


Spatio-temporal Image Analysis with Application to Cancer



Monica Enescu
Christ Church

Supervised by
Professor Julia Schnabel
Professor Michael Chappell

This thesis is submitted to the Department of Engineering Science, University of Oxford, in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Declaration

I declare that this thesis is entirely my own work, and except where otherwise stated, describes my own research.

Monica Enescu
Christ Church
University of Oxford

Abstract

Cancer is one of the main causes of premature death worldwide, leading to almost one in three deaths in the UK. Medical imaging plays a fundamental role in the diagnosis and treatment of cancer patients. This thesis focuses on the assessment of early stage clinical trials of novel cancer treatment, which is performed based on dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) and perfusion computed tomography (pCT). From a clinical viewpoint, the objectives are tumour microenvironment characterisation, which involves the assessment of tumour perfusion and vascularity, and early prediction of response to treatment. From a methodological viewpoint, this thesis is at the interface between pharmacokinetic (PK) modelling and image registration for dynamic images. Major challenges addressed in this work are PK model selection and arterial input function derivation, as well as intra-sequence motion correction for dynamic imaging data. Registration of dynamic imaging data is a particularly challenging problem due to the variety of motion types that appear in these data: bulk patient motion, periodic motion due to breathing, and random motion of small features due to peristalsis. Moreover, contrast-enhanced modalities such as DCE-MRI and pCT pose additional challenges as image intensities change due to contrast inflow. The key contributions of this thesis are the derivation of patient-specific arterial input functions from pCT data, which are employed for performing PK modelling of contemporaneous DCE-MRI scans of the same patients, and the development of an MRF-based discrete optimisation framework for the nonrigid registration of dynamic sequences. While the former was applied on DCE-MRI and pCT images from an early stage trial of new treatment for colorectal cancer, the latter was employed for the motion correction of DCE-MRI images from the same trial. Additionally, this new registration framework was applied to dynamic CT images of the lung to demonstrate the advantage of temporal regularisation in an application where there is periodic motion (i.e. breathing). Experimental validation shows that the proposed registration framework improves over state-of-the-art methods based on continuous optimisation in terms of registration accuracy and computational complexity.

Acknowledgements

First and foremost, I would like to thank my supervisors - Professor Julia Schnabel and Professor Michael Chappell - for their constant guidance and tremendous support. Without Julia, the Oxford DPhil adventure would not have started in the first place. Her support was invaluable in many ways: creating all the necessary conditions for work to go smoothly, giving me the freedom to pursue my research interests, encouraging me to participate at conferences and to publicise my work. She has been an inspiration both professionally and as a human being. I am very grateful to Michael for bringing in his expertise in modelling, and for always asking the right questions, which helped me gain a better understanding of the assumptions behind models and how this is linked with reality. His advice has been extremely helpful in presenting my written work in a clearer way.

I am also thankful to the EPSRC and to Christ Church, for providing the funding that made this DPhil possible, and to Julia for making sure that I never lacked any material or logistical support. I thank our wonderful clinical collaborators at the Churchill hospital - Ricky Sharma, Esme Hill, Jamie Franklin - and our collaborators from the Preclinical Imaging Group - Sean Smart, Veerle Kersemans, Danny Allen- who provided the data and worked closely with us to advance what is known about cancer.

I was lucky to have wonderful colleagues. Firstly, I thank Mattias Heinrich for his great advice with respect to methodological aspects of my thesis. I also thank Manav Bhushan, Amalia Cifor, Ben Irving, and Jesper Kallehauge for many engaging and useful discussions, and for making the lab an enjoyable place. I thank Bartek for his sharp comments, and Thais, Andre and Liam for keeping me company while writing the thesis.

Finally, I thank my family and friends for their support and love, and for being patient when I disappeared in the lab for prolonged periods. Above all, I thank Andrei for everything: joining me in Oxford although he loves the sea, trying to make me have a good work-life balance, but nevertheless accepting that I bring work when we go on holidays, occasionally chaperoning me at conferences, and making daily life so much more fun.

Contents

1	Introduction	1
1.1	The Clinical Problem	1
1.1.1	Colorectal cancer	1
1.1.2	The biology of cancer	3
1.2	Imaging Challenges	4
1.3	Scientific Contributions and Thesis Outline	5
2	Functional Imaging of Cancer	8
2.1	Magnetic Resonance Imaging	9
2.1.1	Basic Principles of MRI	9
2.1.2	Dynamic contrast-enhanced MRI	12
2.2	Computed Tomography	16
2.2.1	Basic Principles of CT	16
2.2.2	Perfusion Computed Tomography	16
2.3	Other Modalities	18
2.4	Modelling of Contrast Agent Uptake	19
2.4.1	Two compartment exchange model (2CX)	20
2.4.2	Adiabatic approximation to tissue homogeneity model (AATH)	21
2.4.3	Distributed parameter model (DP)	22
2.4.4	Common approximations	24
2.4.5	The Arterial Input Function	26
2.5	Discussion	27
3	Registration of Dynamic Images	29
3.1	General Principles of Image Registration	30
3.1.1	Transformation models	31
3.1.2	Similarity measures	35
3.1.3	Optimisation strategies	37

3.1.4	Validation	40
3.2	Registration of Dynamic Image Sequences	42
3.2.1	Respiratory motion modelling from dynamic images	43
3.2.2	Cardiac image registration	44
3.2.3	Dynamic contrast-enhanced image registration	46
4	Personalised Arterial Input Functions	48
4.1	Materials	49
4.1.1	Sonatina clinical trial	49
4.1.2	Image acquisition	50
4.2	Methods	51
4.2.1	Pharmacokinetic modelling	51
4.2.2	Intra-sequence motion correction	53
4.2.3	Derivation of personalised AIFs	54
4.2.4	Pre-processing	56
4.3	Results	57
4.3.1	Model comparison	57
4.3.2	Patient-specific AIFs	59
4.4	Discussion	63
5	An MRF-based Discrete Optimisation Framework for Deformable Registration of Dynamic Images	66
5.1	Introduction to Markov Random Fields	67
5.2	Registration as an MRF labelling problem	69
5.3	Transformation model: Parameterisation and properties	71
5.4	Similarity metric calculation	72
5.5	Inference strategies	72
5.5.1	Tree-reweighed message passing on the full 4D graph	74
5.5.2	Message passing on the reduced 4D graph	75
5.5.3	Diffusion regularisation	77
5.6	Experiments on simulated data	79
5.6.1	Experiments	79
5.6.2	Results	81
5.7	Discussion	83

6	Application to Cancer	85
6.1	DCE-MRI pharmacokinetic estimation and intra-sequence registration for colorectal cancer	85
6.1.1	Experiments	86
6.1.2	Results	88
6.2	Registration of lung 4D CT images	91
6.2.1	Experiments	92
6.2.2	Results	95
6.3	Discussion	100
7	Outlook	102
7.1	Extensions of the DireP framework	102
7.2	Joint modelling of pCT and DCE-MRI	104
7.2.1	Preliminary results	105
7.2.2	Future work	108
7.3	Tumour growth modelling from longitudinal DCE-MRI scans	108
7.3.1	Preliminary results	109
7.3.2	Future work	114
8	Conclusions	118
A	List of publications	122
B	List of abbreviations	124
	Bibliography	127

List of Figures

2.1	DCE-MRI image before and after contrast arrival	13
2.2	pCT image before and after contrast arrival	17
2.3	The 2CX Model	22
2.4	The AATH Model	23
2.5	The DP Model	23
4.1	An example of DCE-MRI derived AIF vs. the pCT derived AIF for Patient 3. $AIF_{DCE-MRI}$ misses the peak of the uptake and has a much smaller height than AIF_{pCT}	56
4.2	T2-weighted MRI to DCE-MRI registration result	57
4.3	Model comparison in terms of tracer kinetic maps	58
4.4	AIC_{min} maps of all patients	60
4.5	Intra-patient variability of the AIFs vs. variability with respect to Orton	61
4.6	AIF variability in terms of peak enhancement and area under the curve	62
4.7	Correlation between PK parameters obtained with population AIF and those obtained with the individual AIFs	63
4.8	K^{trans} maps and histograms before and after treatment with Nelfinavir for a Responder.	64
4.9	K^{trans} maps and histograms before and after treatment with Nelfinavir for a Non-responder.	64
5.1	The Discrete registration and Pharmacokinetic estimation framework (DireP).	70
5.2	TRW-S message passing schedule	75
5.3	The MST-based reduced 4D graph	76
5.4	Lower envelope computation for efficient computation of diffusion regularisation	79
5.5	Voxel plots before and after motion correction	82
5.6	TRE as a function of the motion level	83
6.1	Minimum spanning tree for a colorectal image	87

6.2	Registration result for the Sonatina dataset: Time-cuts	89
6.3	Expert grading of pre-treatment DCE-MRI images before and after registration	90
6.4	Expert grading of mid-treatment DCE-MRI images before and after registration	91
6.5	Expert grading of post-treatment DCE-MRI images before and after registration	92
6.6	Residual error before and after registration for Patient 3 Mid	93
6.7	K^{trans} maps before and after registration with DireP	94
6.8	Minimum spanning tree for a lung image	95
6.9	Registration result for Case 2 of the DIR-lab dataset	97
6.10	Registration result for Case 8 of the DIR-lab dataset	98
6.11	Registration result for 4DCT images of the lung: Time-cuts	99
6.12	Irregularity of landmark trajectories for DIR-lab dataset	100
7.1	New MST-based 4D graphical structure	103
7.2	Ktrans maps for (a)Individual dceMRI estimation (b) Individual pCT estimation (c)Joint modelling	107
7.3	Preclinical pharmacokinetic parameter maps together with corresponding DCE-MRI baseline images.	111
7.4	The correlation between the K^{trans} map and tumour volume change	112

List of Tables

4.1	DCE-MRI median parameter values for the tumour volume for each of the Tofts and 2CTU models. The AIC values are also listed.	59
5.1	Registration results on synthetic data	81
6.1	Motion scoring scale	88
6.2	Residual fitting error for the Sonatina dataset.	90
6.3	Registration results on the DIR-lab dataset. The average target registration error (TRE) between T00-T50 for Cases 1-10 is reported.	96
6.4	Registration results on the DIR-lab dataset. The average target registration error (TRE) for the 75 landmarks in all the volumes for Cases 1-10 is reported.	96
7.1	DICE overlap % for the tumour area	113

Chapter 1

Introduction

1.1 The Clinical Problem

1.1.1 Colorectal cancer

Cancer causes almost one in three deaths in the UK, with 161,823 deaths recorded in 2012 alone. At a global scale, 14.1 million new cases of cancer and 8.2 million deaths were recorded in the same year. Colorectal cancer (CRC) makes up a sizeable part of the problem, being the third most common type of cancer in both men and women after lung and gender specific cancers, and also the third most common cause of cancer related deaths worldwide, with 1.41 million new cases diagnosed and 738,000 deaths recorded in 2012 [1]. Moreover, the incidence is expected to increase to 2.2 million cases over the next two decades, and most of this increase is bound to happen in the developing world, due to demographic trends and changes in lifestyle in conjunction with economic growth and development [75]. Managing this form of cancer is therefore a key healthcare challenge.

