. Microstructural imaging is a promising candidate, by exploiting the behaviour of water molecules within the tissue enabling the microstructural properties of the tissue to be accessed throughout the tumour non-invasively.

2.2 Fundamental Physics of MRI

The following section contains an introduction and overview of the phenomenon of nuclear magnetic resonance (NMR) and its application to medical imaging. To understand the body of work contained in this thesis it is necessary to have a grasp on the fundamental physical principles that underpin MRI. Numerous textbooks have been published that provide an in-depth analysis and exploration of the physics of MR, so this section will be limited to a brief discussion [42–45]. Despite the fact that the fundamental physics of MRI are inherently quantum mechanical by nature, the theory presented in this chapter will not include an in-depth QM description of phenomena. Such analysis can be found in various other literature [46, 47].

2.2.1 Spin and Nuclear Magnetism

In order to explain the physics of MRI, we must first discuss the quantum mechanical property known as *spin*. All fundamental particles possess spin, it is an inherent and measurable property just like mass or charge. Despite the name, spin has nothing to do with particles rotating - the name comes from observations of the way particles interact with magnetic fields; the way in which particles are deflected by a static magnetic field is comparable to the behaviour of a charged, spinning object in classical physics. Instead, spin is related to the quantisation of rotational angular momentum. As with other quantum properties, the spin of a particular physical system can take discrete (quantised) values,

in this case given by:

$$S = \hbar\sqrt{s(s+1)} \quad (2.1)$$

Where \hbar is the reduced Planck constant ($1.054 \times 10^{-34}\text{Js}$) and s is the spin quantum number ($s = \frac{n}{2}$, where n is any non-negative integer). In addition, spin behaves as a vector quantity with both magnitude and direction [48], where the component of spin measured along a particular direction can take values

$$S_i = \hbar s_i, \quad s_i \in \{-s, -(s-1), \dots, s-1, s\} \quad (2.2)$$

Particles with an integer spin are known as *bosons*, while those with half-integer spins are known as *fermions*. All protons, neutrons and electrons are fermions known as “spin-half” particles with $s = \frac{1}{2}$ and so $S = \frac{1}{2}$ or $S = -\frac{1}{2}$. These two spin states are known as “spin up” or “spin down” respectively. In the field of MRI, we are only interested in spin-half nuclei (a particle/nucleus made of an odd number of fermions is known as a composite fermion) such as ^1H , ^{19}F and ^{31}P .

When an individual spin-half particle is placed in a strong external magnetic field (such as that of an MRI scanner), and it’s spin vector directly measured, it is forced to occupy one of the two possible spin states via a process called *Zeeman splitting* (the measurement causes the particle to ‘collapse’ into one of two possible eigenstates). In this case, the particles can more simply be represented as a magnetic dipole that aligns itself parallel or anti-parallel to the magnetic field direction.

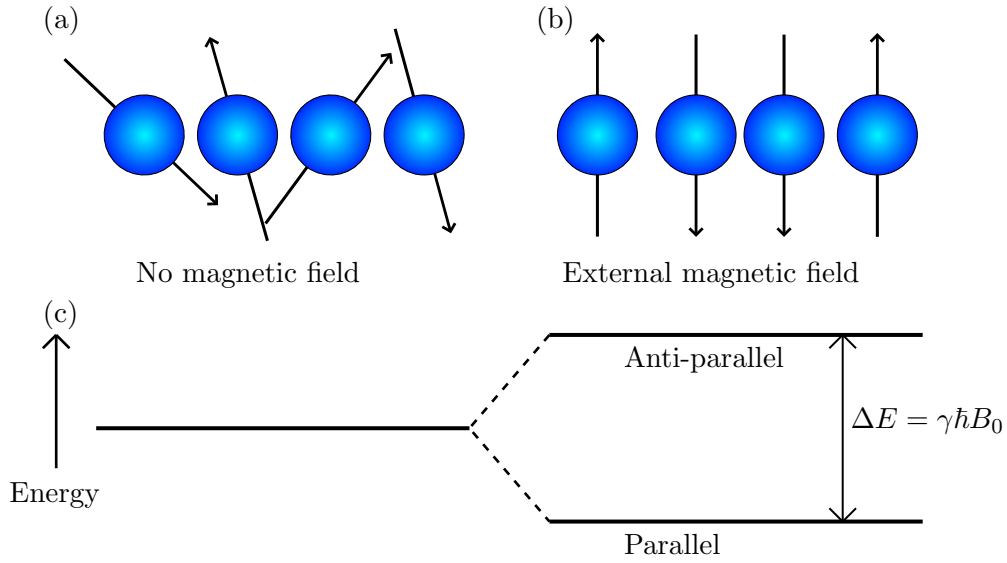


Figure 2.1: A schematic showing nuclei (a) without the influence of an external magnetic field and (b) with an external magnetic field. The nuclei are here classically represented as individual magnetic dipoles that align parallel or anti-parallel to the field. The energy levels associated with the two configurations is shown in (c).

As shown in figure 2.1 there is a difference in the energy levels associated with the parallel and anti-parallel configuration, with the former being energetically more favourable. The energy level of each configuration is given by

$$E = -\boldsymbol{\mu} \cdot \boldsymbol{B}_0 \quad (2.3)$$

Where μ is the magnetic moment associated with the dipole, given by $\mu = \gamma S$ where γ is the *gyromagnetic* ratio which is an inherent property of any nucleus dependent on it's mass and charge. Substituting in for $S = \pm \frac{1}{2}$ for spin-half nuclei we find that the difference in energy between the two configurations is [49]

$$\Delta E = \gamma \hbar B_0 \quad (2.4)$$

From this result, it is possible to derive one of the most fundamental equations associated with MRI - The Larmor equation. It is known that a particle occupying the lower energy state can be excited into the higher energy state by absorption of a photon with an associated energy equal to the difference between the two states. According to the Planck-Einstein relation $E = \hbar\omega$, the energy of a single photon of light is proportional to its frequency. Equating this to [2.4](#) and rearranging gives:

$$\omega = \gamma B_0 \tag{2.5}$$

This result demonstrates the existence of a ‘nuclear resonance’ frequency, referred to as the Larmor frequency (named after Sir Joseph Larmor who discovered it), unique to each particle. This is the basic condition for magnetic resonance - the MR in MRI.

As previously mentioned, the parallel configuration of spins is energetically slightly favourable to the anti-parallel. This means that given a large ensemble of spins (at body temperature) placed in an external magnetic field, the number of spins that adopt the parallel configuration will be very slightly greater. The reason that all spins do not just occupy a lower energy state is due to the fact that large thermal fluctuations tend to oppose the preference for spins to occupy the lower state [\[50\]](#). A popular analogy to illustrate this is the “compasses in a washing machine” analogy. If unperturbed compasses will naturally point north, however this changes if they are placed in the rotating drum of a washing machine. As the compasses bounce around, the needles will no longer all point north, however averaging over all the compasses will show that there is still a slight tendency for them to do so [\[46\]](#).

At thermal equilibrium, the relative number of spins occupying the upper and lower states is predicted by a Boltzmann distribution:

$$\frac{N_{upper}}{N_{lower}} = e^{-\frac{\Delta E}{k_B T}} \quad (2.6)$$

Where N_{upper} and N_{lower} are the populations occupying the upper or lower energy states respectively, k_B is the Boltzmann constant ($k_B = 1.38 \times 10^{-23} \text{m}^2 \text{kg s}^{-2} \text{K}^{-1}$), ΔE is the energy difference between the states and T is the temperature of the system. At 37°C , and an external field strength of 9.4T (typical preclinical MRI B_0 field), equation 2.6 gives $\frac{N_{upper}}{N_{lower}} = 0.999938$. This shows the significance of the thermal fluctuations, and how slight the preference is for spins to occupy the lower energy state. The excess of spins in the lower energy state results in a small net magnetisation vector $\underline{M} = \sum \underline{\mu}_{net} = \gamma \underline{L}$, where \underline{L} is the total angular momentum of the sample. It is this net magnetisation that is measured during NMR or MRI experiments.




Figure 2.2: The distribution of spins within a macroscopic sample, with a spherical distribution of directions, all precess around the principle magnetic field direction \underline{B} at angular velocity ω .

The net magnetisation and static external field interact causing a torque to be exerted on the magnetisation vector:

$$\boldsymbol{\tau} = \mathbf{M} \times \mathbf{B} \quad (2.7)$$

The torque causes the magnetisation vector to precess around the B-field vector. Given that the torque is equal to the rate of change of angular momentum, and that the angular momentum $L = \frac{M}{\gamma}$, an equation for the time-dependent evolution of the net magnetisation vector can be written as

$$\frac{d\mathbf{M}}{dt} = \gamma(\mathbf{M} \times \mathbf{B}) \quad (2.8)$$

Expanding the vector product as a matrix gives

$$\begin{bmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ B_x & B_y & B_z \\ M_x & M_y & M_z \end{bmatrix} = (M_y B_z - M_z B_y)\mathbf{i} + (M_z B_x - M_x B_z)\mathbf{j} + (M_x B_y - M_y B_x)\mathbf{k} \quad (2.9)$$

Conventionally, the static external field B is defined to lie parallel to the z-axis. As B_x and B_y are zero, 2.9 can be separated into the three directions and simplified to

$$\frac{dM_x}{dt} = \gamma M_y B_z, \quad \frac{dM_y}{dt} = \gamma M_x B_z, \quad \frac{dM_z}{dt} = 0 \quad (2.10)$$

For a freely precessing system in the absence of relaxation and RF excitation, these equations have the following solutions:

$$M_x(t) = M_x(0) \cos(\omega t) - M_y(0) \sin(\omega t) \quad (2.11)$$

$$M_y(t) = M_y(0) \cos(\omega t) + M_x(0) \sin(\omega t) \quad (2.12)$$

$$M_z(t) = M_z(0) \quad (2.13)$$

These solutions demonstrate that ignoring the effects of relaxation and excitation, the bulk magnetisation vector \mathbf{M} precesses around the principle magnetic field vector at the Larmor frequency ω (as shown in figure 2.2) with no time-dependent change in the z-component of the magnetisation [51, 52]. However, in the context of MRI the situation is more complex and requires the effects of relaxation to be taken into account.

2.2.2 Relaxation and RF Excitation

MR experiments require the measurement of the net magnetisation vector. As described above, when a system is in equilibrium the net magnetisation vector lies parallel with the static B_0 field. This makes detection of the magnetisation impossible as the static field is many orders of magnitude larger than the magnetisation vector. In order for measurement to be possible, the net magnetisation vector must be tilted into the xy plane perpendicular to the static field. This is achieved using radio-frequency (RF) excitation. Considering an ensemble of spins in their thermodynamic equilibrium state, within a strong homogeneous magnetic field. The net magnetisation vector \mathbf{M} will be aligned with the magnetic field vector \mathbf{B} , and precesses around it at the Larmor frequency given by equation 2.5. The individual spins possess no phase coherence and precess out of phase with each other resulting in zero net transverse magnetisation \mathbf{M}_{xy} .

The sample can be *excited* by irradiating it with a radio frequency (RF) pulse. An RF pulse consists of a rotating magnetic field \mathbf{B}_1 with frequency ω orientated perpendicular to the static field \mathbf{B}_0 that is switched on for a short period of time. It is called an RF pulse because the typical frequency of oscillation of such fields lies within the radio frequency region of the electromagnetic spectrum. Switching to a reference frame that rotates around \mathbf{B}_0 simplifies the analysis, as the total effective magnetic field is stationary in such a frame. If the secondary field \mathbf{B}_1 is applied parallel to the x-axis, the total field is given by

$$\mathbf{B}_{total} = B_1 \mathbf{i} + (B_0 + \frac{\omega}{\gamma}) \mathbf{k} \quad (2.14)$$

where \mathbf{i} and \mathbf{k} are unit vectors along the x and z axis respectively [53]. The magnitude of \mathbf{B}_{total} can be written as

$$B_{total} = \sqrt{B_1^2 + (B_0 + \frac{\omega}{\gamma})^2} = \frac{\sqrt{\omega_1^2 + (\omega_0 - \omega)^2}}{\gamma} \quad (2.15)$$

where $\omega_1 = -\gamma B_1$. In the rotating frame of reference the net magnetisation vector \mathbf{M} precesses around \mathbf{B}_{total} . Looking at figure 2.3, it can clearly be seen that $\sin\theta = \frac{\omega_1}{\sqrt{\omega_1^2 + (\omega_0 - \omega)^2}}$, meaning that when the oscillating frequency of the rotating field is equal to the Larmor frequency ω_0 , the angle between the total effective magnetic field and the z -axis becomes 90° . As the net magnetisation vector precesses around the total effective magnetic field vector, this causes the net magnetisation to tip into the xy plane. By varying the duration and strength of the rotating magnetic field the angle α by which the net magnetisation is tipped can be varied.




Figure 2.3: Diagrams showing the orientations of the magnetic fields viewed from the rotating reference frame. (a) shows the general case where the longitudinal field B_0 and the B_1 field combine to create the total effective magnetic field B_{total} . (b) shows the case when the frequency of the rotating field ω is equal to the Larmor frequency ω_0 , where the B_1 becomes the total effective field. In both cases the net magnetisation vector M precesses around the total effective field vector. In (b) the magnetisation vector tilts by an angle $\alpha = 90^\circ$ into the xy plane.

The rotating field, or RF pulse, is often defined by the angle α that it tips the net magnetisation by. The pulse illustrated in figure 2.3 (b) would be called a 90° pulse, for example. Once the magnetisation vector has been tipped into the xy plane, it can be detected and measured using a radio-frequency coil. However, over time the system will return to it's equilibrium state via a process known as *relaxation*.

When Felix Bloch and Edward Purcell independently discovered the phenomenon of NMR, it was noted that the observed ‘nuclear induction signal’ was transient in nature and would decrease in amplitude over time following excitation [51, 52]. It was Bloch who hypothesised that this was due to the individual spins interacting with one another and their environment in such a way as to return the net magnetisation vector \mathbf{M} to it's equilibrium state parallel to the B_0 field, via a process that he termed *relaxation*. Bloch introduced two time constants, **T1** and **T2** that describe the rates at which the different components of the magnetisation decay.

These relaxation time constants were not derived mathematically from basic principles, but rather were defined in order to explain observed phenomena. **T1** was defined to explain the re-growth of the longitudinal magnetisation caused by the interactions between the spins and the environment, known as ‘*spin-lattice*’ relaxation. **T2** was defined to explain the decay of the transverse component of the magnetisation caused by internuclear interactions, referred to as ‘*spin-spin*’ relaxation. Additional terms are added to 2.10 to include these effects:

$$\frac{dM_x}{dt} = \gamma M_y B_z - \frac{M_x}{T2} \quad (2.16)$$

$$\frac{dM_y}{dt} = \gamma M_x B_z - \frac{M_y}{T2} \quad (2.17)$$

$$\frac{dM_z}{dt} = -\frac{M_z - M_0}{T1} \quad (2.18)$$

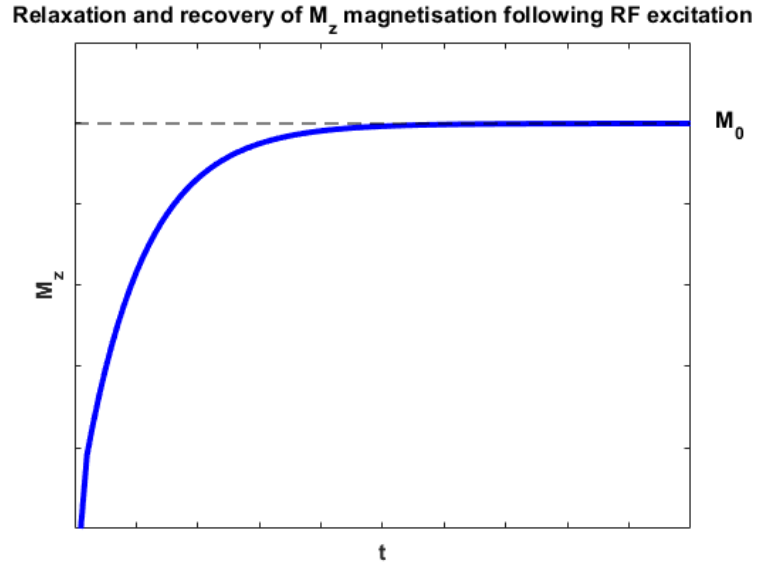


Figure 2.4: Graph showing the characteristic relaxation and recovery of the z-component of the magnetisation following excitation.

Figure 2.4 shows the solution to equation 2.18, describing the evolution of the z-component of the net magnetisation vector over time following excitation by a 90° RF pulse. It can be seen that the recovery of the magnetisation follows an inverse exponential, eventually returning to the equilibrium state with $M_z = M_0$.

In a real-world scenario, the rate of transverse (T2) relaxation can be increased by the presence of B_0 inhomogeneities. If the static field is not perfectly uniform, it can cause spins at different locations within the sample to precess at different frequencies. The difference in precessional frequency combines with the effect of *spin-spin* interaction to cause increased dephasing. This combined effect is referred to as T2*.

2.2.3 Signal Reception and Spatial Encoding

As previously mentioned, the magnetisation of a sample can only be measured perpendicular to the static B_0 field. The detection and measurement is performed using an *RF coil*. The net magnetisation vector of the sample continues to precess around the B_0 field after it has been tilted into the xy plane by a suitable RF pulse. The precessing magnetisation vector produces an oscillating magnetic field, capable of inducing a current via Faraday's law in a suitably-designed radio-frequency coil oriented perpendicular to the static field and placed within the oscillating field. The *emf* (electro-motive force) induced in a coil by the changing magnetic field is given by:

$$\varepsilon(t) = -\frac{d\phi}{dt} \tag{2.19}$$

$$= -\frac{d}{dt} \int \mathbf{M}(t, x) \cdot \mathbf{B}_c(x) dx \tag{2.20}$$

Where ϕ is the flux through the RF coil, and can be thought of as the number of magnetic field lines that are 'cut' by the surface enclosed by the loops of the coil. The quantity B_c is the ratio of the magnetic field produced by the coil to the current in the coil - it is thereby a function of the individual coil used [54]. The induced signal in the coil is known as the FID (Free-Induction Decay). The FID oscillates at the same frequency as the oscillating magnetic field that produced it: the Larmor frequency of the sample [55].




Figure 2.5: A typical Free Induction Decay signal. By Nmr_fid_good_shim.svg: GyroMaggianderivative work: Imalipusram This file was derived from Nmr fid good shim.svg;, CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=30312230>

As can be seen in figure 2.5, the induced current in the RF coil decays in magnitude over time, due to the effects of relaxation described previously.

The signal detected by the RF coil originates from the net magnetisation of the entire sample. In order to produce an image (to achieve the 'I' in MRI) it is necessary to encode spatial information into the MR signal. This is done using *spatial encoding*. Spatial encoding is achieved by the addition of three resistive magnets to the MRI scanner, known as the 'imaging gradients'. Each of the three gradients lies along one of the principle axes

of the scanner, and are thereby sometimes referred to in notational form as G_x , G_y and G_z . The function of each gradient for any given acquisition depends on the orientation of the desired imaging plane relative to the scanner, each gradient performs one of three spatial encoding roles: slice selection, frequency encoding or phase encoding.

Slice selection, as the name would suggest, involves selecting a slice of interest from within the sample. In the vast majority of situations, it is not desirable to excite the entire sample at once. For imaging applications, the excitation must be confined to a small region, or slice. This is achieved by using a slice-selective RF pulse that only excites protons within the desired imaging plane. A linear gradient (usually G_z) is applied, which changes the effective field strength throughout the sample to:

$$B_z(t) = B_0 + G_z(t)z \quad (2.21)$$

The spatially varying magnetic field strength throughout the sample causes the Larmor frequency of the sample to be spatially varying ($\gamma(B_0 - G_z) < \omega_0 < \gamma(B_0 + G_z)$). By switching on this gradient at the same time as the RF pulse, only the region (slice) of the sample where the Larmor frequency matches the frequency of the RF pulse will be excited. The position of the excited slice can be adjusted by changing the centre frequency of the RF pulse. The thickness of the slice can be specified by varying the bandwidth (the frequency range of the pulse) or the strength of the slice-selective gradient [54, 56]. Once a particular slice of the sample has been excited, the further two dimensions of spatial localisation are achieved using frequency and phase encoding.

Frequency encoding is performed using a similar technique to the slice-selection method described above. A linear gradient is applied along one direction (usually the x-axis), causing spins to precess at different frequencies $\omega = \gamma(B_0 + G_x x)$ depending on their

x-position. The linear gradient is applied while the MR signal is being acquired, meaning that the recorded signal contains a range of frequencies. The different frequency components correspond to different positions along the x-axis of the imaging plane.

Phase encoding is used to spatially encode the signal along the final remaining dimension (the y-axis). It is also performed using a linear gradient applied along the y-axis, that similarly causes the spins to precess at a frequency that depends on their position on the y-axis. The gradient is applied for a brief period of time and is then switched off again. When the gradient is switched off, the spins return to precessing at the Larmor frequency ω_0 , but will each have gained or lost a certain ‘phase’ that depends on their position along the y-axis. This process is then repeated a number of times, changing the strength of the phase-encoding gradient with each repetition [57]. After a number of repetitions with different gradient strengths have been acquired, the multiple signals can be spatially decoded and an image of the chosen slice produced.

2.2.4 Basic Pulse Sequences

The combinations of various RF pulses and gradients used to acquire the data needed to form an MR image are known as ‘pulse sequences’. The two main types of pulse sequence used are known as ‘gradient echo’ and ‘spin echo’. These basic sequences form the foundation for many more advanced and more complex pulse sequences and acquisition techniques. The term ‘echo’ refers to the way the detected signal appears after a short period of time following a refocusing RF pulse as the spins re-phase in the xy plane.

Gradient Echo




Figure 2.6: A timing diagram showing a typical gradient echo pulse sequence. The diagram shows the order in which the various RF pulses and gradient waveforms occur during the acquisition. Shown below and labelled is the effect of the aforementioned pulses and gradients on the net magnetisation vector (red) and the individual spins (blue).

Figure 2.6 shows a timing diagram for a gradient echo pulse sequence. The number labels correspond to different events, namely the application of gradients and RF pulses. Timepoint 1 shows the system in equilibrium, with the net magnetisation aligned with the B_0 field. At 2, an α° RF pulse is applied, exciting the system and tilting the net

magnetisation vector into the xy plane. Simultaneously, the slice-selection gradient G_z is applied, to ensure that only the desired slice is excited by the RF pulse. At **3**, the spins are dephased in the xy plane by G_x , referred to as the dephasing gradient. The gradient dephases the spins until the signal is completely null due to complete loss of coherence. Phase encoding is also performed using G_y . At **4**, the polarity of G_x is reversed, causing the spins to rephase. When the area of the rephasing gradient is equal to the dephasing gradient, a signal can be detected (the gradient ‘echo’ **5**). The echo will be somewhat smaller than the initial excited signal, due to $T2^*$ effects. At **6**, the spins continue to dephase naturally due to transverse ($T2$) relaxation.

Spin Echo

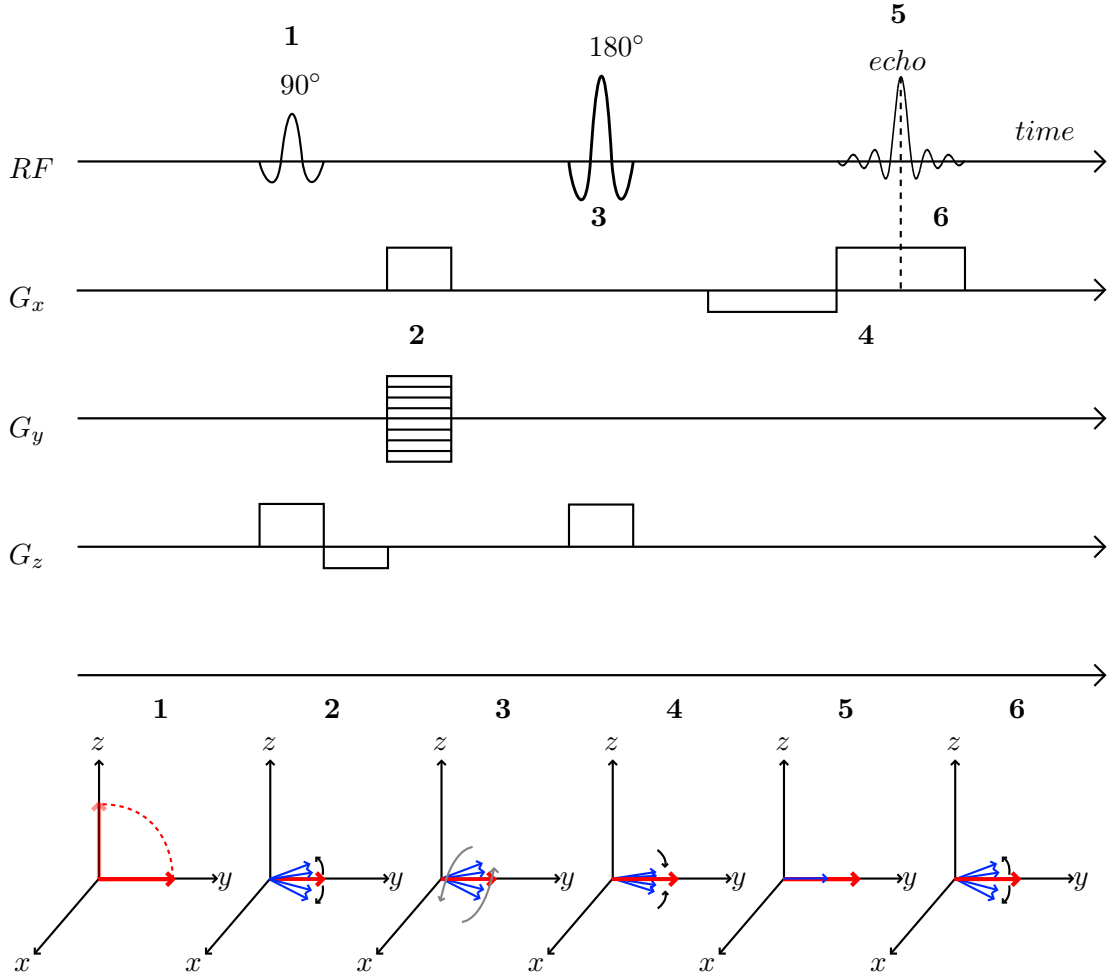


Figure 2.7: A timing diagram showing a typical spin echo pulse sequence. The diagram shows the order in which the various RF pulses and gradient waveforms occur during the acquisition. Shown below and labelled is the effect of the aforementioned pulses and gradients on the net magnetisation vector (red) and the individual spins (blue).

Figure 2.7 shows a timing diagram for a typical spin echo pulse sequence. Different time points are labelled on the diagram, corresponding to the key features of the pulse sequence. Time point 1 shows the application of a 90° RF pulse, the effect of the RF pulse on the

net magnetisation vector is to tilt it entirely into the xy plane. The slice-selection gradient G_z is applied during the RF excitation, to ensure only spins within the desired slice are excited. At time point **2**, the spins begin to dephase in the xy plane due to the applied gradients and T2 relaxation driven by spin-spin interactions. At time point **3**, a 180° RF pulse is applied. This pulse effectively flips the entire spin distribution by 180° , in a manner similar to flipping a pancake. The slice-selection gradient G_z is again applied during RF excitation, as only spins within the desired slice should be influenced by the 180° pulse. Time point **4** shows the spin population shortly after the application of the 180° pulse. The individual spins continue to precess at the same frequencies relative to each other as they did before the 180° pulse. However, instead of continuing to dephase, they now rephase as the whole population has been flipped. Time point **5** shows the formation of a signal ‘echo’ (a ‘spin echo’) caused by the rephasing of spins in the transverse plane. At this point, the signal is acquired using the RF receive coil. During acquisition, the frequency-encoding gradient G_x is applied, to spatially localise the spins along the x-axis. Time point **6** shows that following the signal echo, the spins begin to dephase again, and the signal decays.

It can clearly be seen that there are common features shared between the gradient-echo and spin-echo pulse sequences; both involve initial excitation by an RF-pulse and both utilise the imaging gradients to perform frequency-encoding and phase-encoding. The main difference between the two is the approach used to form the signal echo that is detected. The gradient echo uses the imaging gradients to dephase and then rephase the spins within the transverse plane, whereas the spin echo approach uses a 180° RF pulse to flip the spin population causing them to naturally rephase. Gradient echo sequences are generally faster than spin echo, as acquisition time can be reduced by increasing the strength of the dephasing/rephasing gradients. Spin echo sequences must wait for the spin population to rephase following the 180° pulse. Typical gradient echo repetition times are

of order 3ms, whereas spin echo times are closer to 30ms. Spin echo sequences generally produce images with higher SNR (signal to noise ratio) as $T2^*$ effects are cancelled out by the refocussing pulse.

2.2.5 Image Formation

The signals detected using the RF coil each correspond to an acquisition with certain phase encoding values, with additional spatial information included via frequency encoding. The acquired signals from each phase encoding repetition are in the frequency domain and are stored in a two-dimensional matrix known as k-space. Each row of the k-space matrix corresponds to a sampled signal echo at a single phase encode value, as shown in figure 2.8.

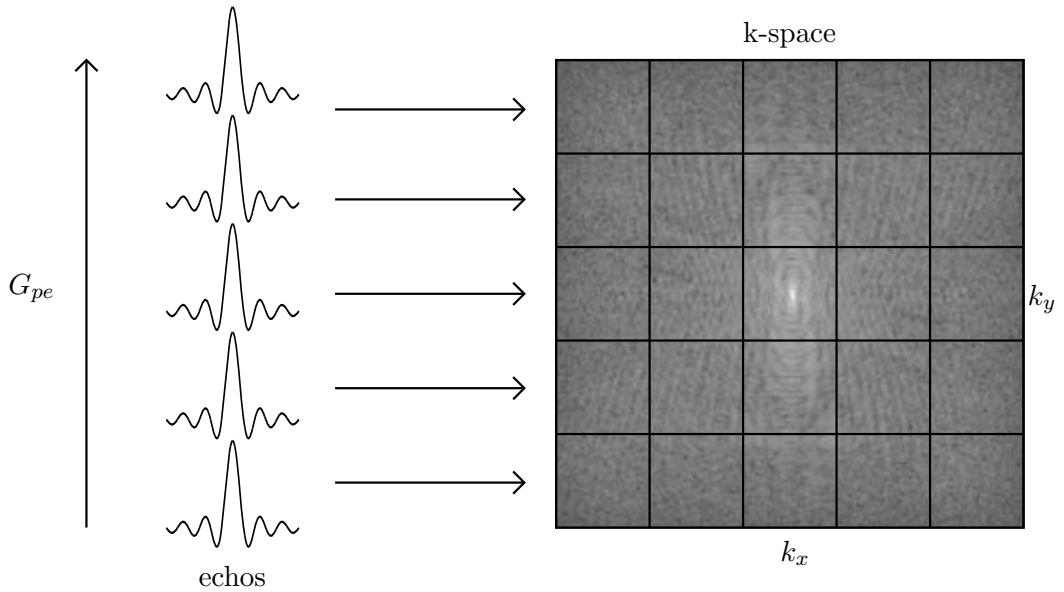


Figure 2.8: The construction of k-space from detected signal echos. Each digitally sampled echo with a particular phase encode value corresponds to a row of the k-space matrix.


A 2D fast Fourier transform is used to convert the k-space matrix from the frequency to the spatial domain, creating complex-type image data with real and imaginary components. These components represent the real and imaginary components of the oscillating circularly polarised magnetic field that was detected by the RF coil. Different information can be extracted from the complex-valued data, depending on the application. The magnitude of the complex data can be calculated by taking the modulus of the real and imaginary components:

$$Mag = \sqrt{Real^2 + Imag^2} \quad (2.22)$$

The images formed by taking the magnitude data are most commonly used for structural imaging. The phase of the complex data can be calculated by:

$$\phi = \arctan\left(\frac{Imag}{Real}\right) \quad (2.23)$$

Specialised pulse sequences can be developed that encode specific information into the phase of the complex data. Thereby phase images can contain information about bulk fluid motion, spatially-varying magnetic susceptibility and static field inhomogeneity. Examples of typical magnitude and phase images are shown in figure [2.9](#).



Magnitude

Phase

Figure 2.9: Typical examples of a magnitude and phase image of a human brain. Images taken from <http://mriquestions.com/making-an-sw-image.html> on 22/05/2018

2.3 Diffusion-Weighted MRI

2.3.1 Basic Overview

One of the limitations of MRI as an imaging modality is the relatively low spatial resolution. Recent advances in the fields of UHF (ultra high field) MR have shown that structural images with a resolution of $100\mu\text{m}$ are possible in the brain [58]. However, even with a resolution of the order $100\mu\text{m}$ conventional MRI is unable to resolve the microstructural details of tissue. To achieve this, it is necessary to take advantage of the fact

that the MR signal can be sensitive to the microscopic random motion of water molecules known as ‘self-diffusion’ or ‘Brownian Motion’ shown in figure 2.10. This motion occurs over length scales much smaller than the resolution of conventional MRI, and comparable to the length scale of the important structures within the tissue microenvironment such as cells and capillaries. Using diffusion-weighted MRI (DW-MRI) makes it possible to sensitise the MRI signal to the tissue microstructure.




Figure 2.10: Illustration of the random thermal motion of a water molecule.

The development of diffusion-weighted MRI gathered momentum in the 1980's to 1990's, with the first diffusion-weighted MR image published in 1985 [20, 59]. The main basis of diffusion-weighted MRI is the aforementioned *Brownian Motion* or *self-diffusion*, a simple

schematic of which is illustrated in figure 2.10. The molecule shown at the starting point at $t = 0$ will undergo a 'random walk', and after a time τ will have a displacement r from the starting position. It should be appreciated that when this effect is propagated over a large system of particles the mean displacement will be 0, as particles will move in all directions. It is therefore more useful to consider the mean squared displacement:

$$\langle r^2 \rangle = 6D\tau \quad (2.24)$$

Where D is the *diffusion coefficient* [60], which is a measure of how freely the particles can diffuse within a medium. The diffusivity is a property of the temperature, viscosity and particle properties and can be calculated by:

$$D = \frac{k_b T}{6\pi\eta R} \quad (2.25)$$

Where k_b is the Boltzmann constant, T is the absolute temperature, η is the fluid viscosity and R is the Particle radius. For a population of diffusing molecules, the course molecular dynamics are given by Fick's second law

$$\frac{\partial \Phi}{\partial t} = D \nabla^2 \Phi \quad (2.26)$$

Where Φ is the particle concentration (mol/liter). The above equation can be solved in one dimension, treating the system as a thin diffusing film to give

$$\Phi(x, t) \propto \frac{1}{\sqrt{Dt}} \exp\left(\frac{-x^2}{4Dt}\right) \quad (2.27)$$

It can be shown that the probability distribution of molecules diffusing in such a system is equivalent to the concentration distribution described above to give:

$$P(x, t) \propto \frac{1}{\sqrt{Dt}} \exp\left(\frac{-x^2}{4Dt}\right) \quad (2.28)$$

Equation 2.28 shows that the probability distribution for particles undergoing free diffusion is Gaussian in nature. Typical values of diffusivity measured in human tissue are of order $3 \times 10^{-9} m^2/s$. This means that over the typical time scale of an MRI acquisition (50ms) a water molecule would be expected to move on average $30\mu m$. In this case, the value of the root mean squared distance is typically smaller than the dimensions of the microstructural features of the tissue. This means that the diffusion of water molecules in tissue is not completely free, and so equation 2.28 is no longer strictly true. However, the interactions between water molecules and their local cellular environment means that the diffusive motion of particles contains information about the tissue microstructure.

Using specially-adapted MR pulse sequences, it is possible to produce images with contrast based on the diffusion of water through the target tissue. Various analysis techniques have been developed that aim to extract parameters that relate to the diffusive motion, and hence the microstructural geometry, of the target tissue.

2.3.2 Acquisition

The most commonly-used pulse sequence for diffusion-weighted imaging is the pulsed-gradient spin-echo (PGSE), first proposed by Stejskal and Tanner in 1965 [61]. The PGSE sequence is formed by the addition of a pair of motion-sensitising gradients to a basic spin-echo sequence, either side of the 180° refocussing pulse. A basic schematic is shown in figure 2.11.




Figure 2.11: A schematic showing the basic pulse sequence structure of the pulsed-gradient spin-echo acquisition

The start of each acquisition is the 90° RF-pulse, which tips the net magnetisation vector into the xy -plane as [previously](#) described. After the 90° RF-pulse, the magnetisation vector is precessing in the transverse plane at the Larmor frequency given by equation [2.5](#). The next event in the sequence is the first of the two diffusion-sensitising gradients, with a gradient strength \mathbf{G} , rise time ϵ and duration δ . While the gradient is 'on', spins will precess faster or slower depending on their position along the axis the gradient is applied in - as the magnetic field they experience will be greater or reduced. For example, if the gradient is applied along the x -axis, the angular frequency in the rotating frame of reference is given by:

$$\omega = \gamma G_x x \quad (2.29)$$

Where ω is the angular frequency in the rotating frame, γ is the gyromagnetic ratio, G_x is the strength of the gradient (mT/m) and x is the position along the x-axis. The gradient is switched off after a time period δ , and the spins revert to precessing at the Larmor frequency. However, they will have acquired a phase-offset ϕ_1 that is directly related to their initial position in the sample. The phase offset is calculated by integrating equation 2.29 over the duration and shape of the gradient waveform G .

$$\phi_1 = \gamma \left(\int_0^{\delta+\epsilon} G(t)x(t)dt \right) \quad (2.30)$$

After the first motion-sensitising gradient is switched off, the 180° RF pulse inverts the phase shift of the spins, then after a time Δ following the onset of the first gradient the second motion-sensitising gradient is applied. The second gradient introduces a phase shift that acts to cancel out the phase shift introduced by the first gradient. The overall net phase shift accumulated is given by

$$\phi_{net} = \phi_1 - \phi_2 = \gamma \left(\int_0^{\delta+\epsilon} G_1(t)x(t)dt - \int_{\Delta}^{\Delta+\delta+\epsilon} G_2(t)x(t)dt \right) \quad (2.31)$$

The effect of the two motion-sensitising gradients and the 180° refocussing pulse is shown in figure 2.12.




Figure 2.12: A schematic showing the effect of the motion sensitising gradients and RF pulses of the PGSE sequence on the phase of a single spin in the xy plane. (a) shows the spin following the application of the 90° RF pulse, tipping the spin into the xy plane. (b) Following the first motion-sensitising gradient, the spin has gained a phase offset ϕ_1 . (c) Following the 180° RF pulse, the phase offset has been inverted. (d) Following the second motion-sensitising gradient, the spin gains a further phase offset ϕ_2 , which acts to cancel out the first phase offset due to the spin phase being flipped by the 180° pulse.

By making the assumptions of a negligibly small rise time ϵ , constant gradient strength and that the spin remains stationary during the application of the gradients equation 2.31 can be simplified to

$$\phi_{net} = \gamma G \delta (x_1 - x_2) \quad (2.32)$$

Where x_1 and x_2 are the particle positions during the first and second motion-sensitising gradients respectively. Equation 2.32 shows that there is a direct relationship between the net phase accumulated by spins and the distance moved between the diffusion gradients. If there is no diffusion, spins will remain in the same position throughout both gradients with $x_1 = x_2$. In this case, the phase shift accumulated by the second gradient will exactly cancel out the phase shift from the first gradient. On the other hand, if the molecules (spins) are diffusing through the tissue they will change position between the two gradients. If this happens, they will experience a different field strength between the second and first gradient. This means the phase shift from the second gradient will not perfectly cancel the first resulting in a net accumulated phase. With a whole population of spins in a sample undergoing the random thermal motion shown in 2.10, the result is a distribution of phase shifts over the whole population. This reduced coherence causes a reduction in magnitude of the transverse magnetisation and an attenuation of the signal echo. The amount of signal loss correlates with the diffusivity of the tissue, and it is by this mechanism that the diffusion-based contrast is generated for DW-MRI.

The theoretical signal amplitude can be calculated in terms of the distribution of phase shifts by

$$S = S_0 \int_{-\infty}^{\infty} P(\phi) \cos(\phi) d\phi \quad (2.33)$$

Where $P(\phi)$ denotes the distribution of phase shifts in the sample, and S_0 is the theoretical signal amplitude in the absence of diffusion weighting [62]. $P(\Phi)$ can be calculated from equation 2.28 to give:

$$P(\Phi) = \frac{1}{\sqrt{2\pi\langle\phi^2\rangle}} \exp\left(-\frac{\phi^2}{2\langle\phi^2\rangle}\right) \quad (2.34)$$

Where the quantity $\langle \phi^2 \rangle$ is the mean square phase shift of the population of spins. Substituting the phase shift distribution into equation 2.33, simplifies the signal equation:

$$S = S_0 \exp\left(-\frac{\langle \phi^2 \rangle}{2}\right) \quad (2.35)$$

Using equations 2.32 and 2.24, the commonly known diffusion signal equation can be derived [61]:

$$S = S_0 \exp(-\gamma^2 G^2 \delta^2 \Delta D) \quad (2.36)$$

$$= S_0 \exp(-bD) \quad (2.37)$$

The above relationship has been simplified using the previously-mentioned short gradient pulse approximation and so does not take into account the effect of the gradient ramp-up and ramp-down periods. In a real-world situation with a real MRI scanner, these effects should not be ignored and equation 2.37 becomes

$$S = S_0 \exp(-\gamma^2 G^2 (\delta^2 (\Delta - \delta/3) - \delta \epsilon^2/6 + \epsilon^3/30) D) \quad (2.38)$$

Where ϵ is the gradient ramp-up/ramp-down time. Equation 2.37 defines the acquisition *b-value*, as suggested by Le Bihan [20]. In diffusion MRI, the b-value is commonly used to describe the degree to which a scan has been diffusion-weighted. Acquisitions with a higher b-value have a stronger diffusion weighting and visa versa. Usually, the b-value is adjusted by changing the diffusion gradient strength G , while keeping all other parameters

constant. With unrestricted diffusion, the equations above still hold, as the probability distribution of molecular displacements is still Gaussian in nature. In this case, the signal is only dependent on the b-value of the scan, regardless of the combination of acquisition parameters that were used. However, when the diffusion is no longer free such as in biological tissue where the microscopic structure influences the motion of water molecules, this is no longer the case. With restricted diffusion, the gradient separation Δ can have a strong influence on the signal, as Δ effectively defines the length the water molecules are able to diffuse during the scan. With larger Δ values, water molecules are more likely to come into contact with the restrictive barriers within the tissue.

Based on the theory set out above, it is clear that there are many factors that can affect the diffusion MRI signal. The take-home message is that the diffusion signal is inherently influenced by the microstructure of the tissue. The acquisition parameters can be adjusted to fine-tune the diffusion acquisition depending on the desired application. It is possible for clinicians to use diffusion-weighted images to infer information about the structure of tissue without any post-processing via qualitative visual inspection. However, the problem with this is that the diffusion-sensitivity is not the only mechanism by which the contrast in the image is generated. Other confounding factors such as different T1 or T2 relaxation times will also affect the amplitude of the measured signal in different tissues. It is therefore often more desirable and more common practise to obtain quantitative information from DW-MRI, using various analysis techniques.




Figure 2.13: Figure showing (left) $b = 500$ and (right) $b = 1000$ images of the brain. Taken from <https://mrimaster.com/characterise%20image%20dwi%20.html>, accessed 30/05/2018

2.3.3 Analysis

2.3.3.1 Apparent Diffusion Coefficient

As previously discussed, in the case of free diffusion the measured diffusion signal is related to the diffusion coefficient D , which in turn is related to the mean squared displacement of a diffusing water molecule over a given time. However, in biological tissue the root mean squared diffusion distance is typically smaller than the dimensions of the microstructural features of the tissue. The effect of this is that the diffusion of the water molecules can be seen as being ‘hindered’ by the tissue. For this reason, the diffusion coefficient D is replaced with the *Apparent Diffusion Coefficient* (ADC) when studying biological tissue using DW-MRI. Equation 2.37 becomes

$$S = S_0 \exp(-b.ADC) \quad (2.39)$$

The value of the ADC is a measure of the degree to which diffusion is hindered in the tissue, which itself can be indicative of the microstructural characteristics. Equation 2.39 above shows that in order to calculate the ADC value from diffusion-MR data it is necessary to acquire a minimum of two separate images with different b-values. A commonly used approach is to acquire a b_0 image (an image without diffusion weighting) and one image with diffusion weighting. An estimate of the ADC value can then be calculated by rearranging 2.39:

$$ADC = \frac{\ln \frac{S_0}{S_1}}{(b_1 - b_0)} \quad (2.40)$$

Though this approach is mathematically correct, it is not commonly used due to its inaccuracy. One cause of this is that the true b-value of the acquisition is difficult to calculate. As previously discussed, the finite gradient ramp-up and ramp-down contribute to the true b-value. In addition, the imaging gradients used to spatially encode the signal also contribute to the b-value of the acquisition [63]. Another source of inaccuracy is the presence of Rician noise. The noise in complex MR data is normally distributed, whereas the magnitude is Rician distributed. The effect of the Rician distribution is that even in the absence of any true MR signal, the expectation value of the rectified signal is greater than zero. Therefore, at low SNR the amplitude of diffusion-weighted MR data is over-estimated, while the amplitude of the non-diffusion-weighted signal is unaffected. Therefore the ADC value calculated from the two diffusion weightings will be *underestimated* [64]. A more commonly-used approach for calculating ADC is to perform non-linear least squares fitting to the unlogged diffusion data using a gradient descent optimisation scheme [65]. A numerical approach such as this is more robust and therefore produces a more reliable value. By calculating the ADC value on a voxel-wise basis it is possible to generate a quantitative image where the contrast is generated solely from the apparent

diffusivity of the underlying tissue, this is often referred to as an 'ADC map' an example of which can be seen in figure 2.14.