As with other types of cancer, primary prevention through lifestyle adjustments, screening, and early diagnosis are amongst the most effective options to improve CRC management. In the case of CRC, screening programs involve any of fecal occult blood tests, colonoscopy, or sigmoidoscopy [14]. To date, the span of such programs is limited to a number of countries in the developed world, enabling only a small number of people

to access these procedures. Once the cancer has developed, a treatment combination is chosen based on tumour staging, which is performed using imaging such as magnetic resonance imaging (MRI), computed tomography (CT), or endorectal ultrasound (ERUS) [4]. These modalities are explained in further detail in Chapter 2. After tumour staging, patients can be classified into four categories: those with superficial tumours, who can be treated with surgery alone; those with operable tumours and a wide circumferential resection margin, who can be treated with a short course of neoadjuvant radiation therapy followed by total mesorectal excision; those with advanced cancer and a close or involved resection margin, who require a long course of chemo-radiotherapy (CRT), and extensive surgery [10], and those with inoperable cancer who qualify for palliative treatment. Surgery remains the principal treatment for colorectal cancer, and 80% of all the patients undergo this procedure [1]. After surgery, the excised tumour tissue is assessed histologically, to determine the nature and extent of the cancer.

However, both CRT and surgery dramatically alter the patient's short and long term quality of life, and involve pain, suffering, and pose a high financial burden on the health-care system. Furthermore, it has been shown that about 25% of the patients who undergo surgery have had a pathological complete response to neoadjuvant CRT, meaning that the surgery could have been avoided [111]. This would be particularly important for patients who undergo sphincter removal and/or a colostomy procedure [1]. At the same time, some patients do not respond to CRT at all. For this subset, at least a few cycles of CRT could be avoided and the patients could be sent straight to surgery. Novel treatment approaches are currently being explored, such as the replacement of standard CRT with short course radiotherapy preceded by drug administration [67]. The main clinical challenge remains to identify individual patient response as early on treatment as possible, and to tailor further therapy accordingly.

Because modifications at a cellular level precede macroscopic tumour growth, changes in the tumour microenvironment during the course of therapy can act as early indicators of treatment outcome. The following section contains an overview of the key biological

processes that shape the tumour microenvironment. These biological processes can be measured non-invasively using medical imaging. Parameters of interest which indicate the underlying tumour biology are extracted from the imaging data using automated image analysis methods. Section 1.2 introduces the principal image analysis challenges related to image biomarker extraction. The proposed contributions which come to address these challenges and the thesis outline are presented in Section 1.3.

1.1.2 The biology of cancer

Cancer is a family of diseases that involve abnormal cell growth, where these cells are able to move and spread to other parts of the body. To date, it is known that cancer is a genetic disorder, caused by successive mutations in a cell's DNA. These mutations enable the cell to divide incessantly, to escape programmed cell death (apoptosis), to become insensitive to growth regulating signals, to induce blood vessel growth (angiogenesis), and to move and invade other tissues (metastasis) [57], [58]. Tumour growth relies on the creation of an aberrant microenvironment around tumour cells. The key processes that shape the tumour microenvironment - angiogenesis, hypoxia, perfusion and metabolism - are further discussed in this section.

As tumour cells proliferate and their density in tissue increases, there is a rise in local demand for oxygen and nutrients. The insufficient amounts of oxygen and glucose are suspected to trigger the secretion of growth factors (i.e. VEGF), which induce the formation of new blood vessels designed to feed the tumour. More specifically, activated endothelial cells migrate from the capillaries into the extravascular extracellular space where they proliferate. This process is termed angiogenesis, and its result is a network of chaotic, fragile and leaky vessels.

The faulty vascular network resulting from tumour angiogenesis leads to a mismatch between the oxygen demand and the delivered supply. This gives rise to regions of mild to severe oxygen deprivation, also known as hypoxia. Tumour hypoxia was shown to correlate with resistance to radiotherapy and chemotherapy [51]. Moreover, the lack of

oxygen leads to enhanced metastatic behaviour [90].

Tumour tissue perfusion governs the amount of nutrients that reach tumour cells, being thus directly related to tumour growth and proliferation. Among the most prominent factors that influence tumour perfusion we note angiogenesis and the extent of the extracellular space. Poor perfusion is also associated with resistance to chemoradiotherapy [66].

As noted by Warburg [149] more than 50 years ago, normal cells rely on the mitochondria for 90% of their energy supply, and use glucose metabolism (glycolysis) for only 10% of their adenosine triphosphate (ATP) production. At the same time, tumour cells use glycolysis for over 50% of their ATP production. This involves that tumour cell metabolism has much higher glucose demand than normal cell metabolism, allowing cells to survive longer in the absence of oxygen.

1.2 Imaging Challenges

Medical imaging is increasingly used for supervising early stage clinical studies of cancer therapies, as it provides an invaluable tool for distinguishing between patients who benefit from the treatment procedures and those who do not. Therapy response is typically assessed using functional imaging modalities such as pCT, DCE-MRI, or dynamic positron emission tomography (dPET) [47]. PET imaging can be used in combination with different radiotracers, according to the physiological property that is being imaged. For instance, FDG-PET is used for measuring the glucose metabolism which is expected to be higher in the perfused parts of tumour tissue and FMISO-PET is currently investigated as a tool for measuring tumour hypoxia. However, PET comes with the disadvantage of high costs, requires the administration of a radioactive tracer, and has poor spatial resolution (3-5mm).

DCE-MRI and pCT are both used for imaging tumour perfusion and permeability, which are indicators of intra-tumour angiogenic activity [97]. These modalities are based

on the injection of a contrast agent and repeatedly imaging the region of interest, over a time course of up to 7 minutes. While pCT has an excellent temporal resolution, lower costs, and is easier to analyse due to the linear relationship between signal enhancement and contrast agent concentration, DCE-MRI has the advantage of involving non-ionising radiation and has a better signal-to-noise ratio.

To extract clinically relevant biomarkers from DCE-MRI or pCT, contrast agent uptake patterns have to be carefully modelled. The modelling process entails quantifying tracer in- and out-flow from the tissue, which is usually done using a PK model [136,141]. PK models are fitted to the data on a voxel-by-voxel basis, yielding parameter maps which are directly related to the underlying tissue physiology. Tracer kinetic modelling is a challenging process, which involves adequate model selection and a strategy to deal with patient and organ motion.

Intra-sequence motion may introduce significant errors to the per-voxel tracer kinetic model fitting, as anatomical features of interest might move to different voxel locations in subsequent volumes. Motion correction for dynamic sequence is typically performed using automated image registration algorithms. Nevertheless, image registration for dynamic data is not a trivial problem, as the sequences can exhibit a variety of motion types: bulk patient motion, periodic motion caused by breathing, and random movements such as peristalsis. Furthermore, additional challenges are posed by contrast inflow, which makes image intensities change with time.

The analysis is further complicated by scanning noise, and in the case of DCE-MRI, by the non-linear relationship between signal enhancement and contrast agent concentration.

1.3 Scientific Contributions and Thesis Outline

The first part of this thesis reviews fundamental concepts in medical image acquisition and image analysis upon which the proposed contributions are based. Chapter 2 describes the common imaging modalities that are used for cancer therapy planning, with a focus

on DCE-MRI and pCT, which are the modalities used in this work. We then review the principal approaches for tracer kinetic modelling, discussing their assumptions and limitations. In Chapter 3, the fundamental principles of image registration are presented and state-of-the art literature concerning the registration of dynamic image sequences is reviewed.

In the second part of the thesis, we discuss a number of contributions which are at the interface between PK modelling and image registration. In Chapter 4, the effect of different design choices in the process of PK modelling is analysed. We propose the derivation of patient-specific arterial input functions (AIFs) from pCT data, and using these for the analysis of contemporaneous DCE-MRI images of the same patients. The issue of selecting the optimal tracer kinetic model for capturing the biological processes of interest is also investigated. The Tofts model, which is widely used in clinical practice, is compared to a second generation model, the two compartment uptake model (2CTU). These methods were tested in the context of a novel cancer drug trial. We attempt to evaluate whether the administered drug, Nelfinavir, has the potential to increase blood flow, acting as a radiosensitiser. To this end, tumour perfusion in patients with rectal cancer is measured using pre- and post-treatment DCE-MRI.

The problem of motion correction for dynamic imaging sequences is investigated in Chapter 5. To address the challenges related to this problem (Section 1.2), we propose a novel framework for nonrigid registration of 4D dynamic sequences. This framework is formulated on a Markov Random Field (MRF), and can be applied across imaging modalities, being suitable for both contrast-enhanced and non contrast-enhanced data. We also formulate a DCE-MRI tailored similarity metric which incorporates PK modelling information, to address the problem of contrast inflow.

Chapter 6 presents two clinical applications of the proposed nonrigid registration framework: the intra-sequence motion correction of DCE-MRI images from colorectal cancer patients and the registration of dynamic CT lung images for respiratory motion modelling. The former is performed together with PK parameter estimation on in-house

data from an early stage clinical trial of novel cancer drugs. The latter is performed on a publicly available dataset (DIR-Lab [25]), which comprises 10 4DCT images of lung cancer that have been acquired for radiotherapy planning.

Chapter 7 provides an outlook of this thesis. Three different projects are identified as potential subsequent areas of research. To establish the potential of these research avenues, initial investigations were also performed. First, we propose improvements to the framework presented in Chapter 5, and discuss its application to the registration and PK modelling of DCE images of the liver. Second, initial work towards the joint modelling of DCE-MRI and pCT is presented. As these modalities image the same biological quantities of interest, we investigate whether the joint modelling of both modalities yields improved parameter maps compared to the individual modelling of each modality. Third, modelling tumour growth from preclinical longitudinal DCE-MRI is performed using image registration. The end goal of this work is quantifying tumour progression and/or subject response to treatment. Chapter 8 concludes this thesis.

Chapter 2

Functional Imaging of Cancer

Imaging plays a fundamental role in the care of cancer patients. As previously explained, an optimal treatment combination is chosen using image-based tumour staging. Furthermore, changes in tissue physiology as a result of treatment can be detected using image derived biomarkers. Depending on the biological quantity to be measured, a range of imaging modalities can be used. While diffusion-weighted magnetic resonance imaging (DW-MRI) can be used to measure cell structure and density, contrast-enhanced modalities such as pCT and DCE-MRI measure tumour perfusion, permeability and angiogenesis. Metabolism can be assessed using dPET, and hypoxia can be quantified using FMISO-PET .

This chapter sets out to describe imaging modalities that are used for cancer therapy planning, with an emphasis on DCE-MRI and pCT which are the modalities used in this work. While the underlying principles of MRI and DCE-MRI are described in Section 2.1, Section 2.2 is dedicated to CT and pCT acquisition. Other cancer imaging modalities are also briefly described (Section 2.3). In Section 2.4, we review tracer kinetic modelling approaches, discussing their assumptions and limitations.

2.1 Magnetic Resonance Imaging

2.1.1 Basic Principles of MRI

Since it was first invented in the 1970s, magnetic resonance imaging has greatly improved, becoming an invaluable tool in clinical practice, for both diagnosis and treatment planning purposes. MRI does not involve ionising radiation, and is particularly useful in oncology due to the good soft tissue contrast which enables differentiation between tumours and normal tissue. While the principles of MRI are briefly described in the following part, an excellent comprehensive description can be found in [60]. Although MRI could be used to scan any element with paramagnetic properties, clinical protocols rely on imaging hydrogen (1H) protons which are most abundant in the human body, as more than 60% of it is water.

Nuclear magnetic resonance relies on the fact that a spinning charged particle (such as 1H) creates an electromagnetic field. From quantum theory we know that these charged particles have specific energy states, whose number depends on a property called the spin quantum number (S). Each 1H proton has two energy states which are aligned in opposite directions, denoted $\frac{1}{2}$, 'spin up', and $-\frac{1}{2}$, 'spin down'. Each proton is spinning about its axis creating a micro-magnetic field. At a macro level, because of the random orientation of individual 1H protons, the micro-magnetic fields cancel each other out yielding a null net magnetisation.

If the protons are placed in an external magnetic field \mathbf{B}_0 , they will align with that field, approximately half pointing in one direction and the rest pointing in the opposite direction. As there are slightly more spins pointing in the direction of \mathbf{B}_0 , the resulting net magnetisation \mathbf{M}_0 (also called longitudinal magnetisation) will point in the direction of \mathbf{B}_0 . The protons also precess about the axis of the external magnetic field (usually the z-axis) at a rate which is given by the Larmor frequency $\omega_0 = \gamma B_0$. Here, γ denotes the gyromagnetic ratio. However, the net magnetisation does not precess, as the individual spins are out of phase with each other. This means there is no net magnetisation in the

x-y plane.

To create an image, radio frequency (RF) pulses at the Larmor frequency are transmitted into the object to be imaged. These pulses flip the longitudinal magnetisation away from the z-axis, by causing the protons to precess in phase about \mathbf{B}_0 , as well as making them to precess about the axis of B_1 , the RF-induced magnetic field. The double precession causes a spiral motion of the net magnetisation from the z-axis into the x-y plane. After the RF pulse is switched off, the RF energy is slowly released in a process called relaxation, and the spins gradually return to alignment with \mathbf{B}_0 . There are two relaxation processes: (1) spin-lattice relaxation, which refers to the return of the net magnetisation along the z-axis, and (2) spin-spin relaxation, which refers to spin dephasing. The signal that is generated during relaxation is recorded by the RF receive coils.

As the received signal has no spatial information, gradient coils are employed for spatial encoding. These gradient coils cause small linear perturbations in the magnetic field, encoding a spatial location into the received signal. There are three types of gradient coils: slice-select G_z , frequency-encoding G_x , and phase-encoding G_y . While G_z encodes slice position and thickness, G_x and G_y encode the spatial locations within a slice.

The spatially encoded signals are stored in k-space, a frequency representation of the image content. While the central part of the k-space contains low frequency information corresponding to major image structures, the outer part of the k-space contains high frequency information corresponding to edges and details. The actual image can be recovered from k-space by computing the inverse Fourier transform.

In the following, several parameters that influence MR image acquisition are discussed, and the MRI signal formation equations are formulated. As briefly described above, to create an MR image, an RF pulse is used to excite the spins, causing them to precess about the external magnetic field. The change in angle of the net magnetisation M_0 caused by the application of the RF pulse is known as the **flip angle** α . After the RF pulse is switched off, the relaxation of spins from a high energy state to a low energy state causes the net magnetisation to return to the direction of \mathbf{B}_0 . This relaxation is an

exponential process which depends on the tissue type. T_1 , the relaxation time constant of the exponential, represents the time needed for 63% of the original magnetisation M_z to be recovered as spins release the energy obtained from the RF pulse back into the surrounding lattice. This constant is also called the **spin-lattice relaxation time**, or the longitudinal relaxation time. For a flip angle $\alpha = 90^\circ$, the spin-lattice relaxation is described by Eq. 2.1.

$$M_z(t) = M_0(1 - e^{-t/T_1}) \quad (2.1)$$

Here, t represents time elapsed and M_0 represents the net magnetisation.

The **spin-spin relaxation** occurs when spins in high and low energy states exchange energy, without losing energy to the surrounding lattice. This represents a loss in transverse magnetisation M_{xy} , which decays exponentially. For a flip angle $\alpha = 90^\circ$, the spin-spin relaxation is described in Eq. 2.2.

$$M_{xy}(t) = M_0 e^{-t/T_2} \quad (2.2)$$

Here, t and M_0 are the same as above. This decay is governed by the T_2 relaxation time constant, which is a measure of the dephasing of the spins. Magnetic field inhomogeneity speeds up the dephasing process at a rate described by T_2^* . Both T_1 and T_2 are inherent properties of tissues, and are therefore fixed for a specific tissue at a given magnetic field strength. T_2 decay occurs 5-10 times faster than T_1 recovery.

A variety of MR image types can be acquired by changing the parameters of the sequence of RF pulses. Unlike T_1 and T_2 , which are tissue parameters, the sequence parameters can be controlled and adjusted by the operator. Apart from the flip angle α , two other essential parameters govern the RF pulse: the **repetition time, TR**, and the **echo time, TE**. To create an image, a sequence of RF pulses needs to be applied. The time interval between successive RF pulses is called the repetition time (TR). The time elapsed between the application of an RF pulse and the measurement of the transverse magnetic field is called the echo time, TE. By combining the effect of all parameters

described above, the MR signal can be written as:

$$S = M_0(1 - e^{-TR/T_1})e^{-TE/T_2} \quad (2.3)$$

The parameters TR and TE can be chosen as to focus on either the measurement of T_1 or the measurement of T_2 . While employing a short TR enhances the T_1 effect, producing what is called a T_1 -weighted image, employing a long TE enhances the T_2 effect, producing a T_2 -weighted image.

The MRI clinical data used in this thesis were acquired using a spoiled gradient echo (SPGR) sequence. This type of sequence involves destroying the transverse magnetisation immediately after recording the measurement at time TE , to ensure that there is no residual magnetisation when the new RF pulse is applied. This is performed using a spoiled gradient echo pulse. The SPGR signal model is described in Eq. 2.4.

$$S = M_0 e^{-TE/T_2^*} \frac{\sin(\alpha)(1 - e^{-TR/T_1})}{1 - \cos(\alpha)e^{-TR/T_1}} \quad (2.4)$$

In the following section, we explain how functional information regarding the tissue of interest can be acquired using an MRI scanner.

2.1.2 Dynamic contrast-enhanced MRI

DCE-MRI enables measurement of tissue physiology and function, alongside structural information. This modality is based on the injection of a paramagnetic contrast agent followed by repeated acquisitions of MRI volumes containing the region of interest. The uptake and wash-out patterns of the contrast agent provide information about the underlying tissue perfusion and vascularity.

The contrast agent acts by shortening the T_1 and T_2 relaxation times of hydrogen atoms in its vicinity, and this effect can be observed as a change in MRI image intensities. More specifically, in regions where contrast is present, the image intensities appear




Figure 2.1: A slice from a DCE-MRI volume before the injection of contrast (left) and after contrast arrival (right). The brighter parts of the right-hand side image indicate the presence of contrast agent.

brighter (Fig. 2.1). The effect of contrast on the tissue relaxation times is described in Eqs. 2.5 and 2.6.

$$\frac{1}{T_1} = \frac{1}{T_{10}} + r_1 C(t) \quad (2.5)$$

$$\frac{1}{T_2^*} = \frac{1}{T_{20}^*} + r_2 C(t) \quad (2.6)$$

Here, T_{10} and T_{20}^* are the relaxation times T_1 and T_2^* prior to the injection of contrast agent, r_1 and r_2 are relaxivity constants that are specific to the agent, and $C(t)$ is the total concentration of contrast agent at time t .

MRI contrast agents approved for clinical use are Gadolinium based, have a low molecular weight, and are soluble in water. These properties enable the contrast molecules to travel through vessels and capillaries, and to diffuse rapidly into the extravascular extracellular space (EES), enabling us to image tissue vasculature and perfusion. The images in this thesis were acquired using ProHanceTM, one of the most widely used contrast agents. This agent was injected intravenously at a rate of 0.1 mmol/kg body weight.