Figure 2.14: Figure showing the ADC map, calculated from the b500 and b1000 images shown in figure 2.13. The map demonstrates how the ADC value relates to the microstructural properties of the underlying tissue. In this case, the ventricles and Cerebrospinal Fluid (CSF) spaces within the brain appear bright as they have a higher ADC value. These regions are fluid-filled, meaning water molecules are free to diffuse. Conversely, the white and grey matter appears darker, as these regions are densely packed with cells and neurons meaning diffusion is restricted and the ADC value is lower. Image taken from <https://mrimaster.com/characterise%20image%20dwi%20.html>, accessed 30/05/2018

Analysis of these maps can give a reasonable representation of changes in diffusion across tissues [66]. This technique has had great success upon translation to the clinic with widespread usage in a variety of disciplines. One of the most well-known applications being the early detection of ischaemic stroke, where the ADC value has shown to provide

a considerable increase in detection sensitivity compared to standard MR techniques. It has been shown that conventional MR sequences (T1 and T2 weighted imaging) can detect an infarct after approximately 6 hours, whereas a decrease in the ADC value is seen after only a few minutes [67, 68]. The ADC has also been shown to have useful applications in oncology where the value has been shown to be lower in malignant tumours than in the surrounding tissue, but recovers again after successful treatment with chemotherapy [69].

One of the main limitations of the ADC is that its value is effected by many physical processes occurring on the microscopic scale of tissue. It is impossible to infer anything about any of these individual processes or properties from the ADC value itself, and as such the ADC can be seen as a representative ‘summary’ of the diffusion-related properties of the tissue microstructure. One such property that influences the ADC value is the presence of perfused blood vessels. Blood flowing in the chaotic network of capillaries can be interpreted as rapid pseudo-diffusion. This will affect the ADC value calculated in that voxel and is a very challenging, if not impossible inverse problem to remove from the overall ADC value [66].

The ADC remains a highly relevant and useful tool in both a clinical and a research setting. It provides a useful insight into the microstructural geometry of tissue, and helps to bridge the gap between the micro and macro length scales. The simplicity of the ADC is one of its main advantages making it easy to implement, but also one of its main limitations. In order to gain a more detailed understanding of the tissue microstructure, it is necessary to apply more sophisticated techniques that allow for the separation of the different factors that contribute to the diffusion signal. This is the main motivating factor behind the development of more advanced models.

2.3.3.2 Diffusion Tensor Imaging

Diffusion Tensor Imaging (DTI) was a technique developed specifically for studying the characteristics of diffusion in the brain. It is commonly known that the white matter tissue within the brain consists of bundles of axons also known as ‘tracts’ [70]. The presence of these tracts cause the underlying tissue microstructure to be highly anisotropic in nature. The basic concept underpinning DTI is that the diffusion of water molecules will be less restricted along the principle direction of the fiber tracts in a given voxel. Previous diffusion imaging experiments had aimed to measure $D_{||}$ and D_{\perp} diffusion components to characterise the anisotropy of the tissue [71, 72]. Using a diffusion tensor to characterise anisotropy was first proposed by Bassler et al [73, 74], where the diffusion probability density function was defined as

$$\rho(x|x_0, \tau) = \frac{1}{\sqrt{|\mathbf{D}(\tau)|(4\pi\tau)^3}} \exp \left[\frac{-(x - x_0)^\tau \mathbf{D}(\tau)(x - x_0)}{4\tau} \right] \quad (2.41)$$

where \mathbf{D} is a 3×3 covariance matrix

$$\begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} \quad (2.42)$$

Where $D_{xx,yy,zz}$ represent the diffusion coefficient measured along the x, y or z-axis respectively. The off-diagonal terms $D_{xy...}$ represent the correlation between the diffusion coefficients along the principle directions. In the special case of isotropic diffusion, the three diagonal components $D_{xx,yy,zz}$ are all equal and the off-diagonal terms are 0 and can therefore be ignored. In the case of anisotropic diffusion, the diagonal components are no

longer equal, and the off-diagonal components can no longer be ignored.

A mathematical technique known as 'eigen decomposition' (or matrix diagonalisation when applied to a square matrix) can be applied to the diffusion tensor matrix in order to extract the eigenvectors and eigenvalues. The three eigenvectors $\epsilon_1, \epsilon_2, \epsilon_3$ correspond to the three principle diffusion directions, and the eigenvalues $\lambda_1, \lambda_2, \lambda_3$ are directly proportional to the root mean squared diffusion displacement in the relative direction. When fitting the diffusion tensor model to diffusion data, the eigenvalues can be constrained in order to capture different diffusion geometries.

The diffusion tensor is often visualised as an ellipsoid, where the size of the ellipsoid along each principle direction is defined by the eigenvalues $\lambda_1, \lambda_2, \lambda_3$ of the covariance matrix, and the orientation of the ellipsoid defined by the corresponding eigenvectors $\epsilon_1, \epsilon_2, \epsilon_3$ (see figure 2.15 [75]).




Figure 2.15: Schematic showing how a diffusion tensor can be visualised as an ellipsoid, with the principle axes direction/dimensions given by the eigenvalues of the diffusion tensor matrix.

Other parameters that may be derived from the calculated diffusion tensor such include the mean diffusivity and fractional anisotropy, which have been demonstrated as effective biomarkers for ischaemic brain injury [76] and Multiple Sclerosis [76]. The mean diffusivity is calculated as the average of the three eigenvalues

$$MD = \frac{(\lambda_1 + \lambda_2 + \lambda_3)}{3} \quad (2.43)$$

and describes the rotationally invariant magnitude of water diffusion within a given voxel (the ‘ADC’ value). Fractional anisotropy aims to describe the degree of diffusion asymmetry in a voxel:

$$FA = \sqrt{\frac{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_1 - \lambda_3)^2}{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)}} \quad (2.44)$$

The fractional anisotropy is a scalar value that varies between 0 and 1, where 0 indicates that the diffusion ellipsoid is spherical (since $\lambda_1 = \lambda_2 = \lambda_3$) and diffusion within the voxel is perfectly isotropic. As the diffusion within a given voxel becomes increasingly anisotropic, the value of FA approaches 1. The mean diffusivity and fractional anisotropy have found widespread application in brain imaging. However, these metrics lack the sensitivity to detect more subtle changes to the tissue. As such, more advanced models are required in order to access tissue metrics that are more sensitive to pathological changes.

2.3.3.3 Intra-Voxel Incoherent Motion

As previously discussed, one of the factors that can affect the value of the ADC in tissue is the perfusion of tissue via blood flow in the capillaries. It has been shown via experiments in an MR phantom that the measured ADC is close to the known value as long as there is no perfusion. In *in-vivo* studies the value is often higher than expected, and this discrepancy is attributed to incoherent motion in the microvasculature (capillary network) [77]. Intra-Voxel Incoherent Motion is a two-compartment model of diffusion proposed by Le Bihan et al [77] in order to account for the effect of pseudo-diffusion on the overall diffusion signal.

The term 'pseudo-diffusion' arises from the pseudo-random orientations of the perfused vessels in the capillary bed, shown in figure 2.16.




Figure 2.16: Image showing the orientations of the capillaries. The image is a zoomed-in section of a segmented vessel network from a subcutaneous LS174T tumour model.

The equation giving the signal attenuation due to diffusion (equation 2.39 is modified when there is flow in the capillaries. The modification takes the form of an additional term, F , that represents the pseudo-diffusion:

$$\frac{S}{S_0} = e^{(-b.ADC)}.F \quad (2.45)$$

Where $F \leq 1$. By making certain assumptions about the geometry of the capillary network the value of F can be calculated.

The capillary network can be modelled as a series of straight segments of length l with an average flow velocity $\langle v \rangle$ [77]. Combining this assumption with a model of biological tissue containing two compartments; a perfusion fraction f of water flowing with the blood in the capillaries and a fraction $(1 - f)$ of static water in the intra/extracellular space. This combination enables a new relationship between the diffusion and detected signal to be derived:

$$\frac{S}{S_0} = (1 - f)e^{(-b.ADC)} + fe^{(-b.D^*)} \quad (2.46)$$

Where D^* is the diffusivity associated with a combination of the pseudo-diffusion and the diffusion of water in whole blood. The increase in complexity of the IVIM model compared

to the ADC means that additional data is required in order to extract values from the model.

Usually, the value of the pseudo-diffusion coefficient D^* is much larger than the diffusion coefficient ADC , therefore at any given diffusion weighting the pseudo-diffusion term decays much faster than the diffusion term. Therefore, at higher b-values the total diffusion signal is dominated by the effects of diffusion and the effects of pseudo-diffusion are negligible. Therefore, in order to estimate the IVIM model parameters it is necessary to acquire diffusion MRI data with low diffusion-weightings (typically $b \leq 150s/mm^2$) as well as high diffusion-weightings ($b \geq 200s/mm^2$).

The most commonly-used method for estimating the IVIM parameter values is by using least-squares fitting. This is a well-documented technique that involves iteratively changing the model parameters in order to minimise the difference between the predicted signal values and the measured values [78].

One of the drawbacks of the IVIM model is the long-running debate over whether or not the perfusion parameters produced by the model correlate with the perfusion values measured via other established techniques such as using intravascular contrast agents [79]. Intravascular tracers measure the perfusion value by detecting the rate at which the tracer is taken up by the tissue, rather than via a direct measurement of the blood flow. The counter-argument to this is that the rate at which any given tracer is taken up is directly dependent on a number of factors, with the blood flow velocity and microvascular geometry included. It can therefore be argued that there should be a degree of correlation between the perfusion parameters measured through IVIM and those measured using contrast agents, though there are examples where this is not the case [80]. Another limitation of the IVIM model is that it has been demonstrated that the parameter values have poor repeatability [81]. A lack of repeatability is a severe limitation as it makes multi-centre

studies very difficult, and may even cast doubt over the results of longitudinal imaging studies. One of the possible reasons for this lack of repeatability is that the low b-value measurements can be corrupted by tissue motion effects. Incoherent motion of the tissue being imaged from processes such as glandular secretion can mimic the signal attenuation associated with microvascular perfusion.

Despite its limitations, the IVIM model has been shown to be useful in both a preclinical and clinical setting for separating the effects of diffusion and microvascular pseudo-diffusion in biological tissue. The increased complexity of the model means it is better able to fit the signal decay curve than the simple mono-exponential approach of the ADC. However, it does not provide a means of effectively characterising the microstructural environment of tissue. The perfusion measurement provided by the IVIM is a useful biomarker in its own right, however it is just one of many possible microenvironmental properties that may be sensitive to the changes caused by new cancer therapies. In order to access these other parameters, more sophisticated models of tissue diffusion are required.

2.3.3.4 Multi-Compartment Models

The ADC and IVIM methods can both provide useful insights into the structure of tissues, and go some way towards characterising their function. However as previously mentioned, the parameters produced by these techniques may be sensitive to multiple microstructural properties that can reduce the sensitivity or specificity of these techniques. Biophysical models aim to provide a simplistic geometrical representation of the underlying target tissue. These generally involved modelling the tissue as discrete ‘compartments’ such as the intracellular and extracellular space. The MR signal corresponding to these models is calculated, then the inverse problem solved to estimate the model parameters that match the measured MR data. These model parameters tend to be analogous to tissue metrics

such as cell radius, intra/extracellular volume fraction etc.

Stanisz’ Optic Nerve

One of the first multi-compartment models was proposed by Stanisz et al [82]. The model aimed to represent the diffusion of water within a bovine optic nerve using three discrete compartments: The axons, glial cells and the extracellular space between them. Diffusion within the axons is represented by a prolate ellipsoid, the glial cells by a sphere and the extracellular diffusion modelled by a hindered isotropic diffusion model. The model also allows for the effect of permeability by allowing exchange between compartments [83]. Due to the complex nature of this model, a large quantity of high-quality data is required in order to fit it. In the study, a sample of bovine optic nerve was scanned *ex-vivo* oriented parallel and perpendicular to the diffusion encoding gradient, and data was acquired with a signal-to-noise (SNR) ratio > 1000 . The model was fitted to this data, and estimates of various tissue properties were calculated including glial cell radius (sphere radius), axon diameter and length (short and long axis of ellipsoids) and both intra- and extracellular diffusivities. However, it is impossible to translate this model to *in-vivo* data as the typically-achievable SNR is much lower (< 50).

Ball and Stick

Behrens et al [84] proposed an adapted version of Stanisz’ model, known as the ‘Ball and Stick’ model, which was intended to serve as the simplest possible model representing diffusion in white matter. The model consists of two compartments with the ‘Ball’ representing the isotropic unrestricted diffusion in the extracellular space and the ‘stick’ representing the directionally-restricted diffusion in the axons. The signal is represented as a weighted linear sum of the two compartmental signal models:

$$S = S_0 \left(f \exp(-bd(\hat{\mathbf{n}} \cdot \hat{\mathbf{g}})) + (1 - f) \exp(-bd) \right) \quad (2.47)$$

Where f is the volume fraction associated with the ‘stick’ compartment (intracellular volume), d is the diffusivity parallel to the stick orientation and within the ball compartment, $\hat{\mathbf{n}}$ is the principle stick orientation and $\hat{\mathbf{g}}$ is the diffusion encoding gradient direction. By adding additional ‘stick’ compartments, the ball and stick model framework can be easily extended in order to model multiple fibre populations within a single voxel [85]. However, the model is unable to access tissue features such as cell/axon radius.

CHARMED

Following on from the development of the ‘ball and stick’ model, a more complex multi-compartment model of diffusion was developed by Assaf et al [86]. The Composite Hindered and Restricted Model of Diffusion (CHARMED) introduces a more accurate portrayal of the intra-axonal tissue structure. The intra-axonal (intracellular) space is represented by either one, or two populations of axons modelled by packed cylinders, with a single orientation per population. The axon diameters are fixed to expected values based on histology. The extracellular space is modelled as anisotropically hindered diffusion, under the assumption that diffusion in the extracellular space is more hindered perpendicular to the axons. Fitting the model to data enables the axon orientation(s), intra/extracellular diffusivities and compartmental volume fractions to be estimated. The main drawback of the model is that due to its increased complexity it requires a large volume of diffusion-weighted data acquired in numerous directions and with multiple diffusion-weightings.

AxCaliber

The CHARMED model was further developed by Assaf et al [87] with the introduction of the AxCaliber technique. AxCaliber extends the CHARMED model by estimating the distribution of axon radii with an assumed orientation. The distribution of axon radii was assumed to follow a Gamma distribution based on previous histological analysis [88]. A large quantity of diffusion-weighted data is required in order to fit the AxCaliber model. The diffusion-encoding direction must be perpendicular to the principle fiber orientation with a combination of different diffusion times and b-values. The AxCaliber technique has been proven to provide a highly accurate estimation of axon diameter in *ex-vivo* tissue samples, and has also been applied *in-vivo* in rats [89]. The main limitations of AxCaliber are the large quantity of data required, and the need for prior knowledge of the principle fibre orientation.

ActiveAx

The problems with the AxCaliber framework were addressed by Alexander et al [90] with the ActiveAx model. The model presents a simpler representation of the intracellular space - as randomly packed cylinders with identical radii and orientation. The extracellular space is represented in the same way as with the CHARMED model, with anisotropically-hindered diffusion. The model features an additional compartment in order to account for partial volume effects arising from the CSF - an isotropic diffusion compartment (a simple mono-exponential). The model is capable of estimating the axon diameter index (correlated with the volume-weighted diameter), fibre orientation and intra/extracellular volume fractions. Despite the fact the model does not attempt to estimate the entire axon diameter distribution, the ActiveAx framework still requires a large volume of high-quality diffusion-weighted image data in order to provide a robust estimate of the model parameters. However, it may be said that requiring a large volume of data is a fairly common limitation within the field of multi-compartment diffusion modelling [91].

NODDI

An extension to the ActiveAx framework was proposed by Zhang et al [92] that aimed to improve the robustness of the model in the presence of orientation dispersion. The intracellular compartment is identical to that of the ActiveAx model, but the orientations of the packed cylinders no longer identical. The orientations instead are assumed to follow a Watson distribution. The model was shown in [92] to be capable of estimating the axon diameter index and orientation dispersion with *in-vivo* DW-MRI data from a human brain. This model formed the foundation of the widely-used Neurite Orientation Dispersion and Density Imaging (NODDI) model developed by Zhang et al [21]. The NODDI framework models the intracellular compartment as sticks with Watson-distributed orientations, meaning the axon radius is not estimated. The relative simplicity of this model means it can be fitted to *in-vivo* data acquired on a clinical scanner with two b-value shells. This has led to the technique being used in a variety of studies [93–96].

VERDICT

All of the models discussed above were developed to measure microstructural features of the brain, or nervous tissue. Compartment models for use in other tissues such as tumour tissue are somewhat lacking. The bi-exponential IVIM technique described previously has been applied numerous times in an attempt to characterise the vascular properties of various types of tumour [97–100]. Although IVIM is able to successfully separate the additional signal from vascular water, the remaining signal corresponding to the intracellular and extracellular spaces is represented by a mono-exponential.

One of the first attempts to extend the use of compartmental/biophysical diffusion models for microstructural characterisation of tumours is VERDICT (Vascular, Extracellular and Restricted Diffusion for Cytometry in Tumours). The VERDICT framework characterizes water diffusion in the intracellular, extracellular and vascular compartments of

tumour tissue. The foundation for the VERDICT model comes from work carried out by Panagiotaki et al in 2012 [6], where a full taxonomy of mathematical diffusion models was evaluated and compared. The results of this in-depth study show that in general models with three separate compartments outperform models with two. A three-compartment model was thereby selected for the VERDICT framework, consisting of separate vascular, extracellular and intracellular compartments. All three compartments contribute to the total calculated diffusion signal, which is given by a volume-weighted linear sum:

$$S = \sum_{i=1}^3 f_i S_i \quad (2.48)$$

Where S_i is the calculated signal from each compartment and f_i is the volume fraction associated with that compartment ($\sum_{i=1}^3 f_i = 1$). An appropriate mathematical model of the diffusion signal must be selected for each compartment, that is suitably representative of the microstructural geometry of the tumour.

The intracellular volume compartment describes the signal resulting from spins diffusing within the cells of the tumour. The diffusion signal from the intracellular compartment is dominated by the effect of restricted diffusion, which must be accounted for by the mathematical model selected. Various models of restricted diffusion have been developed to represent a range of different restriction geometries [6, 62, 101, 102]; including spheres, cylinders, ellipsoids, astrocyinders and astrosticks. A different restricted compartmental geometry can be chosen depending on the expected characteristics of the tumour in order to achieve the most reliable result. As the majority of tumours lack the organised anisotropic structures such as the axons found in the brain, the most often-used restriction geometry is the sphere. The mathematical model representing the signal generated from spins diffusing within a restrictive spherical medium was developed by Murday and Cotts [102].

A three-dimensional diffusion tensor (DT) model (see section 2.3.3.2) is used to model both the extracellular extravascular (EE) and intravascular compartments. The eigenvalues of the tensor model describing the extracellular space are most often constrained to be equal, as the extracellular space is expected to be isotropic. In contrast, when the tensor model representing the vascular pseudo-diffusion is fitted to the data the eigenvalues are left unconstrained, in order to capture the anisotropic nature of the microvascular bed. In more recent work, alternative models have been considered for representing the vascular compartment, such as the ‘astrosticks’ model (see below). The modular nature of the technique gives the flexibility to swap models in and out depending on the application.

Due to the large number of parameters that are fitted to the diffusion data by the VER-DICT model, a greater number of different b-value scans are required in order to achieve a reliable quantification of the tissue microstructure. For pre-clinical *in-vivo* imaging, Panagiotaki et al [9] used a diffusion scheme of 46 different b-values, in 3 orthogonal directions. The scan time (2.5 hours) was therefore considerably longer than a basic diffusion-weighted sequence. However, using an echo-planar acquisition scheme on a clinical scanner it has been possible to reduce the acquisition time to around 15 minutes.

Taxonomy of Model Compartments

The various compartmental models used in the techniques described above were summarised and collated in a study by Panagiotaki et al [6] who compiled a taxonomy of compartmental models. The most commonly-used compartment models are outlined in table 2.1, along with their associated signal models and the tissue compartments they typically represent.

Model Name	Signal Equation	Parameters	Intracellular	Extracellular	Intravascular
Ball	$S = \exp(-b\hat{G}^T D \hat{G}), D = dI$	d	No	Yes	Yes
Sphere	$D = -\frac{1}{b} \ln \left[2 \frac{1 - \cos(\gamma\delta l)}{(\gamma\delta l)^2} + 4(\gamma\delta l)^2 \Sigma e^{(-n^2\pi^2 D \frac{\delta}{l^2})} \frac{1 - (-1)^n \cos(\gamma\delta l)}{(\gamma\delta l)^2 - (n\pi)^2} \right]$	d, R_s	Yes	No	No
Stick	$S = \exp(-bd(\mathbf{n} \cdot \mathbf{G})^2)$	d, θ, ϕ	Yes	No	Yes
Zeppelin	$D = \alpha \mathbf{m} \mathbf{m}^T + \beta I$	$d_{ }, d_{\perp}, \theta, \phi$	No	Yes	No
Cylinder	$S = \exp(L_{ }(\mathbf{G} \cdot \mathbf{n})^2) \exp(L_{\perp}(\mathbf{G} \cdot \mathbf{G} - (\mathbf{G} \cdot \mathbf{n})^2))$	d, θ, ϕ, R	Yes	No	No
Tensor	$D = d_{ } \mathbf{m} \mathbf{m}^T + d_{\perp} \mathbf{n}_{\perp 1} \mathbf{n}_{\perp 1}^T + d_{\perp} \mathbf{n}_{\perp 2} \mathbf{n}_{\perp 2}^T$	$d_{ }, d_{\perp 1}, d_{\perp 2}, \theta, \phi, \alpha$	Yes	Yes	Yes
Astrosticks	$S = \pi^{1/2}/2(2 G (L_{\perp} - L_{ })^{1/2})^{-1} \exp(G ^2 L_{\perp}) \phi(G (L_{\perp} - L_{ })^{1/2}), R = 0$	d	Yes	No	Yes
Astrocylinders	$S = \pi^{1/2}/2(2 G (L_{\perp} - L_{ })^{1/2})^{-1} \exp(G ^2 L_{\perp}) \phi(G (L_{\perp} - L_{ })^{1/2}), R > 0$	d, R_c	Yes	No	Yes
Dot	No attenuation	None	Yes	No	No

Table 2.1: A table showing the formulation of the various compartmental models discussed in the previous section. The mathematical form, fitted parameters and most commonly-represented compartments are shown. S is the diffusion signal, b is the b-value, \hat{G} is the gradient direction, d is the diffusivity, I is the identity tensor, γ is the nuclear gyromagnetic ratio, g is the gradient strength, δ is the gradient duration, l is the coil length, \mathbf{n} is the principle fiber direction, \mathbf{G} is the gradient vector, L is the cylinder radius, R_s is the sphere radius, θ, ϕ are the tensor orientations, $d_{||}$ is the diffusivity parallel to the principle fiber direction, d_{\perp} is the diffusivity perpendicular to the principle fiber direction [6].

2.3.4 Validation and Microstructural Modelling

As shown by the compartmental diffusion models described [above](#), the complexity of diffusion-MRI microstructural imaging techniques is ever-increasing. However, the validation of these techniques via comparison with an established ground truth is essential if they are to be translated into widespread clinical practice. Such a ground truth can be achieved by numerous approaches that can broadly be divided into three categories: histological validation, numerical phantoms and physical phantoms.

2.3.4.1 Histological Validation

One of the most popular methods for validating microstructural imaging methods is by comparison with histology. This enables the tissue in question to be directly visualised and measured in order to establish an accurate ground truth. Histology, also known as *microanatomy* is the study of biological samples using a microscope. Tissue samples are taken, then sliced into thin sections/slices (often using a cryostat). Biological tissue in its unaltered state has very little natural contrast when viewed with a light microscope. Tissue staining is used to give contrast, and specific stains have been developed that target particular features of interest. The most common stains used for studying tissue microstructure are hematoxylin and eosin (H&E). These stain and provide contrast to the cell nuclei and cytoplasm respectively, allowing the tissue microstructure to be visualised.

A large portion of the histological validation of diffusion imaging metrics has been performed on brain-imaging applications. Much of the compartmental models developed for use in the brain focus on recovering the orientation, density or radius of the axon tracts within white matter[[21](#), [84](#), [86](#), [87](#), [90](#)]. As a reflection of this, much of the histological validation work that has been performed has been focused on assessing the accuracy of these methods in measuring such parameters.

Choe et al [103] compared the metrics calculated diffusion metrics such as diffusion tensor orientation, fractional anisotropy and diffusivities from diffusion-weighted data acquired from a monkey brain. They then investigated the correlation between these and similar metrics calculated from light microscopy images of histological slices, similar studies in rats and *ex-vivo* human brain have been performed by other groups [104, 105]. More recently, Gangolli et al [106] performed a similar validation of fibre orientation measurements via q-space imaging by co-registering 3D volumes of histology and diffusion data. This was performed in regions of *ex-vivo* human brain tissue with highly complex white matter structures, as regions with crossing and fanning fibres present a greater challenge when attempting to measure fibre orientation distributions using diffusion MR models. Other recent work has been focused on imaging histology samples with alternative optical imaging techniques such as confocal microscopy [107, 108], optical coherence tomography (OCT) [109] and polarized light imaging [110, 111].

One of the main limitations of histological validation is the extrapolation of the *ex-vivo* imaging conditions and tissue properties to those of *in-vivo*. It has been shown in various studies that the directional dependence and anisotropy of the diffusion signal in white matter is preserved between *in-vivo* and *ex-vivo* [112, 113]. However, it is not definitively known whether other diffusion metrics and tissue parameters can be validated via histology due to the potential microstructural perturbation caused by the fixing and dehydrating steps involved in the sample preparation process. It is also not known whether the same extrapolation between *ex-vivo* and *in-vivo* can be made in other tissues and pathologies such as cancerous tissue.

2.3.4.2 Physical Phantoms

An alternative to histological validation is the use of physical hardware phantoms. Physical phantoms consist of model systems constructed to represent a chosen tissue, whose microstructural properties are known. These systems can then be imaged using a real scanner to assess the accuracy of new DW-MRI techniques. As such, hardware phantoms can be used to bridge the gap between theoretical/computational models and *in-vivo* imaging.

Phantoms can take various forms, depending on the technique they are aiming to validate. They can broadly be split into two main categories: liquid phantoms and structural phantoms. Liquid phantoms are typically MR-visible fluids that have known diffusion properties. Structural phantoms usually consist of MR-visible fluid contained within MR-invisible structures designed to mimic the microstructural properties of certain tissues. These structures restrict the diffusion of water in the same way as the barriers in biological tissue.

The most readily-available and popular liquid phantom is water. It is inert, stable and very easy to obtain. However, the ADC of water is higher than the ADC of *in-vivo* tissues, and is highly temperature dependent [114]. In addition, water has a low viscosity, meaning that large containers of water may produce flow artefacts. One attempt to overcome these issues is to submerge a small container of water within iced water, the ADC of water at $\approx 0^\circ$ is approximately equal to that of living brain tissue and the iced water keeps the temperature stable over long periods of time. This makes the iced-water phantom a useful validation tool for basic diffusion techniques and multi-centre comparisons [115–117]. One limitation of ice-water phantoms however is that they have a single fixed ADC value. This can be overcome using aqueous solutions as phantoms. Solutions of water and albumin [118], water and sucrose [119] and agar gels [120] have been proposed as methods for tuning the ADC value of liquid phantoms. More advanced solutes such as polyvinylpyrrolidone

(PVP) [121] and polyethylene glycol (PEG) [122] have also been shown to be effective. Alternative liquids have also been proposed including various alkanes [123], and water-fat emulsions such as dairy cream [124–126].

Liquid phantoms are useful for validating simple diffusion measurements such as the ADC value. They are usually easy to manufacture and a relatively inexpensive solution. However, they are unable to represent the complex structures found in biological tissue that leads to the restricted nature of diffusion in such tissues. In addition diffusion in liquids is isotropic, meaning liquid phantoms can not be used to validate anisotropic diffusion models such as diffusion tensors or many compartmental models.

Structural phantoms can be constructed that aim to represent a specific feature of tissue. By surrounding diffusing molecules with an MR-invisible material they are able to recreate the effects of anisotropic, hindered and restricted diffusion. Structural phantoms can consist of synthetic phantoms or biological tissue analogues. An example of a biological analogue is the use of asparagus to mimic the anisotropic diffusion properties of white matter or nerve fibres [127]. Multiple synthetic phantoms have been developed that aim to recreate the structure and properties of white matter. Water-filled silica microcapillaries with specific radii have been used to investigate the sensitivity of novel gradient waveform PGSE sequences to small radii, and the effect of changing the frequency of oscillating-gradient spin-echo acquisitions on ADC values [128, 129]. A growing area of research involves the use of various polymers for constructing DW-MRI phantoms. Amongst these, poly(η -caprolactone) (PCL) has been suggested as a promising option. PCL is non-toxic, is easy to process and can be woven into nanofibres via a process known as electrospinning. This has lead to the use of PCL to construct phantoms for validating diffusion-weighted imaging of the heart [130] and brain [131, 132]. A number of structural phantoms have also been developed for applications other than brain imaging. A phantom consisting of polystyrene latex microspheres packed into tubes was used to validate measurements of

pore size and surface/volume ratio using oscillating-gradient spin-echo sequences [133]. In more recent work, McHugh et al [134] used coaxial electrospraying of PCL and PEG to create ‘roughly spherical, micron-sized core-shell polymer “cells”’ as a phantom designed to mimic the cellular structure of tumours.

Physical phantoms (both liquid and structural) are useful tools for validating diffusion-weighted imaging metrics, and have been proven essential for assessing inter-centre variability. The advantages of using physical phantoms is that they take into account the acquisition-specific and scanner-specific effects. However, both liquid and structural phantoms are simple representations of complex biological tissue. They can also be difficult to manufacture, and can suffer from phantom-to-phantom variability.

2.3.4.3 Numerical Phantoms

A numerical phantom is a computer model constructed to represent a physical object/phenomena that enables computerised testing and validation. One of the key advantages of numerical phantoms is that they are highly flexible, allowing for the validation and optimisation of biophysical models with near-infinite combinations of tissue properties and acquisition parameters [135]. Numerical phantoms can broadly be split into 3 categories: Matrix formalisms, finite difference methods and Monte Carlo simulations.

Matrix formalism methods involve solving for the eigenmodes of the diffusion propagator [136]. This allows the expected echo attenuation to be calculated mathematically. However, in practise this is only possible for substrates with simple geometries (such as spherical pores, parallel planes, cylinders) [137]. Solving for the eigenmodes of the propagators for more complex geometries can be very difficult, or impossible. A finite-difference approach aims to solve the Bloch-Torrey equations numerically for substrates with more

complex geometries, including the effects of barrier permeability. Finite difference approaches have been used to investigate the relationship between cellular properties and the ADC value [138]. A finite difference approach has been used in various other studies to predict the relationship between tissue properties and signal attenuation in substrates with varying geometries [139–144]. However, both matrix formalisms and finite difference methods struggle to cope with highly disordered and complex microstructural geometries. They are also unable to take into account additional contrast mechanisms such as exchange and T1/T2 relaxation. By far the most flexible and popular numerical phantom method is Monte Carlo simulations.

Monte Carlo simulations aim to recreate the process of individual spins diffusing through complex substrates, taking into account the interactions between the spins and the permeable/impermeable barriers. The simulated spins, known as ‘walkers’ are initiated within the substrate, and are then propagated through the substrate. The resulting diffusion signal is calculated from the accumulated phase changes of each walker resulting from their changing position relative to simulated diffusion gradients. Walkers are propagated through the substrate following a series of discrete time steps, the length of which is usually of fixed length according to the Einstein diffusion equation in three dimensions:

$$\Delta x = \sqrt{6D\Delta t} \quad (2.49)$$

Where D is the diffusivity and Δt is the length of the time step (also known as the temporal resolution). The direction of the propagation step is chosen randomly, with the walker propagating homogeneously in any direction with equal probability. Once a new step has been calculated a check is performed to assess whether the new step would cause the walker to cross a boundary within the simulation substrate, such as the boundary between the intracellular and extracellular space. If the new step is deemed to cross a

boundary, the interaction between the walker and the boundary is then handled. At a boundary, walkers can either be reflected, or transmitted through the boundary depending on whether or not the simulation includes the effect of permeability. In Monte Carlo, all step generation and boundary interaction processes are driven by random numbers. For a barrier with permeability, the decision to transmit through or be reflected is decided by a random number generator drawing numbers from a uniform distribution. The whole step generation and checking process is repeated until the specified simulation time is reached, at which point the next walker is propagated. Once all of the walkers have been propagated, the diffusion MR signal is generated based on the accumulated phase of each walker. Each walker gains or loses phase based on its time dependant position, and the diffusion gradients specified for the acquisition. A more detailed description of this process will be given in the methods section of chapter 4.

Monte Carlo simulations have been extensively used for validating basic diffusion models as well as diffusion microstructural imaging techniques. Early work investigated the effect of water exchange on the accuracy of biophysical models in various substrate geometries [82, 145, 146]. The vast majority of Monte Carlo modelling studies in recent years have been focussed on modelling diffusion in white matter. Various microstructural measurements of the axons such as axon diameter, axon bending and undulation as well as tractography algorithms have been validated using Monte Carlo simulations [147–152].

One of the main drawbacks of Monte Carlo simulations is the computational expense. Often, Monte Carlo simulations take a significant amount of computational time to run, since a large number of walkers must be propagated, each with a large number of simulation steps and boundary interactions to be calculated. To achieve a reasonable degree of accuracy, a typical Monte Carlo simulation may take minutes, or even hours to complete. Another drawback of Monte Carlo simulations, as will be discussed in the following chapters, is their reliance on simplistic models representing biological tissue. Without real

segmented tissue substrates, simulation substrates instead tend to represent biological tissue as simple geometric shapes such as spheres, cylinders etc. This simplistic representation of tissue limits the ability of Monte Carlo simulations to capture diffusion effects resulting from the complex microstructural geometry of biological tissue.

2.4 Summary

The development of modern cancer therapeutics that can alter the function of tumour tissue without necessarily changing the volume of the tumour has highlighted the need to develop more sensitive non-invasive measures of tumour physiology. One promising candidate is compartmental modelling of diffusion MRI data. Compartmental models such as VERDICT MRI can enable histologic features of tumour tissue to be measured non-invasively, and with full coverage and repeatability. However as the mathematical models involved become increasingly sophisticated and complex, the need for thorough validation becomes increasingly essential. In the following chapters, the capabilities of the VERDICT framework are investigated, and a Monte Carlo validation framework is demonstrated. This is followed by an exploration of the role that machine learning algorithms can play in extracting useful tissue features from diffusion data.

Imaging Microstructural Changes using VERDICT MRI

3.1 Motivation

As previously mentioned, VERDICT MRI is a novel imaging technique that aims to enable access to histologic tissue features non-invasively. The three-compartment signal model is fitted to the signal vs b-value curve from a multi-shell imaging scheme. Fitting the model returns parameters such as the cell radius, compartment volume fractions and associated diffusivities. The aim of the work in this chapter was firstly to demonstrate that the VERDICT model provides a better fit to diffusion data than simple models like the IVIM and ADC, and secondly to show that the parameters are more sensitive to microstructural perturbations than simpler models.

3.2 Introduction

Within the broad scope of cancer research, being able to quickly and accurately quantify the effectiveness of an administered therapy is currently one of the most critical themes of research. In a clinical setting, the current gold standard for assessing therapeutic response is the RECIST criteria, whereby changes in the bulk volume of primary tumour nodes are used to determine whether therapy has been successful [7, 8]. Often the bulk tumour volume is a sensitive predictor of therapeutic response. However, some newly developed cytostatic cancer treatments do not necessarily cause significant changes in tumour volume - meaning such measurements will not be predictive of therapeutic effect [153]. In addition, changes in bulk volume usually take a significant amount of time to occur, with the time between the baseline and follow-up scans usually measured in weeks or months. The risk of tumour metastasis and/or mutation within such time-scales is thereby increased, presenting a greater risk to the patient. This is one of the main motivations for the development of new biomarkers for cancer diagnosis and assessment of treatment response.

As mentioned in the previous chapter, diffusion MRI has been suggested as a promising alternative to bulk volume measurements or invasive techniques such as tumour biopsy and histology. The most commonly-used diffusion MRI technique in a clinical setting is the ADC, which represents the degree to which diffusion within the tissue is hindered by the microstructural environment. Although such simple models like the ADC can infer certain differences in tumour pathophysiology, the exact relationship between the model parameters and the tumour microstructure is equivocal. There are many different factors that could effect the value of the ADC, including the cellularity, cell radius, intra/extracellular diffusivities and vascular properties for example. Recently, work has been focused on developing more complex biophysical models, that aim to separate the effects of these confounding factors allowing histologic tissue features to be accessed non-invasively.

In 2014, Dr Eleftheria Panagiotaki et al developed a novel new diffusion MRI technique for measuring the histologic features of tumour tissue non-invasively called VERDICT (Vascular Extracellular and Restricted Diffusion for Cytometry in Tumours) MRI [9]. The VERDICT model aims to more accurately characterise the tumour microenvironment by employing a multi-compartment modelling approach whereby the tissue is represented as three discrete compartments: the intracellular, extracellular and vascular. Each compartment has it's own associated signal model with parameters describing the geometry of the diffusion propagator. These parameters (including the cell radius, intra/extra/vascular volume fractions and intra/extracellular diffusivities), are intended to correspond to real tissue properties that are potentially useful cancer biomarkers. To date, VERDICT has demonstrated the ability to distinguish between tumour and healthy tissue, between different tumour types and to accurately quantify cell size and density [9, 154] in both prostate and colorectal cancer.

The aim of this chapter is to demonstrate the capabilities of the VERDICT technique in detecting changes in tissue microstructure. Firstly, in detecting the microstructural changes that occur when fixing tissue with formaldehyde. Secondly, in a comparison of the performance of the VERDICT 'BallSphereStick' and 'BallSphereAstrosticks' models for detecting changes induced in a mouse model of glioblastoma treated with Temozolomide.

In order to give context and background, further details on tissue fixation and Temozolomide are given below.

3.2.1 Tissue Fixation

The process of tissue fixation is most commonly involved in the preparation of samples for histological analysis. The purpose of fixation is to halt the post-mortem decay and

degradation of tissue by stopping all chemical and biological processes that occur in the tissue. Changes in osmotic pressure, autolysis and bacterial colonisation can all severely damage and distort the microscopic structure of tissue, making *ex-vivo* analysis of non-fixed tissue futile.

Tissue fixation has extensively been used in the medical imaging field, as fixed tissue samples can be imaged with much longer acquisition times, and without the effects of motion. Typically, fixation is performed by immersion or perfusion of the tissue with a fixative chemical, meaning the fixing process occurs outside-in or inside out respectively. There are a number of different fixative chemicals that can be used. The most prevalent non-aldehyde fixatives are alcohol based such as ethanol, methanol or acetone. Alcohol fixation was commonplace before the discovery of formaldehyde as a method for hardening tissue samples so that they could be examined via microscopy. However, alcohol fixatives are known to cause significant tissue shrinkage, meaning histological measures are unlikely to be reflective of the *in-vivo* tissue [155].

3.2.1.1 Formaldehyde Fixation

By far the most frequently used chemical fixative is formaldehyde. The smallest member of the aldehyde family, formaldehyde is a naturally occurring gaseous organic compound that when dissolved in water creates a substance known as *formalin*. It is important that the pH of the fixative solution is kept within the biological range, ideally neutral (pH 7.0) [12, 156]. In order to achieve this, a buffering compound is added to the formalin (such as phosphate buffered saline, PBS) to create the substance known as Neutral Buffered Formalin (NBF) that is used to fix biological tissue [157].

Formaldehyde fixation works via a process called protein cross-linking, whereby the co-

valent bonds form between the proteins within the tissue, trapping other molecules such as carbohydrates and lipids in the process [158]. The proteins themselves consist of long chains of amino acids (see 3.1).




Figure 3.1: Structure of a typical protein

Formaldehyde readily bonds with amino acids such as Lysine through interaction with the amide group. After a formaldehyde molecule has bonded to an amino acid attached to a protein, it is possible for it to bond to an additional amino acid attached to a second protein. This forms a chemical bridge between the two proteins, thus bonding them together as shown in figure 3.2.




Figure 3.2: Methylene Bridge

Where the CH_2 molecule that links the two proteins is known as a Methylene bridge. When thousands of these Methylene bridges form within tissue, the solubility of the proteins is greatly reduced and the stiffness of the tissue increased.

The effect of formaldehyde fixation on the tissue microstructure has been a topic of frequent investigation, although the majority of studies have focussed on assessing the effect of fixation on various immunohistochemical staining protocols [159–162]. The effect of fixation on the microstructural geometry of tissue has not been studied in the same level of detail, due to the difficulty of quantifying the tissue microstructure both *in-vivo* and *ex-vivo*. Some studies have found the diffusivity of fixed tissue to be significantly lower than *in-vivo*. Changes in cell radius caused by aldehyde-based fixatives have been observed, with one study measuring a small decrease in cell radius [163] and one finding a slight increase in radius [156], however neither of these were statistically significant. It is possible that any changes in cell size and tissue morphology could occur post-mortem due to other factors such as changes in osmotic pressure.

3.2.2 Temozolomide Therapy

Temozolomide is an orally-administered chemotherapy drug, used to treat certain brain cancers. It is used as a primary treatment for glioblastoma, and a secondary treatment for astrocytoma. First synthesised in 1984, the drug is a derivative of Dacarbazine. Temozolomide is an alkylating agent, that works by methylating or alkylating the DNA within cells. The damaged DNA can trigger apoptosis, or stop the cell from replicating [164]. Besides inducing cell death within the tumour, *in-vitro* experiments have shown that glioma cells exposed to temozolomide exhibit significant morphological changes, adopting a fusiform shape and becoming smaller in size [165].

3.3 Study Outlines

As previously mentioned, the main objective of the work in this chapter was to demonstrate VERDICT MRI for measuring changes in tissue microstructure. To this end, two separate studies were performed.

3.3.1 Tissue Fixation Study

This study involved imaging subcutaneously-grown tumours pre- and post-fixation, to assess which VERDICT parameters were sensitive to the changes in tissue structure. Following *in-vivo* imaging with the VERDICT protocol, animals were perfusion and immersion fixed with 4% PFA (Paraformaldehyde). The tumours were then imaged using the VERDICT protocol a second time. The VERDICT 'BallSphereStick' and 'BallSphere' models were then fitted to the *in-vivo* and *ex-vivo* datasets, and the fitted parameters compared. The goal of this study was to assess the performance of the VERDICT MRI framework,

and investigate which of the parameters are sensitive to microstructural changes in the tissue.

3.3.2 Temozolomide Study

The study presented in the latter part of this chapter, involving the administration of temozolomide therapy was highly collaborative in nature. Dr Thomas Roberts was the study leader, performing much of the study design and imaging experiments. I contributed to the work by carrying out the imaging and experiments alongside Dr Roberts, and performing the data analysis.

The data gathered during the study was published in a paper co-authored by myself and Dr Roberts [166]. The paper presents an analysis of the VERDICT technique, comparing the parameters of the VERDICT ‘BallSphereStick’ model to values measured from histology and Optical Projection Tomography. The paper also presents an investigation into the capabilities of the VERDICT model for detecting a response to temozolomide therapy compared to established techniques such as the ADC. The temozolomide study involved imaging mouse-brain glioma cells at multiple time points following administration of temozolomide chemotherapy, to investigate whether the parameters of the VERDICT model are more sensitive markers of treatment response than traditional DW-MRI metrics such as the ADC. Glioma cells were implanted into the brains of mice, and allowed to develop for a number of days. The treatment group were then given temozolomide therapy, while the control group were given a saline bolus. Mice were imaged with VERDICT on the day of therapy/bolus, and then repeat imaging timepoints performed at intervals of 3 days. A single VERDICT model - ‘BallSphereStick’, and ADC were fitted to the data, to investigate whether VERDICT was able to detect therapeutic response at an earlier timepoint than the ADC model.