It is most common to assume a non-linear relationship between MRI signal and contrast agent concentration [143]. By substituting Eqs. 2.5 and 2.6 into the SPGR signal model given by Eq. 2.4, we obtain the expression for the DCE-MRI signal as a function

of contrast agent concentration:

$$S(t) = M_0(t)\sin(\alpha)\frac{1 - e^{-TR/T_{10}-r_1C(t)TR}}{1 - \cos(\alpha)e^{-TR/T_{10}-r_1C(t)TR}} \quad (2.7)$$

where $M_0(t)$ is the equilibrium longitudinal magnetisation given by $g\rho e^{-TE/T_{20}^*-r_2C(t)TE}$, g being the scanner gain and ρ being the proton density. To extract the contrast agent uptake curve, the signal enhancement curve needs to be calculated at every voxel. The signal enhancement $E(t)$ is defined as the ratio between the signal at time t , $S(t)$, and the signal at baseline, $S(0)$: $E(t) = \frac{S(t)}{S(0)}$. Thus, we obtain a non-linear relationship between signal enhancement and contrast agent concentration [5]:

$$E(t) = e^{-r_2C(t)TE}\frac{1 - e^{-P-Q} - \cos(\alpha)(e^{-P} - e^{-P-2Q})}{1 - e^{-P} - \cos(\alpha)(e^{-P-Q} - e^{-P-2Q})} \quad (2.8)$$

where $P = TR/T_{10}$ and $Q = r_1C(t)TR$. Based on Eq. 2.8, the contrast agent concentration can be extracted from the observed signal enhancement, provided that r_1 , r_2 , T_{10} and α , TR and TE are known. Whilst most of these are either acquisition parameters or contrast agent constants, T_{10} , the unenhanced relaxation time of the tissue can only be known with additional measurements. This parameter is typically calculated by the variable flip angle method [38].

Recently, Medved et al. [98] proposed a different approach for extracting contrast agent concentration from the image signal. They compare the signal S in the tissue of interest (tumour) to the signal S_{ref} in a normal reference tissue with known relaxation time T_{1ref} . This yields a simpler, linear relationship between signal intensity $S(t)$ and contrast agent concentration $C(t)$:

$$S(t) = S(0) + rT_{1ref}S_{0ref}C(t) \quad (2.9)$$

Here, $S(0)$ and S_{0ref} are the pre-contrast signals in the tissue of interest and reference tissue, respectively. T_{1ref} is the initial T_1 relaxation time for the reference tissue and its

values are taken from literature.

DCE-MRI has proven to be particularly useful for cancer imaging, as it can discriminate microvasculature changes that occur in tumour tissue. Namely, the many leaky vessels which result from angiogenesis enable a higher and earlier contrast uptake for tumour tissue than for normal tissue. Thus, contrast agent enhancement patterns reflect changes in tumour perfusion, permeability and intra-tumour angiogenic activity [97]. DCE-MRI has been employed in a variety of settings, for differentiating between malignant and benign lesions [108], for non-invasive tumour grading [10], for predicting clinical outcome after chemo-radiotherapy [159], and for testing the effect of drugs in clinical trials of novel therapies [109].

Microvasculature changes can be inferred from the image signal by computing either semi-quantitative or quantitative measures. Semi-quantitative parameters describe signal intensity enhancement and are computed directly on the signal enhancement - time curve. These parameters include bolus arrival time, peak enhancement, area under the curve, time to peak [117], and have the advantage that they are straightforward to compute. However, they do not accurately reflect contrast uptake, are highly dependent on the MRI acquisition protocol, and do not have physiological meaning [114].

Quantitative techniques address these limitations, but they are more difficult to apply. These techniques involve modelling contrast agent uptake kinetics to reveal the underlying tissue physiology. After the contrast agent concentration is estimated from the MRI signal enhancement - time curve (using i.e. Eq. 2.8 or Eq. 2.9), the underlying tissue physiology is recovered by fitting a pharmacokinetic model (detailed in Section 2.4) to the obtained contrast agent uptake curve and computing the model parameters.

2.2 Computed Tomography

2.2.1 Basic Principles of CT

Although the first radiographic image was taken at the end of the 19th century, it was only until the 1970s that 3D X-ray imaging of the human body became possible with the advent of computed tomography (CT). This invention relied on the advances in computer science of the sixties, which led to the construction of the first CT scanner by G. Hounsfield in 1972. Similar to radiographic images, CT involves passing X-ray beams through the human body and recording their attenuation profile which depends on tissue density. However, while radiographs are 2D projection images of all the structures encountered by the beam from the X-ray tube to the detector, CT scans are 3D representations where each picture element represents an individual anatomical structure in the body. This is accomplished by measuring the patient from different directions, with the help of a rotating X-ray tube and detector. The different X-ray projections obtained thus are used for computing the superposition-free images, using a reconstruction algorithm such as filtered backprojection [73]. CT image intensities are measured in Hounsfield Units (H.U.), and are independent of scanner manufacturer or image sequence.

Compared to other modalities, CT scans are very fast, have an excellent dense tissue contrast, a high spatio-temporal resolution, and are low cost. However, these advantages come at the cost of exposing the patient to ionising radiation, which in high doses can cause DNA damage and induce cancer. Moreover, CT provides a modest soft tissue resolution compared to modalities such as MRI. To alleviate the issue of soft tissue contrast which is essential for oncological imaging, perfusion CT was proposed.

2.2.2 Perfusion Computed Tomography

pCT involves the injection of an iodine-based contrast agent, after which the volume of interest is imaged repeatedly over a couple of minutes. The resulting dynamic sequences have high temporal resolution, typically of one volume per second. Similarly to DCE-




Figure 2.2: A slice from a pCT volume before the injection of contrast (left) and after contrast arrival (right). In the right-hand side image, the contrast can be mostly seen in the arteries.

MRI, quantifying the inflow and outflow of contrast agent gives a handle on the underlying tissue haemodynamics. The most prominent applications of pCT are stroke and oncology. In oncological applications, pCT is used for measuring microvasculature changes that occur due to angiogenesis [102]. pCT is widely used in clinical practice due to its simplicity in terms of acquisition and analysis. More specifically, it is a low-cost, widely available modality, which features a linear relationship between contrast agent concentration and signal intensity. Moreover, reproducible data analysis can be performed with available commercial software tools (i.e. Philips CT Perfusion, GE CT Perfusion 4D). Nevertheless, pCT has some major drawbacks such as the high dose of ionising radiation to which patients are exposed [110], the limited field of view and the small enhancement (approx. 10 H.U.). A slice through a pCT volume before and after contrast enhancement is presented in Fig. 2.2.

Assuming that the contrast agent causes a signal intensity change proportional to its concentration, the relationship between contrast agent concentration and signal enhancement can be written as follows [104]:

$$C(t) = k(S(t) - S(0)) \quad (2.10)$$

Here, S_0 is the baseline signal intensity, and $S(t)$ is the signal intensity at time point t . The unknown proportionality constant k does not depend on tissue type.

To quantify tumour microvasculature changes, either semi-quantitative or quantitative analysis can be performed on signal intensity - time curves obtained from pCT. Typical semi-quantitative parameters are peak enhancement, standardised perfusion value, or area under the curve [110]. Quantitative analysis of pCT images can be performed by pharmacokinetic modelling. pCT pharmacokinetic modelling has the same theoretical foundation as modelling of DCE-MRI data, and will be described in the following section.

2.3 Other Modalities

DW-MRI measures the diffusion of water molecules in tissue in a non-invasive way. The rate of diffusion depends on the tissue microenvironment, and is quantified as the apparent diffusion coefficient (ADC). It has been shown that the ADC is inversely correlated to tissue cellularity [92].

PET is a functional imaging modality that enables in-vivo measurement of biological processes in the body. To this end, a radioactive tracer is injected in the patient and the gamma rays emitted due to tracer decays are recorded. Using computerised reconstruction, an image is formed. PET imaging is employed in combination with different radiotracers, according to the physiological property on which the imaging is based. If the chosen radiotracer is FDG, an analogue for glucose, PET imaging can be applied to image cancer metabolism, as glycolysis is expected to be higher in tumour tissue than in normal tissue.

Dynamic PET produces a full image sequence, yielding a spatio-temporal map of tracer concentration. This provides considerably more information about the underlying tissue biology. The per-voxel indicators of blood flow and metabolism can be extracted by means of tracer kinetic modelling.

FMISO, a newer PET tracer, has the potential to map tumour hypoxia. Several

studies have shown the correlation between FMISO accumulation in cells and oxygen deprivation [44, 128].

2.4 Modelling of Contrast Agent Uptake

Modelling of DCE data is a two stage process, composed of independent steps. First, contrast agent-time characteristics have to be derived from the signal enhancement-time curves, for each individual voxel. To this end, the signal model of the respective imaging modality needs to be employed. Second, tissue physiology is recovered by fitting a pharmacokinetic model to the contrast agent uptake curve in order to derive haemodynamic parameters.

All tracer kinetic quantification methods assume the tissue is a linear time invariant system [135]. This means that tissue response is proportional to the injected dose and is independent of the moment of injection. Thus, the contrast agent concentration in tissue, $C_t(t)$, can be described as the convolution between the concentration in the supplying blood vessel, $C_a(t)$, and a residue function $R(t)$ (Eq. 2.11):

$$C_t(t) = F_p R(t) * C_a(t) \quad (2.11)$$

Here, F_p is the plasma flow that carries contrast agent into the tissue. F_p , is defined as the volume of plasma which enters a unit of tissue volume per unit of time, and is measured in ml/min/100ml, or 1/min. The concentration of contrast agent in the supplying vessel is termed the arterial input function (AIF), and its knowledge is essential for accurate tracer kinetic modelling. The residue function, $R(t)$, is the tissue specific response and reflects the distribution of the transit times required for contrast molecules to pass through the tissue [19]. Thus, the mean transit time of particles is defined as the area under $R(t)$, $MTT = \int_0^\infty R(t)dt$.

Without making any further assumptions about the internal structure of the tissue, the parameters F_p and MTT can be derived using mathematic deconvolution. This

approach has gained a lot of popularity, and is implemented in commercial software tools such as GE CT PerfusionTM. Nevertheless, $R(t)$ encodes additional information that can be extracted by further modelling the internal tissue components.

Pharmacokinetic modelling makes assumptions about the different tissue components, such as the intra- and extra-vascular space, and their behaviour to contrast agent passage. By means of pharmacokinetic modelling, additional information about tissue vascularity, vessel permeability and cellularity can be derived, alongside the tissue perfusion given by F_p . Most tissues are composed of cells, interstitial fluid, and capillaries which supply cells with nutrients. We denote the extravascular extracellular volume as v_e , and the intra-vascular plasma volume as v_p . Blood plasma and the interstitial fluid exchange substances via the capillary walls, at a rate given by the permeability surface area product, PS (1/min). Because most contrast agents do not enter in cells, pharmacokinetic models describe only plasma and the EES and the interactions between them.

In the following part, several models that have been used for DCE-MRI and pCT pharmacokinetic analysis are presented, together with their underlying assumptions. Following the classification recently proposed by Sourbron et al. [71], we discuss the two compartment exchange model (2CX), the tissue homogeneity model (TH), the adiabatic approximation to the tissue homogeneity model (AATH) and the distributed parameter model (DP). The most popular particular cases of these general models are also presented.

2.4.1 Two compartment exchange model (2CX)

The two compartment exchange model (2CX) (Fig. 2.3) describes plasma and the EES as two separate compartments with the volumes v_p and v_e respectively. A compartment is defined as a pool of contrast agent which is well mixed (uniform concentration) and weakly connected to its surroundings [143]. Moreover, the contrast agent flux out of a compartment is proportional to the contrast concentration. The concentration in each compartment as a function of time can be described applying the conservation of mass principle (Eqs. 2.12, 2.13). Solving these coupled differential equations, the tissue

response function can be derived.

$$v_p \frac{dC_p}{dt} = F_p C_a(t) - F_p C_p(t) + PSC_e(t) - PSC_p(t) \quad (2.12)$$

$$v_e \frac{dC_e}{dt} = PSC_p(t) - PSC_e(t) \quad (2.13)$$

Here, $C_p(t)$ is the concentration of contrast in plasma, $C_a(t)$ is the concentration in the supplying artery (defined before as the AIF), and $C_e(t)$ is the concentration in the EES. The rate of change of contrast concentration in the plasma is given by Eq. 2.12. This equation states that the change in $C_p(t)$ is equal to inflow from the supplying artery, $F_p C_a(t)$, and the inflow from the EES, $PSC_e(t)$, minus the outflow into the EES, $PSC_p(t)$, and the outflow through the venous outlet, $F_p C_p(t)$. The rate of change of contrast concentration in the EES is given by Eq. 2.13. Again, according to the conservation of mass principle, the change in $C_e(t)$ is equal to the difference between in and outflow into the EES. Recently, the 2CX model has gained more popularity in clinical applications [8, 39]. While being able to provide separate estimates for PS and F_p , the 2CX model is mainly limited by its requirements with respect to data quality: it is only applicable to high temporal resolution data, with a contrast to noise ratio of at least 40 [40, 91].

2.4.2 Adiabatic approximation to tissue homogeneity model (AATH)

The tissue homogeneity model (TH) is different from 2CX model in the way it describes the intra-vascular compartment. While the 2CX model describes plasma as a well mixed-compartment, the TH model describes plasma as a plug flow system (Fig. 2.4). Namely, it is assumed that no mixing occurs at all, but the contrast agent particles travel with the same velocity, like a coin through a tube. This is a plausible assumption given the small diameter of the capillaries. The EES is still viewed as a compartment. As there is no time-domain solution for the differential equation defining the TH model, this approach has

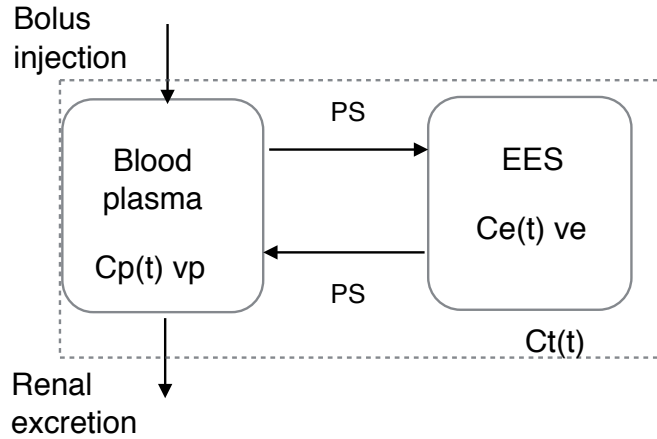


Figure 2.3: *The 2CX Model. Contrast agent is delivered into the plasma (vascular) compartment via the arteries, after a bolus injection. A fraction of the contrast agent then diffuses into EES (tissue compartment) at a rate dependent on the permeability and surface area product of the capillaries. The contrast agent is afterwards eliminated through renal function.*

had limited applicability. To address this issue, the adiabatic approximation to the tissue homogeneity (AATH) model was proposed. The main difference between the two models regards the assumptions about capillary wall permeability. The AATH model assumes the capillary wall is impermeable to contrast agent, and that the exchange between plasma and the EES can only happen at the venous end of the capillary.

2.4.3 Distributed parameter model (DP)

Similarly to the AATH model, the distributed parameter model (DP) describes the intravascular component as a plug flow (Fig. 2.5). Nevertheless, in contrast with the AATH model, the DP model views the EES as a collection of infinitesimal compartments which interact to the closest vessel wall, but do not exchange substances with each other. This assumption has no physiological grounding, but is employed for mathematical convenience as it provides an analytic solution to the residue function.

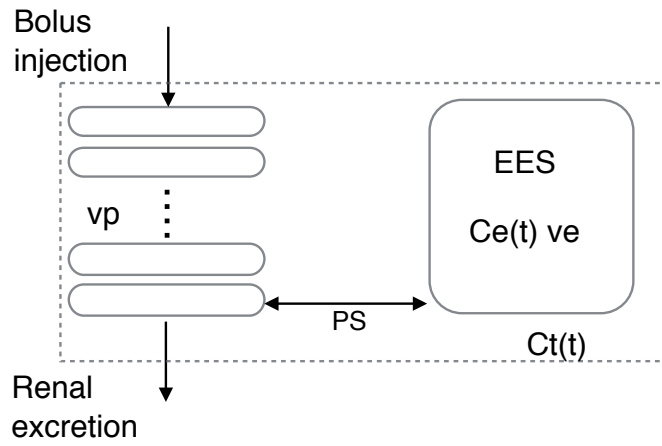


Figure 2.4: *The AATH Model. The plasma component is described as a plug flow system and the EES is represented as a compartment.*




Figure 2.5: *The DP Model. Both the plasma component and the interstitial space are described as a plug flow system.*

2.4.4 Common approximations

One-compartment model. The Tofts model

One of the simplest approximations of the generic models presented above is the one-compartment model, which describes a single well-mixed volume v with in and outflow F :

$$v \frac{dC_t}{dt} = FC_a(t) - FC_t(t) \quad (2.14)$$

In this equation, F represents the flow that transports contrast agent in and out of the volume v . This model is suitable in a number of contexts, and the interpretation of F and v depends on the context in which the model is applied. The first such scenario is a *no exchange regimen*. If the contrast agent extravasation is negligible, the flow F into the compartment is equivalent to the plasma flow F_p , and the volume of the compartment v is the plasma volume v_p . The second such scenario is the *fast exchange regimen*. In this case, the contrast agent extravasates so quickly, that the plasma volume and the EES appear to be a single compartment of volume $v_p + v_e$. F is still equivalent to the plasma flow F_p . Third, assuming the tissue is weakly vascularized, the plasma volume v_p becomes negligible [71]. This is the *extra-vascular regimen*, and has been referred to as the Tofts model [143].

The mathematical formulation of the Tofts model results from Eq. 2.14, by equating F with K^{trans} and assuming the volume of vessels in tissue is negligible:

$$v_e \frac{dC_e(t)}{dt} = K^{trans}(C_a(t) - C_e(t)) \quad (2.15)$$

By solving Eq. 2.15, the form of the Tofts residue function can be derived (Eq. 2.16). The full derivation can be found in Tofts et al. [141].

$$C_t(t) = C_a(t) * K^{trans} e^{-k_{ep}t} \quad (2.16)$$

Here, C_t is the observed tissue response, C_a is the AIF, and the tissue response is given

in terms of model parameters.

The Tofts model is governed by three parameters: the rate constant between EES and blood plasma, k_{ep} , the fractional volume of the EES, v_e , and the volume transfer constant between blood plasma and EES, K^{trans} . However, only two of these parameters are independent, as $k_{ep} = \frac{K^{trans}}{v_e}$. K^{trans} was proven to be the most significant parameter in clinical studies, although its connection with physiology is uncertain. In sclerotic tissue, where trans-endothelial permeability is low, K^{trans} represents permeability surface area product. In tumours, where endothelial permeability is much higher, K^{trans} represents perfusion [142].

Although the standard Tofts model has become very popular in clinical practice, it is important to note that its application is limited to situations where the tissue is weakly vascularized $v_p = 0$. This scenario excludes tumours, which are often highly vascularized due to the process of angiogenesis. However, the usage of this model may be justified for low temporal resolution data, where a second generation model could fail.

Extendend Tofts model

Another commonly used model that can be derived from any of the general cases above is the extended Tofts model. This model differs from the standard Tofts model by accounting for a separate plasma compartment. In this case, the solution becomes:

$$C_t(t) = K^{trans} e^{-\frac{K^{trans}}{v_e} t} * C_a(t) + v_p C_a(t) \quad (2.17)$$

This model is widely used for DCE-MRI analysis, although its underlying assumption is only valid for highly perfused tissue. Recent work has shown that this model is not appropriate for tumours, providing a poor fit [40]. Moreover, if the model is applied to low temporal resolution data, F_p cannot be measured reliably [71].

Tissue uptake model

This model assumes that the concentration of contrast agent in the EES is much smaller than that in plasma, due to large v_e , small PS , short acquisition times, or a combination of the above. This involves that the contrast backflux from the EES is negligible, $PSC_e \approx 0$. Thus, equation 2.13 becomes:

$$v_e \frac{dC_e}{dt} = PSC_p(t) \quad (2.18)$$

The solution is composed of an intravascular part, which washes out exponentially, and an EES part, which accumulates. This model is particularly useful in situations when the acquisition time is too short for the washout of contrast to be determined, or for brain tissue with intact blood-brain barrier, when contrast agent cannot extravasate. The model is also a nice compromise when the data quality required for fitting the 2CX model cannot be attained.

2.4.5 The Arterial Input Function

As mentioned before, accurate estimation of the pharmacokinetic model parameters depends on an accurate measurement of the concentration of contrast agent in blood plasma. The latter is termed the arterial input function (AIF).