In this section of the chapter, an extension of the above-mentioned work is presented, wherein the performance of two different VERDICT models - so called ‘BallSphereStick’ and ‘BallSphereAstrosticks’, for detecting a response to therapy is compared. These models differ in their representation of the vascular volume fraction, with the ‘BallSphereStick’ model representing the diffusion in the vascular volume as an anisotropically-constrained tensor, while the ‘BallSphereAstrosticks’ model represents it as an isotropic diffusion model [6].

3.4 Methods

3.4.1 Tissue Fixation Study

3.4.1.1 Animal Models

All *in-vivo* experiments were carried out in accordance with the local ethical review panel and UK Home Office Animals Scientific Procedures Act 1986 (ASPA) under PPL 70/7309 ‘Targeted Cancer Therapies’. In addition the United Kingdom Co-ordinating Committee on Cancer Research guidelines [167] were adhered to at all times. MF1 nu/nu mice were used due to their limited immune response to human cancer cell-lines. A total of 6 female MF1 mice between 6-8 weeks of age were used for experiments. All animals were kept in UCL BSU facilities in individually ventilated cages (IVCs), with free access to food and water and a 12-hour alternating day/night cycle. After admission to animal housing facilities a 7-day acclimatisation period was observed before the commencement of any experimental work in order to reduce stress. Procedures carried out on the animals included restraint via scruffing, injection of substances (subcutaneous injection of cells) and MRI imaging under anaesthesia. All procedures were performed by a suitably trained

and competent person, with the necessary home office license (Modules 1-4). Daily checks were performed to ensure the health and well-being of the animals, with weights being recorded at least every 2 days. Baseline body-weights were recorded during the 7-day acclimatisation period. Any animal whose weight dropped by 10% of baseline was weighed daily, and any animal whose weight dropped by 15% was terminated via schedule 1. This was not necessary for any of the animals in this study, and all remained healthy until the predetermined study end-point.

3.4.1.2 Tumour Development

The cell-line used for the subcutaneous tumour models was LS174T, a human caucasian colon adenocarcinoma model known to be tumorigenic in nude mice. The cells used were from stocks kept in-house and were handled under sterile conditions. Incubation was carried out in vented flasks at 5% CO_2 , 37°C. For subcutaneous injection a concentration of 5×10^6 cells per 100 μ l of serum free media was prepared then injected into the right flank of the animals. The injection was carried out in a sterile down-flow cabinet to reduce the risk of exposing the animals to external contaminants. Daily checks for tumour growth were carried out using callipers to ensure that none exceeded the maximum permissible volume of 1.5cm³. Tumours took between 17-26 days to develop to a suitable size for imaging.

3.4.1.3 Anaesthesia and Physiological Monitoring

Anaesthesia was induced with 4% isoflurane in 100% oxygen administered at a total volume flow rate of 2L/min. Once appropriate depth of anaesthesia was established, it was maintained using around 1.5% isoflurane in 100% oxygen at a rate of 1 L/min. The concentration of isoflurane was adjusted depending on the breathing rate of the animal.

During *in-vivo* MRI experiments the animals were continuously monitored. Body temperature was monitored using a rectal thermometer (SA instruments, New York, USA), and maintained using a warm water circulation system and warm air blower. Respiration rate was monitored using a neonatal apnoea pad, and maintained by adjusting the isofluorane concentration.

3.4.1.4 Perfusion/Immersion Fixation

Perfusion fixation was performed following a lethal dose of pentobarbital (Euthatal, Merial UK) administered through IP injection. Once full depth of anaesthesia was established (via non-existent paw-pinch reaction) the chest cavity of the animal was opened. A 27G butterfly needle was inserted into the left ventricle of the heart, and an incision made in the right atrium. Isotonic saline was then used to flush the circulatory system. This was followed by 4% paraformaldehyde mixed with phosphate buffered saline. Following perfusion fixation, the whole animals were immersed in 4% paraformaldehyde in order to ensure the tissue was completely fixed. The full perfusion fixation protocol can be found in the [appendix](#).

3.4.1.5 MRI Protocol

In-Vivo MRI

All MR data was acquired using a 9.4T Varian horizontal bore scanner (Varian Inc. Palo Alto, California USA) with a maximum gradient strength of 400mT/m, with 1000mT/m gradient inserts, and a 39mm birdcage RF coil. Mice were anaesthetised as per the method described above then placed onto the scanner bed. The physiological monitoring equip-

ment was then attached, and the respiration rate and temperature checked. The tumour was then secured in place with dental paste that cures to a stiff consistency, to help reduce respiratory motion. The scanner bed was then placed into the bore of the 39mm birdcage coil, which was in turn placed into the bore of the MRI scanner. A GEMS (Gradient Echo Multi Slice) scout sequence (5 slices axial, sagittal and coronal) was performed in order to locate the mouse and tumour and help align the tumour with the isocentre of the magnet. Once the tumour was aligned, a manual shim was performed using a 1D LASER sequence. Following shimming, Fast Spin Echo Multi-Slice (FSEMS) sequences were used to create a slice plan for the VERDICT acquisitions. Firstly, an axial scan with enough slices to cover the entire tumour volume (slice thickness 0.5mm). This was followed by an oblique scan, using the axial scan images to orient the slices so that they bisected the tumour at 90° . The clearest five slices from the oblique scan were used as the slice plan for the VERDICT scans.




Figure 3.3: A schematic showing the PGSE sequence used for the VERDICT acquisitions, with the key sequence parameters labelled.

For VERDICT acquisitions, a PGSE acquisition was used with in plane field-of-view $2.5\text{cm} \times 2.5\text{cm}$, slice thickness 0.5mm and data matrix 64×64 . Diffusion-weighted images were acquired with 46 different b-values, each encoded in 3 orthogonal directions and accompanied by a b_0 image for normalisation. In addition to this, a 42-direction b1000 DTI acquisition was performed. TE was minimised for all acquisitions to maximise the signal-to-noise ratio. The specific combinations of gradient parameters are shown in table 3.1. A corresponding b_0 scan with the same TE as each combination of Δ and G was acquired, in order to normalise the data and account for T2 effects.

δ (ms)	Δ (ms)	G (G/cm)	b-value (s/mm ²)
3	10/20/30/40	3.6	8/16/24/33
3	10/20/30/40	7.2	30/63/97/130
3	10/20/30/40	10.8	68/143/218/293
3	10/20/30/40	14.4	120/254/387/521
3	10/20/30/40	18.0	188/397/606/814
3	10/20/30/40	21.6	270/571/872/1173
3	10/20/30/40	25.2	368/778/1187/1596
3	10/20/30/40	28.8	481/1016/1550/2085
3	10/20/30/40	32.4	609/1285/1962/2639
3	10/20/30/40	36.0	752/1587/2422/3257
10	30/40	4.0	305/420
10	30/40	8.0	1221/1680
10	30/40	12.0	2749/3780

Table 3.1: The gradient parameters used for the VERDICT acquisitions in this study, along with the corresponding b-value for each combination.

This study was attempted twice, during the first attempt no respiration gating was used - making the total scan time 2.5 hours per animal. However, on inspection of the data, some ghosting artefacts were noticed that were likely due to respiratory motion. Steady-state respiration gating was then added to the sequence, that ensures the data is only acquired during the relatively motion-free part of the breathing cycle. This pushed up the total scan time per animal to around 5.5 hours.

Ex-Vivo MRI

Between 20 and 24 days after fixation *in-vivo* imaging, the animals were scanned again. At least 10 days prior to *ex-vivo* imaging the animals were removed from the PFA immersion and washed with de-ionised water. The washing process was repeated 3-4 times before imaging took place, to maximise the diffusion signal. The samples were scanned overnight using the same experimental set-up as for the *in-vivo* scans. The temperature of the *ex-vivo* tissue was raised to 38° by pumping warm water (constant temperature maintained via water bath) through rubber tubing next to the sample, to equal the body temperature of the *in-vivo* subjects. Temperature was monitored using a temperature probe to ensure the temperature was stable prior to commencing scanning. The TR of the scans was increased to 350ms and 6 averages were performed per acquisition. With the same combinations of b-values and directions the total scan time was 8 hours 10 minutes per animal.

3.4.1.6 Model-Fitting

The acquired diffusion data was fitted in accordance with the procedures set out by Panagiotaki et al [6, 9]. The data (entire dataset including 42-direction DTI) was first normalised using the accompanying b_0 acquisitions and then fitted using the iterative maximum likelihood procedure that accounts implemented by the open-source Camino diffusion MRI toolkit [168]. A MATLAB wrapper was implemented to make data handling, plotting and visualisation easier. Camino implements the Levenburg-Marquardt algorithm (LMA) for model fitting. The objective function rewards goodness-of-fit via a chi-squared function, which is modulated by a noise term based on the Rician noise present in MRI data:

$$L_{obj} = \sum_{n=1}^M (\tilde{S}_n(\delta, \Delta, G) - \sqrt{S_n(\delta, \Delta, G)^2 + \sigma^2})^2 \quad (3.1)$$

Where M is the number of observations, S_n is the predicted signal by the model and \tilde{S}_n is the observed signal [6]. The LMA fitting algorithm requires an initial estimate of the parameters that are to be fitted, in order to begin the iterative process. Unlike with some other fitting techniques, these initial estimates can not be chosen randomly, as the performance of the algorithm is sensitive to the presence of local minima. The presence of these local minima depends on the complexity and number of parameters of the model being fitted. To find estimates of the model parameters for the VERDICT models, the toolkit uses a multi-stage approach, where the complexity of the model is built up in stages. For example, when fitting the 'BallSphereStick' model, the toolkit will first fit simpler 'BallSphere' and 'BallStick' models, then use the fitted values of those parameters as the initialisations for the final fitting process. The fitted parameters are also constrained to lie within certain bounds, either depending on the nature of the parameter itself, or based on the physiologically expected values. The constraints applied to the parameters in this study are shown in table 3.2.

Parameter	Constraint
Cell Radius, R	$0.1\mu m \leq R \leq 20.1\mu m$
Pseudo-Diffusion Coefficient P	$P \geq 3.05\mu m^2/ms$
Compartment Volume Fractions	$0 < v < 1$

Table 3.2: The constraints applied to the parameters during the model fitting process.

Models Fitted Four models were fitted to the diffusion-weighted data. The VERDICT 'BallSphereStick' and 'BallSphereAstrosticks' models were both fitted. As described above, these models have identical compartment models for the intracellular and extracellular-extravascular compartments but a different model for the vascular compartment. The

'BallSphereStick' compartment assumes that the vascular signal is anisotropic, whereas the 'BallSphereAstroticks' model allows for isotropic vascular signal. For *ex-vivo* data, the 'BallSphere' mode was used, that lacks a vascular compartment (due to the lack of perfusion in post-mortem tissue). In addition, the ADC model was fitted, as a comparison.

3.4.1.7 Statistical Analysis

All statistical analysis was performed using the GraphPad Prism software (version 6.0). For comparisons between *in-vivo* and *ex-vivo* parameters from the same animal, a two-tailed Mann-Whitney U-test was used. The Mann-Whitney U test is non-parametric, and can be applied on subject-matched datasets. Asterisk notation is used on graphs to denote significant differences between groups.

3.4.2 Temozolomide Study

3.4.2.1 Animal Models

All *in-vivo* experiments were carried out in accordance with the local ethical review panel and UK Home Office Animals Scientific Procedures Act 1986 (ASPA) under PPL 70/7309 'Targeted Cancer Therapies'. In addition the United Kingdom Co-ordinating Committee on Cancer Research guidelines [167] were adhered to at all times. A total of 24 female 8-week old C57BL/6 mice were used for the study. All animals were kept in UCL BSU facilities in individually ventilated cages (IVCs), with free access to water and a 12-hour alternating day/night cycle. A 7-day acclimatisation period was observed before the commencement of any procedures. All procedures were performed by a suitably trained and competent person, with the necessary home office license (Modules 1-4).

3.4.2.2 Tumour Development

The mice were injected with GL261 mouse glioma cells, into the right caudate nucleus. Anaesthesia was induced using 4% isoflurane, mixed with 100% oxygen administered at a flow rate of 2L/min. A stereotactic frame (David Kopf Instrument, Tujunga CA) was used to hold the heads of the mice steady. While in the frame, anaesthesia was maintained using 2% isoflurane, with 100% oxygen at 2L/min. Before the incision was made, the head was sterilised using 4% chlorhexadine. The skull was then exposed by cutting the skin with a sterile scalpel. Coordinates for the burr hole were found using a blunt syringe (1mm anterior, 2mm to the right of the bregma). The burr hole was then created using a 25G needle. A hamilton syringe (75N, 26s/2"/3, 5 μ L), was then used to create a small reservoir within the caudate nucleus, by inserting the syringe 4mm into the tissue then retracting 1mm to form the reservoir. The GL261 cells were then administered in a volume of 2 μ L over 2 minutes, a total of 2×10^4 cells were injected into each mouse. After delivering the cells, the syringe was slowly removed at a rate of 1mm/min. The burr hole was then closed using bone wax, and the wound sutured closed.

3.4.2.3 Anaesthesia and Physiological Monitoring

Anaesthesia was induced with 4% isoflurane in 100% oxygen administered at a total volume flow rate of 2L/min. Once appropriate depth of anaesthesia was established, it was maintained using around 1.5% isoflurane in 100% oxygen at a rate of 1 L/min. The concentration of isoflurane was adjusted depending on the breathing rate of the animal. During *in-vivo* MRI experiments the animals were continuously monitored. Body temperature was monitored using a rectal thermometer (SA instruments, New York, USA), and maintained using a warm water circulation system and warm air blower. Respiration rate was monitored using a neonatal apnoea pad, and maintained by adjusting the isoflurane concentration.

3.4.2.4 Study Design

A total of 24 female 8 week-old C57BL/6 mice were inoculated with mouse glioma cells for this study, using the method described above. Previous work using this model of glioma has shown that the tumours usually take around 2 weeks to develop to a suitable size for imaging. With this in mind, baseline imaging was carried out 13 days after tumour cell inoculation. At this timepoint, the mice were randomly-assigned into two groups: control and therapy. The therapy group were then administered with the first dose of temozolomide mixed with vegetable oil via oral gavage (130mg/kg). The control group were given an oral gavage of vegetable oil with the same volume. Treatments were repeated on the following two consecutive days, bringing the total administered dose to 490mg/kg. Meanwhile, the control group were given two further sham doses of vegetable oil. Further imaging timepoints were performed at 3-day intervals, until the final timepoint 22 days post injection. A schematic of the study timeline is shown in figure 3.4.




Figure 3.4: A schematic of the timeline of the Temozolomide therapy study, showing the tumour inoculation day, imaging and therapy timepoints..

3.4.2.5 MRI Protocol

MRI experiments were performed using a 9.4T Varian horizontal bore scanner (Varian Inc. Palo Alto, California USA) with a maximum gradient strength of 400mT/m, with 600mT/m gradient inserts installed. A 72mm volume RF coil was used as the transmit coil, with a separate 2-channel receiver head coil (RAPID biomedical, Ripmar Germany). Mice were anaesthetised using the protocol described above and then positioned prone onto the scanner bed. An MR-compatible head-holder with a bite-bar was used to hold the mice in position. The head was kept secure and rigid within the holder using plastic ear-bars, with local anaesthetic (Lidocaine) applied to minimise any discomfort. An intra-peritoneal line was inserted, for administration of gadolinium contrast agent. The physiological monitoring equipment was then attached, and the temperature and respiration rate checked. The scanner bed was then inserted into the centre of the RF coil, and into the scanner bore.

After tuning and matching, a GEMS scout sequence was used to ensure that the mouse was correctly positioned within the isocentre of the magnet. Following manual shimming, a structural T2-weighted FSEMS sequence was used to localise the tumour inside the brain. VERDICT data was then acquired using a spin echo sequence with an EPI (Echo Planar Imaging) readout, with a TR of 3 seconds, minimised TE, 3 shots, $20\text{mm} \times 20\text{mm}$ field of view, 0.5mm slice thickness, 5 slices, 2 averages and 64×64 data matrix. The combination of b-values and gradient timings was identical to those used in the tissue fixation study shown in table 3.1, including the 42-direction b-1000 DTI acquisition. Following the acquisition of VERDICT data, contrast-enhanced images were taken. Mice were given intra-peritoneal injections of gadolinium-DTPA (dosage 0.6mmol/kg), which was allowed 10 minutes to circulate before imaging with a T1-weighted spin-echo sequence slice-matched to the VERDICT data.

3.4.2.6 Data Analysis

Regions of interest (ROIs) were manually drawn around the tumours in a slice-by-slice manner, using the *roipoly* function in MATLAB, based on the enhancing regions of the gadolinium-enhanced MR images. The binary masks produced were then applied to the diffusion data, then data analysis was performed on the voxels within these ROIs.

Model fitting was performed using the open-source Camino diffusion MRI toolkit [168], using a MATLAB wrapper. Further details can be found previously in section 3.4.1.6. Two different VERDICT models were fitted to the data, the 'BallSphereStick' and 'BallSphereAstrosticks' models, as well as the ADC. Three parameters from the aforementioned models were selected, the intracellular volume, extracellular volume and cell radius. The performance of corresponding parameters between the two VERDICT models was compared, to see if there was a difference in sensitivity between the two models. The sensitivity of these parameters to the microstructural changes induced by the temozolomide therapy were compared against the ADC. In order to improve the stability of the fitting process, the intracellular and extracellular diffusivities were fixed at $9e - 10m^2s^{-1}$ for both VERDICT models (Similar to the methods of the original VERDICT paper [9]). The fitting process was identical to that of the tissue fixation study 3.4.1.

As a further comparison, the T2-weighted structural scans were segmented in order to give a clinically-relevant estimation of tumour volume. The tumour volume was used as an additional parameter to compare against the ADC and VERDICT parameters to see which is the most sensitive predictor of therapeutic effect.

In addition, to assess and compare the goodness-of-fit achieved by the two VERDICT models, the final value of the objective function minimised during the fitting process was output for each subject. The mean and standard deviations of the values were then calculated so that the performance of the two models can be compared.

3.4.2.7 Statistical Tests

The comparisons made in this study were between control and treated animals imaged at the same timepoint. In each case, a two-tailed Mann-Whitney U test was used to determine statistically significant difference between the two populations. On graphs, a single asterisk is used to denote timepoints where a significant difference was found. All statistical analysis was performed using the GraphPad Prism software (version 6.0).

3.5 Results

3.5.1 Tissue Fixation Study

3.5.1.1 Cell Radius

Figure 3.5 shows the cell radius parameter values produced from fitting the two VERDICT models to the *in-vivo* and *ex-vivo* data. On the left hand side are the results produced by the 'BallSphereStick' and 'BallSphere' models. There was a significant difference between the cell radius parameter fitted to the *in-vivo* and *ex-vivo* data ($P < 0.0001$), with the *in-vivo* radii significantly larger. The mean cell radius parameter from the *in-vivo* fits was $5.18\mu m$, compared to $4.29\mu m$ for the *ex-vivo* fits.

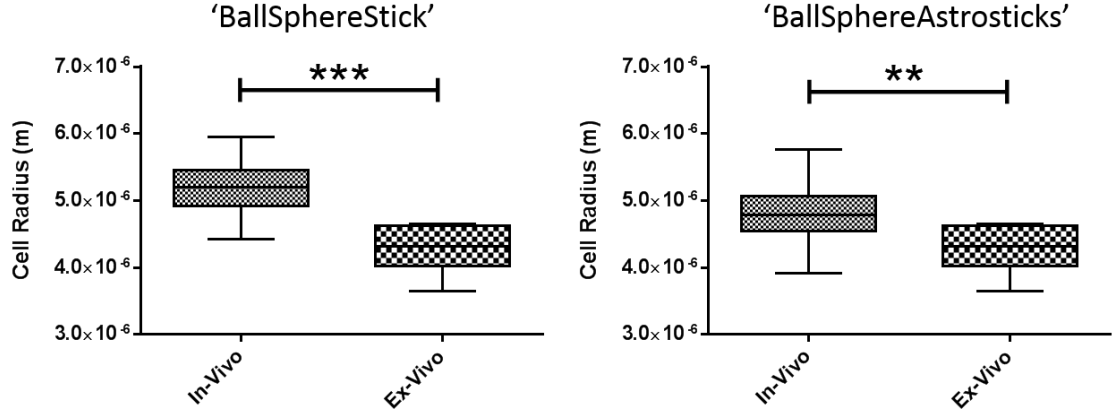


Figure 3.5: Box & Whisker plots showing the values of the cell radius parameter produced by the VERDICT models fitted to the *in-vivo* and *ex-vivo* data. LEFT: 'BallSphereStick' and 'BallSphere' models. RIGHT: 'BallSphereAstrosticks' and 'BallSphere'.

On the right-hand side of figure 3.5 are the results produced by the 'BallSphereAstrosticks' model. A significant difference was found between the cell radius parameter values from the *in-vivo* and *ex-vivo* fits ($P = 0.0011$), with the *in-vivo* radii significantly larger. The mean *in-vivo* cell radius was $4.8\mu\text{m}$, compared to $4.28\mu\text{m}$ *ex-vivo*.

3.5.1.2 Intracellular Volume Fraction

The intracellular volume fraction values produced from fitting the the two VERDICT models to the *in-vivo* and *ex-vivo* data are shown in figure 3.6. The f_{icvf} values produced by the 'BallSphereStick' and 'BallSphere' models are shown on the left-hand side. A significant difference between the *in-vivo* and *ex-vivo* was found ($P < 0.0001$), with larger intracellular volume fractions associated with the *in-vivo* data. The mean *in-vivo* f_{icvf} was 0.498, compared to 0.377 *ex-vivo*.

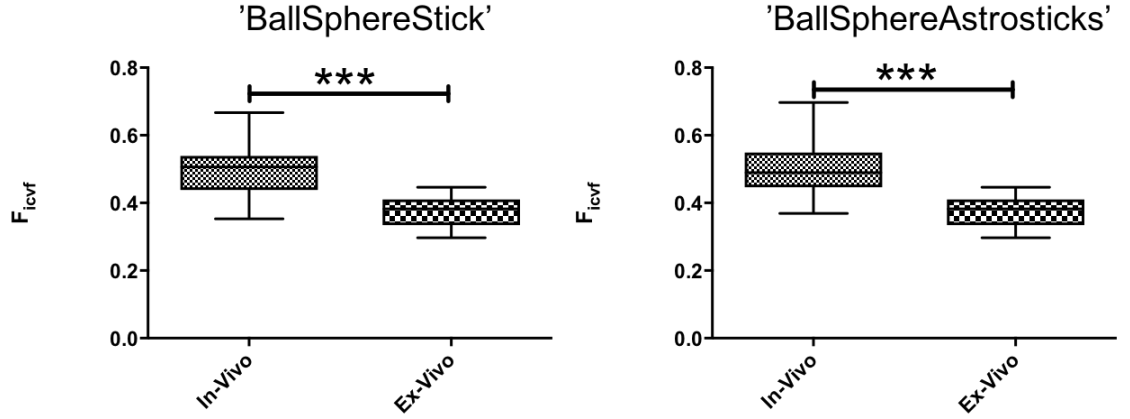


Figure 3.6: Box & Whisker plots showing the values of the intracellular volume fraction produced by the VERDICT models fitted to the *in-vivo* and *ex-vivo* data. LEFT: 'BallSphereStick' and 'BallSphere' models. RIGHT: 'BallSphereAstrosticks' and 'BallSphere'.

The right-hand side of figure 3.6 shows the results achieved by fitting the 'BallSphereAstrosticks' and 'BallSphere' models. Similarly, a significant decrease in intracellular volume was measured between the *in-vivo* and *ex-vivo* data ($P < 0.0001$). The mean intracellular volumes were 0.499 and 0.377 for the *in-vivo* and *ex-vivo* datasets respectively.

3.5.1.3 VERDICT Diffusivity

Figure 3.7 shows the diffusivity values measured using the VERDICT framework. Shown on the left are the results achieved by fitting the 'BallSphereStick' model to the *in-vivo* data, and the 'BallSphere' model to the *ex-vivo* data. A significant decrease in diffusivity was measured between the *in-vivo* and *ex-vivo* data ($P < 0.0001$). The mean values were $1.89 \times 10^{-9} m^2 s^{-1}$ and $6.96 \times 10^{-10} m^2 s^{-1}$ for the *in-vivo* and *ex-vivo* respectively.




Figure 3.7: Box & Whisker plots showing the values of the diffusivity parameter produced by the VERDICT models fitted to the *in-vivo* and *ex-vivo* data. LEFT: 'BallSphereStick' and 'BallSphere' models. RIGHT: 'BallSphereAstrosticks' and 'BallSphere'.

The right-hand side shows the results applying the 'BallSphereAstrosticks' model. Similarly, a significant decrease with $P < 0.0001$ was detected between the *in-vivo* and *ex-vivo* datasets. The mean *in-vivo* diffusivity was $1.32 \times 10^{-9} m^2 s^{-1}$, compared to $6.96 \times 10^{-10} m^2 s^{-1}$ *ex-vivo*.

3.5.1.4 ADC

Figure 3.8 shows the ADC value calculated from the *in-vivo* and *ex-vivo* data. The ADC values were found to be significantly lower in the *ex-vivo* data ($P < 0.0001$).




Figure 3.8: Box & Whisker plots showing the values of the ADC fitted to the *in-vivo* and *ex-vivo* data.

The mean values of ADC were $6.1 \times 10^{-10} m^2 s^{-1}$ and $4.39 \times 10^{-10} m^2 s^{-1}$ *in-vivo* and *ex-vivo* respectively, with standard deviations of $\pm 8.39 \times 10^{-11}$ and $\pm 3.15 \times 10^{-11}$.

3.5.2 Temozolomide Study

3.5.2.1 VERDICT 'BallSphereStick'

Figure 3.9 shows the parameters produced by fitting the VERDICT 'BallSphereStick' model to the data from the control and therapy groups at each study time-point. The results show that the intracellular volume fraction parameter showed no significant difference at the 3 day time-point. A significant difference was found between the control and therapy groups at the 6 and 9-day time-points with p-values of 0.007 and < 0.0001 respectively.

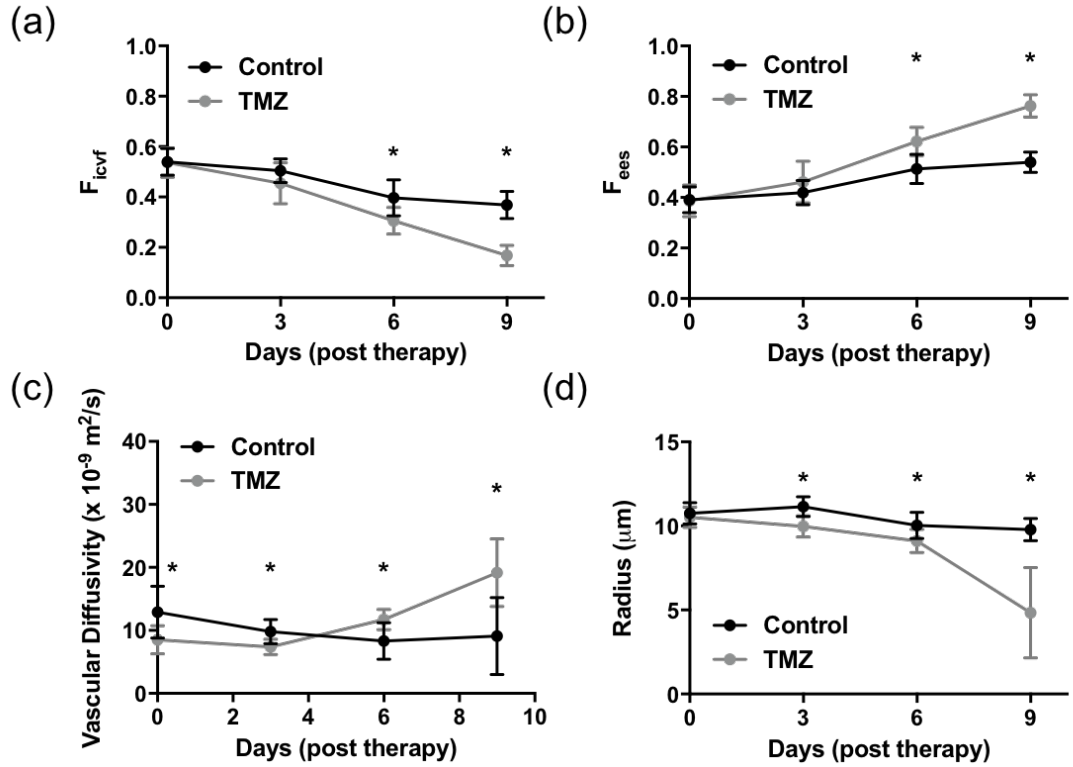


Figure 3.9: The parameter values produced by the VERDICT 'BallSphereStick' model. (a) Intracellular Volume Fraction (b) Extracellular Volume Fraction (c) Vascular Diffusivity (d) Cell Radius. On each plot, the values produced from fitting the model to the control and therapy group data are shown, with error bars corresponding to the standard deviations. Asterisks are used to denote time-points where a significant difference was found between the control and therapy group.

A similar pattern was found for the extracellular volume fraction, with no significant difference found between the control and therapy groups at either the 3 or 6 day post therapy timepoints. However, a significant difference was found after 6 and 9 days, with p-values of 0.0006 and < 0.0001 respectively. The vascular diffusivity parameter presented a significant difference between the control and therapy groups across all four timepoints. At 3 and 6 days post-therapy the vascular diffusivity of the control group was significantly higher (p values 0.005 and 0.003), whereas at 6 and 9 days post therapy the therapy group was significantly higher (p values 0.008 and 0.008). A significant difference was detected

between the cell radius parameter and the 3 ($p = 0.0002$), 6 ($p = 0.0142$) and 9 ($p = 0.0003$) day time points.

3.5.2.2 VERDICT 'BallSphereAstrosticks'

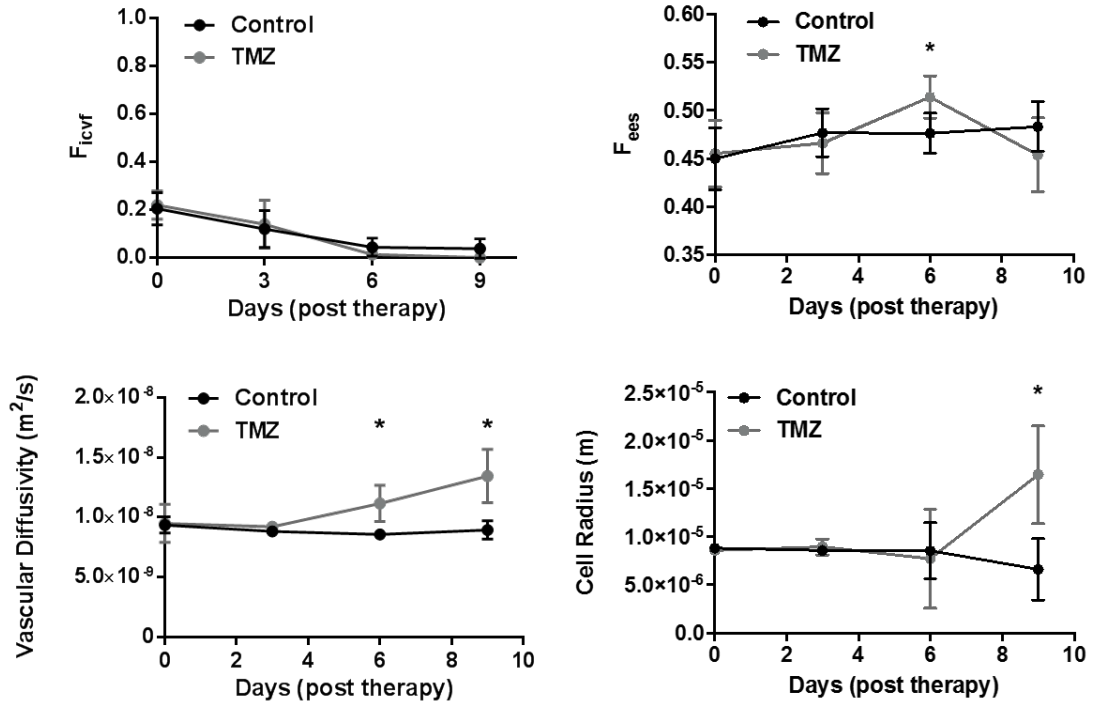


Figure 3.10: The parameter values produced by the VERDICT 'BallSphereAstrosticks' model. (a) Intracellular Volume Fraction (b) Extracellular Volume Fraction (c) Vascular Diffusivity (d) Cell Radius. On each plot, the values produced from fitting the model to the control and therapy group data are shown, with error bars corresponding to the standard deviations. Asterisks are used to denote time-points where a significant difference was found between the control and therapy group.

Figure 3.10 shows the results from the VERDICT 'BallSphereAstrosticks' model (with an anisotropic vascular compartment model). The intracellular volume fraction showed no significant difference between therapy and control at any of the study timepoints, with very small intracellular volume values estimated across all datasets. The extracellular volume

fraction exhibited similar results in that there was no significant difference between the two groups at the 0, 3 or 9-day time points. The therapy group values were significantly higher than the control group at the 6-day time point ($p = 0.0017$). The vascular diffusivity parameter showed no significant difference at the 0 and 3-day time points. After 6 days the therapy group demonstrated a significantly larger vascular diffusivity (mean 1.11×10^{-8} vs 8.56×10^{-9} , $p < 0.0001$). The following 9-day time point showed a similar result, with a higher diffusivity in the therapy group ($p = 0.0005$). There were no significant differences between the cell radius parameters of the two groups up until the 9-day time point. At 9 days post-therapy, the therapy group was significantly larger than the control ($p = 0.002$).

3.5.2.3 ADC




Figure 3.11: The parameter values produced by the ADC model. Asterisks are used to denote time-points where a significant difference was found between the control and therapy group.

Figure 3.11 shows the results from fitting the ADC model to the temozolomide therapy data. There was no significant difference between the control and therapy groups at day 0 or 3 days post therapy. At 6 days post therapy the ADC values of the therapy group were significantly larger than the control group (mean 1.296×10^{-9} vs 1.139×10^{-9} , $p < 0.0001$). This difference increased in magnitude at the 9-day time point (mean 1.456×10^{-9} vs 1.163×10^{-9} , $p < 0.0001$).

Parameter		Days Post Injection							
		0		3		6		9	
		Control	Therapy	Control	Therapy	Control	Therapy	Control	Therapy
F_{icvf}	BSS	0.540 ± 0.053	0.536 ± 0.059	0.504 ± 0.047	0.454 ± 0.081	0.397 ± 0.072	0.306 ± 0.054	0.369 ± 0.054	0.168 ± 0.040
	BSA	0.205 ± 0.067	0.219 ± 0.059	0.120 ± 0.077	0.140 ± 0.100	0.044 ± 0.038	0.013 ± 0.034	0.038 ± 0.041	0.000 ± 0.000
F_{ees}	BSS	0.391 ± 0.051	0.387 ± 0.064	0.419 ± 0.048	0.461 ± 0.082	0.514 ± 0.058	0.621 ± 0.057	0.540 ± 0.041	0.762 ± 0.045
	BSA	0.450 ± 0.032	0.455 ± 0.035	0.477 ± 0.025	0.466 ± 0.032	0.476 ± 0.021	0.514 ± 0.022	0.483 ± 0.026	0.454 ± 0.038
Vascular Diffusivity ($\times 10^{-9}$)	BSS	12.907 ± 4.106	8.495 ± 2.26	9.791 ± 1.920	7.390 ± 1.22	8.333 ± 2.931	11.735 ± 1.59	9.083 ± 6.116	19.164 ± 5.35
	BSA	0.347 ± 0.073	0.478 ± 0.578	0.821 ± 0.049	0.202 ± 0.061	0.565 ± 0.004	1.154 ± 0.513	0.932 ± 0.061	3.434 ± 0.238
Cell Radius	BSS	10.752 ± 0.627	10.516 ± 0.609	11.144 ± 0.590	9.969 ± 0.630	10.029 ± 0.779	9.104 ± 0.694	9.787 ± 0.660	4.831 ± 2.690
	BSA	9 ± 0.052	9 ± 0.097	9 ± 0.013	9 ± 0.031	9 ± 3	8 ± 5	7 ± 3	16 ± 5
ADC		0.996 ± 0.060	0.980 ± 0.062	1.051 ± 0.064	1.066 ± 0.106	1.140 ± 0.050	1.296 ± 0.080	1.163 ± 0.062	1.456 ± 0.036

Table 3.3: Mean and standard deviations of the various fitted parameters at each study time point for the therapy and control groups. For the VERDICT models, the parameter values from the ‘BallSphereStick’ (BSS) and ‘BallSphereAstrosticks’ (BSA) models are shown.

3.5.2.4 Volumetric Measurements

Time (post baseline)	Tumour Volume ($mm^3 \pm \sigma$)	
	Control	Therapy
0	8 ± 1	8 ± 1
3	20 ± 1	19 ± 2
6	47 ± 2	48 ± 4
9	89 ± 7	61 ± 8

Table 3.4: Mean and standard deviations of the volumetric tumour measurements calculated from manual segmentation of the structural (T2-weighted) images of the therapy and control groups, for all time points.

Table 3.4 shows the tumour volume values calculated from the manual segmentations of the T2-weighted images at each time point. At baseline, the therapy and control groups had identical mean volumes. Over the following 6 days the tumours grew almost linearly with time, with control and therapy groups increasing in size at an almost identical rate. At 9 days following baseline, the therapy and control groups diverged, with the tumour growth in the therapy group significantly slowing relative to the control ($p = 0.029$).

3.5.2.5 VERDICT Model Performance Comparison




Figure 3.12: Side-by-side plots showing exemplary VERDICT data from one of the therapy cohort animals. The symbols represent the measured data, and the lines represent the fits from the two VERDICT models.

Figure 3.12 above shows an example VERDICT dataset taken from one of the therapy cohort animals. The left-hand plot shows the fit produced by the ‘BallSphereStick’ model, and the right-hand plot shows the ‘BallSphereAstrosticks’ model. Both models produced a close fit across the full range of data. The ‘BallSphereStick’ model appeared to produce more consistent fits across the different acquisition directions than the ‘BallSphereAstrosticks’ model.

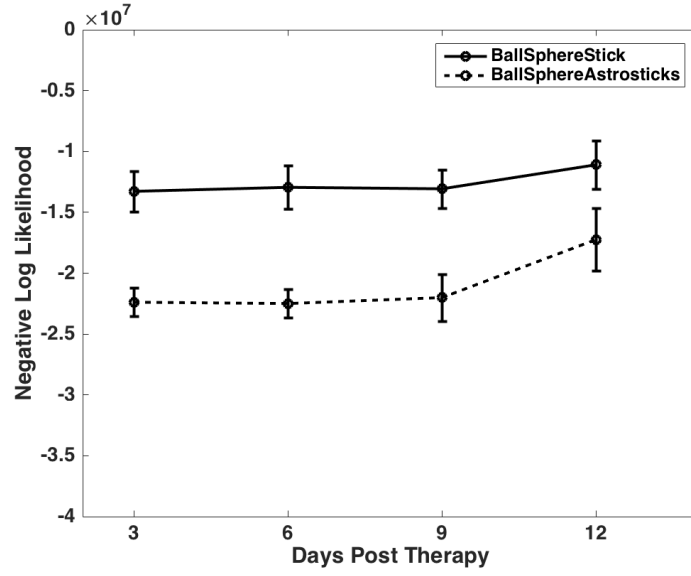


Figure 3.13: **Therapy group:** The final value of the maximum likelihood objective function achieved by the VERDICT ‘BallSphereStick’ (Solid line) and ‘BallSphereAstrosticks’ (Dashed line) models for each imaging time point post therapy, averaged across all study subjects, with standard deviation across the subjects also shown.

Figure 3.13 shows the objective function (maximum likelihood) values achieved during the fitting process for the therapy group data by the ‘BallSphereStick’ and ‘BallSphereAstrosticks’ models. As discussed previously, the objective function value is indicative of how well the model is able to fit the measured data. In this case, a higher value indicates a better fit to the data.

As can be seen looking at Figure 3.13, in general the ‘BallSphereStick’ (BSS) model was able to achieve a closer fit to the data across all the study time-points, with an average value of $-1.2655e7$, vs $-1.9796e7$ for the ‘BallSphereAstrosticks’ (BSA) model. The quality of fit for both models showed little change between the 3rd and 9th day post-therapy, followed by a considerable increase between the 9th and 12th day post-therapy. In particular, the BSS model objective function increased by $1.975e6$ between the 9th and 12th day, while the BSA model objective function increased by $4.676e6$.




Figure 3.14: **Control group:** The final value of the maximum likelihood objective function achieved by the VERDICT ‘BallSphereStick’ (Solid line) and ‘BallSphereAstrosticks’ (Dashed line) models for each imaging time point post therapy for the control dataset, averaged across all study subjects, with standard deviation across the subjects also shown.

Figure 3.14 shows the objective function (maximum likelihood) values achieved during the fitting process for the control group data by the ‘BallSphereStick’ and ‘BallSphereAstrosticks’ models.

The results from the control group were broadly similar to those of the therapy cohort, with the BSS model achieving a closer fit to the data across all study timepoints, with an average value of $-1.342e7$ vs $-2.218e7$ for the BSA model. The values for both models remained broadly consistent throughout the study, without the increase between the 9th and 12th days post therapy seen in the therapy cohort data.

3.6 Discussion

3.6.1 Tissue Fixation Study

The aim of this study was to apply the VERDICT MRI framework to *in-vivo* and fixed tumour tissue, to assess the sensitivity of the fitted parameters to the changes in tissue microstructure attributed to formaldehyde fixation. This was motivated by the clinical need for imaging biomarkers that are sensitive to particular changes in tissue structure caused by therapy, as opposed to the currently-used ADC which is influenced by a broad range of tissue characteristics. To achieve the primary aim of this study, subcutaneous tumours were imaged using the VERDICT acquisition scheme both *in-vivo* and post-fixation. Two VERDICT models ('BallSphereStick' and 'BallSphereAstrosticks'), as well as their *ex-vivo* counterparts were fitted to the data, as well as the ADC. The resulting fitted parameter values were compared to assess which parameters were sensitive.

The cell radius parameters of both VERDICT models were significantly reduced following tissue fixation, this was supported and mirrored by a significant reduction in the intracellular volume fractions. Interpretation of these parameters is difficult. According to the mathematical model, a reduction in the cell radius parameter indicates a reduction of the length scale over which diffusion within the tissue is restricted. This could be attributed to a number of factors. For one, interpreting the cell radius parameter literally would attribute its reduction in value following tissue fixation to a significant reduction of the mean cell radius in the tissue. Such a significant change has not been reported in the literature following histological analysis of fixed tissue - however these studies have not had a reliable means of measuring the cell radius *in-vivo*. The reduction in the cell radius parameter could also be consistent with the additional effects that formaldehyde fixation has on the microstructure of tissue. The formation of a dense matrix of cross-linked

proteins could create a restrictive diffusion environment, trapping the water molecules in position. The absolute values of the fitted cell radius (4-5 microns) are consistent with the literature and previous histological analysis of the LS174T cell line. However, if the diffusion coefficient of the tissue is assumed to be roughly $9 \times 10^{-10} m^2 s^{-1}$, the smallest Δ of the VERDICT acquisition (10ms) would correspond to a diffusion length of $7.34 \mu m$. The fact that this is larger than the fitted radius values introduces a degree of doubt over their absolute accuracy. The diffusivity parameter from both models were also significantly reduced following fixation. This is an unsurprising result, as the *ex-vivo* tissue is so markedly different to the *in-vivo*, and was scanned at a different temperature.

A significant decrease between the ADC values of the *in-vivo* and *ex-vivo* tissue was also found, with a p-value less than 0.0001. The ability to distinguish between living and fixed tumour tissue is not a ground-breaking result in itself. However, the fact that the VERDICT model was able to do so with the same sensitivity as the ADC is a promising result. The decrease in ADC value also supports the VERDICT results, as it is in agreement with the decrease in the cell radius and intracellular volume fraction parameters. The changes induced in the tumour microstructure by the fixation process would likely reduce the radius over which the thermal motion of water molecules is restricted, as discussed above. Such a decrease in restriction radius would likely cause the observed decrease in ADC. The advantage presented by the VERDICT framework in this case is that it provides greater insight into what kinds of microstructural alterations have occurred within the tissue to produce the change in signal.

3.6.2 Temozolomide Study

The aim of this study was to apply the VERDICT MRI framework to a therapy study administering temozolomide to treat a mouse model of glioblastoma, to test whether VERDICT was able to detect therapeutic response at an earlier time point than a the simpler ADC model. Previous work by Panagiotaki et al [9] demonstrated the ability of VERDICT to detect response to gemcitabine chemotherapy in a subcutaneous mouse model of human colon adenocarcinoma. In this current study, glioblastoma tumours were established in 24 mice, then split into therapy and control groups following baseline imaging. Further imaging time points were performed alongside administration of temozolomide therapy. Two separate VERDICT models were fitted to the diffusion data from the therapy and control groups, as well as the ADC. The 'BallSphereStick' (BSS) and 'BallSphereAstrosticks' (BSA) models are identical apart from their vascular compartments (anisotropic for BSS, isotropic for BSA). Four fitted parameters from each VERDICT model were investigated and their performance compared with the ADC and tumour volume measurements.

Of the two VERDICT models, the 'BallSphereStick' performed better with all four fitted parameters detecting significant differences between the therapy and control groups. The intracellular volume fraction decreased more rapidly in the therapy group than in the control. At the 3-day time point, the therapy group intracellular volumes had decreased relative to the control group to a greater degree than those of the ADC or tumour volume measurements. The extracellular volume fraction supported the trend of the intracellular volume, with the therapy group increasing more rapidly relative to the control group and to a greater degree than the ADC and volume measurements. By the 6 day time point a significant difference between therapy and control groups was found. The most sensitive model parameter was the cell radius. By the 3 day time point the cell radius of the therapy group had significantly decreased compared to the control, 3 days earlier than was achieved

using the ADC and 6 days earlier than was achieved using bulk volume measurements. Significant differences were also found at the 6 and 9 day time points.

The parameters of the 'BallSphereAstrosticks' model did not prove as sensitive to the particular microstructural changes. The intracellular volume fraction parameter did not detect a significant difference between groups at any time point. The extracellular volume fraction demonstrated a significant difference at the 6 day time point, but the trend was reversed by the 9-day point making the significant difference at the earlier time point less useful in a clinical setting. The vascular diffusivity was the most reliable at separating the two groups, with significant differences after 6 and 9 days (greater significance at the 9 day point). The cell radius parameter showed no difference between the groups at the 0, 3 and 6 day points. At the 9 day time point, the therapy group values were significantly larger than the controls. Overall, only the vascular diffusivity and cell radius parameters were as sensitive as or more sensitive than the bulk tumour volume measurements. Only the vascular diffusivity parameter proved to be as sensitive as the ADC value as a biomarker for therapy response.

The results demonstrated that the VERDICT framework was able to detect therapeutic response at an earlier time point than the clinically-accepted techniques. In particular, the 'BallSphereStick' model was the most sensitive for determining therapeutic response, with all parameters exhibiting a significant difference between therapy and control groups at the 6-day time point. The cell radius parameter detecting therapeutic response after 3 days post therapy, which supports a decrease in cell size caused by Temozolomide therapy. The 'BallSphereAstrosticks' model was not as effective, with only the vascular diffusivity demonstrating the ability to determine therapeutic effect after 6 days.

In addition, analysis of the fitting performance of both VERDICT models revealed that the 'BallSphereStick' model is able to achieve a closer fit to the *in-vivo* data than the

‘BallSphereAstrosticks’ model. This in itself is perhaps not a hugely surprising result, since the ‘BallSphereStick’ model comprises 7 fitted parameters, compared to the 5 fitted parameters of the ‘BallSphereAstrosticks’ model. Interestingly, towards the latter stages of the study, the performance of the two models converged, as the ‘BallSphereAstrosticks’ model was able to produce a much better fit for the therapy data after the 12th day post-therapy. This effect was not seen in the control data, where the performance of both models remained consistent throughout the study duration. This suggests that the change in model performance is driven by structural changes occurring within the tissue. For example the triggering of apoptosis of the cells within the tumour tissue, resulting in widespread cell death, would result in a more homogeneous and watery tissue which may affect the fitting performance of the two models. Despite this, the results show that the ‘BallSphereStick’ model was able to produce a closer fit to the data, and was more sensitive to microstructural changes induced in the tissue caused by the administration of temozolomide therapy.

One limitation of this study, that could be addressed in future work, is the lack of a comparative measure of perfusion. One option for future analysis would be to make use of the contrast-enhanced acquisitions to obtain an independent measure of perfusion fraction. Dynamic contrast-enhanced MRI has previously been applied to measure perfusion volume and blood flow metrics in tumours and in the brain [169, 170].

In summary, this study compared the parameters produced by the VERDICT MRI framework to the clinically-accepted MRI standards for assessing brain tumour therapy response. This chapter builds upon the work presented in Roberts et al [166], which demonstrated that VERDICT is more sensitive to microstructural changes caused by the administration of temozolomide therapy than conventional techniques. The results presented in this study confirmed by a parallel analysis that the VERDICT method provides increased sensitivity for detecting response to the Temozolomide therapy in a mouse model of glioblastoma,

with differences between therapy and controls detected with VERDICT 6 days before changes in the bulk volume measurements. Further, the work in this chapter expanded on the analysis presented in [166] by comparing the performance of the two commonly-used VERDICT models - ‘BallSphereSticks’ and ‘BallSphereAstrosticks’, both in terms of their sensitivity to microstructural tissue alterations and in terms of the quality-of-fit achieved by both models. The results demonstrated that the ‘BallSphereStick’ model is both highly sensitive (more so than the ‘BallSphereAstrosticks’ model), and produces a closer fit to the data. The work presented in the following chapters undertakes a thorough validation of both of these models, to investigate how their parameters correlate with tissue features, and interrogate their sensitivity to experimental variables such as signal-to-noise.

3.7 Chapter Summary

The aim of this chapter was to assess the capability of the VERDICT framework to detect changes in tissue microstructure. Firstly by determining whether the VERDICT parameters can quantify changes in tumour microstructure induced by tissue fixation and secondly by demonstrating that the parameters can serve as sensitive biomarkers for tumour therapy response. The tissue fixation study successfully demonstrated the utility of the VERDICT framework, and showed that the VERDICT parameters were able to quantify the changes in cell radius, intracellular volume and diffusivity between the *in-vivo* and fixed tissue. The same level of significance was found in the parameter changes for the VERDICT models and the ADC. The temozolomide chemotherapy study presented a more in-depth investigation into the potential clinical utility of the VERDICT framework. The results discussed above showed that in this case the VERDICT model (in particular the cell radius parameter of the ‘BallSphereStick’ model) was able to detect a significant response to chemotherapy at an earlier time point (3 days post therapy vs 6 days) than the

ADC and considerably earlier than the conventional bulk volume measurements. Further, a comparison of the performance of two commonly-used VERDICT models ('BallSphereStick' and 'BallSphereAstrosticks') was presented, showing that the 'BallSphereStick' model was the more sensitive, and produced a better fit to the data. These results are a positive indication that the VERDICT framework may provide useful non-invasive cancer biomarkers by providing a quantitative insight into the tumour microstructure. However, one of the key challenges associated with the development of a complex technique such as VERDICT is validation of the model parameters. Exploring the link between the fitted parameter values and the actual tissue microstructure is vital if the technique is to become further established in the clinic. The following chapter presents the development of a Monte Carlo modelling framework designed to perform *in-silico* validation of VERDICT.

Development and Application of a Monte Carlo Model of Tissue Diffusion and Flow

4.1 Motivation

As previously mentioned, compartmental models such as VERDICT have the potential to unlock access to histologic tissue features through mathematical modelling of the diffusion MR signal. These complex models necessitate thorough validation to determine their sensitivity and specificity to the tissue property they are trying to measure. Various studies have focussed on validating diffusion models using numerical phantoms - namely Monte Carlo models of diffusion. Monte Carlo models are by far the most popular approach for modelling diffusion, as although they are computationally very expensive they allow for the greatest flexibility in modelling diffusion in a wide-range of tissue geometries. However, to date numerical approaches have widely neglected the effect of pseudo diffusion arising from water molecules diffusing in the intravascular space. Only simple, non-branching vascular structures have been implemented, generating synthetic signal deterministically. It has been recognized that Monte Carlo methods incorporating the effect of water diffusion in realistic intravascular spaces is a gap in the field that should be filled [135, 171].

The aim of this chapter was firstly to develop a Monte Carlo simulation framework that recreates the diffusion process of water in biological tissue, including water diffusion in a realistic vascular structure, and generates synthetic diffusion MRI data based on this motion. Secondly, to apply said framework as a validation tool for the mathematical models used by the VERDICT technique and their associated parameters.

4.2 Introduction

Simple models such as the ADC have been successfully employed in the clinic for detecting tumours and assessing response to tumour therapy. However, they are limited by the fact that there are various factors that can affect their fitted parameters. This means it is difficult to ascertain what physiological change has caused the corresponding change in parameter value(s). Models such as VERDICT look to overcome this limitation by using mathematical models that aim to reflect the structure of the underlying tissue. As previously discussed in chapter 2 [2](#) a similar approach has been successfully applied to quantitative imaging of the brain, with multi-compartment models such as NODDI able to extract histologic features of white matter such as axon diameter and inter/extracellular volume fractions.

Along with the development of more sophisticated or complex diffusion models such as VERDICT comes an increased necessity for thorough validation. It is vital to quantify the relationship between the fitted parameters and the actual tissue microstructure. Previous work on validation has been performed via a variety of methods including comparison with histology, physical phantoms (including hardware and *ex-vivo* tissue samples) and numerical methods. Each of these methods has its own associated advantages and disadvantages, as was discussed in chapter 2 [2](#). Histology is a popular validation technique as it allows for the direct visualisation of the imaged tissue sample. However, tissue samples

must be removed from the subject via surgical resection, biopsy or *ex-vivo* dissection. This is usually followed by a series of sample preparation steps, that may alter the structure of the tissue itself. The necessity for removing the samples from the subjects also makes it very difficult to perform subject-matched longitudinal studies. Phantoms provide an alternative to histological validation, by providing a tangible yet simplistic representation of the target tissue microstructure, where the corresponding properties of the phantom can be finely adjusted and measured. Phantoms are however limited by their simplistic nature, and in addition can be difficult and costly to manufacture. The most popular alternative to physical hardware phantoms and histology is numerical validation.

As discussed in chapter 2 [2](#), numerical phantoms are computer models that are constructed to represent a physical object or phenomenon that allows the detailed testing and validation thereof. In this case the physical object is biological tissue, more specifically the tumour microstructural environment. The computational nature of such phantoms imparts a great deal of flexibility, allowing for validation of biophysical models in tissue with near-infinite combinations of physical properties and acquisition parameters.

The aim of the work presented in this chapter was firstly to develop a Monte Carlo modelling framework that simulates the motion of water molecules diffusing within tumour tissue and generates synthetic diffusion MRI data from said motion. Secondly, to apply the aforementioned modelling framework as a validation tool for the mathematical models of the VERDICT MRI technique, as well as validating the fitted parameter values of other popular diffusion models such as the ADC and IVIM.

4.3 Methods

4.3.1 Animal Models

A total of 5 female MF1 nu/nu mice were used to gather the *in-vivo* DW-MRI data. MF1 mice were used due to the fact that they are immunocompromised. Animals were kept in UCL BSU facilities in individually ventilated cages (IVCs), with free access to water and food and an artificial 12-hour alternating day/night cycle. A 7-day acclimatisation period before performing any procedures was implemented in order to help reduce any unnecessary stress on the animals. Procedures carried out after this period included restraint via scruffing, subcutaneous injection of substances, and MRI under anaesthesia. All *in-vivo* experiments were carried out in accordance with the local ethical review panel and UK Home Office Animals Scientific Procedures Act 1986 (ASPA) under PPL 70/8421 'Imaging of Cell Therapy in Tumour Models'. The United Kingdom Co-ordinating Committee on Cancer Research guidelines were also adhered to at all times. All procedures were performed by a suitably trained and competent person with the necessary Home Office license (PIL, modules 1-4).

The cell-line used for the development of subcutaneous tumours was LS174T, a well-studied human caucasian colon adenocarcinoma model known to be tumorigenic in MF1 nu/nu mice. The cells were grown from stocks kept in-house, and were incubated in vented flasks (5 CO₂, 37° C). A concentration of 5×10^6 cells per 100 μ l of serum-free media was prepared for each injections. The cells were injected into the right flank of each mouse, then the mice were monitored daily to check for tumour growth. Tumours were measured using callipers to ensure that none exceeded the permitted 1.5cm³ volume. On average, tumours took 20 days to develop to a suitable size for imaging.

4.3.2 Data Acquisition

All MR data was acquired using a 9.4T Varian 20cm horizontal bore scanner (Varian Inc. Palo Alto, CA, USA) with a maximum gradient strength of 400mT/m, and a 39mm birdcage RF coil. Mice were anaesthetised and placed in the bore of the RF coil, with the tumour close to the RF centre. The RF coil was then tuned and matched to maximise the detected signal. In order to confirm the optimal positioning of the animal in the scanner a gradient-echo scout sequence was used to check that the tumour was properly centred within the RF coil. Manual shimming was then performed to ensure that the B_0 field was as homogeneous within the sample as possible. Fast Spin-Echo Multi-Slice sequences were used to create a slice plan for the VERDICT acquisition. Consecutive slices through the tumour were chosen that were free of skin while still maintaining the thickest cross-sectional area of tumour tissue.

VERDICT data was acquired using a Pulsed-Gradient Spin-Echo sequence, as previously discussed. A full dataset consists of 46 different diffusion weightings, each applied in the 3 principle directions. Gradient separations of $\Delta = 10, 20, 30, 40ms$ were used. Gradient durations of $\delta = 3ms$ were used for all Δ , and $\delta = 10ms$ in addition for $\Delta = 30, 40ms$. For each timing combination with $\delta = 3ms$, the gradient amplitude was stepped from 40 to 400mT/m in steps of 40mT/m. For timing combinations with $\delta = 10ms$ the maximum gradient strength was 120mT/m. The in-plane FOV was $2.5cm \times 2.5cm$ with a 64×64 data matrix, and the slice thickness was $0.5mm$. Steady-state respiration gating was implemented in order to reduce the effect of motion artefacts. The tumours were also given a generous coating of dental paste, in order to further reduce motion during scans.

4.3.3 Monte Carlo Simulation Framework

The following section outlines the development and methods of the Monte Carlo simulation framework that forms a significant portion of this PhD project as a body of work. The development of the framework can be split into three main sections; the generation of the simulation substrate, the development of the simulation dynamics, and the generation of the synthetic diffusion MR data. Each of these sections have involved various stages of development. Where possible, each of these stages will be presented and discussed.

4.3.3.1 Substrate Generation

The substrate of the Monte Carlo simulation framework represents the environment through which the simulated walker molecules are propagated. Within the context of this project, the aim is to create a simulation substrate that recreates the microstructural environment of tumour tissue as closely as possible. Biological tissue can broadly be described as consisting of three discrete elements, interchangeably referred to as ‘compartments’. Therefore, the development of the substrate generation element of the project addressed each compartment separately, although the intracellular and extracellular spaces are intrinsically linked.

Intracellular/Extracellular Compartments

The first compartment is the intracellular space, and this represents the space within the cells of the tissue. In reality, these cells can take a variety of forms depending on the tissue.




Figure 4.1: Left: H&E stained histology of a giant cell tumour (bone). Right: H&E stained histology of a neurofibroma tumour. By Nephron - Own work, CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=14831566>

Figure 4.1 shows two examples of the different cellular morphologies found within tumour tissue. The left-hand side image shows a histology slice from a giant cell bone tumour, while on the right is a histology slice from a neurofibroma tumour. The difference between the cellular morphologies of these two tumours is apparent. The giant cell tumour exhibits large ellipsoid-shaped cells, while the cells of the neurofibroma tumour are much more elongate in nature. These differences are reflective of the healthy tissues from which the tumours originate. The cells themselves are packed together in a chaotic fashion, separated by the extracellular space. Typical intracellular volume fractions can range from 0.2 up to around 0.7 depending on the tumour type [166, 172].

The aim when developing the intracellular compartment of the simulation substrate was to create a representation of the cellular microstructure of tumours that was simple to implement, while trying to capture as much of the complexity of the structure as possible. The MATLAB function *tissuemodel.m* was written to generate the simulation substrates, and includes the code for creating the intracellular, extracellular and, if required, the vascular tissue compartments. In the early development stages the simulations were limited to two compartment substrates without simulated vasculature.

The *tissuemodel.m* function generates the intra/extracellular volume fraction using a sphere-packing algorithm. The function takes the following user inputs:

- Intracellular Volume Fraction
- Extracellular Volume Fraction
- Vascular Volume Fraction
- X-dimension
- Y dimension
- Z dimension
- Diffusivity
- Cell Radius Parameters

These parameters characterise the properties of the target simulation substrate, and are used by the subroutines within the *tissuemodel.m* function. The cell radius parameters define the distribution of cell radius values within the simulation substrate/voxel (for example the mean and standard deviation)

Intracellular and Extracellular Volume Compartments

The first step in generating the intracellular compartment is to calculate the volume of the substrate. This is simply calculated as the product of the three user-specified inputs:

$$V_{target} = X_{dim} \times Y_{dim} \times Z_{dim} \tag{4.1}$$

This is then used to determine the target intracellular volume:

$$V_{ic} = f_{icvf} V_{target} \quad (4.2)$$

A random cell-packing algorithm is then implemented. A flow-diagram of the algorithm is shown in figure 4.2. After initialising the loop parameters the first random cell is generated. In particular, a loop parameter that tracks the current intracellular volume of the substrate is initialised (set to 0). Simulated cells within the substrate are represented parametrically by a centre coordinate and a radius.

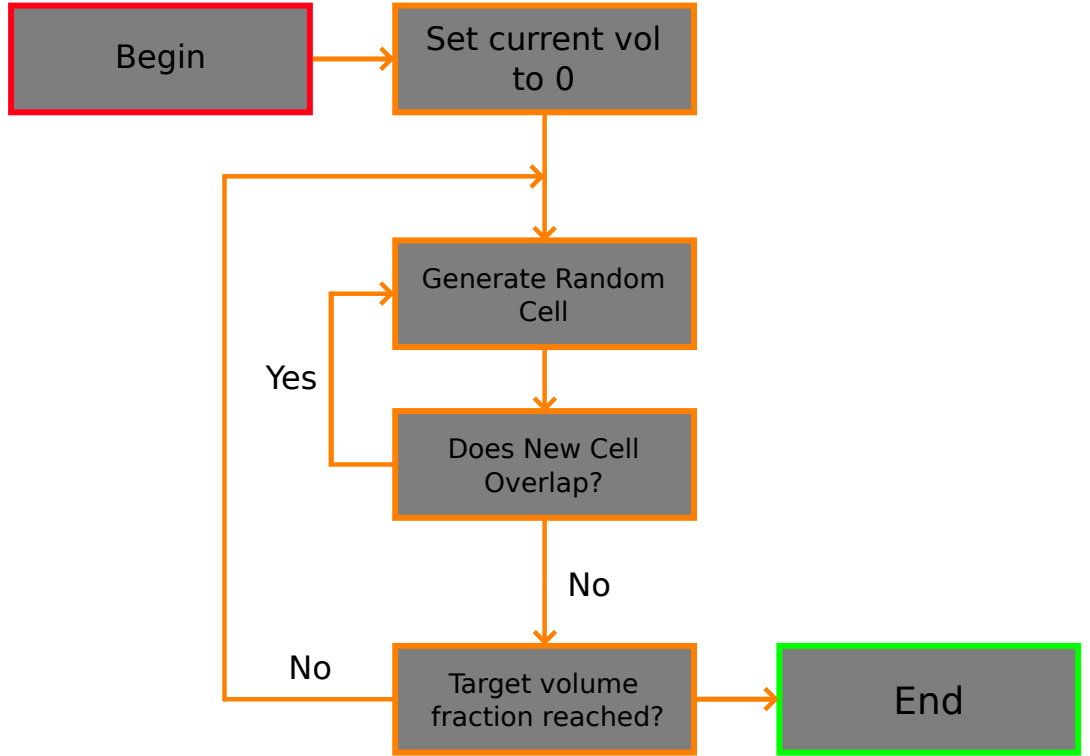


Figure 4.2: Flow diagram showing the algorithm used to pack spherical cells into the tissue substrate.

Random numbers between 0 and the size of the substrate along each cartesian coordinate direction are generated. These three random numbers are used as the centre coordinate of the cell. The radius of the cell is then calculated. The method for doing so depends on the inputs provided to the *tissuemodel.m* function. The most commonly-used inputs are a mean radius and standard deviation for the substrate cell population. In this case, the cell radius is selected by sampling a normal distribution with the mean and standard deviation specified. The current intracellular volume parameter is then updated:

$$V_{icurrent} = V_{icurrent} + V_{cell} \quad (4.3)$$

When generating subsequent cells, the algorithm includes an additional step that involves checking whether the new proposed cell would overlap any of the existing cells. After the centre coordinate and radius is selected a separate function checks the proposed cell against each of the previously generated cells to check if the proposed cell overlaps. An illustration of this process is shown in figure [4.3](#).




Figure 4.3: An illustration of the process of checking proposed cells for overlap.

If the proposed cell overlaps a previously generated cell, it is rejected and a new cell is proposed using the methods described above. Otherwise, the new cell is accepted and added to the simulation substrate. The volume of the new cell is then added to the current intracellular volume fraction. At the end of each loop iteration (generating one new cell), the current intracellular volume is compared with the target value. If the current volume is greater than or equal to the target value then the process completes. Otherwise, the algorithm will continue.




Figure 4.4: An example output from the sphere packing algorithm described above. The parameters used to generate the substrate shown were mean radius $10\mu m$, standard deviation $5\mu m$. A large standard deviation was only used in order to give a clearer illustration of the variation in cell radius that can be achieved using this algorithm.

Figure 4.4 shows an example substrate produced using the sphere-packing algorithm. In this case, the substrate produced comprises two compartments: the intracellular and extracellular space. The extracellular space is defined by the tight spaces in between and around the packed spheres, in which the extracellular water molecules will diffuse. As previously discussed, the majority of previous work in Monte Carlo simulations of biological tissue have been limited to this two-compartment approach. The main goal of this project was to extend this framework to account for the effect of water molecules diffusing/flowing within the blood vessel and capillary network.

Vascular Compartment - Part 1

The most challenging aspect of this project was developing a realistic representation of the vascular structure. The aim was to capture both the complex structure of the capillary bed and large feeder vessels, as well as the functional properties. The most important functional property is the speed at which water molecules flow through the vessels as this has the greatest influence on the resulting DW-MRI signal.

Part 1 of the development of the vascular substrate began with a random network generating algorithm, that joining together a randomly-generated cloud of seed points (vertices) using a 3D Minimum Spanning Tree (MST) algorithm.

The *tissuemodel.m* function is modified with an additional input ‘vesseldensity’, which is a three-element vector that indicates the number of vessel vertex coordinates to be placed along each of the substrate dimensions. It is therefore analogous to the density of the resulting vascular network.

The first step of the network generation process is to generate a random cloud of points within the simulation substrates, evenly distributed along each of the three dimensions. The number of X, Y and Z coordinates are generated according to the vessel density parameter. These points will form the vertices of the vascular network, which will be joined together using the MST algorithm. The algorithm was implemented using the TREES toolbox developed by Hermann Cuntz et al [173]. The toolbox was originally created to generate and analyse neuronal branching structures, however for the purposes of this project, the implementation of the MST algorithm proved very efficient for generating simulated vascular structures.

The toolbox implements a greedy version of the MST algorithm. Firstly, a root point is chosen (normally the origin of the simulation voxel $[0,0,0]$). The closest point to the any part of the tree is then selected and connected to the tree according to two conditions: 1. The euclidean distance between the selected point and the point on the tree. 2. The resulting path length from the point to the root point of the tree. A the cost function used to select how to connect each subsequent point is a balanced sum of these two factors, weighted accordingly by the ‘balancing factor’. A low balancing factor ($bf \approx 0.1$) will result in a highly compartmentalised tree structure, while a high balancing factor ($bf \approx 1$) will result in a direct connection between each vertex and the root. Figure 4.5 shows the resulting connected vessel structure.




Figure 4.5: A random vessel network generated using random seed points joined together using a MST algorithm. The figure shows a network produced using 1000 vertices ($vesseldensity = [1000, 1000, 1000]$, with root vertex at $(0, 0, 0)$).

Once the vessel structure has been generated, the convex hull is calculated using the *hull_tree.m* function. The function defines a high-resolution three-dimensional grid of points over the vessel structure, and then calculates the closest tree-point to each point in

the grid, and stores the distance for each point. The MATLAB *isosurface.m* function is then used to define a 3D contour that surrounds the vessel structure at a specified distance (corresponding to the vessel radius). The 3D contour is then converted into a 3D mesh of faces and vertices, that form the outer wall of the simulated blood vessels. An example of the resulting simulated vasculature is shown in figure 4.6.




Figure 4.6: The resulting simulated vessel network produced from the vessel structure of figure 4.5. The vessel radius in this case was specified as $5\mu m$ throughout the network.

The methods for generating a simulated vessel network presented above allow for a great deal of flexibility. By tuning the parameters used, vessel networks with a wide range of properties can be created. The framework also allows for multiple networks to be generated within a single simulation substrate:




Figure 4.7: Multiple vessel networks generated within a single simulation substrate, using the methods described in part 1.

However, this vessel-generation approach was abandoned when a more physiologically realistic approach became available using real segmented tumour vessel networks. The second approach is described below, which is the approach used for the rest of the work presented in this chapter.

Vascular Compartment - Part 2

It became apparent during the development of the first method for generated simulated vessel structures that calculating realistic flow-velocities within the resulting networks would be a challenge. In addition, the tree-shaped structures are not necessarily representative of the vascular morphology of real-world tumours. With this in mind, a new approach for creating a realistic vascular was pursued.




Figure 4.8: Images showing the various processing stages involved in extracting the vessel structures for use in the simulation framework. (a) A surgically resected LS174T human colon adenocarcinoma tumour. The tumour was grown subcutaneously on the right flank of an immunocompromised mouse. (b) The resected tumour following optical clearing using the benzyl alcohol, benzyl benzoate (BABB) protocol [3]. (c) OPT image of the cleared tumour, resolution $5-10\mu m$. (d) Segmentation produced by Gaussian-subtraction, Frangi filtering, thresholding and skeletonisation. (e) Schematic representing the conversion of the segmented data into a 3D spatial graph format. (f) Resulting 3D graphical representation of the full tumour vascular network.

The vessel structures are generated from optical projection tomography (OPT) data from real subcutaneous tumours. During the perfusion fixation process by which the animals are culled, a systemic injection of fluorescently conjugated lectin is administered which binds to the walls of the blood vessels. The tumours are then resected and optically

cleared so that they can be imaged with the OPT system [3, 174]. By using an optical filter that matches the fluorescence wavelength of the fluorophore, it is possible to capture a 3D image of the blood vessel network of the tumour, down to a resolution of $5 - 10\mu m$ [175]. The fluorescent signal from the vessels is segmented from the background noise by subtracting a 3D Gaussian-filtered image ($25 \times 25 \times 25$ kernel), then applying a multiscale Frangi filter to enhance the vessel structures [176]. The medial axis transform (MAT) of the segmented vessels is then used to extract a skeletonised version of the structure. The result of the skeletonisation process is then converted into a graph format (nodes and segments), which can be used in the simulation to dictate the directions of walkers diffusing within the vascular volume. An example of a fully-segmented vascular structure, converted into graph format is shown in figure 4.9.




Figure 4.9: Segmented OPT data from an entire subcutaneous tumour, skeletonised and converted into graph format. The vessel network has been overlaid on structural MR data of the subcutaneous tumour on which the OPT imaging was performed.

Once the vessel structure has been extracted using the processing framework described above, the next step is to estimate the expected steady-state blood flow values throughout the network. These blood flow values are necessary for determining the motion of the walkers propagating through the vessel network during the simulation.

Blood flow values were estimated using the REAListic Numerical Image-based Modelling of biologicAl Tissue substratEs (REANIMATE) framework developed by D’Esposito, Sweeney et al [4]. The framework (shown in figure 4.10)

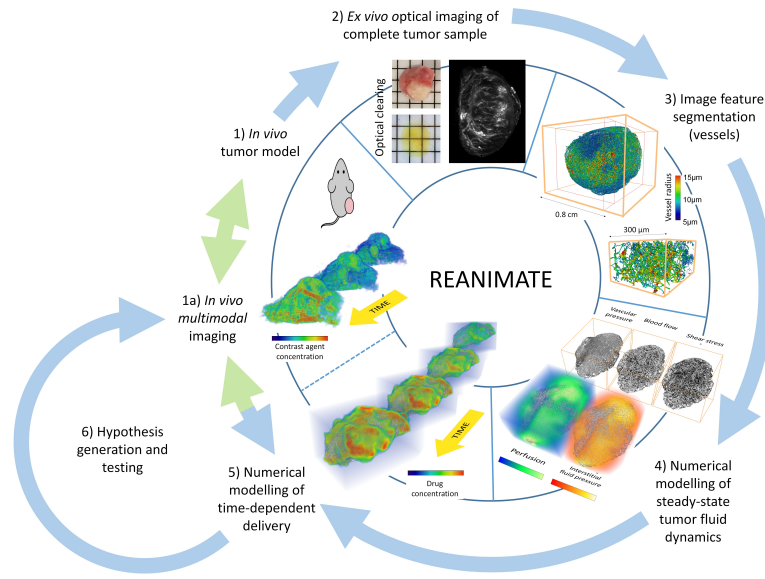


Figure 4.10: Illustration of the REANIMATE modelling pipeline. The pipeline includes the complete process for estimating the steady-state blood flow values for complete vessel networks extracted from real tumours [4]

The framework implements a discrete network model for estimating the steady state volumetric blood flow throughout the network (at figure 4.10 stage (4)). The discrete network model calculates flow values according to Poiseuille’s law:

$$F_{volume} = \frac{P_1 - P_2}{R} \quad (4.4)$$

Where $P_1 - P_2$ represents the pressure drop over the vessel segment, and R is the 'resistance to flow' which is dependant on the viscosity of whole blood and the dimensions of the segment:

$$R = \frac{8\eta L}{\pi r^4} \quad (4.5)$$

Where η is the blood viscosity, L and r are the length and radius of the vessel segment respectively. By providing the boundary conditions to the simulation in the form of the pressure drop at one of the inlet vessels, the resultant volumetric blood flow rate can be calculated over the entire tumour network (see 4.11). Unknown boundary conditions are estimated based on minimising the deviation of network pressures and wall shear stresses from specific target values [177] The volumetric flow rate can be converted into a blood flow velocity using the radius of the vessel segment.

Voxel-sized portions are cropped from the whole tumour network, and used to form the vascular compartment of the simulation substrate. The simulation substrate is then packed with virtual cells using the sphere-packing algorithm described previously. An example of a completed simulation substrate is shown in figure 4.11.

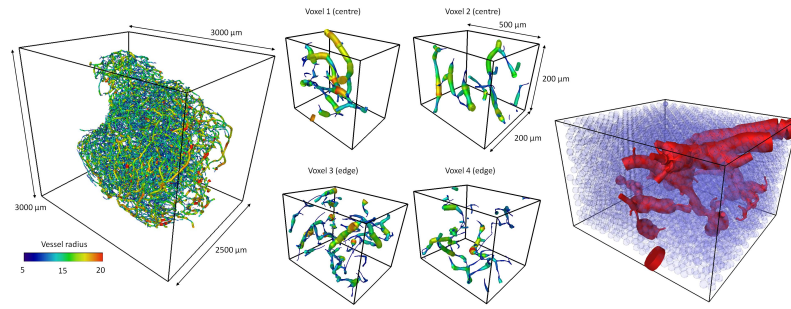


Figure 4.11: Left: A full simulated tumour vessel network, produced using the REANIMATE framework. Middle: 4 examples of voxel-sized portions cropped from the complete tumour network, two from the middle of the tumour and two from the edge. Right: A complete simulation substrate comprising spherical cells packed around the vessel network from one of the example voxels.

Now that the simulation substrate is complete, the next step in the simulation framework is to implement the simulation dynamics, that simulate how water molecules interact with the representation of the tumour microstructure.

4.3.3.2 Simulation Dynamics

Walkers in the intracellular or extracellular space are propagated through the simulation substrate following a random-walk procedure. Walkers within the vessel network are propagated through the network at the blood flow velocity which is dependent on their current vessel segment. The simulation dynamics are implemented in the *tissuediffusion.m* function, written in the MATLAB programming language. The basic structure of the code takes inspiration from the ‘MCML’ code developed by Lihong Wang and Stephen Jaques [178]. The method involves propagating one walker at a time, calculating the trajectory step-by-step through the simulation substrate based until the elapsed simulated time reaches the time duration specified by the user (specified to be at least as long as a DW-MR acquisition). The total number of steps taken by each walker is determined by the

time resolution of the simulation, specified by the user. The time resolution corresponds to the length of time covered by a single step. For example, a typical duration of 0.2seconds with a time resolution of 1×10^{-4} seconds would result in each walker propagating 2000 steps through the substrate. The function takes the following inputs:

- Substrate structure - The output of the *tissuemodel.m* function.
- Numwalkers - The number of walkers to be propagated.
- Duration - The desired duration of the simulation.
- Temporal resolution - Time per step.
- Initialisation - Specifies the characteristics of the initial walker positions.
- Poolsize - The number of CPU cores to utilise for parallel processing.
- Nanremoveflag - A binary flag determining whether to scan the output for NaN values (legacy code from when things were behaving strangely)

The first step is to initialise the walker within the substrate. The positions at which the walkers are initialised are specified by the user and parsed to the function via the ‘initialisation’ variable described above. There are a number of possible options including:

- ‘random_full’ - Walker positions are chosen at random within the substrate, can be within intracellular, extracellular or vascular compartment, with relative numbers of walkers in each compartment corresponding to the volume fractions of the substrate.
- ‘random_quartile’ - Most common option. Walker positions are chosen at random within the central region defined by the inner quartiles along each dimension. Walkers can be within intracellular, extracellular or vascular compartment, with relative

numbers of walkers in each compartment corresponding to the volume fractions of the substrate.

- ‘random_cells’ - Walkers are initialised at the centre coordinate of a random cell.
- ‘fixed_cells’ - Walkers are initialised at the centre coordinate of cells, in the order they were generated (repeatable).
- ‘vasc’ - Walkers are initialised at a random vessel node within the vascular structure.
- ‘vasc_single_seed’ - Walkers are initialised at a specified vessel node within the vascular structure.

After initialising the remaining walker properties such as elapsed time and current step size the initialised walker is propagated through the substrate. The initial walker propagation direction is determined by the *scatter.m* function that performs both the initial direction calculation and all subsequent direction determinations based on the walker’s compartment and position within the substrate. The propagation direction of the walker is represented by directional cosines ux , uy , uz , that represent the angles between the propagation vector and the three coordinate axes.

If the walker is not in the vascular compartment, the new direction is chosen randomly and can take any value. The directional cosines are first assigned random values between $-1 \leq u \leq 1$. The values are then scaled to ensure that the sum of the squared directional cosines is equal to 1. Accordingly:

$$ux_{temp} \in U(-1, 1)$$

$$uy_{temp} \in U(-1, 1)$$

$$uz_{temp} \in U(-1, 1)$$

$$\begin{aligned} ux &= \sqrt{\left\| \frac{ux_{temp}}{(\|ux_{temp}\| + \|uy_{temp}\| + \|uz_{temp}\|)} \right\|} \times \text{sign}\left(\frac{ux_{temp}}{(\|ux_{temp}\| + \|uy_{temp}\| + \|uz_{temp}\|)}\right) \\ uy &= \sqrt{\left\| \frac{uy_{temp}}{(\|ux_{temp}\| + \|uy_{temp}\| + \|uz_{temp}\|)} \right\|} \times \text{sign}\left(\frac{uy_{temp}}{(\|ux_{temp}\| + \|uy_{temp}\| + \|uz_{temp}\|)}\right) \\ uz &= \sqrt{\left\| \frac{uz_{temp}}{(\|ux_{temp}\| + \|uy_{temp}\| + \|uz_{temp}\|)} \right\|} \times \text{sign}\left(\frac{uz_{temp}}{(\|ux_{temp}\| + \|uy_{temp}\| + \|uz_{temp}\|)}\right) \end{aligned}$$

The step size is then calculated using the Einstein diffusion equation:

$$s = \sqrt{6Dt_s} \quad (4.6)$$

Where D is the diffusivity of the substrate (assigned during the substrate generation), and t_s is the time per step (or ‘temporal resolution’) of the simulation.

If the walker is in the vascular compartment, the *scatter.m* function chooses the new direction based on the walker’s current position within the vessel network. The next vertex is then determined based on the closest network vertex downstream of the current vertex. The flow direction at the current position in the network is calculated by the difference in pressure between the current vertex and those either side. A positive pressure difference indicates the direction of flow. In the case where the vessel network is bifurcating (the

walker is currently at the bifurcation point), the next vertex is determined probabilistically based on the relative pressure drops between the current vertex and the vertices of the two bifurcations. The step size for walkers in the vascular compartment is calculated by determining the blood flow velocity based on the vessel radius and the volumetric flow rate at the current position.

Once the propagation direction and step size have been determined, the next step is to update the walker position. However, before this is done it must be determined whether or not the proposed step size and direction would cause the walker to hit a boundary within the tissue. The three boundaries considered are the cell walls, blood vessel walls and the edge of the simulation substrate. Each is handled in a simplistic manner: if the proposed step would cause the walker to cross any of the above boundaries, the step is abandoned and a new step generated.

If it is determined that the proposed step will not cause the walker to cross any boundaries within the tissue, then the walker position is updated as follows:

$$walker.x = walker.x + (walker.s * walker.ux)$$

$$walker.y = walker.y + (walker.s * walker.uy)$$

$$walker.z = walker.z + (walker.s * walker.uz)$$

After updating the walker position, the elapsed time of the simulation is increased by the time per step (according to the temporal resolution). A check is then performed to ensure that the time elapsed is less than the desired simulation duration. If the elapsed time is less than the desired duration, the process is repeated generating a new step and updating the walker position. If the elapsed time is greater than the desired duration, the walker

is marked as ‘dead’ and propagation finishes. After each walker has been propagated, the walker trajectory is stored. The trajectory data contains a record of the walker position, time elapsed, step number and step length for every step taken during the simulation.

Once the propagation of a single walker through the substrate has completed, the process repeats for the next walker and so on until the specified number of walkers has been propagated through the substrate. The computational cost of running Monte Carlo simulations can be very high; a typical simulation can involve propagating tens of thousands of walkers, resulting in long waiting times. To reduce the amount of time taken per simulation, the code makes use of parallel processing. The MATLAB parallel-processing toolbox allows for multiple walkers to be propagated on different CPU cores simultaneously, meaning the time taken to run a simulation can be decreased by a factor depending on the number of CPU cores available. On an 8-core processor, this allows a simulation propagating 10,000 walkers to be completed in 7 minutes, as opposed to almost an hour when run on a single core.

When all walkers have been propagated, the function exits and returns the large matrix containing the trajectory data for each walker, as well as statistics about the simulation including time taken, and the number of walkers propagated in each volume compartment. The trajectory data is then used to generate synthetic DW-MR data.

4.3.3.3 Diffusion Data Synthesis

Synthetic data is generated using the *scan* function of the Camino diffusion MRI toolkit [168]. The toolkit calculates the diffusion signal based on the positions of the walkers within the substrate at each time point, for different acquisition parameters specified by the user.

The walker positions are read into the *scan* function in the form of a trajectory file (.traj file). The data contained in the trajectory file comprises the position history of every simulated walker at every time step of the simulation, stored in big-endian binary format. Understandably given the nature of the data that they contain, trajectory files can be very large - possibly tens of gigabytes per file.

A matlab function *WriteTraj.m* was written to convert the trajectory data output from the random-walk simulation into a .traj file. The function starts by re-ordering the trajectory data. Due to the walker-by-walker propagation approach of the simulation the trajectory data output is ordered by walker index first, followed by time step and position. The .traj format requires that the trajectory data is ordered by time step first, followed by walker index and position. After re-ordering the data the function initiates a for-loop that iterates through each time step of the trajectory data and writes the walker index and other data to the .traj file as follows:

- Time Step - 64-bit double.
- Walker Index - 32-bit int.
- x-coordinate - 64-bit double.
- y-coordinate - 64-bit double.
- z-coordinate - 64-bit double.

The resulting .traj file is then parsed to the *scan* function along with a ‘scheme’ file containing the acquisition parameters for the DW-MR acquisition protocol. For example, for a 220-acquisition VERDICT scheme, the scheme file contains 220 rows, each containing the gradient strength, direction, duration and separation and echo time corresponding to each acquisition.

After parsing the trajectory and acquisition scheme to the *scan* function, the synthetic data is generated by calculating the acquired phase offset from each walker along each gradient direction by:

$$\begin{aligned}\phi_x &= \gamma \int_0^\delta G_x(t)x(t)dt \\ \phi_y &= \gamma \int_0^\delta G_y(t)y(t)dt \\ \phi_z &= \gamma \int_0^\delta G_z(t)z(t)dt\end{aligned}$$

Where γ is the gyromagnetic ratio, G_{xyz} is the time-dependent gradient strength along the x,y, and z-direction and $x(t)$, $y(t)$ and $z(t)$ are the time dependent walker positions in the x, y and z-direction. The total signal for each synthetic acquisition is then calculated by:

$$S(G, \Delta, \delta) = S(0, \Delta, \delta) \int_{-\infty}^{\infty} P(\phi) \cos(\phi) d\phi \quad (4.7)$$

Where $S(0, \Delta, \delta)$ is the signal calculated with no diffusion weighting and $P(\phi)$ is the walker phase distribution [179].

4.3.4 Model Fitting

Model fitting was performed using the Camino toolkit, see section 3.4.1.6 for details. The VERDICT ‘BallSphereSticks’ and ‘BallSphereAstrosticks’ were both fitted to the synthetic data generated using the simulation framework described above, along with the ADC and IVIM models.

4.3.5 Verification

The simulation framework itself was validated by generating synthetic DW-MR signal using simple substrate geometries for which analytical models for the DW-MR signal have also been derived. This is a popular approach for validating Monte Carlo models [135]. The *tissuemodel.m* function was altered to generate a substrate with either uniformly-packed or randomly-packed parallel cylinders. This substrate was used alongside the original packed spheres for validation purposes.

Callaghan [137] derived the exact solutions for DW-MR signals arising from parallel cylinders and spheres with impermeable, perfectly-reflecting walls. The analytically-calculated diffusion signals for different acquisition parameters can be plotted against the signal generated from Monte Carlo simulations run on corresponding substrate geometries as a way of validating the simulation dynamics.

Two validation experiments were performed, using signal calculated using two different analytical models. Firstly, the *tissuemodel.m* code was modified to generate a substrate comprising parallel packed cylinders with user-specified radius. The cylinders were either arranged using cylindrical packing (a pseudo two-dimensional version of the spherical packing algorithm), or were regularly arranged on a grid. For validation purposes, a cylinder radius of $8\mu m$ was chosen, along with an intracellular volume fraction of 0.6, and diffusivity of $10 \times 10^{-9} m^2 s^{-1}$. Data was generated using a typical VERDICT acquisition scheme of 220 acquisitions, with 46 b-values and 3 directions. The synthetically and analytically-generated data was then normalised, and plotted in the order of increasing b-value. The average relative difference between the two Monte Carlo models and the analytical model were calculated.

The process outlined above was then repeated using a simulation substrate comprising packed spheres, with a sphere radius of $8\mu m$, intracellular volume fraction of 0.4 and diffusivity $10 \times 10^{-9} m^2 s^{-1}$.

4.3.6 Application to Validation of Diffusion Models

The Monte Carlo simulation framework described above was applied as a tool for validating mathematical diffusion models, and investigating the sensitivity of their associated parameters to different substrate characteristics. The parameters of four diffusion models were validated; the ADC, IVIM and VERDICT ‘BallSphereStick’ (BSS) and ‘BallSphere-Astrosticks’ (BSA) models.

4.3.6.1 VERDICT Validation

The primary objective of the simulation framework was to be apply the framework as a tool for validating the mathematical models of the VERDICT MRI technique.

The first step of the validation process was to run multiple simulations, varying the parameters of the simulation substrate. The substrates were generated using a real blood vessel network captured from an LS174T adenocarcinoma tumour grown subcutaneously on the flank of an immuno-compromised mouse, simulated using the REANIMATE methods described previously. Four voxel-sized portions of the full network were cropped out to form the vascular volume compartment of the simulation substrates, shown in figure [4.11](#).

Voxels 1 and 2 were taken from the centre of the tumour, and voxels 3 and 4 were taken from the edge of the tumour. Summary vascular parameters of each voxel, and the network

as a whole, are shown in table 4.1.