Ideally, the concentration of contrast agent in plasma should be measured by arterial blood sampling. Although very accurate, this procedure is highly invasive and therefore not recommendable for routine clinical practice. A common approach is to use a functional form of the AIF, whose parameters were previously derived based on a population of individually measured AIFs [70, 112, 116, 150]. For instance, Orton et al. [112] fit an exponential form to a population averaged AIF (Eq. 2.19). The model is summarized by its parameters, A_1, A_2, μ_1, μ_2 , and can be used in further studies where the AIF is difficult to compute.

$$C_a(t) = A_1 t e^{-\mu_1 t} + A_2 (e^{-\mu_2 t} - e^{-\mu_1 t}) \quad (2.19)$$

However, the AIF varies between individuals and visits due to variation in physiologi-

cal parameters such as cardiac output, vascular tone, renal function, and also in external parameters such as injection timing. Moreover, it has been shown that using a standard population-averaged AIF for all patient scans at all time-points can lead to large errors in PK parameter values [118]. Measuring patient-specific AIFs is therefore important for accurate PK modelling [124].

Directly extracting the AIF based on the DCE-MRI signal enhancement curve has several disadvantages. First, there is a non-linear relationship between signal enhancement and contrast agent concentration. Second, most DCE-MRI sequences are acquired with low temporal resolution (9.5s between subsequent volumes for the data used in this study), which makes the AIF very difficult to recover. Third, the inherent artifacts of MR imaging (i.e. $T2^*$ effects at higher concentrations or $B1$ -field inhomogeneities) also affect AIF measurement [84]. Phase-based AIFs [84] have been recently proposed to address several of these issues related to signal magnitude AIFs. They do not, however, tackle the challenges posed by low temporal resolution data, as their approach was validated on high temporal resolution MRI data (2.4 second intervals). The issue of low temporal resolution was addressed by pre-bolus (or dual-bolus) studies [86], which involve injecting a small dose of contrast agent, followed by imaging with a high temporal resolution (2s) before the main sequence is acquired. This approach still does not account for the non-linear relationship between image signal and contrast concentration, and also extends the length of the imaging session by 8 – 10 minutes.

2.5 Discussion

In this chapter, we have explained DCE-MRI and pCT data acquisition (Sections 2.1, 2.2), these being among the most important modalities used for cancer therapy monitoring. Both DCE-MRI and pCT measure tumour microvasculature changes that result from angiogenesis. Both modalities involve the injection of a contrast agent and repeated

acquisitions of MRI or CT volumes containing the region of interest. The underlying tissue changes can be quantified by pharmacokinetic analysis, which was described in Section 2.4. Classic models such as 2CX, AATH and DP were presented, together with their most common approximations. The choice of AIF, which is important for accurate pharmacokinetic modelling, was also discussed.

Generally speaking, the ideal model provides a good fit to the data and has as little parameters as possible. To differentiate between equally fitting models, prior knowledge about the problem to be solved is often included. For the case of colorectal tumours, the tissue is expected to have many leaky vessels which are formed through angiogenesis, interspersed with regions of necrosis where the supply of nutrients is limited.

The historic approach to PK modelling is the Tofts model. Although its underlying assumption, $v_p = 0$ is probably not adequate for colorectal tumours, it has the advantage of being applicable to low temporal resolution data, has a small number of parameters (two) and is very popular in clinical practice. In Chapter 4, we show the results of applying the Tofts model to colorectal cancer. This model was also compared with a higher order model, namely the tissue uptake model.

As there is a lot of variability between patients' physiological properties, the use of individualised arterial input functions is recommended. In Chapter 4 we elaborate on this idea and propose to extract patient specific AIFs from contemporaneous pCT data. These AIFs are then employed for pharmacokinetic modelling of DCE-MRI.

Chapter 3

Registration of Dynamic Images

Image registration is one of the main fields of research in medical image analysis. Its importance stems from the need to combine or compare information from different medical images, either from the same patient, or across a group of patients, to improve diagnosis and treatment. Classic examples of image registration problems include multimodal image fusion, longitudinal registration, registration of dynamic sequences, registration for treatment planning, etc [54]. The first involves aligning different imaging modalities to combine complementary information of the same patient. Examples include MR/CT registration for improved treatment planning, or MR/PET (or CT/PET) registration, which helps combine structural and functional information. Longitudinal registration is used for comparing series of images of the same subject, taken over a longer period of time. This is particularly useful to monitor changes and the effectiveness of treatment in applications such as neurodegenerative diseases and cancer. The third example involves tracking and correcting for motion in dynamic image sequences. While tracking motion is used to assess the spatio-temporal behaviour of an organ such as the lungs or the heart, correcting for motion enables accurate functional analysis of dynamic images such as DCE-MRI or pCT. Lastly, surgery and radiotherapy often involve the overlay of high-resolution images acquired before the procedure, with low-resolution images used during the intervention. In the following, we present the general principles of image registration, and then review

the most relevant literature concerning the registration of dynamic images, which is one of the focus points of this thesis.

3.1 General Principles of Image Registration

An image I can be defined as a mapping $I : X \mapsto \mathbb{R}$ from the ND grid X to the space of allowed intensity values. Image registration is defined as the process of finding the optimal transformation T that maps a source image I_s onto a target image I_t , to achieve correspondence. Depending on the application, the user might be interested in structural correspondence (i.e. same patient, before and after treatment), functional correspondence (i.e. aligning equivalent regions of the brain for a group of subjects), or structural-functional correspondence (i.e. overlapping functional information over a structural image) [35]. The transformation T deforms the underlying grid into a new grid. At any point $\mathbf{x} \in X$ in the ND grid, the transformation can be defined as $T(\mathbf{x}) = \mathbf{x} + \mathbf{u}$ where \mathbf{u} is the ND vector of displacements. For a $3D$ image, (u, v, w) would be the displacements on the x-, y- and z-axis, respectively. A generic registration algorithm contains the following three components:

- a transformation model T , which imposes a prior on the allowed deformations based on assumptions about patient motion, image acquisition, tissue characteristics. A transformation model is typically characterised by its number of parameters, or 'degrees of freedom'.
- a cost function C , which measures how well the images are aligned for a given transformation T , at some point in the optimisation process. Most often, a cost function contains a (dis)similarity term and a regularisation term. The dissimilarity metric can be based on image intensities, geometric or anatomical features, or high level information i.e. segmentations. The term dissimilarity is adopted, as this cost is minimised during registration. The regularisation term is used to penalise deformations which are unrealistic.

- an optimisation strategy, to find the transformation \hat{T} which minimises the cost function C . This is typically done in an iterative manner.

In mathematical terms, the cost function can be written as the combination between the dissimilarity cost C_D and the regularisation cost C_R , weighted by a factor λ :

$$C(T) = C_D(I_t, T(I_s)) + \lambda C_R(T) \quad (3.1)$$

To summarise, the solution to the image registration problem is the optimal transformation \hat{T} , where $\hat{T} = \arg \min C(T)$. In the following, the three building blocks of image registration are discussed, together with key methodological contributions. A comprehensive review of methodological developments is given by Sotiras et al. [133]. We then give some considerations about validation of registration algorithms.

3.1.1 Transformation models

According to the number of degrees of freedom, transformations can be classified as linear or non-linear (non-rigid). Linear transformations can be defined by a transformation matrix, and allow for image translations, rotations, scaling and shearing. The simplest linear transformation is the rigid body transformation, which, for a $3D$ image, has 6 degrees of freedom (3 represent translations along the axes, and 3 represent the rotation angles around the axes). A transformation that allows for rigid movement, together with scaling and shearing, is called an affine transformation and has 12 degrees of freedom. Although robust, linear transformation models are often insufficient for capturing the complex and localised deformation that is present in medical images.

Non-linear transformation models come to address this issue, as they can feature tens to million of degrees of freedom. The main challenges in deformable image registration reside in the problem being under-constrained, non-convex and non-linear. Often, the choice of a particular deformation model involves a compromise between computational efficiency and the accuracy of the description. Following the classification proposed

by Holden [69] and Sotiras et al. [133], (i) transformations inspired by physical models, (ii) transformations inspired from interpolation and approximation theory, and (iii) knowledge-based transformations, we briefly review some of the most important transformation models.

Among deformations inspired by physical models, we note elastic registration, viscous fluid registration, diffusion registration and flows of diffeomorphisms. These approaches are usually formulated in a variational setting, which involves adding a regularisation term to the cost function, such as in Eq. 3.1, and minimizing the cost function with respect to the displacement field \mathbf{u} . The derivation is performed in the infinite-dimensional continuous setting. The optimum can be found by solving the Euler-Lagrange partial differential equation (PDE) (Eq. 3.2), which corresponds to the derivative of the cost function with respect to the displacement.

$$\lambda A_R \mathbf{u} + f_D(\mathbf{u}) = 0 \quad (3.2)$$

Here, $\lambda A_R \mathbf{u}$ corresponds to $\nabla C_R(\mathbf{u})$, and $f_D(\mathbf{u})$, which is also called the driving force, corresponds to $\nabla C_D(\mathbf{u})$. The regularisation operator is generically denoted as A_R and the formula is valid for linear regularisers.

Elastic registration [20,48], the oldest approach for non-rigid registration, assumes the images deform like an elastic body. For this algorithm, the driving force $f_D(\mathbf{u})$ depends on the image similarity metric, and the regularisation imposes elasticity constraints. The main limitation of this registration method is that it cannot capture large deformations due to the elastic regularisation term.

Coming to address this limitation, the viscous flow model [21,30,36] describes the deformation acting on the image as the flow of a viscous fluid. In contrast with elastic registration which operates on displacements, the viscous flow model operates on velocities. This enables the model to capture large deformations. An important limitation of this approach is its computational complexity.

Diffusion models employ the gradient of the deformations as a regularisation penalty. In this case, Eq. 3.2 becomes $\Delta \mathbf{u} + f(\mathbf{u}) = 0$. The corresponding regularisation term for a N dimensional image is given by Eq. 3.3.

$$C_R(\mathbf{u}) = \sum_{i=1}^N |\nabla u_i|^2 \quad (3.3)$$

A popular approach that uses diffusion regularisation is the Demons framework [140]. This framework has appeared from the need of a more efficient deformable registration method, contrasting with the back-then state-of-the art methods that involved the simulation of physical processes (i.e. elastic, fluid registration). Demons registration is a generic algorithm with many different variants. Its main characteristic is the decoupling between the calculation of the driving force and the regularisation, which are performed as separate steps. In the first step, the force field is calculated as to minimise the dissimilarity cost (Eq. 3.4). In the regularisation step, the force field is smoothed by convolution with a Gaussian kernel G_σ (Eq. 3.6).

$$\mathbf{f} = \frac{(I_t - I_s(T))\nabla I_s(T)}{(I_t - I_s(T))^2 + \|\nabla I_s(T)\|^2} \quad (3.4)$$

$$\mathbf{v} = \tau \mathbf{f} \quad (3.5)$$

$$\mathbf{u}_{i+1} = G_\sigma * (\mathbf{u}_i + \mathbf{v}) \quad (3.6)$$

Here, the transformation is modelled as a dense displacement field which is added to the identity transform: $T(\mathbf{x}) = Id + \mathbf{u}$. At each iteration, the current displacement \mathbf{u}_i is incremented by an update value \mathbf{v} , which is computed based on a set of forces \mathbf{f} . The updated displacement field is then smoothed by Gaussian regularisation. Filtering with a Gaussian kernel can be seen as an approximation of the diffusion regularisation in Eq. 3.3.

In its initial formulation [140], the demons algorithm was purely heuristic. Since its initial formulation, several works [121, 147] have shown that the force term (3.4) is an approximation to an update step of a gradient descent minimisation of the sum of

squared differences (SSD). SSD will be discussed in more detail in Section 3.1.2, and gradient descent optimisation will be described in Section 3.1.3. The demons algorithm was revisited many times and a number of extensions and generalizations were proposed. According to Zikic [160], these revisions can be grouped into three categories: extensions to incorporate any similarity measure [23, 105] and any regularizer, extensions to make transformations invertible (diffeomorphic) [146, 147], and generalization to establish a correspondence with fluid registration [21], making the algorithm suitable for computing large deformations.

Among deformations motivated by mathematical models, we note thin plate spline registration and B-spline registration. Instead of calculating a displacement field at each voxel, this set of methods proposes representing the transformation model as a function with a low number of parameters. The regularisation is implicitly enforced by the chosen parametrisation. Thin plate splines (TPS) [16] is the most known registration algorithm from the radial basis functions registration family. Its name is given from the analogy with a bending sheet of metal. Whereas the main advantage of this method is the ability to interpolate a deformation field from an irregular set of control points, its main shortcoming is the global support. In other words, the displacement value at one control point influences the values of the points in the entire image.

One of the most popular registration approaches, the Free Form Deformation (FFD) algorithm [131] provides a transformation model with local support by using cubic B-splines as a basis function. A rectangular grid of uniformly spaced control points is superimposed on the image, and displacements are estimated at each control point. Given the displacement at the control points p , the dense displacement field $\mathbf{u}(\mathbf{x})$ can be found using:

$$\mathbf{u}(\mathbf{x}) = \sum_{i=0}^3 \sum_{j=0}^3 \sum_{k=0}^3 B_i(u) B_j(v) B_k(w) p_{m+i, n+j, o+k} \quad (3.7)$$

Here, δ_x , δ_y and δ_z represent the control point spacing, $m = \lfloor \frac{x}{\delta_x} \rfloor - 1$, $n = \lfloor \frac{y}{\delta_y} \rfloor - 1$, $o = \lfloor \frac{z}{\delta_z} \rfloor - 1$, and B_l represents the l th order B-spline function. Up to the 3rd order, the values

for B_i are: $B_0(u) = (1 - u)^3/6$, $B_1(u) = (3u^3 - 6u^2 + 4)/6$, $B_2(u) = (-3u^3 + 3u^2 + 3u + 1)$ and $B_3(u) = u^3/6$.

3.1.2 Similarity measures

There is a plethora of similarity metrics that can be used for image registration. The choice of similarity measure depends on the assumptions about the relationship between the intensity values of the images to be registered, and on assumptions about image noise. SSD, which is the simplest similarity metric, expects an identity relation between the intensities in the two images (Eq. 3.8). This assumption holds for mono-modal registration, when the two images differ only by Gaussian noise.

$$SSD = \frac{1}{N} \sum_{\mathbf{x} \in \Omega} (I_t(\mathbf{x}) - I_s(T(\mathbf{x})))^2 \quad (3.8)$$

Here, Ω is the image domain, and \mathbf{x} represents the voxel coordinate vector.

If we assume a linear relationship between image intensities, normalised cross-correlation (NCC) is a suitable similarity metric. Generally, NCC is a statistical measure of the linearity of the relationship between two random variables, these being image intensities for the registration problem. The NCC between target image I_t and source image I_s is defined as:

$$NCC = \frac{\sum_{\mathbf{x} \in \Omega} (I_t(\mathbf{x}) - \mu_t)(I_s(T(\mathbf{x})) - \mu_s)}{\sqrt{\sum_{\mathbf{x} \in \Omega} (I_t(\mathbf{x}) - \mu_t)^2 \sum_{\mathbf{x} \in \Omega} (I_s(T(\mathbf{x})) - \mu_s)^2}} \quad (3.9)$$

Here, Ω is the image domain, and μ_t , μ_s are the means of I_t , I_s , respectively. NCC is particularly useful for mono-modal images that were acquired with different intensity windowing.

When registering images from different modalities (multi-modal registration), the assumption of a linear relationship between image intensities is no longer valid. Typically, this problem can be addressed either by using similarity measures from information theory, which assume a statistical relationship between image intensities, or by reducing the

problem to a single-modal one. The most widely known information theoretic similarity metric, mutual information (MI) [93] describes how well an image is explained by the other. MI is defined based on entropy, which is a measure of the amount of information in an image. Let $P(I_t(\mathbf{x}))$ the marginal probability of a voxel \mathbf{x} in the target image having intensity $I_t(\mathbf{x})$, $P(I_s(\mathbf{x}))$ the marginal probability of a voxel \mathbf{x} in the source image having intensity $I_s(\mathbf{x})$, and $P(I_t(\mathbf{x}), I_s(\mathbf{x}))$ the joint probability of a voxel \mathbf{x} having intensity $I_t(\mathbf{x})$ in the target image and intensity $I_s(\mathbf{x})$ in the source image. The marginal entropy is defined as:

$$H(I) = - \sum_{\mathbf{x} \in \Omega} P(I(\mathbf{x})) \log P(I(\mathbf{x})) \quad (3.10)$$

This entropy is considered maximal if all the intensities have equal probability. The joint entropy is defined as:

$$H(I_t, I_s) = - \sum_{\mathbf{x} \in \Omega} \sum_{\mathbf{x} \in \Omega} P(I_t(\mathbf{x}), I_s(\mathbf{x})) \log P(I_t(\mathbf{x}), I_s(\mathbf{x})) \quad (3.11)$$

Using MI as a similarity metric relies on the assumption that the images are aligned when the joint entropy is lowest. The MI between the target image I_t and the source image I_s is defined in Eq. 3.12:

$$MI(I_t, I_s) = H(I_t) + H(I_s) - H(I_t, I_s) \quad (3.12)$$

Since it was introduced to image registration, mutual information has been widely used for the registration of images whose intensities are linked by a statistical relationship. Its normalised variant, NMI [137] involves dividing MI by the marginal entropy of the target image, to eliminate the effect of the changing image overlap on the similarity metric.

Another information theoretic measure, the Kullback-Leibler Divergence (KLD) [31] is based on comparing joint intensity distributions. Images are aligned by minimising the distance between the observed distribution and a reference distribution, which is learnt from correctly aligned images.

The information theoretic measures presented above are based on single voxel joint probability models. As a consequence, they have the drawback of being insensitive to changes in the position of the voxels. Later developments addressed this issue, as they involved adding local context to the similarity metric [89].

Another approach to deal with multi-modal registration is to reduce the problem to a mono-modal one, by mapping both modalities in a common space. Such an approach was recently proposed by Heinrich et al. [63], who calculate a descriptor based on similarities between neighbouring patches. Once the descriptor image is constructed for each modality, the difference between descriptors can be taken as a similarity metric.

3.1.3 Optimisation strategies

The optimal transformation that minimises the cost function can be found using an optimisation or inference strategy. Depending on the nature of the variable to be inferred, optimisation methods can be classified as (1) continuous or (2) discrete. Continuous optimisation methods are the traditional approach to image registration. They operate on real valued variables, and can typically be applied when the cost function is differentiable. These types of methods estimate the optimal set of parameters by successively updating the solution. A generic continuous optimisation algorithm is presented in Algorithm 1. Here, \mathbf{u} is the parameter vector of the transformation, \mathbf{u}_0 is the initialisation value, $\hat{\mathbf{u}}$ is the

Algorithm 1 Generic optimisation algorithm

- 1: *initialize*: $\hat{\mathbf{u}} \leftarrow \mathbf{u}_0$
 - 2: **repeat**
 - 3: *compute*: d_k, α_k
 - 4: *update*: $\mathbf{u}_{k+1} \leftarrow \mathbf{u}_k + \alpha_k d_k$
 - 5: **until** *convergence*
-

optimal value, k is the iteration number, α_k is the step size, and d_k is the search direction. The step size and the search direction depend on the chosen optimisation method. Some of the most prominent continuous registration approaches are: Steepest Descent, Conjugate Gradient, Gauss-Newton, Quasi-Newton (i.e. BFGS) and Stochastic Gradient Descent.

A detailed description of these methods, together with their implementations can be found in [125]. In the following, we briefly describe each of them, commenting on their advantages and disadvantages.

The Steepest Descent algorithm search direction is the negative of the gradient. The gradient of the cost function (Eq. 3.1) is denoted by $\mathbf{g}_k = \nabla C(\mathbf{u})$. While this method is fairly easy to implement and compute, its main limitation is the slow convergence. Steepest Descent was often employed for the registration problem [11, 131]. In order to accelerate convergence, Conjugate Gradient proposes a search direction that is conjugate to the direction at the previous step. Newton’s method proposes using both gradient and curvature information, but it has seldom been used in practice as it involves calculating the Hessian matrix of the cost function, \mathbf{H} , at each iteration. Several alternatives to the direct calculation of the Hessian matrix have been proposed. Gauss-Newton optimisation, which can be used if the registration is formulated as non-linear least squares minimisation, proposes approximating the Hessian with $\mathbf{H} \approx 2\mathbf{J}^T\mathbf{J}$. Here, \mathbf{J} denotes the Jacobian and the higher order derivatives have been ignored. This method is particularly useful for single-modal image registration, where the similarity metric can be formulated as minimising the sum of squared residuals. Examples of works that have used Gauss-Newton optimisation are [146, 147]. Quasi-Newton methods such as the Broyden Fletcher Goldfarb Shanno (BFGS) algorithm [125] approximate the Hessian matrix by using information from previous iterations. An application of this approach can be found in [89]. Addressing cases when the deterministic gradient is difficult to calculate, researchers have proposed Stochastic Gradient Descent. For this method, the update rule is based on an approximation of the true gradient.