	Whole Network	Voxel 1 (Centre)	Voxel 2 (Centre)	Voxel 3 (Edge)	Voxel 4 (Edge)
Blood Volume Fraction	0.028	0.025	0.017	0.080	0.026
Mean Vessel Radius	6.96 ± 4.53	5.6 ± 4.2	5.6 ± 4.1	8.0 ± 5.2	7.6 ± 4.1
Mean Blood Flow	4.82 ± 34.8	0.07 ± 0.40	0.15 ± 0.89	22.9 ± 59.0	9.43 ± 25.60

Table 4.1: Table showing the summary vascular parameters of each voxel cropped from the network, as well as the overall parameters of the network as a whole.

Diffusion-weighted MR data was then generated from each simulation run, in accordance with the methods described above. A scheme file representing a typical *in-vivo* VERDICT 46 b-value, 3-direction acquisition was used to generate the synthetic data with 220 acquisitions in total.

After generating the synthetic data, the two VERDICT models were fitted to the data using the Camino toolkit. The model parameters corresponding to the fits were then plotted against the corresponding substrate parameter (for example, the cell radius parameter plotted against the cell radius of the simulation substrate). A positive correlation between the substrate parameter and fitted parameter value would indicate that the model parameters are truly representative of the microstructural properties that they aim to measure.

A line of identity was added to the correlation plots, as this represents the ideal case where the substrate parameter and the fitted parameter correlate 1-1. Linear regression was then performed on the correlation data to fit linear lines of best fit. ANCOVA was then used to compare the gradients of the resulting linear models with the line of identity to evaluate whether they are statistically significantly different.

The substrate parameters were varied according to table 4.2 below. As discussed previously, the intracellular walker fraction and vascular walker fraction represent the proportionate number of walkers propagated through the intracellular and vascular spaces respectively.

Parameter Name	Range	Step Size
Cell Radius	$5\mu m - 15\mu m$	$1\mu m$
Intracellular Walker (Volume) Fraction	0 - 1	0.05
Vascular Walker (Volume) Fraction	0 - 1	0.05

Table 4.2: Table showing the ranges over which the substrate parameters were varied.

As well as correlating the substrate parameters with their corresponding model parameter (i.e. Intracellular walker fraction with intracellular volume fraction), each model parameter was also plotted against the other varied substrate parameters. This was done to investigate whether any of the VERDICT model parameters are cross-correlated with other substrate parameters. The presence of such cross-correlation could adversely affect the usefulness of the parameters as biomarkers for detecting changes in the microstructural properties of tissue.

4.3.6.2 ADC and IVIM

The validation process above was then repeated for two diffusion models that are widely used in a clinical setting: the ADC and IVIM models. This was in order to investigate how their associated fitted parameters are correlated with different characteristics of the tissue microstructural environment. The same batches of simulations were run as for the VERDICT model validation experiment, varying the substrate cell radius, intracellular walker fraction and vascular walker fraction. The two models were then fitted to the resulting synthetic signal, and the correlation between each fitted model parameter and

the varied substrate parameter were plotted. The ADC has a single fitted parameter of the same name. The IVIM model has 3 main parameters: D (Diffusion), D^* (Pseudo-diffusion) and F_v (Pseudo-diffusion volume fraction).

When fitting the ADC model, b-values below 100 were omitted from the fitting process. This is routinely performed in clinical applications of the ADC model [180].

4.3.6.3 Effect of Blood Flow

The results of the VERDICT validation experiments prompted a further investigation into the effect of the mean blood flow on the resulting correlation between vascular walker fraction and the fitted value of the vascular volume fraction. This was done by repeating the methods of the VERDICT validation experiment above for the vascular volume fraction parameter a number of times. For each repeat, the blood flow values of the substrate vessel network were multiplied by a scalar factor. The scalar factor was varied between 0 and 2, in intervals of 0.1.

The substrates were generated from a single selected voxel (voxel 3) using the same methods described previously, but with the blood flow velocities scaled up or down by the constant scaling factor. The simulations were then run, increasing the vascular walker fraction (the relative number of walkers propagated through the vascular space, analogous to the vascular volume) from 0 to 1. The synthetic signal was then generated for each run, and the two VERDICT models were fitted to the resulting data. The resulting fitted values of vascular volume fraction were plotted against the specified values of vascular walker fraction, for each set of simulations run with increasing mean blood flow velocity.

4.3.6.4 Noise Sensitivity

A further aspect of diffusion modelling that can be explored using a Monte Carlo validation approach is the robustness of the mathematical models to noise. Mathematical models that can measure intricate details of biological tissue are rendered useless in a clinical setting if the performance of the models is too sensitive to a decrease in SNR of the MR signal.

To assess the robustness of the VERDICT models to signal SNR, signals were generated from the same simulation substrate with varying SNR values. The generated signals were then fitted with the two VERDICT models, and the accuracy of the resulting cell radius parameter for each SNR value was compared. The cell radius parameter in particular was chosen as this is the main parameter that has been investigated as a potentially promising biomarker for cancer imaging.

A substrate was generated with $8\mu m$ cell radius, and other substrate parameters assigned typical values. High SNR signal ($SNR \approx 100$) was generated by propagating a large number of walkers through the substrate (500,000 walkers, resulting in a .traj file of 50Gb in size). The resulting signal was then corrupted with Rician noise of varied amplitude, to produce signals with SNR values of 50, 40, 30, 20, 17.5, 15, 12.5 and 10. Multiple signals were generated for each SNR value.

The two VERDICT models were then fitted to the synthetic data. The fitting process itself is repeated multiple times for each signal, initialising the model parameters at different starting values to reduce the chances of the minimisation process finding a local minimum. For each signal and for each SNR, the difference between the fitted cell radius value and the actual value was calculated then a mean value calculated accross all the signals at each SNR value.

The resulting accuracy vs SNR data can then be plotted to investigate how the models behave as the SNR decreases, and assess the overall robustness of the models at clinically realistic SNR levels.

4.4 Results

In the following section, the results from the various simulation-based experiments outlined above will be presented.

4.4.1 Validation of Monte Carlo Framework

The first validation experiment performed was to generate synthetic diffusion signal from a substrate comprising packed parallel cylinders. Synthetic signal was generated using the Monte Carlo framework of the current project and the framework included with the Camino diffusion MRI toolkit. Further signal for comparison was then calculated using a corresponding analytical model. The three signal curves are shown in figure [4.12](#).




Figure 4.12: Validation of the Monte Carlo simulation framework performed by generating synthetic diffusion-weighted signal from a substrate comprising parallel cylinders (radius $8\mu m$), then comparing the synthetic diffusion-weighted signal with signal generated using a corresponding analytical model. Also shown is synthetic signal generated using the Monte Carlo simulation tool included in the Camino diffusion toolkit, using a parallel cylinders substrate (radius $8\mu m$).

Looking at figure 4.12, it can clearly be appreciated that both the Monte Carlo simulation frameworks produce synthetic signal that is closely aligned with the signal calculated from the analytical model. Further quantitative analysis showed that the signal produced by the presently-developed Monte Carlo simulation framework differed from the analytical model by an average of 3.9%, compared to 8.4% for the Camino toolkit.

The validation experiment was then repeated for a simulation substrate comprising packed spheres with radius $8\mu m$. The corresponding three signal curves are shown below in figure 4.13.




Figure 4.13: Validation of the Monte Carlo simulation framework performed by generating synthetic diffusion-weighted signal from a substrate comprising packed spheres (radius $8\mu m$), then comparing the synthetic diffusion-weighted signal with signal generated using a corresponding analytical model. Also shown is synthetic signal generated using the Monte Carlo simulation tool included in the Camino diffusion toolkit, using a spherical mesh substrate (radius $8\mu m$).

The results shown above show an even closer relationship between the two Monte Carlo models and the analytical gold standard. The presently-developed Monte Carlo framework differed from the analytical model by an average of 4.8%, compared to 3.2% for the Camino toolkit.

In both validation experiments, the Monte Carlo models differed from the analytical model to a greater extent at higher b-values.

4.4.2 VERDICT Validation

4.4.2.1 Intracellular Walker (Volume) Fraction

The results of the intracellular volume fraction validation experiments are shown below in figure 4.14.

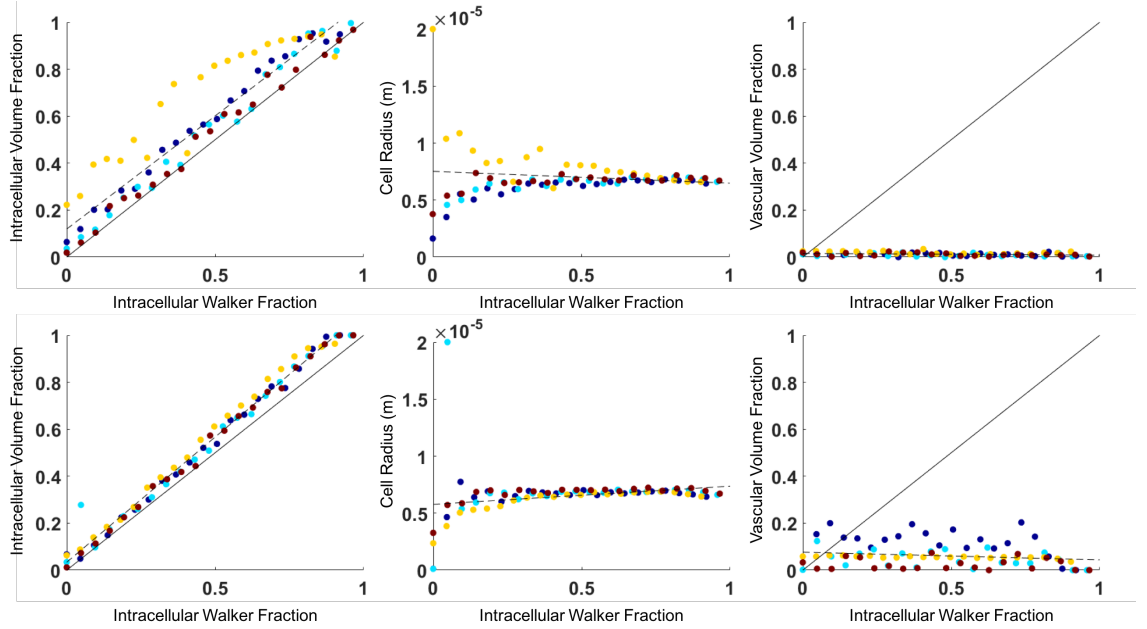


Figure 4.14: Parameter correlation plots for the intracellular walker fraction plotted against fitted model parameters produced using the VERDICT ‘BallSphereStick’ model (top row) and ‘BallSphereAstrosticks’ model (bottom row). The plots show the ground truth values (x-axis) taken from the simulation, plotted against the fitted parameter values produced by the VERDICT models. From left to right: Intracellular walker fraction vs intracellular volume fraction, intracellular walker fraction vs cell radius, intracellular walker fraction vs vascular volume fraction. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The solid black line shows the 1-1 correlation line. The dashed black line is the line of best fit fitted to the average of the four voxels.

Shown above in the top row of figure 4.14 are the results from the ‘BallSphereStick’ (BSS) model (with anisotropic vascular diffusion compartment).

All four voxels show a correlation between intracellular walker fraction and the fitted intracellular volume fraction. ANCOVA revealed that the best fit line gradients of three of the four voxels (1, 2 and 4) were not significantly different to the 1-1 model ($p < 0.001$). Voxel 3 exhibited a weaker correlation than the other 3 voxels, and diverges both from the best case and the other three voxels. Despite this, the best fit line of the mean data does not differ from the ideal best case ($p < 0.001$), but slightly overestimated the intracellular volume fraction by an average of 0.11. There was no correlation found between intracellular walker fraction and cell radius or vascular volume fraction.

The bottom row of figure 4.14 shows the results from the ‘BallSphereAstrosticks’ (BSA) model, with an isotropic vascular diffusion compartment.

All four voxels exhibit a strong correlation between the intracellular walker fraction and the fitted intracellular volume fraction. ANCOVA revealed that all four voxels achieved a correlation that was not significantly different to the 1-1 model ($p < 0.01$). The same was true for the best fit line of the mean data, which was statistically identical to the 1-1 model at the 1% significance level. Similarly to the BSS model, the BSA results slightly overestimate the intracellular volume fraction, by an average of 0.06. Again, there was no correlation between the intracellular walker fraction and either the fitted cell radius parameter or the fitted vascular volume parameter.

4.4.2.2 Cell Radius

The results of the cell radius validation experiments are shown below in figure 4.15.

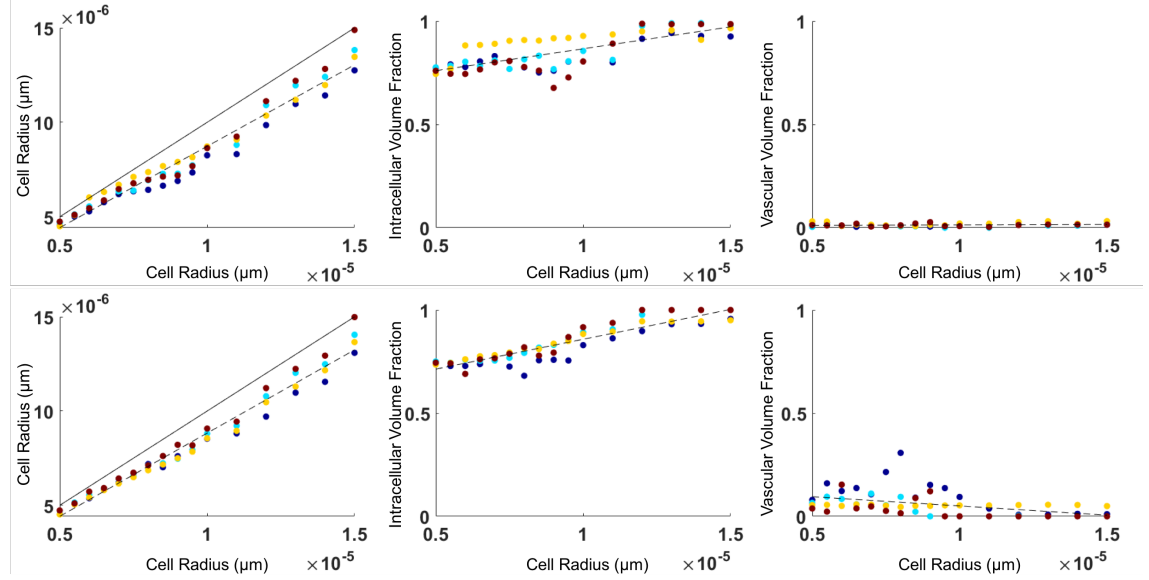


Figure 4.15: Parameter correlation plots for the substrate cell radius plotted against fitted model parameters produced using the VERDICT ‘BallSphereStick’ model (top row) and ‘BallSphereAstrosticks’ model (bottom row). The plots show the ground truth values (x-axis) taken from the simulation, plotted against the fitted parameter values produced by the VERDICT models. From left to right: Substrate cell radius vs fitted cell radius, substrate cell radius vs intracellular volume fraction, substrate cell radius vs vascular volume fraction. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The solid black line shows the 1-1 correlation line. The dashed black line is the line of best fit fitted to the average of the four voxels.

The top row of figure 4.15 shows the results from the BSS model. All four voxels achieved a strong correlation between the substrate cell radius and the fitted cell radius parameter (mean $R = 0.99$, $p < 0.00001$). ANCOVA further revealed that voxels 2 and 4 were not significantly different to the 1-1 best case model ($p = 0.08996$, $p = 0.0757$ respectively). However, a significant difference was found between the mean data (dashed black line)

and the ideal case. The fitted cell radius parameter underestimated the actual substrate cell radius by an average of $1.12\mu m$. There was no significant correlation found between the substrate cell radius and the fitted intracellular volume fraction or the fitted vascular volume fraction.

The bottom row of figure 4.15 shows the results from the BSA model. There was no significant difference between the cell radius (substrate) vs cell radius (fitted) correlations achieved with the BSS and BSA models. Moreover, there was again no significant correlation found between the substrate cell radius and the fitted intracellular volume fraction and vascular volume fraction parameters. Similarly to the BSS model, the BSA model underestimated the cell radius value by an average of $1.11\mu m$.

4.4.2.3 Vascular Walker (Volume) Fraction

The results of the vascular walker fraction validation experiments are shown in figure 4.16.

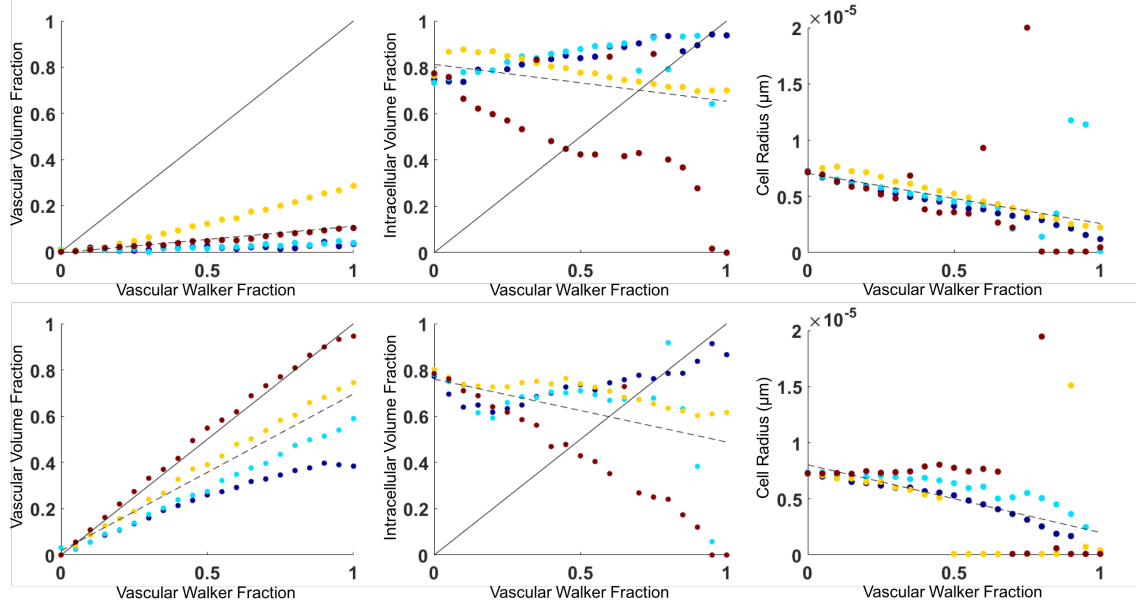


Figure 4.16: Parameter correlation plots for the substrate vascular walker fraction plotted against fitted model parameters produced using the VERDICT ‘BallSphereStick’ model (top row) and ‘BallSphereAstrosticks’ model (bottom row). The plots show the ground truth values (x-axis) taken from the simulation, plotted against the fitted parameter values produced by the VERDICT models. From left to right: Vascular walker fraction vs fitted vascular volume fraction, vascular walker fraction vs intracellular volume fraction, vascular walker fraction vs fitted cell radius. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The solid black line shows the 1-1 correlation line. The dashed black line is the line of best fit fitted to the average of the four voxels.

The top row of figure 4.16 shows the results from the BSS model. Each of the four voxels showed a correlation between the substrate vascular walker (volume) fraction and the fitted vascular volume fraction value, though the correlations were weak. As can be seen looking at the top-left plot, all four voxels significantly underestimated the vascular volume of the substrates. This result was confirmed by ANCOVA, which unsurprisingly found that all of

the voxels demonstrated a correlation significantly different to the 1-1 ideal model. There was also a noticeable variation in the correlation gradients between each voxel, with voxel 3 (orange) demonstrating a more positive correlation than voxels 1 and 2. There was no correlation found between the vascular walker fraction and the fitted intracellular volume fraction. Surprisingly, there was a slight negative correlation found between the vascular walker fraction and the fitted cell radius value.

The bottom row of figure 4.16 shows the results from the BSA model. These results were broadly consistent with those of the BSS model, but the correlations between the substrate vascular walker fraction and fitted vascular volume parameters were much stronger. The correlation demonstrated by voxel 4 was statistically identical to the 1-1 ideal case ($p < 0.0001$). The remaining three voxels were strongly correlated but were significantly different to the 1-1 ideal model. Similarly to the BSS results, there was a noticeable variation in the correlation slopes between the voxels. Again, there was no correlation between the vascular walker fraction and the fitted intracellular volume fraction parameter. Meanwhile, there was no significant correlation between the vascular walker fraction and fitted cell radius.

4.4.2.4 ADC and IVIM models

As previously mentioned, the methods of the VERDICT model validation experiment for which the results are presented above were repeated for the ADC and IVIM models. Both of these, and in particular the ADC, are widely used in a clinical setting so understanding what tissue characteristics can potentially influence their parameter values is of key importance.

The results for the ADC model are shown below in figure 4.17.




Figure 4.17: Parameter correlation plots for the ADC model. From left to right: Cell Radius vs ADC, Intracellular Walker Fraction vs ADC, Vascular Walker Fraction vs ADC. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The dashed black line is the line of best fit fitted to the average of the four voxels.

The left-hand plot of figure 4.17 shows the correlation between the ADC value and substrate cell radius. The ADC increased showed a strong correlation with substrate cell radius for all four substrate voxels (mean $R = 0.91$). There was little variation between the four voxels, with ANCOVA showing the correlation gradients of all four voxels were not statistically significantly different (Significance level 0.05). The opposite was true for the intracellular walker fraction vs ADC, where a strong negative correlation was found for all four substrate voxels (mean $R = -0.94$). As with the cell radius results, there was little variation between the four voxels, with no significant difference between the correlation gradients of the four voxels. The right-hand plot of figure 4.17 shows the results for the vascular walker fraction. All four voxels demonstrated a strong negative correlation between vascular walker fraction and ADC (mean $R = 0.86$). As can be seen, there was a large degree of variation between the four voxels. There was no significant difference between the correlation gradients of voxels 1 and 2. Voxels 3 and 4 had much shallower correlation gradients.

The results for the IVIM model are shown below in figure 4.18.

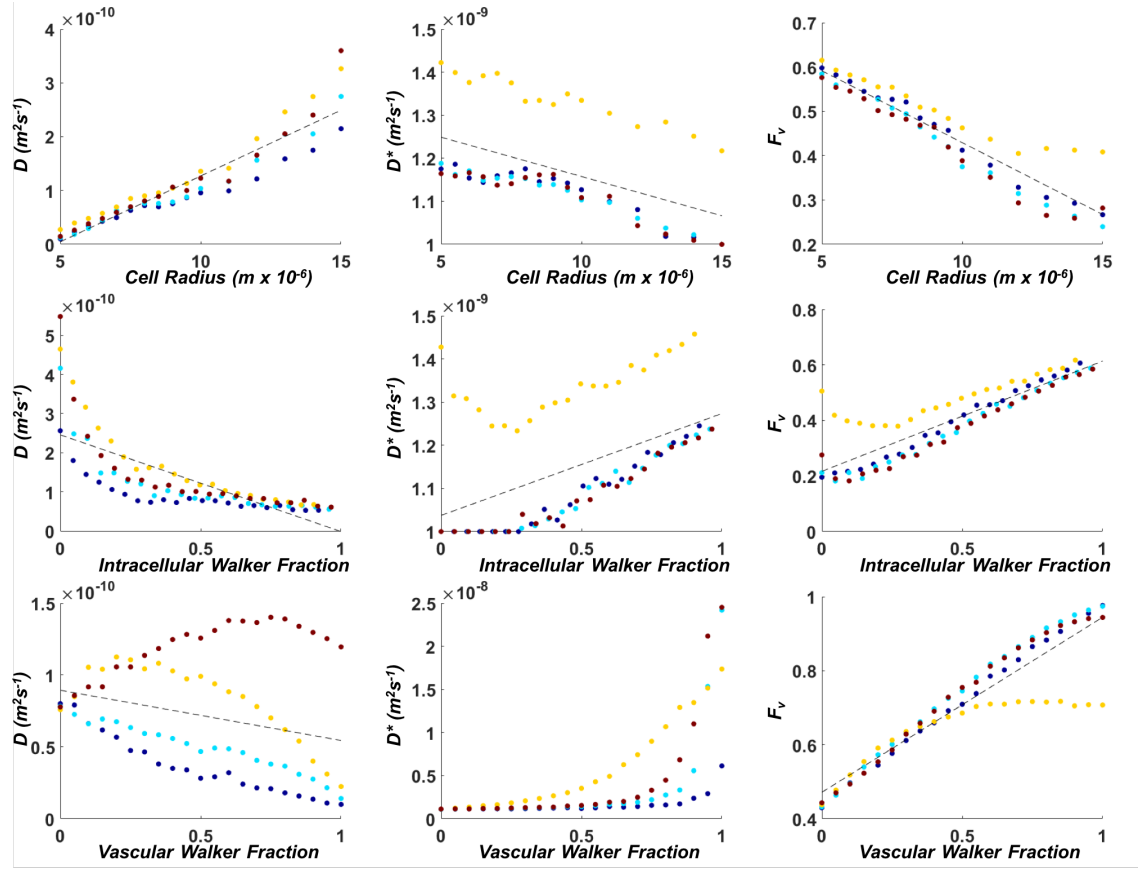


Figure 4.18: Parameter correlation plots for the IVIM model. Top row from left to right: substrate cell radius vs D , substrate cell radius vs D^* , substrate cell radius vs pseudo-diffusion volume. Middle row from left to right: intracellular walker fraction vs D , intracellular walker fraction vs D^* , intracellular walker fraction vs pseudo-diffusion volume. Bottom row from left to right: vascular walker fraction vs D , vascular walker fraction vs D^* , vascular walker fraction vs pseudo-diffusion volume. The four different colours represent the four voxels cropped out of the vessel network (Dark blue - voxel 1, light blue - voxel 2, orange - voxel 3, red - voxel 4). The dashed black line is the line of best fit fitted to the average of the four voxels.

The top row of figure 4.18 shows the results for the experiments with varied substrate cell radius. The left-hand side plot shows the correlation between substrate cell radius and diffusivity. A strong positive correlation was found for all four voxels (mean $R = 0.86$).

There was a slight variation between the four voxels, with ANCOVA revealing voxels 3 and 4 exhibiting a steeper correlation than voxels 1 and 2, though a significant difference was only found between voxels 1 vs 3 and 1 vs 4 (significance level 0.05). The middle plot shows the correlation between substrate cell radius and pseudo-diffusion coefficient (D^*). A strong negative correlation was found for all four voxels (mean $R = 0.92$). There was no significant different difference between the correlation gradients of the four voxels. However, as can be seen the pseudo-diffusion coefficient values for voxel 3 were considerably higher than the values from voxels 1, 2 and 4. The right-hand plot shows the correlation between substrate cell radius and psuedo-diffusion volume fraction. A strong negative correlation was found for all 4 voxels (mean $R = 0.88$). There was a significant difference between the correlations of voxel 3 and voxels 1, 2, and 4.

The middle row of figure 4.18 shows the results for the substrate intracellular walker fraction. A negative correlation can be seen between the intracellular walker fraction and diffusivity. It appears that the relationship between the two is not linear in nature, with the diffusivity dropping rapidly between an intracellualr walker fraction of 0 and 0.3, then a more gradual decrease between 0.3 and 1. Visual inspection of the plot shows that all four voxels demonstrated a broadly similar trend. The middle plot shows the correlation between intracellular walker fraction and pseudo-diffusion coefficient (D^*). All four voxels exhibit a positive correlation. Similar to the diffusivity results, the results from all four voxels show a different trend between intracellular walker fraction values from 0 to 0.3, compared to between 0.3 and 1. Voxels 1, 2 and 4 showed very similar results, with no significance between their correlation gradients between intracellular walker fraction values of 0.3 and 1. The D^* values for voxel 3 were significantly different, though the correlation gradient for intracellular walker fraction values above 0.3 was statistically indistinguishable from the other three voxels.

The bottom row of figure 4.18 shows the results for the substrate vascular walker fraction. The left-hand side plot shows the correlation between vascular walker fraction and diffusivity. As can clearly be seen, there was a significant degree of variation between the results from the four simulation substrate voxels. Voxels 1 and 2 exhibited a strong negative correlation ($R = 0.89$ and $R = 0.87$ respectively). Meanwhile, voxels 3 and 4 exhibited a more complex relationship. The diffusivity values for voxel 3 increased slightly for vascular walker fractions below 0.25, then decreased for values above 0.25. A similar relationship was found for voxel 4, but with the diffusivity values peaking for a vascular walker fraction of 0.75. The middle plot shows the correlation between the substrate vascular walker fraction and the pseudo-diffusion coefficient (D^*). A positive, non-linear correlation was exhibited by all four substrate voxels. The results of voxels 1, 2, and 4 were broadly consistent with a gradual increase in (D^*) up to a vascular fraction of around 0.8, then an almost exponential increase for vascular fractions above 0.8. Voxel 3 showed a steeper increase for lower vascular fractions, but with a closer to linear increase in (D^*) from vascular fractions of around 0.6 and above. The right-hand plot shows the correlation between substrate vascular walker fraction and pseudo-diffusion volume. The results show a strong positive correlation for all four voxels up to a vascular walker fraction of 0.5 (mean $R = 0.96$), with ANCOVA no significant difference between the correlation gradients of all voxels. Above a vascular volume fraction, the results of voxels 1, 2, and 4 remain the same, while the results of voxel 3 diverge and the pseudo-diffusion values are roughly constant with increasing vascular fraction.

4.4.3 Effect of Blood Flow Velocity

The results of the vascular volume fraction validation experiment shown above prompted a further investigation into the possible cause of the variation in correlation gradient observed. It was hypothesised that the mean blood flow velocity within the vessel network

could by some mechanism affect the fitted value of vascular volume fraction, and hence the correlation achieved by the VERDICT models. The results of the blood flow velocity experiments are shown below in figure 4.19.

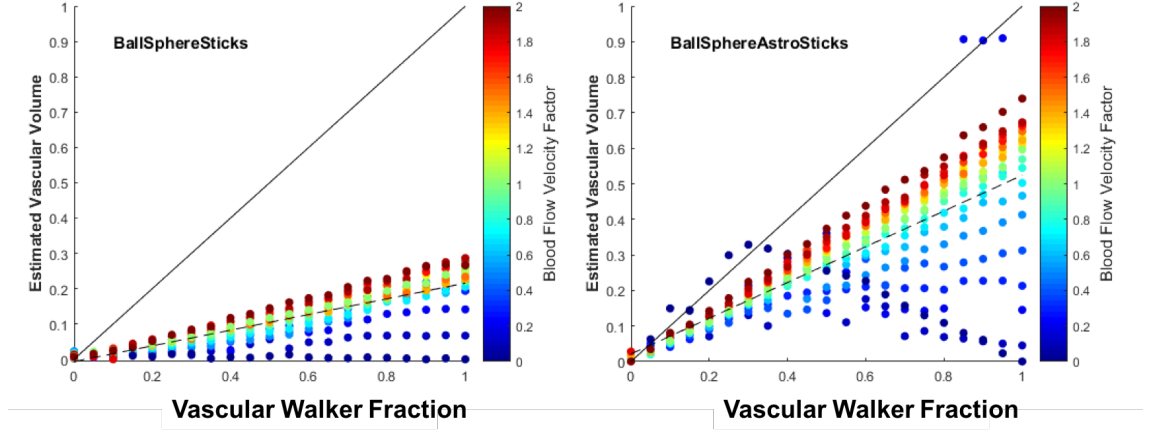


Figure 4.19: Parameter correlation plots for the substrate vascular walker fraction plotted against the fitted vascular volume fraction parameter produced using the VERDICT ‘BallSphereStick’ model (left) and ‘BallSphereAstroSticks’ model (right) for substrates with scaled blood flow velocities. The substrates were all generated from voxel 3, from the VERDICT validation experiments previously described. The colours represent the blood flow velocity scaling factor (Blue = 0, Red = 2)

As shown in figure 4.19, the results achieved by the VERDICT BSS (left) and BSA (right) models show similar trends. The slope of the correlation between the substrate vascular walker fraction and the fitted value of vascular volume fraction changes as a function of the blood flow velocity. In particular, the slope of correlation *increased* as a function of increasing blood flow velocity. As previously described, the correlations achieved by the BSA model were closer to the 1-1 ideal model than those of the BSS model. In both models, when the blood flow scaling factor became close to zero, the behaviour of the vascular volume parameter became less predictable. This can be seen more clearly in the results from the BSA model, where the relationship between vascular walker fraction and fitted vascular volume fraction becomes noticeably more random at lower blood flow values (shown in dark blue).

4.4.4 Noise Sensitivity

As discussed previously, a further aspect of diffusion modelling that can be tested using the developed Monte Carlo framework is the sensitivity of diffusion models to signal noise. The results of the noise sensitivity investigation experiments are shown below in figures 4.20 and 4.21.

Figure 4.20 below shows the results produced by the VERDICT BSS model.

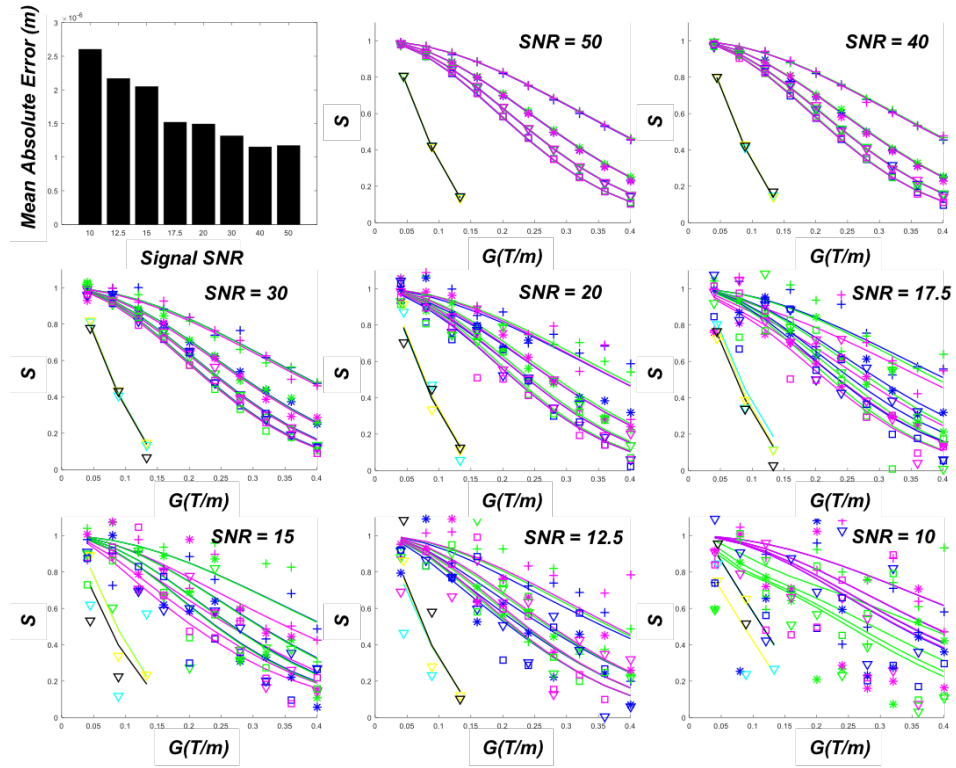


Figure 4.20: SNR sensitivity results produced by the VERDICT ‘BallSphereStick’ (BSS) model. Top left: A bar graph showing the signal SNR value plotted against the mean absolute error of the fitted cell radius parameter. Remaining plots: Show the fits of the VERDICT BSS model to the synthetic data. Each plot shows gradient strength plotted against signal amplitude for the synthetic VERDICT data corrupted with various levels of noise. The points (stars, squares, triangles) show the synthetic data and the lines show the analytically calculated signal calculated from the parameters of the BSS model fitted to the data.

The top-left hand plot of figure 4.20 shows the synthetic (simulated) signal SNR plotted against the mean absolute error. To re-cap, the ‘mean absolute error’ is the average amount by which the fitted cell radius model parameter overestimates or underestimates the substrate cell radius. The performance of the BSS model was stable from an SNR of 50 (error $1.15 \times 10^{-6}m$) down to an SNR of 17.5, with an increase in the mean absolute error of $0.34 \times 10^{-6}m$ between the two. Below an SNR of 17.5, the mean error increased, up to $2.58 \times 10^{-6}m$ at a signal SNR of 10. A qualitative analysis of the fits appears to support this result, with model appearing to fit the data well until the SNR reaches drops down as far as 17.5. At SNR levels of 15 and below, the synthetic (simulated) data appears almost randomly scattered, and the BSS model is unable to produce a sensible fit.

Figure 4.21 below shows the results produced by the VERDICT BSA model.

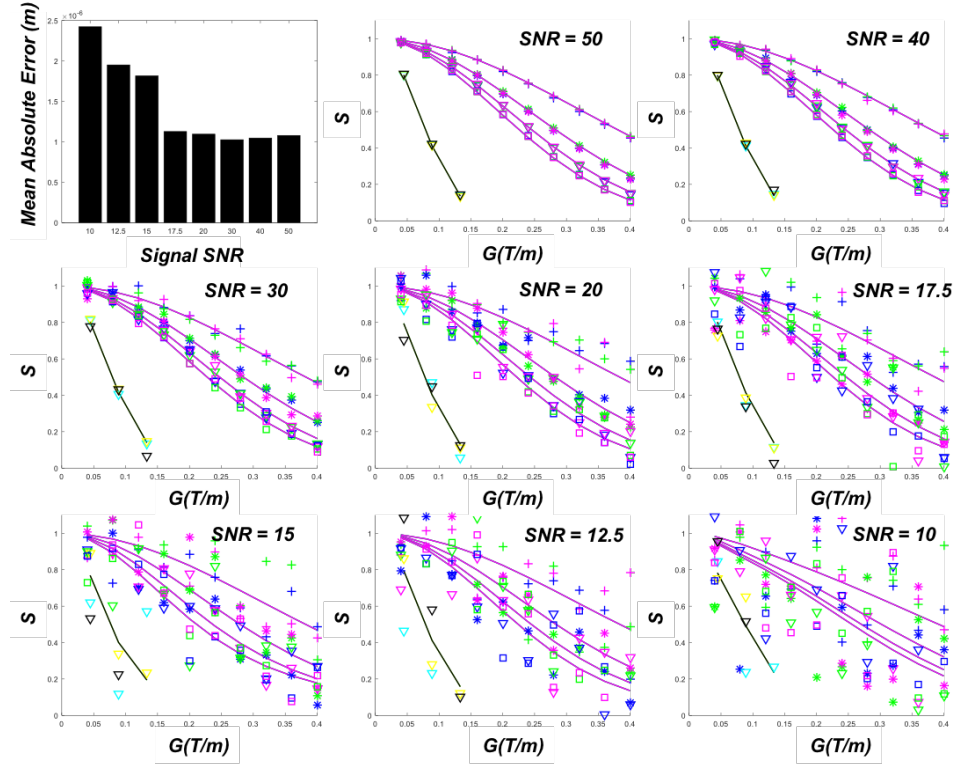


Figure 4.21: SNR sensitivity results produced by the VERDICT ‘BallSphereAstrosticks’ (BSA) model. Top left: A bar graph showing the signal SNR value plotted against the mean absolute error of the fitted cell radius parameter. Remaining plots: Show the fits of the VERDICT BSA model to the synthetic data. Each plot shows gradient strength plotted against signal amplitude for the synthetic VERDICT data corrupted with various levels of noise. The points (stars, squares, triangles) show the synthetic data and the lines show the analytically calculated signal calculated from the parameters of the BSA model fitted to the data.

The results achieved by the BSA model were broadly consistent with those of the BSS model. The top left hand plot of figure 4.21 shows the synthetic (simulated) signal SNR plotted against the mean absolute error. As with the BSS model, the performance of the BSA model was stable from an SNR of 50 (error $1.07 \times 10^{-6}m$) down to an SNR of 17.5. The increase in error was considerably smaller than that of the BSS model at $0.02 \times 10^{-6}m$. Below an SNR of 17.5, there was a step-change in the mean absolute error,

which was $1.81 \times 10^{-6}m$ at an SNR of 15. The error continued to increase to $2.46 \times 10^{-6}m$ at an SNR of 10. Overall, the mean error achieved by the BSA model was lower than that of the BSS model by an average of $0.24 \times 10^{-6}m$. A qualitative analysis of the fits shows the same result as that of the BSS model, with the BSA model capable of providing a good fit to the synthetic data down to an SNR of 17.5, after which the data loses any recognisable structure and the performance of the BSA model deteriorates.

4.5 Discussion

The aim of this body of work was to develop a Monte Carlo model of water diffusion in biological tissue that includes the effect of water diffusing/flowing in the microvascular network. Further, to apply the developed model to a validation study of the mathematical models that form an intrinsic part of the VERDICT diffusion MRI framework. This was motivated by the need for thorough validation of diffusion models, in particular for diffusion models that are applied in tissues that are well perfused, such as those found in tumours. In these tissues, the effect of water molecules diffusing/flowing in the microvasculature is much more significant. It was therefore appreciated that a validation method that takes this effect into account would help in the translation of emerging diffusion MRI techniques into the clinical setting.

4.5.1 Validation of Monte Carlo Framework

The purpose of the first set of simulation experiments was to validate the Monte Carlo framework itself, to ensure that the synthetic signal generated was representative of the chosen substrate geometry. The first validation experiment involved generating signal from a substrate comprising uniformly-packed parallel cylinders with zero permeability.

The synthetic signal generated by the Monte Carlo framework developed for this work was compared with signal calculated using an analytical signal model of parallel cylinders, and signal generated using an established Monte Carlo diffusion simulation framework included in the Camino diffusion toolkit. The results from this experiment demonstrated that the Monte Carlo framework generated synthetic signal that was consistent with the signal calculated using the analytical model. Furthermore, the developed framework outperformed the Camino toolkit diffusion simulation framework, in that the signal generated by the developed framework was closer on average to the analytical model, with an average 3.9% difference vs 8.4% for the Camino toolkit.

In the second experiment, the process was repeated but in a simulation substrate comprising packed spheres with zero permeability. The results from this experiment show that both the developed framework and the Camino toolkit were closely matched to the analytical model with average differences of 4.8% and 3.2% respectively. The results further demonstrated that the difference in performance between the developed framework and the Camino toolkit was smaller. In this case, the Camino toolkit outperformed the developed framework with the generated signal being closer on average to the analytical model.

In both experiments, the signal generated from the developed framework was greater in magnitude than those of both the analytical model and the Camino toolkit. One possible explanation for this is the way in which the developed framework handles boundary interactions. As previously discussed, if a generate step would cause a walker to cross a boundary the step is rejected and a new step generated. This method is discussed in the comparison paper by Xing et al [181]. Handling boundary interactions in this way can cause the boundaries themselves to possess an effective ‘thickness’, wherein walkers can be reflected from the boundaries from a maximum of 1 step length away. In the substrates used for the validation, this would shrink the effective size of the cylinders or spheres in

which the walkers were diffusing. The average signal loss for those populations of walkers would then be lower, causing the slight increase in overall signal seen.

Overall however, the results of the validation experiments demonstrate that the diffusion characteristics of the developed Monte Carlo framework are accurate, and the framework itself produces synthetic signal that is consistent with the gold standard analytical case and an established MC modelling framework. However, one of the limitations of this validation work was the lack of validation for the microvascular flow aspect of the simulation. The challenges of validating flow-sensitive measurements have been discussed previously [135], with no currently available physical phantoms providing a suitably controllable and reproducible solution. The results from the VERDICT parameter validation sections of this chapter go some way towards validation by demonstrating that the microvascular flow characteristics of the simulation are at least correlated with flow-sensitive parameters fitted by the VERDICT models. Future work on the validation of this framework should focus on more thoroughly and explicitly validating the flow element of the framework.

4.5.2 VERDICT Validation

The purpose of the VERDICT validation experiments was to apply the developed simulation framework as a tool for validating the compartmental diffusion models of the VERDICT technique. In addition, the performance of two popular VERDICT models - the ‘BallSphereStick’ and ‘BallSphereAstrosticks’ models were compared.

Some of the experiments were focussed on validating the intracellular volume fraction parameter. The results showed that the fitted value of intracellular volume fraction was strongly correlated with the intracellular walker fraction of the simulation for both models, with an almost 1-1 exact correspondence between the measured and actual values. The

‘BallSphereAstrosticks’ performed slightly more consistently across the different simulation substrates, particularly for correlations between intracellular walker fraction and intracellular volume fraction. Overall, correlations of both models were virtually identical. Furthermore, both models demonstrated that there is no apparent correlation between either the substrate cell radius or vascular walker fraction and the fitted intracellular volume fraction. These results are encouraging, as they suggest that the intracellular volume fraction parameter of both VERDICT models is sensitive to changes in the substrate intracellular walker fraction, but is not influenced by changes in the substrate cell radius or vascular walker fraction.

Another VERDICT parameter that has shown promise as a biomarker for cancer development and therapy response is the cell radius measurement. The results showed that the measured value of cell radius was strongly correlated with the substrate cell radius value, with close to a 1-1 correspondence found. Meanwhile there was no correlation found between the intracellular walker fraction and the cell radius. However, a weak negative correlation between the vascular walker fraction and cell radius was found. There was reassuringly almost no noticeable difference in performance between the ‘BallSphereStick’ and ‘BallSphereAstrosticks’ models, showing that both are equally sensitive to changes in the substrate cell radius, and equally specific.

One possible reason for this is a significant number of vessel segments within the vascular substrates had very low flow velocities. A current limitation of the simulation framework is the fact that walkers propagating in the vascular compartment do so at the blood-flow velocity, neglecting the effects of diffusion. In segments where the blood flow velocity is very low, this could cause the walkers to become almost stationary, rather than following a conventional random walk within the blood vessel as would be expected. The limited motion of these near - stationary walkers in the vascular compartment could appear the same as restricted diffusion of walkers within the cells.

Some of the *in-vivo* work in recent years has suggested that the VERDICT parameters associated with the vascular compartment can behave unpredictably, and are less reliable. This was one of the main motivating factors for developing the diffusion and flow validation framework presented in this chapter. The results outlined above were a very positive indication of the potential of the vascular compartment parameters. Both VERDICT models exhibited a strong positive correlation between the vascular volume fraction and the corresponding substrate vascular volume fraction. This was coupled with a lack of correlation with the substrate cell radius or intracellular volume fraction for both models. Although the two VERDICT models exhibited similar behaviours in terms of achieving strong correlations, the vascular volume values of the BSA model were generally higher. Furthermore, the correlation slopes of the BSA model were steeper. One possible reason for this is the different way in which the BSS and BSA models fit the diffusion signal from the vascular compartment. As previously discussed, the BSS model assumes the vascular signal is anisotropic, whereas the BSA model makes no assumptions of anisotropy. Because of this the isotropic vascular diffusion model of the BSA model can achieve a better fit, and this is then reflected in higher vascular volume values.

This is the first time that the vascular model parameters have been properly validated in such a way, and suggests that these parameters may have some potential for use as disease biomarkers. However, the significant variation between the different subvoxels used as the vascular substrate warranted further investigation.

Another limitation of the simulation framework is that the simulation dynamics do not account for the permeability of membranes between the intracellular/extracellular space, or intravascular/extracellular space. In real biological tissue, water molecules can permeate through cell membranes and vessel walls. Vessel wall permeability is a particularly significant effect in tumour tissue, where blood vessels are known to become ‘leaky’ [182]. A possible effect of this on the synthetic data is a higher vascular volume fraction than

would be expected, since the simulated water molecules flowing in the blood vessels cannot diffuse out into the extracellular space.

4.5.3 Other Diffusion Models

Unsurprisingly, the ADC value demonstrated strong correlations with all of the substrate parameters. This is not surprising as the ADC is a summary value that is meant to represent the general freedom with which water in tissue can diffuse. As discussed, the value of the ADC is influenced by many factors, a statement which is supported by the experimental results presented.

The results achieved by the IVIM model were similar, with all three model parameters influenced by the changes in value of all three varied substrate parameters. Promisingly, the perfusion fraction (pseudo-diffusion volume) was strongly positively correlated with the vascular walker fraction. However, its value was also influenced by the intracellular walker fraction and substrate cell radius.

The fact that the ADC value, and the values of the IVIM model parameters are influenced by a combination of substrate properties suggests that their values may be more difficult to interpret. In a clinical setting, the sensitivity to multiple factors makes the models less reliable. Despite the fact that both models are widely applied in a clinical setting, the results of this chapter provide further evidence that the absolute values of their parameters should be interpreted with caution. This result further encourages and supports the development of more advanced diffusion models that aim to separate out the various contributions to the overall diffusion signal.

4.5.4 Effect of Blood Flow Velocity

In view of the results of the validation experiments focussed on the vascular compartment parameters, further simulations were run to investigate the observed variability in vascular volume correlation between the different simulation substrates.

The results showed that the blood flow velocity within the substrate has a significant effect on the correlation between vascular walker fraction and the fitted value of vascular volume. Faster blood flow velocities caused an increase in vascular volume fraction values. On the other hand, when mean blood flow velocities approached zero, the behaviour of the model became unstable and there was no strong correlation between vascular walker fraction and vascular volume fraction.

The exact mechanism behind this is unclear. However, one possible explanation links back to the significant number of vessel segments within the segments with very low blood flow velocity values. Due to the fact that the simulation dynamics do not include the effect of diffusion for walkers within the vasculature, when the blood flow velocity approaches zero the behaviour of the water molecules becomes unrealistic. When the blood flow velocity in a given segment is zero, the walker will remain stationary rather than undergoing a random walk as would be expected. This could explain why, when the blood flow velocity factor was small, the behaviour of the model became unstable since some of the segments within the vessel network would have blood flow velocities of effectively zero. When the blood flow velocity of the substrate is increased by a constant factor, these velocity values increase and the motion of the walkers is then fast enough to be fitted by the vascular compartment models. This could then be reflected in an increase of the calculated vascular volume fraction, and an increase of the correlation slope general. Likewise when the blood flow velocity became extremely low, the motion of the walkers within the vascular compartment

became indistinguishable from diffusion, making it difficult for the vascular models to fit.

The results of the increased blood flow velocity experiments are not conclusive proof that the observed effect would occur when fitting *in-vivo* data. Multiplying all of the blood flow velocity values throughout the vessel network is unlikely to give a realistic representation of how the flow velocity values would change throughout a vessel network given an increase or decrease of blood flow. A more realistic approach could be to increase the pressure drop boundary conditions of the finite element blood flow model, to generate a whole blood vessel network with increased overall blood flow.

Despite the limitations of the current simulation framework, the results still provide a valuable insight into the behaviour of the VERDICT diffusion models. Promisingly, the vascular volume fraction always exhibited a strong correlation with the vascular walker fraction. Changes in blood flow velocity may affect the overall fitted values, yet the strong correlation remains present.

4.5.5 Noise Sensitivity

The final set of simulation experiments performed as an application of the developed Monte Carlo diffusion and flow framework were focussed on assessing the sensitivity of the VERDICT models to signal noise.

The signal-to-noise ratio (SNR) is a value that reflects the magnitude of the MR signal compared to the magnitude of the background noise. A low value of SNR indicates that a signal is very "noisy" i.e. that the magnitudes of the useful signal and noise are similar. Typically, SNR values are higher for MR scanners with a greater B_0 field strength. Pre-clinical scanners with a 9.4T field strength can achieve SNR values of around 50 for a typical pulsed-gradient spin echo sequence. Meanwhile, clinical scanners often have

a weaker static field (typically around 3T as I write this), though higher-field strength scanners are becoming increasingly popular. Typical clinical SNR values are difficult to come by, though estimates of ~ 25 are common.

It is crucial that diffusion models are robust to noise if they are to be applied in a clinical setting. If a diffusion model performs well at the high SNR values found in pre-clinical imaging, but then cannot achieve comparable results on a clinical scanner, the translation of the model into clinical use becomes a more significant problem.

The ‘BallSphereStick’ model experienced a slight decrease in accuracy between a high SNR of 50 and low SNR of 17.5. However, the performance of the BSS model only appeared to drop off below an SNR of 17.5. The ‘BallSphereAstrosticks’ model performed better, showing no significant decrease in accuracy from an SNR of 50 to 17.5, with a significant decrease in accuracy below an SNR of 17.5.

These results suggest that the performance of both VERDICT models is stable at SNR values typically encountered during both pre-clinical and clinical imaging. The results were further supported by the plots of the fits to the synthetic data, showing that the models were able to achieve good fits to the data above an SNR of 15. On the other hand, the results also suggest that care should be taken when adjusting imaging parameters in order to reduce acquisition times at the cost of signal.

Despite the promising results presented in this chapter, future work should focus on further investigation into the robustness of the VERDICT models to noise. Further experiments could look into the accuracy of additional parameters, and quantitatively assess the quality of the fits to the synthetic data. Additional work could also investigate other potential VERDICT models, with diffusion compartments taken from the taxonomy by Panagiotaki et al [6].

4.6 Chapter Summary

The aim of this chapter was to firstly to develop a Monte Carlo simulation framework that includes the effects of diffusion and flow. The simulation aims to recreate the motion of water molecules within both the intravascular and extravascular spaces. Secondly, to apply the developed Monte Carlo framework as a tool for validating various mathematical models of diffusion including the ADC, IVIM and VERDICT models.

The first half of the chapter was focused on the development of the simulation framework, which itself was split into three main sections and represents a significant portion of this thesis. The three main aspects of the framework development were the substrate generation, simulation dynamics and signal generation. The results presented in this chapter demonstrate that the resulting simulation framework is a valuable tool for investigating the performance and behaviour of diffusion models in detail.

The application of the simulation framework as a validation tool successfully demonstrated that the parameters produced by the VERDICT model are sensitive to the microstructural properties that they are intended to represent. Furthermore, analysis of the sensitivity to noise showed that the VERDICT models are highly robust and perform well over the SNR ranges typically encountered during pre-clinical or clinical imaging.

The approaches demonstrated here have a number of limitations. The simple spherical representation of the cells in the simulation substrate is clearly a vast oversimplification of the intracellular and extracellular space of actual tumour tissue. Walkers propagating within the simulated cells are restricted to move within a certain distance in any direction. The perfectly isotropic restriction within the cellular space is not representative of tumour tissue where the cells can have more chaotic shapes, and the intracellular space can exhibit anisotropically-restricted diffusion akin to that found in brain tissue. This could affect the performance of the models used to fit to the data, which both assume isotropically-restricted intracellular diffusion. Another limitation caused by the simplistic representation of cells as packed spheres is that it is impossible to achieve realistic intracellular volume fractions of the simulation substrate. Instead, as discussed previously the higher intracellular volume fraction is achieved by propagating more walkers through the intracellular space. However, this does not truly recreate the conditions that water molecules would experience in real tumour tissue. For walkers diffusing in the extracellular space, their motion will be much less hindered by the packed spheres since these cannot be packed as closely together as real cells. The effect this may have on the resulting synthetic signal, particularly for acquisitions with long diffusion times, is an artificially higher extracellular diffusion coefficient.

The simplistic substrate model could be further developed by introducing non-spherical cell shapes. The introduction of 3D high-resolution datasets such as those produced using High Resolution Episcopic Microscopy (HREM) presents the opportunity to generate

cellular substrates that are based on segmented images of real tumour tissue. This would further enable higher intracellular volume fractions to be reached without the need for the more artificial approach used during this work. Realistic cellular substrates, coupled with the real vascular structures implemented in this work would represent a significant step towards a truly realistic simulation substrate for tumour tissue. Aside from these geometrical considerations, the effects of membrane permeabilities, T2 values and different diffusion coefficients between the different compartments could also be explored. Various aspects of the walker dynamics could also be investigated, such as interactions between the walkers themselves, and more realistic interactions between the walkers and tissue membranes.

A further significant limitation of this study is how to thoroughly validate the microvascular flow aspects of the walker dynamics. The fact that increasing the fraction of walkers flowing in the vessels is reflected in an increase in the vascular volume fraction of the various relevant diffusion models is reassuring. However, if Monte Carlo simulations of well-perfused tumour tissue are to be explored further future work must focus on proper physical validation of the motion of the walkers within the blood vessels.

All the limitations outlined above will affect the accuracy of the generated synthetic signal, making it less representative of diffusion-MRI signal from real tissue samples. These limitations will therefore affect the results achieved in the following chapter, where the synthetic data is used to train a neural network to learn tissue parameters from diffusion data. The limitations of the simulation framework will limit the performance of the neural network when the trained network is applied to real diffusion MRI data.

Despite the limitations outlined above, the simulation framework presented in this chapter has already produced some valuable insights into the behaviours and performances of diffusion models. The approach used here is a useful foundation upon which future work

should build, and has highlighted the areas where there is significant room for additional research such as development of realistic biological tissue simulation substrates.

A further potentially powerful application for the synthetic signal generated using simulation frameworks such as the one presented in this chapter is the field of machine learning. This is an area of research that has seen an explosion in popularity in recent years, with ever-increasingly complex techniques being developed at an ever-increasing rate. Techniques such as neural networks and deep learning have enormous potential to revolutionise medicine (as well as every aspect of our day-to-day lives). One thing that the vast majority of machine learning methods have in common is their voracious appetite for vast quantities of training data. Therefore, being able to generate synthetic MR data from substrates with near-infinite combinations of microstructural and biophysical properties could prove vital to future work in these fields. An initial exploration of the potential application of synthetic data to machine learning is presented in the following chapter.

Application of Machine Learning to Real and Synthetic MR Data

5.1 Brief Overview of Machine Learning

Over the last 20 years, human beings have moved into the age of so-called ‘Big Data’. This refers to the huge amounts of digital information that is being produced every second around the world. According to the Economist in 2011 from the dawn of time until 2005, it is estimated that human beings have produced around 130 Exabytes of data. That’s 1.3×10^{20} bytes worth of books, films, music and other stored information [183]. However, just 5 years later that number had increased to roughly 1227 Exabytes, and is predicted to reach 40000 Exabytes by 2020. Currently, the majority of the data generated is unused. Machine learning is an area of technology that aims to harness this data as a resource and put it to use in a wide range of applications.

The term ‘*Machine learning*’ is broadly defined by the ability of a machine to alter its own structure, program or data, based on its own inputs or external stimuli, in such a manner that its expected future performance improves. More narrowly, machine learning is a term commonly used to refer to algorithms within the field of computer science that perform complex tasks by iteratively changing parameters of their associated models based

on continuous assessment of their performance. These tasks can cover a broad range of disciplines and include recognising objects within still images or videos, diagnosing diseases based on medical images, predicting the long-term behaviours of financial markets, detecting speech or controlling robots.

Algorithms that fall under the umbrella term of machine learning can be further divided into three main categories: ‘Supervised learning’, ‘Unsupervised learning’ and ‘Reinforcement learning’. Supervised learning algorithms construct mathematical models that predict the mapping function that maps the inputs and outputs of a pre-labelled dataset referred to as a ‘training set’. Popular types of supervised learning algorithm are classification algorithms and regression algorithms. Classification algorithms are characterised by their outputs being categories such as classifying incoming emails as “spam” or “not spam”, or classifying voxels of MRI data as “healthy” or “disease” [184]. Examples of classification algorithms include: Support Vector Machines, Logistic Regression, K-Nearest Neighbours, Kernel SVM, Naive Bayes, Decision Tree Classification, Random Forest Classification. Regression algorithms are characterised by their outputs comprising values such as finding the relationship between input “square feet” and “house price” or “MRI signal” and “diffusivity”. Examples of regression algorithms include: Simple Linear Regression, Multiple Linear Regression, Polynomial Regression, Support Vector Regression, Decision Tree Regression, Random Forest Regression.

Unsupervised learning algorithms construct mathematical models that aim to represent a dataset that contains only inputs and no corresponding outputs. The goal of an unsupervised learning algorithm is to find underlying structure within the data. Popular types of unsupervised learning algorithm are clustering algorithms and association algorithms. Clustering algorithms aim to discover inherent groups within the input data, such as grouping pixels by colour value. Examples of clustering algorithms include: K-means, Fuzzy C-means and Hierarchical Clustering. Association algorithms aim to find rules that

link different inputs of the dataset, such as determining that customers who buy one type product being more likely to buy another type of product. Examples of association algorithms include apriori and Eclat intuition. Reinforcement learning algorithms are designed to improve their performance of a particular task based on continuous feedback in the form of rewards/penalties for good/bad decisions. Examples of reinforcement learning algorithm include: Upper Confidence Boundary and Thomson Sampling. Popular embodiments of reinforcement learning applications are algorithms that learn to play games such as chess, video games or more recently the AlphaGo algorithm developed by Google Deepmind that beat the world's best at the gold standard of strategy games - 'Go' [185].

As outlined above, there are various different algorithms associated with each of the categories discussed above. A detailed analysis of each of those algorithms is beyond the scope of this thesis. However, there is further type of algorithm that can fit into any of the categories above and can be applied to almost any problem. The name of that algorithm: Artificial Neural Networks (ANNs). The work presented in this chapter was focused on utilising an artificial neural network to learn microstructural tissue parameters directly from synthetic diffusion MRI signal, then apply the trained neural network to real *in-vivo* diffusion data. In view of this, a more detailed analysis of neural networks is presented below.

5.2 Artificial Neural Networks

5.2.1 Brief Historical Background

The theoretical underpinnings of neural networks can arguably be traced back to the work of Bain [186] and James [187] in the late 19th Century. Both men hypothesised that

interactions between individual neurons within the brain were responsible for thoughts and actions. In particular, Bain theorised that different subsets of neurons were responsible for individual thoughts and actions, and that the physical links between the neurons within these subsets became stronger over time as a function of their use. Until the 20th century, the theory behind the interaction between neurons in the brain was mostly biological in nature. It was McCulloch and Pitts [188] who first developed a mathematical algorithmic representation of neural networks. The theory, named *threshold logic*, aimed to capture broad concepts from specific perceptions and a simple mathematical representation of the biological neuron.

The computational representations of the individual biological neurons are often referred to as ‘computing elements’, ‘units’ or ‘nodes’. In accordance with the current state of the art, ‘node’ will be used throughout this chapter. Put simply, nodes can be thought of as individual functions, taking inputs $x_{1,2,...,n}$ and producing an output value.




Figure 5.1: A graphical representation of the most simple embodiment of a neural network unit. Shown are Inputs x_1, x_2, \dots, x_n , and output function f [5].

Generally, it was considered advantageous to use a primitive function at each node, that takes a single input, and produces a single output. With this in mind, nodes are split in half - an integration function and an activation function.




Figure 5.2: A graphical representation of a neural network node, showing the integration function g and activation function f [5].

The role of the integration function is to output a single value reflective of the n inputs to the node. The activation function then takes the output value of the integration function and maps it to a single output that reflects the degree to which the node is ‘active’.

The *threshold logic* model developed by McCulloch and Pitts implemented a very simplistic node architecture. The inputs and outputs to each node are purely binary. The nodes aim to recreate the behaviour of human neurons by including both *excitatory* inputs and *inhibitory* inputs. Each node also has an associated threshold θ . The output of each node is determined by a comparison of the sum of the excitatory inputs $x_1 + x_2 + \dots + x_n$ with the threshold value θ . If the sum exceeds the threshold, the node is ‘active’ and the output of the node is 1, otherwise the output is 0. If any of the inhibitory inputs to the node are 1, then the node is deactivated and the output becomes 0. Despite their basic architecture, networks comprising multiple McCulloch/Pitts nodes are capable of synthesising any Boolean logic function. This makes them a powerful computing tool [5].

Following on from the work of McCulloch and Pitts, in the 1950s and 1960s Frank Rosenblatt developed the *perceptron* [189]. Similarly to the nodes developed by McCulloch and Pitts, perceptrons take multiple inputs and produce a single, binary output. However, Rosenblatt introduced the concept of weighted inputs, whereby each input to the percep-

tron has an associated weight. The output is determined by a comparison of the *weighted* sum of the inputs and a predetermined threshold. If the weighted sum is less than or equal to the threshold, the output is 0. Otherwise, the output is 1. To draw a direct comparison with the McCulloch/Pitts model, the ‘integration function’ of the perceptron is the weighted sum of the inputs while the ‘activation function’ is a step function with a value of 1 above the predetermined threshold and zero below. Alternatively, the activation function can be defined using the dot product of a vector of weights and a vector of inputs. Bernard Widrow adapted the threshold approach to use an equivalent ‘bias’ value [190]. The threshold value is converted into a bias, which is added to the result of the dot product. If the resulting value is less than 0 the perceptron is less than 0, and 1 otherwise.

Taken alone, a perceptron is a basic mathematical model of a simple decision-making process. Their full potential is only realised when they are combined together. Perceptrons (and other types of neurons/nodes, as discussed), can be combined to form interlinked networks with multiple layers that are capable of making more subtle decisions based on multiple inputs.




Figure 5.3: A schematic showing a network comprising multiple layers of perceptrons.

Shown in figure 5.3 is an example of a typical layer of perceptrons comprising multiple layers. The inputs are fed in to the first layer, with each perceptron in the layer activating

or remaining dormant based on the weighted sum of the inputs. The outputs from the first layer are then fed into the perceptrons in the second layer as weighted inputs. Finally, the output of the network is determined by the weighted sum of the outputs from the second layer. The multi-layer structure allows the network to make complex and detailed decisions.

The basic structure shown in figure 5.3 is universally familiar, as it is the basic structure adopted by the majority of recently developed neural networks. The problem that remains is how to choose the values for the weights and thresholds/biases for each neuron and input? For a small network solving a simple problem, it may be straightforward to select the weights and biases by hand. However, when networks get significantly larger, this becomes infeasible. The solution is to introduce *learning algorithms*. These are what give the field of Machine Learning its name, and allow the weights and biases of a neural network to be tuned automatically, allowing the network to ‘learn’ how to solve a given problem.

However, a problem arises when applying learning algorithms to a network comprising perceptron neurons. Namely, a small change to the weights on the inputs to a particular neuron may not cause any change to the output of the neuron, since this is either a 0 or 1. On the other hand, a small change may be enough to make be enough to change the output from a 0 to a 1, or visa versa, in turn causing significant changes to the behaviours of consequent neurons. The problem with this is that a small change to the inputs to the network could cause a large change in the output. To overcome this, computer scientists introduced the sigmoid neuron.

5.2.2 Sigmoid Neurons

To avoid the difficulty encountered with perceptrons, a more sophisticated model of neuron is required. One of the most popular of these is the *sigmoid* neuron. The structure of sigmoid neuron is similar to that of a perceptron, but the inputs and outputs are no longer binary. The inputs can take any value between 0 and 1, and the output of the neuron is given by:

$$s_{out} = \sigma(w \cdot x + b) \quad (5.1)$$

Where w is a column vector of the input weights, x is a vector of the inputs and b is the ‘bias’ on the neuron (how easily the neuron becomes activated). The σ represents the ‘sigmoid’ function, from which the form of neuron gets its name. The sigmoid function is represented by:

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (5.2)$$

The sigmoid function may look more complex than the simple threshold model implemented by McCulloch and Pitts, but on inspection the shape of the two functions are in fact fairly similar:

The sigmoid function equation can be combined with [5.1](#) to give the function that computes the output from the neuron for any given set of inputs, weights and bias:

$$s_{out} = \frac{1}{1 + e^{-\sum_{i=1} w_i x_i - b}} \quad (5.3)$$




Figure 5.4: A graph showing a comparison between a typical threshold (dotted line) and corresponding sigmoidal (solid line) activation function.

As can be appreciated looking at figure 5.4, the behaviour of the sigmoid function at very large positive or negative numbers is the same as the threshold function, i.e. it outputs a 1 or a 0. At smaller magnitudes however, the behaviour of the two functions diverges. The key difference being that the sigmoid function outputs values between 0 and 1, as opposed to 0 or 1 exclusively. The crucial advantage of the sigmoid function over the threshold function for applications to learning algorithms is the **smoothness** of the function. The reason that a smooth function is advantageous is that having a small change in output resulting from a small change in input allows the network to *learn*.

5.2.3 Neural Network Structure

The individual neurons can be seen as the building blocks, that are combined to form the overall structure of a neural network. As previously mentioned, neural networks are formed of multiple layers of neurons. The dimensions of each layer can be chosen arbitrarily by the user, according to the problem to be solved. The neurons in each layer are joined to neurons in adjacent layers via connections, each with their own associated weights. In most cases, data passes through the layers of the network in one direction only, which is known as a *feedforward* network. Alternative configurations such as Recurrent Neural Networks (RNNs) allow for time-varying activations and feedback loops [191–193], but a discussion of these types of network is beyond the scope of this preliminary chapter.




Figure 5.5: An illustration showing the fully-connected representation of a typical feedforward neural network, comprising an input layer with 10 nodes, two 9-node hidden layers and a 5-node output layer.

An example of a typical neural network architecture is shown in figure 5.5. The representation shown is a typical way of illustrating neural network structures known as a ‘fully-connected representation’ - due to the fact that all nodes are shown with all their possible connections to all consequent nodes. Networks are commonly illustrated with a left-to-right data flow, with nodes within the same layer aligned vertically to form columns of nodes. The left-most layer shown in figure 5.5 is the input layer. As the name would suggest, this layer takes the data input to the network - for example a list of parameters characterising a house. The right-most layer is the output layer, which concordantly

represents the output values from the network for example the predicted house price, or classification of an object in an image. Between the input and output layers are the ‘hidden’ layers. The term ‘hidden’ denotes the fact that the activation values of the neurons in the hidden layers are usually just passed to the consecutive layers, or to the output layer.

The properties of an individual neural network can generally be described by five parameters [194]:

- Size: The total number of nodes.
- Width: The number of nodes in a given layer.
- Depth: The number of layers in the network.
- Capacity: The types of function that can be represented by the network.
- Architecture: The arrangement of the layers in the network, and the nodes in the layers.

The dimensions of the individual layers, and the number of layers of the network can all be chosen by the user. There is, at the time of writing of this thesis, no accepted general rule-of-thumb for choosing the number of layers or the dimensions. However, there are some generally accepted guidelines for choosing the architecture of a neural network.

The first of these is using experimentation. By defining a range of possible depths, widths and re-training the network with every possible combination of these within the specified ranges, then evaluating which combination gives the best model performance. This is an expensive method in terms of computational time, but will reliably determine which architecture is best for a specific dataset and problem.

In some cases, it may be possible to use intuition to choose the network architecture. If the user has a great deal of understanding of the problem to be solved, it may be possible to choose the depth and width of the network to reflect this. In general, if it is appreciated that a given problem may be expressed in terms of a number of underlying factors, which in turn may be expressed in terms of simpler underlying factors and so on, such a problem may be more suited to a neural network with greater depth [195].

Goodfellow et al [195], asserted that in the majority of problems, the performance of a neural network model is at least loosely correlated with the network depth, i.e. the deeper the network, the better it performs. It is worth noting however, that using a deeper network results in a significant increase in training time for any given dataset.

Another possible method for making an initial selection of network architecture is to use previously published examples. If the problem to be solved is one that has been solved previously, it may be possible to use previous examples as a prior when choosing initial values for the network depth and properties.

Once the architecture for the network has been selected, the network must somehow *learn* how to solve the problem.

5.2.4 How Neural Networks ‘Learn’

As mentioned above, once the number of layers, number of nodes etc. of the neural network have been selected, the network must now somehow *learn* how to solve a given problem - by tuning the parameters of the network to figure out the function that maps the inputs to the correct outputs.

The first ingredient in the learning process is the dataset. The dataset not only defines the problem to be solved, but also provides the data on which the network is to train. Typically, the data is split into two discrete sets: the ‘training set’ and the ‘test set’. A large portion of the overall dataset is reserved for the training of the model - allowing the model to tune and adjust the weights and biases of the model to improve the model performance. Then a small portion of the dataset is held back, and used to evaluate the performance of the model. The reason for doing this is to test the model performance on data that hasn’t been seen before during the training process. Neural networks are not intuitive - given a set of data to train on they will learn the optimal function that maps the inputs of that data to the output. However, this is not the desired solution, as any dataset is typically only a sample of the whole population of possible inputs, therefore the performance of the model should only be evaluated on data that hasn’t yet been seen by the model during the training process.

For supervised learning models, such as those used during the work presented here, each instance of the dataset comprises a number of inputs and the corresponding *correct* outputs. In order to measure the performance of the model for any given set of model parameters, a single value is calculated that reflects the difference between all of the outputs calculated by the model and the correct values from the training set. This is done by computing the ‘Cost Function’.

5.2.4.1 Cost Function

As discussed above, the cost function is a measure of how well a neural network has performed with respect to the training dataset and the outputs calculated by the network. The value of the cost function is typically therefore influenced by the expected output values, correct output values, as well as in some cases the parameters of the network

itself. Importantly, the cost function calculates a single value, not a vector of values corresponding to all instances of the training set. This single value is reflective of the *overall* performance.

There are many different types of Cost Function, but in order to be compatible with the most common learning algorithm ‘Backpropagation’ (discussed [below](#)), cost functions must adhere to two key criteria[196]:

1. The cost function must be able to be expressed as an average value over all the entries of the training set:

$$C = \frac{1}{n} \sum_{x=1} C_x \quad (5.4)$$

This is necessary, as it allows the gradient of the cost function to be calculated, which is key to training the network using ‘Gradient Descent’ (as discussed [below](#)).

2. The cost function must not depend on neuron output values from any layer in the neural network **except** the output layer. This is due to the way in which ‘Backpropagation’ works by calculating the weight updates on the output layer, then propagating these back through all the layers of the network.

The choice of cost function is very much dependent on the problem to be solved. For regression problems, one of the most commonly-used functions is the *Mean Squared Error* (MSE):

$$C(w, b) = \frac{1}{2n} \sum_{x=1} \|a_x - y(x)\|^2 \quad (5.5)$$

Where w is the collective weights of the network, b is the collective biases, n is the number of instances in the training set, a_x is the calculated outputs from the model, and $y(x)$

is the correct output from the training set. The $|||$ notation indicates the modulus of difference values - which is always positive. The value of the MSE is thereby always positive, and becomes very small when the difference between the calculated and actual outputs becomes small, and very large when the difference is large [196].

The aim of the learning algorithm will therefore be to tune the parameters of the network such that the difference between the calculated outputs and correct outputs is as small as possible. The goal is therefore to *minimize* the cost function.

The problem to be solved is how to work out what changes need to be made to the network parameters in order to minimize the cost function. The algorithm that solves this problem is known as ‘Gradient Descent’.

5.2.4.2 Gradient Descent

For convex functions with few input variables, finding the minimum point is possible by differentiating the function and finding the points where the differential is zero. However, the cost function of a neural network typically has thousands, millions or even billions of inputs - the weights and biases of the neurons in the network. This means that an explicit calculation of the cost function minimum is impossible.

A more flexible approach is to choose an initial set of input values at random, and determine the direction in which those inputs should step in order to decrease the value of the cost function. Specifically, by calculating the *gradient* of the function at a given point: and stepping each variable in the direction in which the gradient is most negative.

To achieve this, for a given step, a good starting point is to investigate how the value of the cost function changes for a given change of the input variables. This is given by:

$$\Delta C \approx \frac{\partial C}{\partial w_1} \Delta w_1 + \frac{\partial C}{\partial b_1} \Delta b_1 + \frac{\partial C}{\partial w_2} \Delta w_2 + \frac{\partial C}{\partial b_2} \Delta b_2 + \dots + \frac{\partial C}{\partial w_n} \Delta w_n + \frac{\partial C}{\partial b_n} \Delta b_n \quad (5.6)$$

It can be appreciated, looking at equation 5.6 that the changes in weights given by Δw_1 , Δw_2 etc. are the changes that we wish to make to the weights (and also biases) in order to decrease the cost function - they are what we need to calculate. It can also be seen that the value of each of these will depend on the partial derivative of the cost function with respect to that input variable [196].

Equation 5.6 may be rewritten using vector notation as:

$$\Delta C \approx \nabla C \cdot \Delta w \quad (5.7)$$

Where ∇C is the vector of partial derivatives of the cost function:

$$\nabla C = \left(\frac{\partial C}{\partial w_1}, \frac{\partial C}{\partial b_1}, \dots, \frac{\partial C}{\partial w_n}, \frac{\partial C}{\partial b_n} \right) \quad (5.8)$$

Therefore it is possible to choose a value for Δw such that ΔC is always negative, for instance:

$$\Delta w = -\alpha \frac{\partial C}{\partial w} \quad (5.9)$$

Where α is a scalar multiplier used to define the length of step taken during the gradient descent algorithm, and is known as the ‘learning rate’. When substituted into 5.7, and for the purposes of simplicity assuming that the biases are constant ($\frac{\partial C}{\partial w} = \nabla C$) this gives:

$$\Delta C \approx -\alpha \nabla C \cdot \nabla C \quad (5.10)$$

$$\approx -\alpha \|\nabla C\|^2 \quad (5.11)$$

Where because $\|\nabla C\|^2$ is always positive, ΔC is always negative. Therefore, if the changes to the weights (and biases) are made according to equation 5.9, the cost function is guaranteed to decrease. This holds true when the biases are not constant, as ∇C denotes the partial derivatives with respect to each weight and bias.

The gradient descent algorithm then repeats the process, calculating the cost function gradient, and updating the weights/biases in order to ‘step’ in the direction in which that gradient is most negative. The algorithm continues, until it finds a minimum of the cost function. Figure 5.6 shows an illustration of the gradient descent algorithm, minimizing a simple cost function with a single parameter.




Figure 5.6: An illustration showing the steps of the gradient descent algorithm on a simplified function with a single parameter. Starting at a random parameter value, the algorithm determines the gradient of the cost function at that point, and then updates the parameter according to equation 5.9, then repeats the process until the minimum is found.

An advantage provided by updating the weights/biases according to equation 5.9 is that the length of the step taken is proportional to the gradient of the cost function. This means the algorithm takes smaller steps when the cost function approaches a minimum. As a result, it is less likely that the algorithm will ‘over-shoot’ the minimum.

However, there is a further problem that can be encountered when applying gradient descent to cost functions with many parameters. The shape of the cost function may no longer be the simple, smooth curve such as that shown in figure 5.6. Even though the cost function is a quadratic function of each weight/bias, the high-dimensionality of the

function can cause it to become non-convex, resulting in multiple minima, into which the gradient descent algorithm may fall.




Figure 5.7: An illustration of the gradient descent algorithm being applied to a non-convex function. When the algorithm initialises at the left-hand point, it converges to a local minimum. When the algorithm initialises on the right-hand point, it finds the global minimum.

Figure 5.7 illustrates the problem presented by high-dimensionality cost functions. There is a strong possibility that the algorithm will find a false, local minimum rather than the true global minimum.

One way to increase the chance of finding the true global minimum is simply to run the gradient descent algorithm multiple times, initialising at different combinations of parameters each time. Although this approach becomes less effective on larger networks, as the sheer number of different combinations of network parameters makes it impossible to sample a significant portion of the whole parameter space.

A more popular method for improving the performance of the standard gradient descent algorithm is called **Stochastic Gradient Descent** (SGD). The SGD algorithm differs from the general gradient descent algorithm in that instead of calculating the cost function based on the average of all the training set instances (this is sometimes referred to as Batch Gradient Descent, (BGD)), the cost function is calculated on a single instance at a time. The weights/biases are then updated based on this single instance, then the algorithm repeats. The advantages of SGD are that it is faster to run (despite involving a larger number of steps, the memory requirements are much lighter), and importantly, the fluctuations in cost function value are greater. This additional ‘noise’ in the trajectory of the cost function means it is more likely to avoid local minima, and find the true global minimum. One disadvantage of the SGD algorithm, is that it (by definition) is non-deterministic. Whereas the BGD algorithm is deterministic. A possible compromise between BGD and SGD is known as ‘Mini-Batch Gradient Descent’ (mBGD), where a small ‘batch’ of training set instances are used to calculate the cost function and update the weights at each step. In this case, equation 5.9 becomes:

$$\Delta w = -\frac{\alpha}{m} \sum_{j=1}^m \frac{\partial C_j}{\partial w} \quad (5.12)$$

Where the sum is over the m instances within the mini-batch.

The last vital step in implementing gradient descent for training a neural network is how to calculate the gradient of the cost function ∇C , in order to determine the changes to the weights and biases. For functions with few variables, this is a trivial problem that can be solved using basic calculus. For neural network cost functions with thousands or millions of weights and biases however, this is infeasible. To accomplish this, an algorithm known as ‘Backpropagation’ can be implemented [196].

5.2.4.3 Backpropagation

The backpropagation algorithm was originally developed in the 1960's and 1970's. In 1970, Seppo Linnainma published a method for differentiating connected networks of differentiable functions [197], which Paul Werbos then adapted and applied to artificial neural networks in 1982 [198]. However, the algorithm only gained widespread popularity in 1986 following the publication of a famous paper by David Rumelhart, Geoffrey Hinton, and Ronald Williams [199]. They found that the backpropagation algorithm provided significant advantages in a variety of neural network architectures, vastly speeding up the learning process [196]. Neural network research tailed off in the 1990s and early 2000s, due to the limitations of the computing hardware of the time. In the 2010s, neural network research exploded once again, with the gradient descent and backpropagation algorithms providing the backbone to the vast majority of the state of the art research.

At a fundamental level, the backpropagation algorithm is a method for finding the partial derivatives of the cost function with respect to both the weights w and biases b of a network:

$$\frac{\partial C}{\partial w}, \quad \frac{\partial C}{\partial b} \tag{5.13}$$

These expressions represent the way the weights and biases of the network influence the cost function, and in turn the performance of the network.

For the sake of conceptual simplicity, consider a network with a single neuron in each layer:



Figure 5.8: A simple network with one neuron in each layer, with the layer number indicated by L . Each neuron has an associated activation, a , and an associated bias, b . The neurons in consecutive layers are connected by weights w , which belong to the layer to which they connect.

The network shown in figure 5.8 contains 4 layers, the input, two hidden layers, and the output. Each layer L consists of a single neuron, with an associated activation a^L and bias b^L . The consecutive layers are connected by weights w^L , which by convention are assigned to the layer to which they connect. The activation function of each neuron in this network is the sigmoid function.

The cost function of the network above therefore is a function of the three weights and three biases of the network:

$$C \rightarrow C(w^{L-2}, b^{L-2}, w^{L-1}, b^{L-1}, w^L, b^L) \quad (5.14)$$

The value of this cost function for a single training instance is determined by the activation of the output layer and the desired output from a single instance of the training set:

$$C_0(w^{L-2}, b^{L-2}, \dots) = (a^L - y)^2 \quad (5.15)$$

Where y is the desired output from the training set. The activation of the neuron in the

final layer is determined by the weighted sum of the activations in the previous layer and the bias on the output neuron, and the activation function of the neuron:

$$a^L = \sigma(w^L a^{L-1} + b^L) \quad (5.16)$$

$$= \sigma(z^L) \quad (5.17)$$

For notational simplicity, the weighted sum of activations and the biases are combined into a single variable z^L , which is known as the ‘input’ to the neuron. Given that $C_0 \propto a^L$, and in turn $a^L \propto z^L$ and Z^L is a function of w^L , a^{L-1} and b^L . It is possible to use the chain rule to write an expression for $\frac{\partial C}{\partial w^L}$ and $\frac{\partial C}{\partial b^L}$:

$$\frac{\partial C}{\partial w^L} = \frac{\partial z^L}{\partial w^L} \frac{\partial a^L}{\partial z^L} \frac{\partial C_0}{\partial a^L} \quad (5.18)$$

Where:

$$\frac{\partial C_0}{\partial a^L} = 2(a^L - y) \quad (5.19)$$

$$\frac{\partial a^L}{\partial z^L} = \sigma'(z^L) \quad (5.20)$$

$$\frac{\partial z^L}{\partial w^L} = a^{L-1} \quad (5.21)$$

Therefore, for a single training instance:

$$\frac{\partial C_0}{\partial w^L} = a^{L-1} \sigma'(z^L) 2(a^L - y) \quad (5.22)$$

A similar equation can also be found for the biases:

$$\frac{\partial C_0}{\partial b^L} = \frac{\partial z^L}{\partial b^L} \frac{\partial a^L}{\partial z^L} \frac{\partial C_0}{\partial a^L} \quad (5.23)$$

$$= \frac{\partial a^L}{\partial z^L} \frac{\partial C_0}{\partial a^L} \quad (5.24)$$

$$= \sigma'(z^L) 2(a^L - y) \quad (5.25)$$

Where the quantity $\frac{\partial z^L}{\partial b^L} = 1$. The above expressions can then be averaged over all training instances:

$$\frac{\partial C}{\partial w^L} = \frac{1}{n} \sum_{k=0}^{n-1} \frac{\partial C_k}{\partial w^L} \quad (5.26)$$

$$\frac{\partial C}{\partial b^L} = \frac{1}{n} \sum_{k=0}^{n-1} \frac{\partial C_k}{\partial b^L} \quad (5.27)$$

These two expressions determine the rate of change of the cost function with respect to any particular weight or any particular bias within the network. These results can then be expanded to networks with multiple neurons, where the cost function is given by equation 5.5. Equation 5.18 then becomes:

$$\frac{\partial C}{\partial w_{jk}^l} = \frac{\partial z_j^l}{\partial w_{jk}^l} \frac{\partial a_j^l}{\partial z_j^l} \frac{\partial C_0}{\partial a_j^l} \quad (5.28)$$

Where w_{jk}^l indicates the weight connecting the k^{th} neuron in the $l - 1^{th}$ layer to the j^{th} neuron in the l^{th} layer. Where

$$\frac{\partial a_j^l}{\partial z_j^l} = \sigma'(z_j^l) \quad (5.29)$$

$$\frac{\partial z_j^l}{\partial w_{jk}^l} = a_j^{l-1} \quad (5.30)$$

And a similar approach using the chain rule can be used to find a relation for $\frac{\partial C}{\partial a_j^L}$:

$$\frac{\partial C}{\partial a_j^L} = \sum_{k=0}^{n_{l+1}-1} w_{jk}^{l+1} \sigma'(z_j^{l+1}) \frac{\partial C}{\partial a_j^{l+1}} \quad (5.31)$$

Combining equations 5.28 and 5.31 gives the rate of change of the cost function with respect to any given weight in the network. Importantly, it can be seen from equation 5.31 that the gradient of the cost function with respect to a given weight in a given layer is dependant on the activations in the next layer (in a **forward** direction). This is where the idea of ‘Backpropagation’ comes from.

When implemented, the backpropagation algorithm starts at the output layer, then works backwards through the remaining layers calculating the partial derivatives of the cost function with respect to each weight and bias throughout the network. The weights and biases in the network are then updated according to the gradient descent algorithm. This then allows the weights and biases of the network to be updated such that the cost function decreases and the network performance improves.

The process then repeats for the next training instance or batch of instances until all of the training data has been used. Once the end of the training data is reached, the training process is said to have completed one ‘epoch’. The whole process is then repeated, starting from the first training instance/batch, for as many epochs as required to fully train the network.

5.2.5 Application to Medical Imaging

In recent years, the medical imaging sector has become an increasingly-ideal candidate for the application of neural networks. The majority of work to date has been focussed on a number of key areas: image segmentation, image registration and, predominantly, computer-aided diagnosis or detection (CAD) [200].

Image segmentation is a key part of the image processing pipeline for a wide variety of image analysis techniques. Often, image segmentation is used to extract key regions of interest from the whole image field-of-view. Segmentation may also be used to extract volumes from whole image stacks. Traditional image segmentation techniques usually rely on edge detection by convolving the 2D images with a feature detection kernel such as a Sobel filter. More advanced edge detection techniques such as the Canny edge detection algorithm use combinations of filtering, non-maximum suppression and hysteresis to improve the detection of edges in the image [201]. Other image segmentation techniques work by grouping pixels with similar intensities.

Various groups have applied neural networks to the problem of image segmentation. Lin [202] applied a form of Recurrent Neural Network called a Hopfield Network to the segmentation of MRI images. Kobashi et al [203] used a feed-forward ANN to classify volume regions of Magnetic Resonance Angiography (MRA) data. Milletari et al [204] trained a

CNN to segment the prostate from MRI data, including generating additional synthetic training instances by applying non-linear transformations to existing training data. More recent work has moved toward the use of Deep Learning techniques. Liu et al [205] used a Deep Neural Network (DNN) to segment the prostate from CT images.

Image registration involves the application of rigid, affine and non-rigid transformations to images to align them into a common space. Typically this is performed as part of an integrated image analysis pipeline, where subsequent analysis steps require the different images to be mapped into the same coordinate system. Image registration is of particular importance when comparing images produced using different modalities, such as MRI and X-ray CT. Traditional image registration algorithms typically involve applying iterative transformations to a ‘floating’ image in order to align it with a ‘fixed’ image. The transformations are usually calculated using a gradient descent algorithm that aims to minimise a cost function that is dependent on the correlation of regional intensities or image features.

A large volume of research focussed on the application of machine learning algorithms to image registration have applied the ‘Self Organising Maps’ (SOM) algorithm, which is a type of ANN that implements competitive learning - an alternative to gradient descent and backpropagation in its training. Shang et al [206] used a SOM to calculate the first principle direction and centroid of CT and MRI images in order to align them. Other more recent approaches have shown that convolutional networks may be well suited to image registration problems. For example, de Vos et al [207] used a CNN-based approach to calculate the spatial transformation parameters for rapid alignment of Cardiac CINE-MR images.

One of the key areas where neural networks have shown great promise is in the field of computer-aided detection/diagnosis (CAD). CAD algorithms aim to assist clinicians in the detection of abnormalities and pathologies in medical image data. There have been

numerous studies demonstrating the ability of various neural network architectures in the detection of tumours in the breast [208], brain [209–211], gut [212] and even eye [213].

There has, to date, been very limited research into the application of neural networks (or machine learning algorithms in general) to the quantification of tissue microstructure. The main reason for this is the lack of reliable ground truth data with which to train a model. Golkov et al [214] trained neural networks on human diffusion MR data, using parameters (Kurtosis, and NODDI parameters) measured using conventional model fitting as ground truth values for the model training set. They found that the trained networks were then able to measure the same parameters as the fitted models, but were able to do so on diffusion data that was much more sparse. This enabled a 12-fold decrease in scan time, while still achieving good agreement with the fitted models that were fitted to data from the full acquisitions. Nedjati-Gilani et al [215] used synthetic data generated from Monte Carlo simulations of water diffusing in the intra and extra-axonal spaces to train a random forest machine learning model. The trained model was able to extract microstructural parameters from diffusion data for which there exists no explicit mathematical model (such as water residence time in the axons).

The preliminary work presented in this chapter takes a similar approach to that of [215], using the Monte Carlo simulation framework developed in the previous chapters to generate synthetic diffusion-MR data, on which to train a neural network. The aim being to train the neural network to extract microstructural tissue properties such as cell radius from *in-vivo* diffusion data without needing to fit complex mathematical models.

5.3 Study Overview

The aim of the work presented in this chapter was to investigate the potential applications of artificial neural networks in learning to extract microstructural features from diffusion-weighted MR data.

To achieve this, the Monte Carlo simulation framework presented in the previous chapter was used to generate synthetic diffusion data from simulation substrates with various properties. The synthetic data and ground truth parameters from the simulations were then used as the training dataset for the neural network. The trained neural network was then used to calculate microstructural parameters from *in-vivo* diffusion data of LS174T tumours.

The work presented in this chapter is preliminary in nature, and is intended to show the potential for applying a combination of synthetic data generation and machine learning within the field of diffusion MRI modelling, and MRI in general.

5.3.1 Methods

5.3.1.1 In-Vivo Data

In-Vivo diffusion MRI data was captured from a subcutaneous tumour grown on the flank of a female MF1 nu/nu mouse. The tumour cell line used was LS174T - a human caucasian colon adenocarcinoma model known to be tumorigenic in nude mice. All experimental procedures and methods were the same as those used for the tissue fixation study in [chapter 3](#). The tumours were scanned with the full VERDICT acquisition, with 46 b-value and 3 directions, 220 acquisitions in total. One slice of *in-vivo* data was selected as the test dataset, to which the trained neural network was applied.

The *in-vivo* data was then fitted with the VERDICT 'BallSphereAstrosticks' model, using the same fitting procedures detailed in chapters 3 and 4. The fitting was performed on a voxel-wise basis, as opposed to fitting the averaged signal from an ROI.

5.3.1.2 Synthetic Data

A total of 3676 sets of synthetic diffusion MR data were generated using the Monte Carlo simulation framework presented previously. Random-walk simulations were run on simulation substrates with varied substrate parameters. The substrate parameters were varied as follows:

Parameter Name	Range
Cell Radius	$5\mu m - 15\mu m$
Intracellular Walker (Volume) Fraction	0 - 1
Vascular Walker (Volume) Fraction	0 - 1
Extracellular Walker (Volume) Fraction	0 - 1
Blood Flow Velocity Factor	$\times 0 - \times 2$

Each of the five parameters listed above were selected by via a random sampling within the corresponding parameter range. The remaining substrate parameters were kept constant, with diffusivity fixed at $9 \times 10^{-10} m^2 s^{-1}$.

The synthetic data was then generated from the walker trajectories of each simulation. The data was generated using the full VERDICT acquisition scheme - 46 b-value, 3-directions with 220 acquisitions in total. The resulting dataset comprised 3676 signal vectors, each with a length of 220.

To create the training and testing dataset, these signal vectors were appended with the corresponding substrate parameters (cell radius, intracellular walker fraction, vascular walker fraction, extracellular walker fraction and blood flow velocity factor). The signal

vectors make up the input to the network, with the parameters forming the target values for evaluating the network performance and computing the cost function during network training. The total dataset was then split into the training and testing datasets, with a ratio of 80:20.

To improve the performance of the neural network, feature scaling was applied to the inputs and outputs, this normalises the inputs and outputs to lie within common ranges. This is especially important for the outputs, as the cell radius values are around six orders of magnitude smaller than the volume fractions.

5.3.1.3 Artificial Neural Network Development

The artificial neural network was developed using the Keras neural network API, running with a Tensorflow backend on Python 3.5.

The width of the network hidden layers was chosen as the average between the width of the input layer (220 neurons) and the output layer (4 neurons). This is a popular method for selecting network width, as outlined in the book ‘Introduction to neural networks in Java’ by Jeff Heaton [216].

As discussed [previously](#), at the time of writing this thesis there is no accepted method for choosing the number of layers of the neural network. The basic motivation when selecting the number of layers for the network (the number of hidden layers) is a trade-off between training accuracy and training time. Increasing the number of hidden layers has been shown, empirically, to improve the training accuracy of the network [195]. A neural network with a greater number of layers is able to extract a greater number of features from the data, to a certain extent. Too many layers can increase the chances of ‘overfitting’ the data, which usually manifests in the form of good training dataset performance and poor

test dataset performance. Furthermore, increasing the number of layers in the network increases the amount of computing time required to train it.

For the present study, the number of network layers was determined through systematic experimentation. The network was repeatedly trained for 50000 epochs on the synthetic data while iterating the number of hidden layers, and the final test dataset accuracy recorded for each network.




Figure 5.9: Plot of number of network hidden layers vs final test dataset error. Each network was trained on 2941 instances of synthetically generated VERDICT MRI data, and tested on 735 instances of synthetic data. Each network comprised an input layer of 220 neurons, an output layer of 4 neurons, and a number of hidden layers each with 112 neurons. Each network was trained for 50000 epochs.

Figure 5.9 shows the final test dataset error achieved by varying the number of hidden layers in the network. The error value drops sharply between 1 and 3 hidden layers, from 0.723 to 0.1077. For networks with 3 or more hidden layers the error decreases more gradually, dropping to 0.0746 with 5 hidden layers. For networks with 5 hidden layers or above, the error remains almost constant, decreasing only slightly to 0.0653 for 8 hidden

layers. In view of this, a neural network with 5 hidden layers was selected to provide a balance between performance and training time.

The chosen architecture of the network is shown below in figure 5.10. As described previously, the input layer contained 220 nodes, corresponding to the 220 acquisitions of the synthetic diffusion data. The network comprised 6 layers (not including the input), with 5 hidden layers and an output layer. The width of the hidden layers was 112, and the output layer had a width of 4. Even though there were 5 varied parameters in the training data, one of the volume fractions can be removed to avoid overfitting, due to the fact the 3 volume fractions must sum to 1. The width of the hidden layers was chosen as the average of the input layer and output layer widths.




Figure 5.10: A schematic showing the neural network architecture that was trained on synthetic diffusion MRI data. The network has 5 layers (not including the input layer), with dimensions 108, 108, 108, 108, 4. The input layer comprised 220 nodes.

The activation function chosen for each node of the network was the ‘ReLU’ function, or ‘Rectified Linear’ function. This is a popular activation function for large networks as it avoids issues associated with the sigmoid activation function such as saturation [195].

The network was trained on the training dataset using the ‘Adam’ optimisation algorithm, implemented in the Keras neural network API. The Adam algorithm is an extension of the stochastic gradient descent algorithm, that assigns a separate learning rate parameter α to each of the nodes in the network, and allows each learning rate to be adjusted during the learning process. The network was trained over 50000 epochs, with the training process sped up using GPU parallelisation with the CUDA toolkit 9.0.

5.3.1.4 Applying the Trained Network to In-Vivo Data

Once the network had been trained on the synthetic data, the trained network was applied to the *in-vivo* VERDICT data from the LS174T tumours. The resulting parameters from the neural network were used to generate parameter maps, which were then compared to parameter maps generated by conventional fitting of the VERDICT ‘BallSphereAstro-sticks’ model.

5.4 Results

In the following section, the results from the various simulation-based experiments outlined above will be presented. Each learned parameter is presented separately.

5.4.1 Intracellular Volume Fraction




Figure 5.11: A side-by-side comparison showing the intracellular volume fraction parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data.

Figure 5.11 above shows a side-by-side comparison between the f_{icvf} parameter maps produced by model fitting vs by using the trained neural network. The left-hand side figure shows the parameter map from the model fitting method. The f_{icvf} values are spread across the full range from 0 to 1. In the bottom-right portion of the slice, many of the parameter values hit the upper bound during the fitting process. The average

f_{icvf} value from the model fitted map was 0.64. The right-hand side figure shows the parameter map from the trained neural network. The f_{icvf} values range from 0.23 to 0.5, with a mean value of 0.37. Similarly to the fitted case, the bottom-right portion of the slice shows slightly higher values. Looking at both figures, some common features can be seen between the two parameter maps. In the upper-right portion of the slice, a region of higher f_{icvf} values in the fitted map corresponds closely to a region of lower f_{icvf} in the learned parameter map.




Figure 5.12: Histogram showing the f_{icvf} value distributions for both the fitted and learned results.

Figure 5.12 above shows the f_{icvf} value distributions for both the fitted and learned maps. The values achieved by the model fitting method followed a much broader distri-

bution than the learned values. Ignoring the voxels in which the fitting method hit the upper bound, the median value for the fitted method was 0.57, vs 0.37 for the learned case.

5.4.2 Extracellular Volume Fraction




Figure 5.13: A side-by-side comparison showing the extracellular volume fraction parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data.

Figure 5.13 shows the side by side comparison for the extracellular volume fraction parameter (f_{ees}). The left figure again shows the results from the model fitting approach. It can clearly be seen that a large proportion of the voxels within the slice hit the lower parameter bound during the fitting process and therefore returned a value of 0. The values ranged from 0 to 0.79. The mean value of the fitted f_{ees} map was 0.17, or when ignoring the voxels that hit the lower bounds 0.31. It is difficult to discern any appreciable structure within the fitted map, though there does appear to be some higher extracellular

volume fraction values closer to the centre of the slice. The right hand figure shows the f_{ees} parameter map produced by the trained neural network. The learned map appears more homogeneous than the fitted map. The values in the map range from 0.14 to 0.49. The mean value of the learned f_{ees} map was 0.39. Similarly to the fitted results, the learned map shows some higher values of f_{ees} towards the middle of the slice, though the difference is minor.




Figure 5.14: Histogram showing the f_{ees} value distributions for both the fitted and learned results.

The histograms of the two maps show that, similarly as for the intracellular volume fraction parameter, the f_{ees} values for the model fitting method follow a much broader distribution than those of the neural network. Interestingly, the distribution of f_{ees} values for the neural network method appears to contain two distinct peaks at ≈ 0.35 and ≈ 0.45 .

5.4.3 Vascular Volume Fraction




Figure 5.15: A side-by-side comparison showing the vascular volume fraction parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data.

The comparison between the parameter maps for the vascular volume fraction produced similar results to those of the extracellular volume fraction. The left-hand plot of figure 5.15 shows the f_{vasc} parameter map produced via model fitting. Similarly to the extracellular volume fraction results, the fitting process hit the lower bound in a large proportion of the voxels within the slice. The fitted vascular volume fraction values ranged from 0 to 0.99, with a mean value of 0.19 (which increased to 0.32 when ignoring the 0-value). Generally, it appears that the fitted vascular volume values were lower towards the central portion of the image slice. The right-hand plot shows the f_{vasc} parameter map produced by the neural network. Again, the map is much more homogeneous in nature, with values ranging from 0.12 to 0.42, with a mean value of 0.23. The map produced by the neural network also showed lower values of f_{vasc} towards the central portion of the image slice.




Figure 5.16: Histogram showing the f_{vasc} value distributions for both the fitted and learned results.

The histograms of the fitted and learned maps show the same trend as the other parameters, with the fitted values following a much broader distribution than the learned values. Similarly to the extracellular volume fraction results, the distribution of values from the neural network method exhibited two peaks at ≈ 0.18 and ≈ 0.34 . In addition, ignoring the fitted values that hit the lower bound during the fitting process, the distribution of values from the model fitting method also exhibited two peaks at ≈ 0.08 and ≈ 0.46 .

5.4.4 Cell Radius




Figure 5.17: A side-by-side comparison showing the cell radius parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data.

The cell radius parameter maps shown above in figure 5.17 show a correspondence between the fitted map and the learned map. The fitted values ranged from $0.1 \times 10^{-6}m$ to $17.5 \times 10^{-6}m$, while the learned values ranged from $7.86 \times 10^{-6}m$ to $12.5 \times 10^{-6}m$. The mean fitted value was $6.4 \times 10^{-6}m$ vs $10.96 \times 10^{-6}m$ for the learned value. Overall therefore, the values achieved by the model fitting process were lower than those achieved by the neural network. However, a visual analysis of the two maps shows common features in both. Both maps show a region of higher cell radius in the lower middle portion of the slice. In addition, both maps show a patch of lower cell radius values in the upper-right portion of the slice.

5.4.5 Blood Flow Parameters




Figure 5.18: A side-by-side comparison showing the blood-flow related parameter maps produced by: LEFT - Fitting the VERDICT BallSphereAstroSticks model, RIGHT - Output from neural network trained on synthetic data. The left-hand plot shows the vascular diffusion coefficient D_v , and the right-hand plot is the mean blood flow velocity output from the neural network (cms^{-1}).

The left-hand plot of figure 5.18 shows the vascular diffusion coefficient parameter map produced using the model fitting method. The D_v parameter did not perform in a stable manner, with many voxels hitting both the upper and lower bounds during the fitting process. The right-hand plot shows the mean blood flow velocity parameter output from the trained neural network. The blood flow velocity values range from 0 to $4 \times 10^{-3} cms^{-1}$. Looking at the velocity map, a clear trend can be seen showing higher blood flow velocities near the edge of the tumour, and lower velocities in the interior.

5.5 Discussion

The purpose of the preliminary study presented above was to investigate the possibility of applying a neural network trained on synthetically-generated diffusion MRI data to real-world *in-vivo* data, to extract microstructural features without the need for an explicit model-fitting approach.

The results presented above showed the potential for a machine-learning based approach to microstructural imaging using diffusion MRI. The trained neural network was able to produce parameter maps of the intracellular, extracellular and vascular volume fraction, as well as the cell radius and mean blood flow velocity parameters. The volume fraction maps showed similarities between the fitted and learned results, though the maps produced by the neural network generally contained a much smaller range of values compared to those of the model fitting approach. Notably, the vascular volume fraction parameter output from the neural network performed with much greater stability than that of the model fitting approach, producing a believable vascular volume map while the model fitting process failed in a large proportion of the slice voxels. The machine-learning approach therefore showed promise in as far as it was able to produce sensible parameter maps for the volume fractions, but further investigation is needed to ensure that the absolute volume fraction values are realistic and consistent.

The cell radius map produced by the neural network showed a strong resemblance to that produced by the model fitting procedure, with corresponding regions of high and low cell radius. However, the absolute radius values were generally higher, and again followed a much narrower distribution than those of the model-fitting case. This is a similar case to the results for the volume fraction parameters. The reason for the difference in absolute value, and decrease in the range of values is unknown. Future work should focus on

validating the actual parameter values using histology and alternative *in-vivo* imaging techniques such as two-photon imaging.

A promising result was the ability of the neural network to output a parameter map of the mean blood-flow velocity throughout the image slice. On the other hand, the model-fitting approach was unable to produce a parameter map of the vascular diffusion coefficient, with many of the voxels hitting the upper bound during the fitting process. Promisingly, the voxels in which the neural network output a low blood flow velocity correspond to the voxels in which the network also output a lower vascular volume fraction, and higher intracellular volume fraction and cell radius. Future work is needed to validate these flow measurements, using flow-sensitive MR techniques such as phase contrast MRA (Magnetic Resonance Angiography), or alternative flow-sensitive *in-vivo* imaging techniques.

However, the results presented in this chapter are merely an initial proof-of-concept for the application of synthetic data and machine learning to the field of diffusion MRI. One of the major limitations of this study was the small quantity of training data. Although close to 4000 instances of synthetic data were generated in order to train the neural network, this still represents a significant undersampling of the parameter space. With five parameters being randomly selected per simulation run, to generate synthetic data with every combination of parameters sampling each parameter 10 times would require 10000 simulation runs. This limited parameter sampling may therefore explain why the trained neural network appeared to output more discrete, rather than continuous values. In order to properly train the network, it is likely that orders of magnitude more training data would be required. It may be possible to make use of computer clusters, or web-distributed computing services, to generate sufficient samples of training data to properly train the network.

The results achieved for the mean blood flow and vascular volume fraction parameters

may have been influenced by the blood flow dynamics of the simulation framework. As discussed previously, when the blood flow velocity of a vessel segment approaches zero, the simulation dynamics become less realistic due to the fact that the simulation neglects the effects of diffusion within the vasculature. This may have effected the neural network, having been trained on data synthesized from the simulation framework. In particular, the vascular contribution of synthesized data from simulations with low blood flow velocity would appear more like stationary water. The trained neural network may therefore incorrectly categorise highly restricted diffusion as low blood flow velocity.

Furthermore, the vascular volume fraction parameter may have been influenced by the fact that the simulation dynamics neglect the effects of membrane permeability. This is unrealistic since in real tumour tissue, with high vessel permeability, the water molecules in the vessels are able to pass through the permeable vessel walls and into the extracellular space. Neglecting this effect in the simulation, and hence the synthetic data used as the training data may effect the performance of the neural network for the vascular parameters.

5.6 Chapter Summary

The aim of this chapter was to investigate a novel neural-network based approach for extracting microstructural tissue features from diffusion MRI data without the need for complex diffusion compartment models or an explicit model-fitting process. One of the key novel features of this approach is the use of synthetically-generated diffusion MRI data to train the neural network. Therefore, the key question to be answered was: Can a neural network trained on synthetic data extract useful features from real *in-vivo* data?

The first part of the experiment was the development of the neural network. The theory behind neural networks and how they ‘learn’ was set out in the [artificial neural networks](#) section. The training dataset for the network was then generated using the Monte Carlo

simulation framework developed in chapter 4. The network dimensions and parameters were selected based on the dimensions of the training data, and the network was then trained on the synthetic data.

The neural network was then applied to *in-vivo* data from real-world tumours, and the parameter maps produced using the conventional model fitting approach and the machine learning based approach were compared. The results showed that the machine learning based approach was able to generate sensible parameter maps for the parameters of the VERDICT BallSphereAstroSticks model. However, despite the evident correspondence between the learned and fitted parameters, the absolute values of the learned parameters were consistently different to the fitted parameters. The reason for this discrepancy is not known, and future work should focus on validation of the absolute parameter values output from the network. One promising result was the ability of the network to output parameters relating to the blood flow in the tissue. This has never been attempted previously due to the challenges of generating synthetic data that includes the effects of diffusion and flow, or compiling enough ground truth blood flow measurements to create a training dataset of suitable size.

The work presented in this chapter is only intended to demonstrate the potential for a machine-learning based approach to microstructural imaging, and as such the methods presented herein have many associated limitations and shortcomings. The lack of validation of the neural network outputs when applied to the *in-vivo* data means it is difficult to assess whether the network performed better or worse than the conventional model fitting based approach. The other limitations discussed above relating to the limited number of synthetic datasets, and in chapter 4 in relation to the simplistic nature of the monte carlo simulation framework itself must also be taken into account when assessing the results of this chapter. However, despite these limitations, the work presented demonstrates the potential role that synthetic data and machine learning could play in the field of diffusion

MRI.

Future work should focus on properly validating the neural network output parameters using histology and other flow-sensitive *in-vivo* imaging techniques. However, overall the results presented above show that machine-learning has huge potential in the field of microstructural imaging, and the wider field cancer imaging in general.

Final Summary and Conclusions

In the UK alone, it is estimated that there are currently between 2.5 and 3 million people living with diagnosed cancer [217]. This figure is expected to rise to 4 million by the year 2030. As the social and economic burden of cancer continues to grow, it is imperative that new therapies and treatments continue to be developed. Alongside the development of these novel therapies, it is crucial that new diagnostic and prognostic methods are developed that are sensitive to the pathophysiological changes caused by newly-developed therapies. To this end, medical imaging techniques play a vital role in providing effective biomarkers for disease progression and therapeutic response. One of the main advantages of medical imaging based methods is that they are non-invasive and repeatable, compared to conventional methods that require a tumour biopsy to be performed. MRI has come to the forefront of cancer imaging, due to its excellent soft tissue contrast, and ability to be sensitised to different tissue properties. In the field of cancer imaging, diffusion-weighted MRI has gained particular success, due to its sensitivity to the microstructure of tissue. Many hallmarks of cancer are related to the changes in cellular structure of the tissue. Diffusion-weighted MRI has therefore cemented its place as a powerful tool for quantifying and studying cancer. Its utility has been greatly improved by the introduction of advanced modelling and analysis techniques such as VERDICT and NODDI, and many others that build on the success of the established techniques such as IVIM or ADC. These new techniques have demonstrated great potential for enabling access to a wealth

of information about the microstructure of tissue without the need for invasive biopsy. These new metrics such as the cell radius and compartmental volume fractions could prove to be invaluable biomarkers for both disease prognosis and therapy response. The overall aim of the research in this thesis was to aid in the development of novel non-invasive biomarkers for cancer therapy. In particular, to aid in the validation of advanced diffusion modelling techniques such as VERDICT. In addition, to investigate the potential for machine-learning based techniques that may soon revolutionise the field of microstructural diffusion imaging. This chapter summarises the studies carried out throughout the thesis, discusses some of the successes and limitations of each study and outlines potential avenues for future development.

In chapter 3, a novel diffusion-weighted MRI method for measuring microstructural properties of tumour tissue, known as Vascular, Extracellular and Restricted Diffusion for Cytometry in Tumours (VERDICT) was used to detect changes in tumour microstructure caused by tissue fixation and administration of temozolomide therapy. The VERDICT technique was able to quantify changes in cell radius, intracellular volume and diffusivity between the *in-vivo* and fixed tissue. The same level of significance was found in the parameter changes for the VERDICT models and the conventional ADC. The VERDICT technique was also able to detect changes in the cell radius parameter in response to the temozolomide therapy at an earlier time point than either the ADC, or bulk volume measurements. The studies were therefore successful in demonstrating the capabilities of the VERDICT framework, and the utility of the parameters associated with it as cancer therapy biomarkers. The results of the experiments also showed that the performance of the VERDICT technique is influenced by the choice of compartment models. The parameters of the ‘BallSphereAstroSticks’ model (with an isotropic vascular compartment) were generally not as sensitive to the changes in tissue microstructure as those of the ‘BallSphereStick’ model (anisotropic vascular compartment). This was likely due to the

reduced number of fitted parameters in the ‘BallSphereAstroSticks’ model, which also called into question whether the ‘BallSphereStick’ model may be overfitting the data. An informative future study would be to investigate whether the early detection of microstructural changes using VERDICT corresponds to a better long-term prognosis. This would help to answer the question as to whether the VERDICT parameters can truly be considered useful biomarkers for therapy response.

In chapter 4, the development of a Monte Carlo simulation that incorporates the effects of diffusion and flow was presented. Following this, the framework was applied as a tool for validating the mathematical models of the VERDICT technique. In the first part of the chapter, the underlying methods for the simulation framework were presented in three main parts: the creation of the substrate, the development of the walker dynamics and the generation of synthetic data. The simulation was then validated by comparing the synthetic signal generated by the framework to the gold standard signal calculated using analytical models of diffusion in parallel cylinders and packed spheres. The results showed that the framework was able to generate realistic synthetic MRI signal compared to the analytical gold standard. In addition, the simulation framework performed on par with another publicly available simulation framework. One of the main limitations of this part of the study was the lack of validation of the flow dynamics. Future work in this area should aim to validate the dual flow and diffusion dynamics, using a suitable phantom or by developing an analytical solution. Additional limitations are related to the simplistic nature of the simulation substrate, for example, representing the cells as simple non-permeable spheres. Further development of the simulation framework should also focus on creating more advanced and realistic substrates, preferably based on real biological tumour tissue. In the second part of the chapter, the simulation framework was applied as a tool for validating the mathematical models of the VERDICT framework, as well as other clinically relevant models such as ADC and IVIM. The validation was performed

by generating signal from simulation substrates with varied substrate parameters such as cell radius and tissue volume fractions. The results showed that the parameters of both VERDICT models were generally well correlated with their intended microstructural property, and generally did not correlate with other substrate parameters. The promising results suggest that the parameters of the VERDICT framework may indeed be able to provide reliable access to microstructural tissue properties. However, due to the simplistic nature of the simulation framework, it is difficult to predict whether the model parameters truly correlate with the microstructural tissue features they are intended to measure. Further developments to the Monte Carlo simulation framework relating to the complexity and sophistication of the tissue substrate may assist in drawing more concrete conclusions. In addition, the results for the clinically-accepted models showed that their parameters tended to correlate with a wider range of substrate parameters, thus making it difficult to determine, without prior knowledge, which substrate parameter caused a given change in the model parameter value. This result suggests that, when they are properly validated, more complex compartmental diffusion models such as VERDICT may be beneficial in a clinical setting. Future work should focus on the development of more sophisticated *in-silico* validation methods, combined with in-depth histological validation.

In chapter 5, an investigation into the potential applications for machine learning within the field of diffusion MRI modelling was presented. The aim was to investigate whether a neural network could replace the explicit model-fitting process and output useful microstructural tissue metrics based solely on the raw signal. The Monte Carlo simulation framework from chapter 4 was used to generate a large training dataset on which to train the neural network. The trained network was then applied to real world *in-vivo* data of subcutaneously grown tumours, to generate parameter maps of ‘learned’ tissue features that broadly correspond to the parameters of the VERDICT ‘BallSphereAstro-Sticks’ model. A comparison of the ‘learned’ vs fitted parameter maps showed that the

machine-learning approach was able to generate sensible parameter maps, that loosely correspond to those generated through model fitting. However, the absolute values measured by the neural network were consistently different to the fitted values. The study successfully demonstrated the huge potential for machine learning in the field of microstructural imaging, and further outlined an additional useful application of *in-silico* diffusion modelling techniques for the generation of large training datasets for machine learning models. There is a huge amount of further work to be done in this field, and endless avenues to be explored. One such avenue is the use of other machine learning algorithms such as random forrests, that enable the ‘importance’ of each model input to be measured. This could allow sequences to be fine-tuned to remove unnecessary measurements, and help to reduce scan time without sacrificing accuracy.

To conclude, in this thesis the capabilities of advanced diffusion models such as VERDICT have been demonstrated, and validated using a sophisticated simulation framework. In addition, the potential applications of machine learning as an alternative to explicit diffusion modelling were investigated. As an increasing number of advanced cancer therapies become more readily available, the ongoing development, validation and exploration of new ways to understand and measure the microstructural environment of tumours is crucial. It is my belief that diffusion-MRI will continue to play a vital role in cancer imaging, due to its inherent sensitivity to the tissue microstructure. It is therefore my hope that the work presented in this thesis can benefit the future development of diffusion MRI analysis techniques.

Bibliography

- [1] B Stewart, Christopher P Wild, et al. World cancer report 2014. *World*, 2016.
- [2] World Health Organisation. Cancer. <http://www.who.int/mediacentre/factsheets/fs297/en/>, 2015. [Online; accessed 15-11-2016].
- [3] Angela D’Esposito, Daniil Nikitichev, Adrien Desjardins, Simon Walker-Samuel, and Mark F Lythgoe. Quantification of light attenuation in optically cleared mouse brains. *Journal of biomedical optics*, 20(8):80503, aug 2015.
- [4] Angela D’Esposito, Paul Sweeney, Morium Ali, Magdy Saleh, Rajiv Ramasawmy, Thomas A Roberts, Adrien Desjardins, Mark F Lythgoe, R Barbara Pedley, Rebecca Shipley, and Simon Walker-Samuel. Combining optical imaging of cleared tissue with mathematical modelling to predict drug delivery and therapeutic response. *bioRxiv*, page 219865, nov 2017.
- [5] Raúl Rojas. Neural networks-a systematic introduction springer-verlag. *New York*, 1996.
- [6] Eleftheria Panagiotaki, Torben Schneider, Bernard Siow, Matt G Hall, Mark F Lythgoe, and Daniel C Alexander. Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison. *NeuroImage*, 59(3):2241–54, feb 2012.

- [7] Patrick Therasse, Susan G. Arbuck, Elizabeth A. Eisenhauer, Jantien Wanders, Richard S. Kaplan, Larry Rubinstein, Jaap Verweij, Martine Van Glabbeke, Allan T. van Oosterom, Michael C. Christian, and Steve G. Gwyther. New Guidelines to Evaluate the Response to Treatment in Solid Tumors. *JNCI: Journal of the National Cancer Institute*, 92(3):205–216, feb 2000.
- [8] E A Eisenhauer, P Therasse, J Bogaerts, L H Schwartz, D Sargent, R Ford, J Dancey, S Arbuck, S Gwyther, M Mooney, L Rubinstein, L Shankar, L Dodd, R Kaplan, D Lacombe, and J Verweij. New response evaluation criteria in solid tumours: Revised RECIST guideline (version 1.1). *European journal of cancer (Oxford, England : 1990)*, 2009.
- [9] Eletheria Panagiotaki, Simon Walker-Samuel, Bernard Siow, S Peter Johnson, Vineeth Rajkumar, R Barbara Pedley, Mark F Lythgoe, and Daniel C Alexander. Non-invasive quantification of solid tumor microstructure using VERDICT MRI. *Cancer research*, 74(7):1902–12, apr 2014.
- [10] John F. R. Kerr, Clay M. Winterford, and Brian V. Harmon. Apoptosis. Its significance in cancer and cancer Therapy. *Cancer*, 73(8):2013–2026, apr 1994.
- [11] American Cancer Society. Testing Biopsy and Cytology Specimens for Cancer. <http://www.cancer.org/acs/groups/cid/documents/webcontent/003185-pdf.pdf>, 2015. [Online; accessed 18-11-2016].
- [12] Cecil H Fox, Frank B Johnson, John Whiting, and Peter P Roller. Formaldehyde fixation. *J histochem Cytochem*, 33(8):845–853, 1985.
- [13] John W Gillespie, Carolyn J M Best, Verena E Bichsel, Kristina A Cole, Susan F Greenhut, Stephen M Hewitt, Mamoun Ahram, Yvonne B Gathright, Maria J Merino, Robert L Strausberg, Jonathan I Epstein, Stanley R Hamilton, Gallya Gannot, Galina V Baibakova, Valerie S Calvert, Michael J Flaig, Rodrigo F Chuaqui,

- Judi C Herring, John Pfeifer, Emmanuel F Petricoin, W Marston Linehan, Paul H Duray, G Steven Bova, and Michael R Emmert-Buck. Evaluation of non-formalin tissue fixation for molecular profiling studies. *The American journal of pathology*, 160(2):449–57, feb 2002.
- [14] Joon W. Choi and Diane K. Stoecker. Effects of Fixation on Cell Volume of Marine Planktonic Protozoa. *Appl. Envir. Microbiol.*, 55(7):1761–1765, jul 1989.
- [15] S Richardson, B Siow, A M Batchelor, M F Lythgoe, and D C Alexander. A viable isolated tissue system: a tool for detailed MR measurements and controlled perturbation in physiologically stable tissue. *Magnetic resonance in medicine*, 69(6):1603–10, jun 2013.
- [16] A Afaq, A Andreou, and D M Koh. Diffusion-weighted magnetic resonance imaging for tumour response assessment: why, when and how? *Cancer imaging : the official publication of the International Cancer Imaging Society*, 10 Spec no:S179–88, jan 2010.
- [17] Miho I Yoshikawa, Shozo Ohsumi, Shigenori Sugata, Masaaki Kataoka, Shigemitsu Takashima, Teruhito Mochizuki, Hirohiko Ikura, and Yutaka Imai. Relation between cancer cellularity and apparent diffusion coefficient values using diffusion-weighted magnetic resonance imaging in breast cancer. *Radiation medicine*, 26(4):222–6, may 2008.
- [18] M Eis, T Els, and M HoehnBerlage. High resolution quantitative relaxation and diffusion MRI of three different experimental brain tumors in rat. *Magnetic resonance in ...*, 1995.
- [19] Z. Rumboldt, D.L.A. Camacho, D. Lake, C.T. Welsh, and M. Castillo. Apparent Diffusion Coefficients for Differentiation of Cerebellar Tumors in Children. *AJNR Am. J. Neuroradiol.*, 27(6):1362–1369, jun 2006.

- [20] Denis Le Bihan and E. Breton. Imagerie de diffusion in-vivo par résonance magnétique nucléaire. *Comptes-Rendus de l'Académie des Sciences*, 93(5):27–34, dec 1985.
- [21] Hui Zhang, Torben Schneider, Claudia A Wheeler-Kingshott, and Daniel C Alexander. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage*, 61(4):1000–16, jul 2012.
- [22] D. C. ALEXANDER. Multiple-Fiber Reconstruction Algorithms for Diffusion MRI. *Annals of the New York Academy of Sciences*, 1064(1):113–133, dec 2005.
- [23] Kurt G. Schilling, Vaibhav Janve, Yurui Gao, Iwona Stepniewska, Bennett A. Landman, and Adam W. Anderson. Histological validation of diffusion MRI fiber orientation distributions and dispersion. *NeuroImage*, 165:200–221, jan 2018.
- [24] Qiuyun Fan, Aapo Nummenmaa, Barbara Wichtmann, Thomas Witzel, Choukri Mekkaoui, Walter Schneider, Lawrence L. Wald, and Susie Y. Huang. Validation of diffusion MRI estimates of compartment size and volume fraction in a biomimetic brain phantom using a human MRI scanner with 300mT/m maximum gradient strength. *NeuroImage*, jan 2018.
- [25] Cyril Poupon, Bernard Rieul, Irina Kezele, Muriel Perrin, Fabrice Poupon, and Jean-François Mangin. New diffusion phantoms dedicated to the study and validation of high-angular-resolution diffusion imaging (HARDI) models. *Magnetic Resonance in Medicine*, 60(6):1276–1283, dec 2008.
- [26] Muriel Perrin, Cyril Poupon, Bernard Rieul, Patrick Leroux, André Constantinesco, Jean-François Mangin, and Denis LeBihan. Validation of q-ball imaging with a diffusion fibre-crossing phantom on a clinical scanner. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1457), 2005.

- [27] Koen M Marien, Valerie Croons, Yannick Waumans, Ellen Sluydts, Stefanie De Schepper, Luc Andries, Wim Waelput, Erik Fransen, Peter B Vermeulen, Mark M Kockx, et al. Development and validation of a histological method to measure microvessel density in whole-slide images of cancer tissue. *PLoS One*, 11(9):e0161496, 2016.
- [28] Els Fieremans, Yves De Deene, Steven Delputte, Mahir S Ozdemir, Yves D’Asseler, Jelle Vlassenbroeck, Karel Deblaere, Eric Achten, and Ignace Lemahieu. Simulation and experimental verification of the diffusion in an anisotropic fiber phantom. *Journal of magnetic resonance (San Diego, Calif. : 1997)*, 190(2):189–99, feb 2008.
- [29] Christian Beaulieu and Peter S. Allen. Water diffusion in the giant axon of the squid: Implications for diffusion-weighted MRI of the nervous system. *Magnetic Resonance in Medicine*, 32(5):579–583, nov 1994.
- [30] Evangelia I. Zacharaki, Sumei Wang, Sanjeev Chawla, Dong Soo Yoo, Ronald Wolf, Elias R. Melhem, and Christos Davatzikos. Classification of brain tumor type and grade using MRI texture and shape in a machine learning scheme. *Magnetic Resonance in Medicine*, 62(6):1609–1618, dec 2009.
- [31] Gemma L. Nedjati-Gilani, Torben Schneider, Matt G. Hall, Niamh Cawley, Ioana Hill, Olga Ciccarelli, Ivana Drobnjak, Claudia A.M. Gandini Wheeler-Kingshott, and Daniel C. Alexander. Machine learning based compartment models with permeability for white matter microstructure imaging. *NeuroImage*, 150:119–135, apr 2017.
- [32] John Del Gaizo, Neda Mofrad, Jens H Jensen, David Clark, Russell Glenn, Joseph Helpert, and Leonardo Bonilha. Using machine learning to classify temporal lobe epilepsy based on diffusion MRI. *Brain and behavior*, 7(10):e00801, oct 2017.
- [33] David M. Schnyer, Peter C. Clasen, Christopher Gonzalez, and Christopher G. Beev-

- ers. Evaluating the diagnostic utility of applying a machine learning algorithm to diffusion tensor MRI measures in individuals with major depressive disorder. *Psychiatry Research: Neuroimaging*, 264:1–9, jun 2017.
- [34] Maonan Wang, Jingzhou Zhao, Lishen Zhang, Fang Wei, Yu Lian, Yingfeng Wu, Zhaojian Gong, Shanshan Zhang, Jianda Zhou, Ke Cao, et al. Role of tumor microenvironment in tumorigenesis. *Journal of Cancer*, 8(5):761, 2017.
- [35] Alison C Lloyd. The regulation of cell size. *Cell*, 154(6):1194–1205, 2013.
- [36] Alecsandru Ioan Baba and Cornel Câtoi. Tumor cell morphology. *Comparative Oncology*, 2007.
- [37] JC Mottram. Observations on the morphology of tumour cells. *The Journal of Pathology and Bacteriology*, 54(4):511–514, 1942.
- [38] Rakesh K Jain. Transport of molecules in the tumor interstitium: a review. *Cancer research*, 47(12):3039–3051, 1987.
- [39] Fabienne Danhier, Olivier Feron, and Véronique Préat. To exploit the tumor microenvironment: passive and active tumor targeting of nanocarriers for anti-cancer drug delivery. *Journal of controlled release*, 148(2):135–146, 2010.
- [40] Lýdia Vargová, Aleš Homola, Josef Zámečník, Michal Tichý, Vladimír Beneš, and Eva Syková. Diffusion parameters of the extracellular space in human gliomas. *Glia*, 42(1):77–88, 2003.
- [41] MJ Duffy. The biochemistry of metastasis. In *Advances in clinical chemistry*, volume 32, pages 135–166. Elsevier, 1996.
- [42] Ray H. Hashemi, William G. Bradley, and Christopher J. Lisanti. *MRI : the basics*. Lippincott Williams & Wilkins, 2010.

- [43] Donald W. McRobbie. MRI: From Picture to Proton. *Radiology*, 232(2):474–474, aug 2004.
- [44] Matt A. Bernstein, Kevin Franklin. King, and Xiaohong Joe. Zhou. *Handbook of MRI pulse sequences*. Academic Press, 2004.
- [45] Robert W. Brown, Yu-Chung N. Cheng, E. Mark Haacke, Michael R. Thompson, and Ramesh Venkatesan. *Magnetic resonance imaging : physical principles and sequence design*. John Wiley & Sons, 2014.
- [46] Lars G Hanson. Is Quantum Mechanics necessary for understanding Magnetic Resonance? *Concepts in Magnetic Resonance Part A: An Educational Journal*, 32(5):329–340, 2008.
- [47] Malcolm H. Levitt. *Spin dynamics : basics of nuclear magnetic resonance*. John Wiley & Sons, 2001.
- [48] Marinus T. Vlaardingerbroek and Jacques A. den Boer. *Magnetic Resonance Imaging*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.
- [49] S Mansson and A in Bjornerud. Physical principles of medical imaging by nuclear magnetic resonance. *The chemistry of contrast agents in medical magnetic resonance imaging*. Chichester: John Wiley and Sons, Ltd, pages 1–43, 2001.
- [50] James Keeler. NMR and energy levels, 2004.
- [51] Felix Bloch. Nuclear induction. *Physical review*, 70(7-8):460, 1946.
- [52] Edward M Purcell, H Co Torrey, and Robert V Pound. Resonance absorption by nuclear magnetic moments in a solid. *Physical review*, 69(1-2):37, 1946.
- [53] A. Abragam. *The principles of nuclear magnetism*. Clarendon Press, 1983.

- [54] W.S. Hinshaw and A.H. Lent. An introduction to NMR imaging: From the Bloch equation to the imaging equation. *Proceedings of the IEEE*, 71(3):338–350, 1983.
- [55] R Brown. Signal Detection Concepts. In *Magnetic Resonance Imaging*, pages 95–111. John Wiley & Sons, Inc., Hoboken, NJ, USA, may 2014.
- [56] Robert W. Brown, Yu-Chung N. Cheng, E. Mark Haacke, Michael R. Thompson, and Ramesh Venkatesan, editors. *Magnetic Resonance Imaging*. John Wiley & Sons Ltd, Chichester, UK, apr 2014.
- [57] W A Edelstein, J M S Hutchison, G Johnson, and T Redpath. Spin warp NMR imaging and applications to human whole-body imaging. *Physics in Medicine and Biology*, 25(4):751–756, jul 1980.
- [58] P M Robitaille, A M Abduljalil, and A Kangarlu. Ultra high resolution imaging of the human head at 8 tesla: 2K x 2K for Y2K. *Journal of computer assisted tomography*, 24(1):2–8, 2000.
- [59] Klaus-Dietmar Merboldt, Wolfgang Hanicke, and Jens Frahm. Self-diffusion NMR imaging using stimulated echoes. *Journal of Magnetic Resonance (1969)*, 64(3):479–486, oct 1985.
- [60] Albert Einstein. *Investigations on the Theory of the Brownian Movement*. Courier Corporation, 1956.
- [61] E. O. Stejskal and J. E. Tanner. Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient. *The Journal of Chemical Physics*, 42(1):288, jul 1965.
- [62] C. H. Neuman. Spin echo of spins diffusing in a bounded medium. *The Journal of Chemical Physics*, 60(11):4508, August 1974.

- [63] D Le Bihan, E Breton, D Lallemand, P Grenier, E Cabanis, and M Laval-Jeantet. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–7, nov 1986.
- [64] Derek K Jones and Peter J Basser. squashing peanuts and smashing pumpkins: how noise distorts diffusion-weighted mr data. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 52(5):979–993, 2004.
- [65] Derek K. Jones and Mara Cercignani. Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR in Biomedicine*, 23(7):803–820, sep 2010.
- [66] Denis Le Bihan. Apparent diffusion coefficient and beyond: what diffusion MR imaging can tell us about tissue structure. *Radiology*, 268(2):318–22, aug 2013.
- [67] Ashok Srinivasan, Mayank Goyal, Faisal Al Azri, and Cheemun Lum. State-of-the-Art Imaging of Acute Stroke. *RadioGraphics*, 26(suppl.1):S75–S95, oct 2006.
- [68] Laura M. Allen, Anton N. Hasso, Jason Handwerker, and Hamed Farid. Sequence-specific MR Imaging Findings That Are Useful in Dating Ischemic Stroke. *RadioGraphics*, 32(5):1285–1297, sep 2012.
- [69] Michael A Jacobs, Ronald Ouwerkerk, Kyle Petrowski, and Katarzyna J Macura. Diffusion-weighted imaging with apparent diffusion coefficient mapping and spectroscopy in prostate cancer. *Topics in magnetic resonance imaging : TMRI*, 19(6):261–72, dec 2008.
- [70] Hal Blumenfeld. *Neuroanatomy Through Clinical Cases (Blumenfeld, Neuroanatomy through Clinical Cases)*. Sinauer Associates, 2002.
- [71] M E Moseley, Y Cohen, J Kucharczyk, J Mintorovitch, H S Asgari, M F Wendland,

- J Tsuruda, and D Norman. Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system. *Radiology*, 176(2):439–445, aug 1990.
- [72] T L Chenevert, J A Brunberg, and J G Pipe. Anisotropic diffusion in human white matter: demonstration with MR techniques in vivo. *Radiology*, 177(2):401–405, nov 1990.
- [73] P.J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1):259–267, jan 1994.
- [74] P J Basser, J Mattiello, and D LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of magnetic resonance. Series B*, 103(3):247–54, mar 1994.
- [75] Paul Malloy, Stephen Correia, Glenn Stebbins, and David H. Laidlaw. Neuroimaging of White Matter in Aging and Dementia. *The Clinical Neuropsychologist*, 21(1):73–109, jan 2007.
- [76] Christopher H. Sotak. The role of diffusion tensor imaging in the evaluation of ischemic brain injury - a review. *NMR in Biomedicine*, 15(7-8):561–569, nov 2002.
- [77] D Le Bihan, E Breton, D Lallemand, M L Aubin, J Vignaud, and M Laval-Jeantet. Separation of diffusion and perfusion in intravoxel incoherent motion MR imaging. *Radiology*, 168(2):497–505, aug 1988.
- [78] Chen Ye, Daoyun Xu, Yongbin Qin, Lihui Wang, Rongpin Wang, Wuchao Li, Zixiang Kuai, and Yuemin Zhu. Estimation of intravoxel incoherent motion parameters using low b-values. *PloS one*, 14(2):e0211911, 2019.
- [79] R. Mark Henkelman. Does IVIM measure classical perfusion? *Magnetic Resonance in Medicine*, 16(3):470–475, nov 2005.

- [80] Denis Le Bihan and Robert Turner. The capillary network: a link between ivim and classical perfusion. *Magnetic Resonance in Medicine*, 27(1):171–178, sep 1992.
- [81] A. Andreou, D. M. Koh, D. J. Collins, M. Blackledge, T. Wallace, M. O. Leach, and M. R. Orton. Measurement reproducibility of perfusion fraction and pseudodiffusion coefficient derived by intravoxel incoherent motion diffusion-weighted MR imaging in normal liver and metastases. *European Radiology*, 23(2):428–434, feb 2013.
- [82] Greg J. Stanisz, Graham A. Wright, R. Mark Henkelman, and Aaron Szafer. An analytical model of restricted diffusion in bovine optic nerve. *Magnetic Resonance in Medicine*, 37(1):103–111, jan 1997.
- [83] JÖRG KÄRGER, HARRY PFEIFER, and WILFRIED HEINK. Principles and Application of Self-Diffusion Measurements by Nuclear Magnetic Resonance. *Advances in Magnetic and Optical Resonance*, 12:1–89, jan 1988.
- [84] T E J Behrens, M W Woolrich, M Jenkinson, H Johansen-Berg, R G Nunes, S Clare, P M Matthews, J M Brady, and S M Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 50(5):1077–88, nov 2003.
- [85] Timothy EJ Behrens, H Johansen Berg, Saad Jbabdi, Matthew FS Rushworth, and Mark W Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage*, 34(1):144–155, 2007.
- [86] Yaniv Assaf, Raisa Z Freidlin, Gustavo K Rohde, and Peter J Basser. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic resonance in medicine*, 52(5):965–78, nov 2004.
- [87] Yaniv Assaf, Tamar Blumenfeld-Katzir, Yossi Yovel, and Peter J Basser. Axc caliber:

- a method for measuring axon diameter distribution from diffusion mri. *Magnetic resonance in medicine*, 59(6):1347–1354, 2008.
- [88] Francisco Aboitiz, Arnold B Scheibel, Robin S Fisher, and Eran Zaidel. Fiber composition of the human corpus callosum. *Brain research*, 598(1-2):143–153, 1992.
 - [89] Daniel Barazany, Peter J Basser, and Yaniv Assaf. In vivo measurement of axon diameter distribution in the corpus callosum of rat brain. *Brain*, 132(5):1210–1220, 2009.
 - [90] Daniel C. Alexander. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine*, 60(2):439–448, aug 2008.
 - [91] H Zhang, D Barazany, Y Assaf, HM Lundell, DC Alexander, and T Dyrby. A comparative study of axon diameter imaging techniques using diffusion mri. In *Journal of the International Society for Magnetic Resonance in Medicine*. ISMRM, 2011.
 - [92] Hui Zhang, Penny L Hubbard, Geoff JM Parker, and Daniel C Alexander. Axon diameter mapping in the presence of orientation dispersion with diffusion mri. *Neuroimage*, 56(3):1301–1315, 2011.
 - [93] Gavin P Winston, Caroline Micallef, Mark R Symms, Daniel C Alexander, John S Duncan, and Hui Zhang. Advanced diffusion imaging sequences could aid assessing patients with focal cortical dysplasia and epilepsy. *Epilepsy research*, 108(2):336–339, 2014.
 - [94] Inge Timmers, Hui Zhang, Matteo Bastiani, Bernadette M Jansma, Alard Roebroek, and M Estela Rubio-Gozalbo. White matter microstructure pathology in classic galactosemia revealed by neurite orientation dispersion and density imaging. *Journal of inherited metabolic disease*, 38(2):295–304, 2015.

- [95] Nicolas Kunz, Hui Zhang, Lana Vasung, Kieran R O'brien, Yaniv Assaf, François Lazeyras, Daniel C Alexander, and Petra S Hüppi. Assessing white matter microstructure of the newborn with multi-shell diffusion mri and biophysical compartment models. *Neuroimage*, 96:288–299, 2014.
- [96] Zach Eaton-Rosen, Andrew Melbourne, Eliza Orasanu, M Jorge Cardoso, Marc Modat, Alan Bainbridge, Giles S Kendall, Nicola J Robertson, Neil Marlow, and Sebastien Ourselin. Longitudinal measurement of the developing grey matter in preterm subjects using multi-modal mri. *NeuroImage*, 111:580–589, 2015.
- [97] EE Sigmund, GY Cho, S Kim, M Finn, M Moccaldi, JH Jensen, DK Sodickson, JD Goldberg, S Formenti, and L Moy. Intravoxel incoherent motion imaging of tumor microenvironment in locally advanced breast cancer. *Magnetic resonance in medicine*, 65(5):1437–1447, 2011.
- [98] SF Riches, K Hawtin, EM Charles-Edwards, and NM De Souza. Diffusion-weighted imaging of the prostate and rectal wall: comparison of biexponential and monoexponential modelled diffusion and associated perfusion coefficients. *NMR in Biomedicine*, 22(3):318–325, 2009.
- [99] Miriam Klau, Andreas Lemke, Katharina Grünberg, Dirk Simon, Thomas J Re, Mortiz N Wente, Frederik B Laun, Hans-Ulrich Kauczor, Stefan Delorme, Lars Grenacher, et al. Intravoxel incoherent motion mri for the differentiation between mass forming chronic pancreatitis and pancreatic carcinoma. *Investigative radiology*, 46(1):57–63, 2011.
- [100] Hersh Chandarana, Stella K Kang, Samson Wong, Henry Rusinek, Jeff L Zhang, Shigeki Arizono, William C Huang, Jonathan Melamed, James S Babb, Edgar F Suan, et al. Diffusion-weighted intravoxel incoherent motion imaging of renal tumors with histopathologic correlation. *Investigative radiology*, 47(12):688–696, 2012.

- [101] Paul T. Callaghan. *Principles of Nuclear Magnetic Resonance Microscopy*. Clarendon Press, 1993.
- [102] J. S. Murday. Self-Diffusion Coefficient of Liquid Lithium. *The Journal of Chemical Physics*, 48(11):4938, September 1968.
- [103] Choe A. S., Stepniewska I., Colvin D. C., Ding Z., and Anderson A. W. Validation of diffusion tensor mri in the central nervous system using light microscopy: quantitative comparison of fiber properties. *NMR in Biomedicine*, 25(7):900–908, 2012.
- [104] Trygve B Leergaard, Nathan S White, Alex De Crespigny, Ingeborg Bolstad, Helen D’Arceuil, Jan G Bjaalie, and Anders M Dale. Quantitative histological validation of diffusion mri fiber orientation distributions in the rat brain. *PloS one*, 5(1):e8595, 2010.
- [105] Arne Seehaus, Alard Roebroek, Matteo Bastiani, Lúcia Fonseca, Hansjürgen Bratzke, Nicolás Lori, Anna Vilanova, Rainer Goebel, and Ralf Galuske. Histological validation of high-resolution dti in human post mortem tissue. *Frontiers in neuroanatomy*, 9:98, 2015.
- [106] Mihika Gangolli, Laurena Holleran, Joong Hee Kim, Thor D Stein, Victor Alvarez, Ann C McKee, and David L Brody. Quantitative validation of a nonlinear histology-mri coregistration method using generalized q-sampling imaging in complex human cortical white matter. *NeuroImage*, 153:152–167, 2017.
- [107] Sune Nørhøj Jespersen, Lindsey A Leigland, Anda Cornea, and Christopher D Kroenke. Determination of axonal and dendritic orientation distributions within the developing cerebral cortex by diffusion tensor imaging. *IEEE transactions on medical imaging*, 31(1):16–32, 2012.

- [108] Kurt Schilling, Vaibhav Janve, Yurui Gao, Iwona Stepniewska, Bennett A Landman, and Adam W Anderson. Comparison of 3d orientation distribution functions measured with confocal microscopy and diffusion mri. *Neuroimage*, 129:185–197, 2016.
- [109] Hui Wang, Christophe Lenglet, and Taner Akkin. Structure tensor analysis of serial optical coherence scanner images for mapping fiber orientations and tractography in the brain. *Journal of biomedical optics*, 20(3):036003, 2015.
- [110] Markus Axer, Sven Strohmer, David Gräbel, Oliver Bückner, Melanie Dohmen, Julia Reckfort, Karl Zilles, and Katrin Amunts. Estimating fiber orientation distribution functions in 3d-polarized light imaging. *Frontiers in neuroanatomy*, 10:40, 2016.
- [111] Jeroen Mollink, Michiel Kleinnijenhuis, Anne-Marie van Cappellen van Walsum, Stamatios N Sotiropoulos, Michiel Cottaar, Christopher Mirfin, Mattias P Heinrich, Mark Jenkinson, Menuka Pallegage-Gamarallage, Olaf Ansorge, et al. Evaluating fibre orientation dispersion in white matter: comparison of diffusion mri, histology and polarized light imaging. *Neuroimage*, 157:561–574, 2017.
- [112] Tim B Dyrby, William FC Baaré, Daniel C Alexander, Jacob Jelsing, Ellen Garde, and Lise V Sogaard. An ex vivo imaging pipeline for producing high-quality and high-resolution diffusion-weighted imaging datasets. *Human brain mapping*, 32(4):544–563, 2011.
- [113] Jennifer A McNab, Saâd Jbabdi, Sean CL Deoni, Gwenaëlle Douaud, Timothy EJ Behrens, and Karla L Miller. High resolution diffusion-weighted imaging in fixed human brain using diffusion-weighted steady state free precession. *Neuroimage*, 46(3):775–785, 2009.
- [114] Manfred Holz, Stefan R Heil, and Antonio Sacco. Temperature-dependent self-diffusion coefficients of water and six selected molecular liquids for calibration in ac-

- curate 1h nmr pfg measurements. *Physical Chemistry Chemical Physics*, 2(20):4740–4742, 2000.
- [115] Dariya Malyarenko, Craig J Galbán, Frank J Londy, Charles R Meyer, Timothy D Johnson, Alnawaz Rehemtulla, Brian D Ross, and Thomas L Chenevert. Multi-system repeatability and reproducibility of apparent diffusion coefficient measurement using an ice-water phantom. *Journal of Magnetic Resonance Imaging*, 37(5):1238–1246, 2013.
- [116] Thomas L Chenevert, Craig J Galbán, Marko K Ivancevic, Susan E Rohrer, Frank J Londy, Thomas C Kwee, Charles R Meyer, Timothy D Johnson, Alnawaz Rehemtulla, and Brian D Ross. Diffusion coefficient measurement using a temperature-controlled fluid for quality control in multicenter studies. *Journal of Magnetic Resonance Imaging*, 34(4):983–987, 2011.
- [117] Dariya I Malyarenko, David Newitt, Lisa J. Wilmes, Alina Tudorica, Karl G Helmer, Lori R Arlinghaus, Michael A Jacobs, Guido Jajamovich, Bachir Taouli, Thomas E Yankeelov, et al. Demonstration of nonlinearity bias in the measurement of the apparent diffusion coefficient in multicenter trials. *Magnetic resonance in medicine*, 75(3):1312–1323, 2016.
- [118] Minoru Fukuzaki, Nobuhiro Miura, Naoki Shinyashiki, Daisaku Kurita, Sumie Shioya, Munetaka Haida, and Satoru Mashimo. Comparison of water relaxation time in serum albumin solution using nuclear magnetic resonance and time domain reflectometry. *The Journal of Physical Chemistry*, 99(1):431–435, 1995.
- [119] Ioannis Lavdas, Kevin C Behan, Annie Papadaki, Donald W McRobbie, and Eric O Aboagye. A phantom for diffusion-weighted mri (dw-mri). *Journal of Magnetic Resonance Imaging*, 38(1):173–179, 2013.
- [120] Hans J Laubach, Peter M Jakob, Karl O Loevblad, Alison E Baird, Maria Picone

- Bovo, Robert R Edelman, and Steven Warach. A phantom for diffusion-weighted imaging of acute stroke. *Journal of Magnetic Resonance Imaging*, 8(6):1349–1354, 1998.
- [121] C Pierpaoli, J Sarlls, U Nevo, PJ Basser, and F Horkay. Polyvinylpyrrolidone (pvp) water solutions as isotropic phantoms for diffusion mri studies. In *Proc Intl Soc Magn Reson Med*, volume 17, page 1414, 2009.
- [122] Ryohei Matsuya, Masahiro Kuroda, Yoshitsugu Matsumoto, Hirokazu Kato, Hide-nobu Matsuzaki, Junichi Asaumi, Jun Murakami, Kazunori Katashima, Masakazu Ashida, Takanori Sasaki, et al. A new phantom using polyethylene glycol as an apparent diffusion coefficient standard for mr imaging. *International journal of oncology*, 35(4):893–900, 2009.
- [123] PS Tofts, D Lloyd, CA Clark, GJ Barker, GJM Parker, P McConville, C Baldock, and JM Pope. Test liquids for quantitative mri measurements of self-diffusion coefficient in vivo. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 43(3):368–374, 2000.
- [124] Z Ababneh, M Haque, SE Maier, and RV Mulkern. Dairy cream as a phantom material for biexponential diffusion decay. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 17(2):95–100, 2004.
- [125] Craig Jones, Alex MacKay, and Brian Rutt. Bi-exponential t2 decay in dairy cream phantoms. *Magnetic resonance imaging*, 16(1):83–85, 1998.
- [126] Els Fieremans, Antonio Pires, and Jens H Jensen. A simple isotropic phantom for diffusional kurtosis imaging. *Magnetic resonance in medicine*, 68(2):537–542, 2012.
- [127] Jimmy Lätt, Markus Nilsson, Anna Rydhög, Ronnie Wirestam, Freddy Ståhlberg, and Sara Brockstedt. Effects of restricted diffusion in a biological phantom: a q-

- space diffusion mri study of asparagus stems at a 3t clinical scanner. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 20(4):213, 2007.
- [128] Bernard Siow, Ivana Drobnjak, Aritrick Chatterjee, Mark F Lythgoe, and Daniel C Alexander. Estimation of pore size in a microstructure phantom using the optimised gradient waveform diffusion weighted nmr sequence. *Journal of Magnetic Resonance*, 214:51–60, 2012.
- [129] Hua Li, John C Gore, and Junzhong Xu. Fast and robust measurement of microstructural dimensions using temporal diffusion spectroscopy. *Journal of magnetic resonance*, 242:4–9, 2014.
- [130] Irvin Teh, Feng-Lei Zhou, Penny L Hubbard Cristinacce, Geoffrey JM Parker, and Jürgen E Schneider. Biomimetic phantom for cardiac diffusion mri. *Journal of Magnetic Resonance Imaging*, 43(3):594–600, 2016.
- [131] Q Ye Allen, Penny L Hubbard Cristinacce, Feng-Lei Zhou, Ziyang Yin, Geoff JM Parker, and Richard L Magin. Diffusion tensor mri phantom exhibits anomalous diffusion. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, pages 746–749. IEEE, 2014.
- [132] Penny L Hubbard, Feng-Lei Zhou, Stephen J Eichhorn, and Geoffrey JM Parker. Biomimetic phantom for the validation of diffusion magnetic resonance imaging. *Magnetic resonance in medicine*, 73(1):299–305, 2015.
- [133] Edward C Parsons Jr, Mark D Does, and John C Gore. Temporal diffusion spectroscopy: theory and implementation in restricted systems using oscillating gradients. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 55(1):75–84, 2006.
- [134] Damien J McHugh, Fenglei Zhou, Penny L Hubbard Cristinacce, Josephine H Naish, and Geoffrey JM Parker. Ground truth for diffusion mri in cancer: a model-based

- investigation of a novel tissue-mimetic material. In *International Conference on Information Processing in Medical Imaging*, pages 179–190. Springer, 2015.
- [135] Els Fieremans and Hong-Hsi Lee. Physical and numerical phantoms for the validation of brain microstructural mri: A cookbook. *NeuroImage*, 2018.
 - [136] Paul T Callaghan. A simple matrix formalism for spin echo analysis of restricted diffusion under generalized gradient waveforms. *Journal of Magnetic Resonance*, 129(1):74–84, 1997.
 - [137] Paul T Callaghan. Pulsed-gradient spin-echo nmr for planar, cylindrical, and spherical pores under conditions of wall relaxation. *Journal of magnetic resonance, Series A*, 113(1):53–59, 1995.
 - [138] Kevin D Harkins, Jean-Philippe Galons, Timothy W Secomb, and Theodore P Trouard. Assessment of the effects of cellular tissue properties on adc measurements by numerical simulation of water diffusion. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 62(6):1414–1422, 2009.
 - [139] Scott N Hwang, Chih-Liang Chin, Felix W Wehrli, and David B Hackney. An image-based finite difference model for simulating restricted diffusion. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 50(2):373–382, 2003.
 - [140] Håkan Hagslätt, Bengt Jönsson, Magnus Nydén, and Olle Söderman. Predictions of pulsed field gradient nmr echo-decays for molecules diffusing in various restrictive geometries. simulations of diffusion propagators based on a finite element method. *Journal of Magnetic Resonance*, 161(2):138–147, 2003.
 - [141] Leandro Beltrachini, Zeike A Taylor, and Alejandro F Frangi. A parametric finite

- element solution of the generalised bloch–torrey equation for arbitrary domains. *Journal of Magnetic Resonance*, 259:126–134, 2015.
- [142] Jing-Rebecca Li, Hang Tuan Nguyen, Dang Van Nguyen, Housseem Haddar, Julien Coatléven, and Denis Le Bihan. Numerical study of a macroscopic finite pulse model of the diffusion mri signal. *Journal of Magnetic Resonance*, 248:54–65, 2014.
 - [143] Benjamin F Moroney, Timothy Stait-Gardner, Bahman Ghadirian, Nirbhay N Yadav, and William S Price. Numerical analysis of nmr diffusion measurements in the short gradient pulse limit. *Journal of Magnetic Resonance*, 234:165–175, 2013.
 - [144] Dang Van Nguyen, Jing-Rebecca Li, Denis Grebenkov, and Denis Le Bihan. A finite elements method to solve the bloch–torrey equation applied to diffusion magnetic resonance imaging. *Journal of Computational Physics*, 263:283–302, 2014.
 - [145] Aaron Szafer, Jianhui Zhong, and John C Gore. Theoretical model for water diffusion in tissues. *Magnetic resonance in medicine*, 33(5):697–712, 1995.
 - [146] J Chetley Ford and David B Hackney. Numerical model for calculation of apparent diffusion coefficients (adc) in permeable cylinderscomparison with measured adc in spinal cord white matter. *Magnetic resonance in medicine*, 37(3):387–394, 1997.
 - [147] Daniel C Alexander, Penny L Hubbard, Matt G Hall, Elizabeth A Moore, Maurice Ptito, Geoff JM Parker, and Tim B Dyrby. Orientationally invariant indices of axon diameter and density from diffusion mri. *Neuroimage*, 52(4):1374–1389, 2010.
 - [148] Kévin Ginsburger, Fabrice Poupon, Justine Beaujoin, Delphine Estournet, Felix Matuschke, Jean-François Mangin, Markus Axer, and Cyril Poupon. Improving the realism of white matter numerical phantoms: A step toward a better understanding of the influence of structural disorders in diffusion mri. *Frontiers in Physics*, 6:12, 2018.

- [149] Markus Nilsson, Jimmy Lätt, Freddy Ståhlberg, Danielle van Westen, and Håkan Hagslätt. The importance of axonal undulation in diffusion mr measurements: a monte carlo simulation study. *NMR in Biomedicine*, 25(5):795–805, 2012.
- [150] Cyril Poupon, Bernard Rieul, Irina Kezele, Muriel Perrin, Fabrice Poupon, and Jean-François Mangin. New diffusion phantoms dedicated to the study and validation of high-angular-resolution diffusion imaging (hardi) models. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 60(6):1276–1283, 2008.
- [151] Alonso Ramirez-Manzanares, Philip A Cook, Matt Hall, Manzar Ashtari, and James C Gee. Resolving axon fiber crossings at clinical b-values: An evaluation study. *Medical physics*, 38(9):5239–5253, 2011.
- [152] Chun-Hung Yeh, Benoit Schmitt, Denis Le Bihan, Jing-Rebecca Li-Schlittgen, Ching-Po Lin, and Cyril Poupon. Diffusion microscopist simulator: a general monte carlo simulation system for diffusion magnetic resonance imaging. *PloS one*, 8(10):e76626, 2013.
- [153] Rabiya S Tuma. Sometimes size doesn’t matter: reevaluating recist and tumor response rate endpoints. *Journal of the National Cancer Institute*, 98(18):1272–1274, 2006.
- [154] Eleftheria Panagiotaki, Rachel W Chan, Nikolaos Dikaio, Hashim U Ahmed, James O’Callaghan, Alex Freeman, David Atkinson, Shonit Punwani, David J Hawkes, and Daniel C Alexander. Microstructural characterization of normal and malignant human prostate tissue with vascular, extracellular, and restricted diffusion for cytometry in tumours magnetic resonance imaging. *Investigative radiology*, 50(4):218–227, 2015.
- [155] Hendrik F van Essen, Marina AM Verdaasdonk, Sabrina M Elshof, Roel A de Weger,

- and Paul J van Diest. Alcohol based tissue fixation as an alternative for formaldehyde: influence on immunohistochemistry. *Journal of clinical pathology*, pages jcp–2010, 2010.
- [156] John Randal Baker. *Principles of biological microtechnique*. Methuen;Wiley, 1958.
- [157] W E Grizzle. Special symposium: fixation and tissue processing models. *Biotechnic & histochemistry : official publication of the Biological Stain Commission*, 84(5):185–93, October 2009.
- [158] Rooban Thavarajah, Vidya Kazhiyur Mudimbaimannar, Joshua Elizabeth, Umadevi Krishnamohan Rao, and Kannan Ranganathan. Chemical and physical basics of routine formaldehyde fixation. *Journal of oral and maxillofacial pathology: JOMFP*, 16(3):400, 2012.
- [159] M. M. Arnold, Sudhir Srivastava, Jerry Fredenburgh, Cecil R. Stockard, Russell B. Myers, and William E. Grizzle. Effects of Fixation and Tissue Processing on Immunohistochemical Demonstration of Specific Antigens. *Biotechnic and Histochemistry*, July 2009.
- [160] RW Dapson. Macromolecular changes caused by formalin fixation and antigen retrieval. *Biotechnic & Histochemistry*, May 2015.
- [161] D Otali, C R Stockard, D K Oelschlager, W Wan, U Manne, S A Watts, and W E Grizzle. Combined effects of formalin fixation and tissue processing on immunorecognition. *Biotechnic & histochemistry : official publication of the Biological Stain Commission*, 84(5):223–47, October 2009.
- [162] Mythily Srinivasan, Daniel Sedmak, and Scott Jewell. Effect of fixatives and tissue processing on the content and integrity of nucleic acids. *The American journal of pathology*, 161(6):1961–71, December 2002.

- [163] Debra Gusnard and Robert H. Kirschner. Cell and organelle shrinkage during preparation for scanning electron microscopy: effects of fixation, dehydration and critical point drying. *Journal of Microscopy*, 110(1):51–57, May 1977.
- [164] W Günther, E Pawlak, R Damasceno, H Arnold, and AJ Terzis. Temozolomide induces apoptosis and senescence in glioma cells cultured as multicellular spheroids. *British journal of cancer*, 88(3):463, 2003.
- [165] Wei Shen, Jun-An Hu, and Jie-Sheng Zheng. Mechanism of temozolomide-induced antitumour effects on glioma cells. *Journal of International Medical Research*, 42(1):164–172, 2014.
- [166] Tom A Roberts, Harpreet Hyare, Giulia Agliardi, Ben Hipwell, Angela d’Esposito, Andrada Ianus, James O Breen-Norris, Rajiv Ramasawmy, Valerie Taylor, David Atkinson, et al. Quantitation of brain tumour microstructure response to temozolomide therapy using non-invasive verdict mri. *bioRxiv*, page 182675, 2017.
- [167] P Workman, E O Aboagye, F Balkwill, A Balmain, G Bruder, D J Chaplin, J A Double, J Everitt, D A H Farningham, M J Glennie, L R Kelland, V Robinson, I J Stratford, G M Tozer, S Watson, S R Wedge, and S A Eccles. Guidelines for the welfare and use of animals in cancer research. *British journal of cancer*, 102(11):1555–77, May 2010.
- [168] PA Cook, Y Bai, SKKS Nedjati-Gilani, KK Seunarine, MG Hall, GJ Parker, and DC Alexander. Camino: open-source diffusion-mri reconstruction and processing. In *14th scientific meeting of the international society for magnetic resonance in medicine*, volume 2759. Seattle WA, USA, 2006.
- [169] S Pahernik, J Griebel, A Botzlar, T Gneiting, M Brandl, M Dellian, and AE Goetz. Quantitative imaging of tumour blood flow by contrast-enhanced magnetic resonance imaging. *British journal of cancer*, 85(11):1655–1663, 2001.

- [170] Glyn Johnson, Stephan G Wetzel, Soonmee Cha, James Babb, and Paul S Tofts. Measuring blood volume and vascular transfer constant from dynamic, t-weighted contrast-enhanced mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 51(5):961–968, 2004.
- [171] Louis Gagnon, Sava Sakadžić, Frédéric Lesage, Joseph J Musacchia, Joël Lefebvre, Qianqian Fang, Meryem A Yücel, Karleyton C Evans, Emiri T Mandeville, Jülien Cohen-Adad, et al. Quantifying the microvascular origin of bold-fmri from first principles with two-photon microscopy and an oxygen-sensitive nanoprobe. *Journal of Neuroscience*, 35(8):3663–3675, 2015.
- [172] Ulrich Pilatus, Hyunsuk Shim, Dmitri Artemov, Delphine Davis, Peter CM Van Zijl, and Jerry D Glickson. Intracellular volume and apparent diffusion constants of perfused cancer cell cultures, as measured by nmr. *Magnetic resonance in medicine*, 37(6):825–832, 1997.
- [173] Hermann Cuntz, Friedrich Forstner, Alexander Borst, and Michael Häusser. One rule to grow them all: a general theory of neuronal branching and its practical application. *PLoS computational biology*, 6(8):e1000877, 2010.
- [174] Mark Oldham, Harshad Sakhalkar, Tim Oliver, Ying Min Wang, John Kirpatrick, Yiting Cao, Cristian Badea, G. Allan Johnson, and Mark Dewhurst. Three-dimensional imaging of xenograft tumors using optical computed and emission tomography. *Medical Physics*, 33(9):3193, sep 2006.
- [175] Cemil Kirbas and Francis Quek. A Review of Vessel Extraction Techniques and Algorithms. *ACM Computing Surveys (CSUR)*, 2003.
- [176] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale vessel enhancement filtering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 130–137. Springer, 1998.

- [177] Brendan C Fry, Jack Lee, Nicolas P Smith, and Timothy W Secomb. Estimation of blood flow rates in large microvascular networks. *Microcirculation*, 19(6):530–538, 2012.
- [178] Lihong Wang, Steven L. Jacques, and Liqiong Zheng. MCMLMonte Carlo modeling of light transport in multi-layered tissues. *Computer Methods and Programs in Biomedicine*, 47(2):131–146, 1995.
- [179] Matt G. Hall and Daniel C. Alexander. Convergence and Parameter Choice for Monte-Carlo Simulations of Diffusion MRI. *IEEE transactions on medical imaging*, 28(9):1354–1364, 2009.
- [180] Anwar R Padhani, Guoying Liu, Dow Mu-Koh, Thomas L Chenevert, Harriet C Thoeny, Taro Takahara, Andrew Dzik-Jurasz, Brian D Ross, Marc Van Cauteren, David Collins, et al. Diffusion-weighted magnetic resonance imaging as a cancer biomarker: consensus and recommendations. *Neoplasia*, 11(2):102–125, 2009.
- [181] Haoyang Xing, Fang Lin, Qizhu Wu, and Qiyong Gong. Investigation of different boundary treatment methods in monte-carlo simulations of diffusion nmr. *Magnetic resonance in medicine*, 70(4):1167–1172, 2013.
- [182] Sandy Azzi, Jagoda K Hebda, and Julie Gavard. Vascular permeability and drug delivery in cancers. *Frontiers in oncology*, 3:211, 2013.
- [183] L Siegele. Welcome to the yotta world, the world in 2012. *The Economist*, 122, 2012.
- [184] Jay H Lee, Joohyun Shin, and Matthew J Realff. Machine learning: Overview of the recent progresses and implications for the process systems engineering field. *Computers & Chemical Engineering*, 114:111–121, 2018.

- [185] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- [186] Alexander Bain. *Mind and body: The theories of their relation*, volume 4. Henry S. King, 1873.
- [187] William James. *The principles of psychology*, volume 1. Read Books Ltd, 2013.
- [188] Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, 1943.
- [189] Frank Rosenblatt. Principles of neurodynamics. perceptrons and the theory of brain mechanisms. Technical report, CORNELL AERONAUTICAL LAB INC BUFFALO NY, 1961.
- [190] Bernard Widrow and Marcian E Hoff. Adaptive switching circuits. Technical report, STANFORD UNIV CA STANFORD ELECTRONICS LABS, 1960.
- [191] Xiangang Li and Xihong Wu. Constructing long short-term memory based deep recurrent neural networks for large vocabulary speech recognition. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4520–4524. IEEE, 2015.
- [192] Alex Graves, Marcus Liwicki, Santiago Fernández, Roman Bertolami, Horst Bunke, and Jürgen Schmidhuber. A novel connectionist system for unconstrained handwriting recognition. *IEEE transactions on pattern analysis and machine intelligence*, 31(5):855–868, 2009.
- [193] Haşim Sak, Andrew Senior, and Françoise Beaufays. Long short-term memory re-

- current neural network architectures for large scale acoustic modeling. In *Fifteenth annual conference of the international speech communication association*, 2014.
- [194] How to configure the number of layers and nodes in a neural network, Dec 2018.
- [195] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [196] Michael A Nielsen. *Neural networks and deep learning*, volume 25. Determination press USA, 2015.
- [197] Seppo Linnainmaa. Taylor expansion of the accumulated rounding error. *BIT Numerical Mathematics*, 16(2):146–160, 1976.
- [198] Paul Werbos. Beyond regression:” new tools for prediction and analysis in the behavioral sciences. *Ph. D. dissertation, Harvard University*, 1974.
- [199] David E Rumelhart, Geoffrey E Hinton, Ronald J Williams, et al. Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1, 1988.
- [200] Jianmin Jiang, P Trundle, and Jinchang Ren. Medical image analysis with artificial neural networks. *Computerized Medical Imaging and Graphics*, 34(8):617–631, 2010.
- [201] John Canny. A computational approach to edge detection. In *Readings in computer vision*, pages 184–203. Elsevier, 1987.
- [202] Jzau-Sheng Lin. Segmentation of medical images through a penalized fuzzy hopfield network with moments preservation. *Journal of The Chinese Institute of Engineers*, 23(5):633–643, 2000.
- [203] Syoji Kobashi, Naotake Kamiura, Yutaka Hata, and Fujio Miyawaki. Volume-quantization-based neural network approach to 3d mr angiography image segmentation. *Image and Vision Computing*, 19(4):185–193, 2001.

- [204] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 565–571. IEEE, 2016.
- [205] Chang Liu, Stephen J Gardner, Ning Wen, Mohamed A Elshaikh, Farzan Siddiqui, Benjamin Movsas, and Indrin J Chetty. Automatic segmentation of the prostate on ct images using deep neural networks (dnn). *International Journal of Radiation Oncology* Biology* Physics*, 2019.
- [206] Lifeng Shang, Jian Cheng Lv, and Zhang Yi. Rigid medical image registration using pca neural network. *Neurocomputing*, 69(13-15):1717–1722, 2006.
- [207] Bob D de Vos, Floris F Berendsen, Max A Viergever, Marius Staring, and Ivana Išgum. End-to-end unsupervised deformable image registration with a convolutional neural network. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 204–212. Springer, 2017.
- [208] Jun Ge, Berkman Sahiner, Lubomir M Hadjiiski, Heang-Ping Chan, Jun Wei, Mark A Helvie, and Chuan Zhou. Computer aided detection of clusters of microcalcifications on full field digital mammograms. *Medical Physics*, 33(8):2975–2988, 2006.
- [209] AR Kavitha, C Chellamuthu, and Kavin Rupa. An efficient approach for brain tumour detection based on modified region growing and neural network in mri images. In *2012 international conference on Computing, Electronics and Electrical Technologies (ICCEET)*, pages 1087–1095. IEEE, 2012.
- [210] Hao Dong, Guang Yang, Fangde Liu, Yuanhan Mo, and Yike Guo. Automatic brain tumor detection and segmentation using u-net based fully convolutional networks. In *annual conference on medical image understanding and analysis*, pages 506–517. Springer, 2017.

- [211] Suchita Goswami and Lalit Kumar P Bhaiya. Brain tumour detection using unsupervised learning based neural network. In *2013 International Conference on Communication Systems and Network Technologies*, pages 573–577. IEEE, 2013.
- [212] Toshiaki Hirasawa, Kazuharu Aoyama, Tetsuya Tanimoto, Soichiro Ishihara, Satoki Shichijo, Tsuyoshi Ozawa, Tatsuya Ohnishi, Mitsuhiro Fujishiro, Keigo Matsuo, Junko Fujisaki, et al. Application of artificial intelligence using a convolutional neural network for detecting gastric cancer in endoscopic images. *Gastric Cancer*, 21(4):653–660, 2018.
- [213] Kamil Dimililer, Yoney Kirsal Ever, and Haithm Ratemi. Intelligent eye tumour detection system. *Procedia Computer Science*, 102:325–332, 2016.
- [214] Vladimir Golkov, Alexey Dosovitskiy, Jonathan I Sperl, Marion I Menzel, Michael Czisch, Philipp Sämann, Thomas Brox, and Daniel Cremers. Q-space deep learning: twelve-fold shorter and model-free diffusion mri scans. *IEEE transactions on medical imaging*, 35(5):1344–1351, 2016.
- [215] Gemma L Nedjati-Gilani, Torben Schneider, Matt G Hall, Niamh Cawley, Ioana Hill, Olga Ciccarelli, Ivana Drobnjak, Claudia AM Gandini Wheeler-Kingshott, and Daniel C Alexander. Machine learning based compartment models with permeability for white matter microstructure imaging. *NeuroImage*, 150:119–135, 2017.
- [216] Jeff Heaton. *Introduction to neural networks with Java*. Heaton Research, Inc., 2008.
- [217] J Maddams, M Utey, and H Møller. Projections of cancer prevalence in the united kingdom, 2010–2040. *British journal of cancer*, 107(7):1195–1202, 2012.
- [218] Tim B. Dyrby, Lise V. Sgaard, Matt G. Hall, Maurice Ptito, and Daniel. C. Alexander. Contrast and stability of the axon diameter index from microstructure imaging with diffusion MRI. *Magnetic Resonance in Medicine*, 70(3):711–721, sep 2013.

- [219] Joseph Larmor. Lxiii. on the theory of the magnetic influence on spectra; and on the radiation from moving ions. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(271):503–512, 1897.
- [220] National Cancer Institute. Types of Treatment. <https://www.cancer.gov/about-cancer/treatment/types>, 2015. [Online; Accessed 15-11-2016].
- [221] E A Eisenhauer, P Therasse, J Bogaerts, L H Schwartz, D Sargent, R Ford, J Dancey, S Arbuck, S Gwyther, M Mooney, L Rubinstein, L Shankar, L Dodd, R Kaplan, D Lacombe, and J Verweij. New response evaluation criteria in solid tumours: revised RECIST guideline (version 1.1). *European journal of cancer (Oxford, England : 1990)*, 45(2):228–47, jan 2009.
- [222] Donald W, Elizabeth A McRobbie, Martin J Moore, Martin R Graves, and Prince. *MRI From Picture to Proton 2nd Edition*. Cambridge University Press, The Edinburgh Building, Cambridge CB2 8RU, UK, 2006.
- [223] Daniel Stucht, K Appu Danishad, Peter Schulze, Frank Godenschweger, Maxim Zaitsev, and Oliver Speck. Highest Resolution In Vivo Human Brain MRI Using Prospective Motion Correction. *PloS one*, 10(7):e0133921, 2015.
- [224] E. Hahn. Spin Echoes. *Physical Review*, 80(4):580–594, nov 1950.
- [225] H. Carr and E. Purcell. Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments. *Physical Review*, 94(3):630–638, may 1954.
- [226] Uran Ferizi, Torben Schneider, Eleftheria Panagiotaki, Gemma Nedjati-Gilani, Hui Zhang, Claudia A M Wheeler-Kingshott, and Daniel C Alexander. A ranking of diffusion MRI compartment models with in vivo human brain data. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 72(6):1785–92, dec 2014.

- [227] Yaniv Assaf, Tamar Blumenfeld-Katzir, Yossi Yovel, and Peter J Basser. AxCaliber: a method for measuring axon diameter distribution from diffusion MRI. *Magnetic resonance in medicine*, 59(6):1347–54, jun 2008.
- [228] Masayoshi Ikeda and Yoshinori Oka. The relationship between nerve conduction velocity and fiber morphology during peripheral nerve regeneration. *Brain and behavior*, 2(4):382–90, jul 2012.
- [229] J M Ritchie. On the relation between fibre diameter and conduction velocity in myelinated nerve fibres. *Proceedings of the Royal Society of London. Series B, Biological sciences*, 217(1206):29–35, dec 1982.
- [230] H. C. Torrey. Bloch Equations with Diffusion Terms. *Physical Review*, 104(3):563–565, nov 1956.
- [231] Kevin D. Harkins, Jean-Philippe Galons, Timothy W. Secomb, and Theodore P. Trouard. Assessment of the effects of cellular tissue properties on ADC measurements by numerical simulation of water diffusion. *Magnetic Resonance in Medicine*, 62(6):1414–1422, dec 2009.
- [232] Junzhong Xu, Mark D Does, and John C Gore. Dependence of temporal diffusion spectra on microstructural properties of biological tissues. *Magnetic resonance imaging*, 29(3):380–90, apr 2011.
- [233] Paul T Callaghan. *Translational dynamics and magnetic resonance: principles of pulsed gradient spin echo NMR*. Oxford University Press, 2011.
- [234] J C Ford and D B Hackney. Numerical model for calculation of apparent diffusion coefficients (ADC) in permeable cylinders—comparison with measured ADC in spinal cord white matter. *Magnetic resonance in medicine*, 37(3):387–94, mar 1997.

- [235] Chunlei Liu, Roland Bammer, Burak Acar, and Michael E. Moseley. Characterizing non-gaussian diffusion by using generalized diffusion tensors. *Magnetic Resonance in Medicine*, 51(5):924–937, may 2004.
- [236] MATLAB. *version 9.0.0.341360 (R2016a)*. The MathWorks Inc., Natick, Massachusetts, 2016.
- [237] Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial & Applied Mathematics*, 11(2):431–441, 1963.
- [238] Simon Walker-Samuel, Matthew R Orton, Lesley D McPhail, and Simon P Robinson. Robust Estimation of the Apparent Diffusion Coefficient (ADC) in Heterogeneous Solid Tumours. *Magnetic resonance in medicine*, 62:420–429, 2009.
- [239] Andrew L Alexander, Jee Eun Lee, Mariana Lazar, and Aaron S Field. Diffusion tensor imaging of the brain. *Neurotherapeutics : the journal of the American Society for Experimental NeuroTherapeutics*, 4(3):316–29, jul 2007.
- [240] José M Soares, Paulo Marques, Victor Alves, and Nuno Sousa. A hitchhiker’s guide to diffusion tensor imaging. *Frontiers in neuroscience*, 7:31, 2013.
- [241] A. Castriota-Scanderbeg, U. Sabatini, F. Fasano, R. Floris, L. Fraracci, Di M. Mario, U. Nocentini, and C. Caltagirone. Diffusion of water in large demyelinating lesions: a follow-up study. *Neuroradiology*, 44(9):764–767, sep 2002.
- [242] Tanguy Duval, Blanche Perraud, Manh-Tung Vuong, N Lopez Rios, Nikola Stikov, and Julien Cohen-Adad. Validation of quantitative mri metrics using full slice histology with automatic axon segmentation. In *Proceedings of the 24th Annual Meeting of ISMRM (Singapore)*, 2016.

- [243] Junzhong Xu, Mark D Does, and John C Gore. Sensitivity of mr diffusion measurements to variations in intracellular structure: effects of nuclear size. *Magnetic resonance in Medicine*, 61(4):828–833, 2009.
- [244] Paddy J Slator, Jana Hutter, Laura McCabe, Ana Dos Santos Gomes, Anthony N Price, Eleftheria Panagiotaki, Mary A Rutherford, Joseph V Hajnal, and Daniel C Alexander. Placenta microstructure and microcirculation imaging with diffusion mri. *Magnetic resonance in medicine*, 80(2):756–766, 2018.
- [245] Els Fieremans, Yves De Deene, Steven Delputte, Mahir S Özdemir, Eric Achten, and Ignace Lemahieu. The design of anisotropic diffusion phantoms for the validation of diffusion weighted magnetic resonance imaging. *Physics in Medicine & Biology*, 53(19):5405, 2008.
- [246] Els Fieremans, Yves De Deene, Steven Delputte, Mahir S Özdemir, Yves DAsseler, Jelle Vlassenbroeck, Karel Deblaere, Eric Achten, and Ignace Lemahieu. Simulation and experimental verification of the diffusion in an anisotropic fiber phantom. *Journal of magnetic resonance*, 190(2):189–199, 2008.
- [247] R Lorenz, ME Bellemann, J Hennig, and KA Il’yasov. Anisotropic phantoms for quantitative diffusion tensor imaging and fiber-tracking validation. *Applied Magnetic Resonance*, 33(4):419, 2008.
- [248] Pim Pullens, Alard Roebroek, and Rainer Goebel. Ground truth hardware phantoms for validation of diffusion-weighted mri applications. *Journal of Magnetic Resonance Imaging*, 32(2):482–488, 2010.
- [249] Nathan Yanasak and Jerry Allison. Use of capillaries in the construction of an mri phantom for the assessment of diffusion tensor imaging: demonstration of performance. *Magnetic resonance imaging*, 24(10):1349–1361, 2006.

- [250] Michal E Komlosh, Evren Özarslan, Martin J Lizak, Ferenc Horkay, Vincent Schram, Noam Shemesh, Yoram Cohen, and Peter J Basser. Pore diameter mapping using double pulsed-field gradient mri and its validation using a novel glass capillary array phantom. *Journal of magnetic resonance*, 208(1):128–135, 2011.
- [251] Darya Morozov, Leah Bar, Nir Sochen, and Yoram Cohen. Modeling of the diffusion mr signal in calibrated model systems and nerves. *NMR in Biomedicine*, 26(12):1787–1795, 2013.
- [252] Damien J McHugh, Feng-Lei Zhou, Ian Wimpenny, Gowsihan Poologasundarampillai, Josephine H Naish, Penny L Hubbard Cristinacce, and Geoffrey JM Parker. A biomimetic tumor tissue phantom for validating diffusion-weighted mri measurements. *Magnetic resonance in medicine*, 80(1):147–158, 2018.
- [253] Jonathan I Epstein, William C Allsbrook Jr, Mahul B Amin, Lars L Egevad, ISUP Grading Committee, et al. The 2005 international society of urological pathology (isup) consensus conference on gleason grading of prostatic carcinoma. *The American journal of surgical pathology*, 29(9):1228–1242, 2005.
- [254] Guidance on the use of temozolomide for the treatment of recurrent malignant glioma (brain cancer) — Guidance and guidelines — NICE.
- [255] Charles S Sherrington. Experiments in examination of the peripheral distribution of the fibres of the posterior roots of some spinal nerves. part ii. *Philosophical Transactions of the Royal Society of London. Series B, Containing Papers of a Biological Character*, 190:45–186, 1898.
- [256] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012.
- [257] Russell Reed and Robert J MarksII. *Neural smithing: supervised learning in feed-forward artificial neural networks*. Mit Press, 1999.

Perfusion-Fixation Protocol

Materials:

- 0.07ml sodium pentobarbital (Euthatal, obtained from Merial UK) diluted with heparinized NaCL up to 2ml.
- 0.9% saline solution.
- 4% PFA (Paraformaldehyde).
- Perfusion Pump.
- Heparin.
- 27G butterfly needle.
- Scissors, 2 clamps, round-ended tweezers.
- 50ml falcon tube.

Methods:

- Terminally anaesthetise the animal via IP (intraperitoneal) injection of Euthatal.
- Check for depth of anaesthesia using paw-pinch reaction, and secure animal in prone position on collection tray.
- Open the chest with a V-shaped incision using scissors.
- Use scissors to cut ribs, and clamp xiphoid process and ribs back.
- Carefully puncture diaphragm, will cause lungs to collapse.

- Manoeuvre beating heart into position to enable easy access to left ventricle.
- Start perfusion pump and insert butterfly needle into left ventricle (2mm penetration to prevent perfusion of pulmonary circuit). Snip right atrium to allow blood to drain.
- Perfuse with around 25ml saline ($37^{\circ}C$) at a rate of 3ml/min. Pour distilled water on torso to remove blood and help prevent clotting.
- Once the circulatory system has cleared, perfuse with 30ml PFA at 3ml/min.