Several ways of choosing the step size α_k have been proposed: it can be constant $\alpha_k = \alpha$, slowly decreasing with the iteration number $\alpha_k = \alpha/k^n$, or it can be chosen to minimise the cost function for the current transformation parameters $\alpha_k = \operatorname{argmin}_\alpha C(\mathbf{u}_k + \alpha d_k)$. A comparison of the methods presented above was performed by Klein et al. [79] using a B-spline transformation model and MI as a similarity metric.

However, most continuous optimisation approaches require computing the gradient of a cost function. This can be limiting, as not all the cost functions are easy to differentiate or even differentiable. Moreover, these approaches are highly dependent on the initialisation and can get trapped in local optima if they are gradient based. To address these limitations, discrete optimisation methods operate on variables which can take values in a discrete set. They are robust to initialisation and perform global searches over the quantised space of possible parameter values. Their main limitation is the tradeoff between the size of the search space and the computational efficiency.

Discrete optimisation is typically formulated as a Markov Random Field (MRF) labelling problem [50]. For deformable image registration purposes, a graph is defined in which the nodes represent voxels or groups of voxels, and the edges connect voxels with similar anatomical features and spatial proximity. For every node, there is a set of labels which represent possible discrete displacements of the source image volume with respect to the target image volume. Finding the optimum displacement at each voxel equates to finding the labelling that minimises the MRF energy function.

Discrete optimisation approaches generally fall in one of three categories: 1) graph-cuts, 2) message passing, and 3) linear programming. Graph-cuts methods rely on the max-flow/min-cut algorithm, which states that the maximum flow between the sink node and the source node is equal to the minimum cut that separates the two nodes. Originally, graph cuts could only solve the binary labelling problem, but they have been extended to multiple labels [18]. α -expansion [18], one of the most popular graph-cuts algorithm, performs multi-labelling by iteratively applying binary graph cuts. The main limitation of this approach is its applicability solely to submodular functions [81]. This involves that the truncated quadratic regularisation, which is a popular regularisation strategy in image registration, cannot be used in conjunction with graph-cuts.

Message passing approaches involve message exchanges between nodes according to a set schedule. Their main advantage is the ability to cope with non-submodular functions. A message is a vector of the same size as the label set, and it represents the belief at a

particular node regarding each solution. On a chain or tree-like graph, belief propagation methods are equivalent to dynamic programming and give an exact solution. Loopy belief propagation (LBP) [45], the variant of the algorithm for graphs containing loops, was shown to have good convergence properties. An improved message passing schedule was proposed in tree-reweighted message passing algorithm (TRW-S) [80]. This algorithm also involves linear programming (LP) relaxation of the original problem. Instead of minimising the energy of the original problem, which may be hard, TRW-S maximises the lower bound on this energy. In particular, the original graph is decomposed as a convex combination of trees. Belief propagation is performed for each tree, and the results are averaged.

The Fast Primal-Dual (FastPD) algorithm [82] is also based on LP relaxation. The solutions are derived so that the primal-dual gap decreases. This approach shows improved performance with respect to α -expansion and has been used extensively for image registration [50]. However, its advantages come at the cost of greatly increased memory requirements. Excellent comparative studies of the different discrete optimisation methods can be found in [74, 138].

When discussing the performance of an optimisation algorithm, aspects such as convergence, the number of iterations, the cost per iteration, and the memory footprint are usually discussed.

3.1.4 Validation

Evaluating the quality of a registration algorithm is a key task when developing new methodologies. Validation enables evaluating the algorithm’s performance, understanding its strengths and weaknesses, and benchmarking it against other algorithms. Often, validation is a very challenging problem, as there is no ground truth. To date, many evaluation criteria have been proposed. Examples can range from evaluating the quality of the resulting transformation, over calculating the distance between corresponding anatomical landmarks, to higher level clinical metrics.

Rohlfing [129], who tests several validation metrics in the context of a mock registration algorithm, has shown that some metrics can yield completely unreliable results. His algorithm, a completely useless registration tool (CURT), produces a meaningless transformation but manages to outperform some state-of-the-art registration algorithms in terms of similarity-based surrogate metrics. Thus, similarity-based metrics should be avoided for evaluating registration accuracy.

The most widely accepted method of evaluating registration accuracy is based on identifying corresponding anatomical landmarks on the images to be registered. The metric employed to compare the degree of mismatch between corresponding landmarks before and after registration is target registration error (TRE). The TRE is defined as the Euclidian distance between corresponding points in the image pair. For a given landmark \mathbf{p} in the source image, and the corresponding landmark \mathbf{q} in the target image, the TRE is given by:

$$TRE = \|T(\mathbf{p}) - \mathbf{q}\| \quad (3.13)$$

Another approach to validating a registration algorithm involves evaluating the quality of the resulting transformation. The most common way of doing that involves calculating the Jacobian determinant $Jac(\mathbf{u})$ of the resulting transformation $T = Id + \mathbf{u}$:

$$Jac(\mathbf{u}) = \begin{bmatrix} 1 + \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} & \frac{\partial u}{\partial z} \\ \frac{\partial v}{\partial x} & 1 + \frac{\partial v}{\partial y} & \frac{\partial v}{\partial z} \\ \frac{\partial w}{\partial x} & \frac{\partial w}{\partial y} & 1 + \frac{\partial w}{\partial z} \end{bmatrix} \quad (3.14)$$

The Jacobian at position (x, y, z) represents the volume change caused by the transformation at that position, giving a measure of local deformation properties. While $Jac(\mathbf{u}) < 1$ represents tissue compression, $Jac(\mathbf{u}) > 1$ represents expansion. Negative Jacobian values represent the disappearance of the corresponding voxels, also known as tissue folding. It is generally thought that this should be avoided for medical images, due to the assumption of tissue incompressibility. The fraction of voxels with negative Jacobian values is

often used as a surrogate measure for registration quality.

If the proposed registration algorithm is used to better discriminate between two categories of patients, e.g. responders and non-responders to cancer treatment, then the outcome data of the respective clinical trial can be used as a secondary measure of registration quality. For instance, in Bhushan et al. [12] image registration was used to improve prediction of response to chemo-radiotherapy in colorectal cancer patients. To this end, the Kolmogorov-Smirnoff (KS) distance between pre- and post-therapy PK parameter distributions was calculated, and was used to classify each patient as a responder or a non-responder. A registration algorithm is considered effective if it increases the separation between responders and non-responders in terms of KS distance. A similar approach is used in Chapter 4, where we looked at the shift between pre- and post-therapy PK parameter distributions and correlated this with clinical outcomes. It should be noted that a better correlation with the clinical outcome does not necessarily imply increased accuracy of the proposed registration algorithm. Nevertheless, in clinical applications the goal to extract clinically relevant information, which does not always coincide with increased registration accuracy. For this reason, secondary validation measures are often the best option.

3.2 Registration of Dynamic Image Sequences

Dynamic imaging data is becoming increasingly important for a number of applications such as quantification of cardiac function, respiratory motion modelling prior to image guided interventions, or quantitative analysis of DCE-MRI or pCT. For all these applications, intra-sequence registration of the dynamic images is a key challenge. In addition to patient and organ motion, dynamic images present extra challenges due to both small and large motion, and potentially contrast inflow. In this section, we review some of the most important work concerning dynamic data registration. Contrast enhanced and non-contrast enhanced data are discussed separately, as the former presents additional

challenges.

3.2.1 Respiratory motion modelling from dynamic images

Respiratory motion estimation is important for a number of clinical applications, including pulmonary function evaluation, in order to diagnose and understand breathing disorders such as COPD, or radiotherapy planning for lung cancer patients. Nevertheless, estimating respiratory motion is a challenging problem, due to the complex lung motion which includes sliding motion and large movement of small features, and due to reduced image quality associated with low radiation dose acquisition protocols. Image registration methods that estimate respiratory motion in an automated way have been proposed since the early 2000s [96]. In this section, we focus on respiratory motion modelling approaches that rely on 4D image registration or that enforce temporally smooth transformations in some way. For a comprehensive overview of respiratory motion models, we refer the reader to the recent review paper of McClelland et al. [96].

Castillo et al. [24] use a 4D optical flow method which is based on trajectory modelling for the registration of 4D CT lung images. The inference is solved using the Levenberg-Marquardt algorithm. Their algorithm was tested on the DIR-lab dataset [25], yielding an average registration accuracy of 1.25mm. Metz et al. [101] proposed using a $2D + \text{time}$ or $3D + \text{time}$ FFD B-spline model for performing non-rigid registration of dynamic images. Their method uses a similarity metric that minimises intensity variance over time, and the inference is performed using stochastic gradient descent with an adaptive step size. A number of different applications were tackled, including $4D$ CT images of the lung (DIR-lab dataset [25], POPI dataset [145]) and $4D$ CT angiography images. Several variants of the algorithm were investigated: registration with respect to a reference timepoint, consecutive timepoint registration, groupwise registration, and the groupwise approach with applying additional temporal smoothness constraints. The groupwise approach was shown to outperform the other variants. The proposed method achieves temporally smooth deformations. However, this comes at the cost of reduced registration accuracy compared

to the best published results on the DIR-lab and POPI data [24, 72]. This work is also limited by the use of lung segmentations as an input to the registration algorithm. Besides involving extra computations, this approach only achieves registration inside the lungs. A B-spline transformation model was also used by Vandemeulebroucke et al. [144] and by Yigitsoy et al. [158]. Yigitsoy et al. [158] address the problem of bias towards the selected reference image by employing a similarity metric which involves calculating mean squared intensity differences between all possible pairs of images. Recently, Wu et al. [153] proposed the alignment of all image volumes to a high resolution group-mean image. Temporal smoothness of the transformation is enforced using virtual fibres that connect temporally corresponding points.

3.2.2 Cardiac image registration

Dynamic images of the heart offer information about cardiac anatomy and function, enabling early diagnosis and treatment of cardiac disease. Image registration is essential to derive information about cardiac function, enabling the estimation of cardiac deformation and strain. The literature generally falls in one of two application categories: intra-sequence image registration to quantify motion and strain, which could be used to identify diseased patients, and inter-sequence spatio-temporal sequence matching, for the comparison of cardiac motion across subjects, or the construction of a motion atlas. Inter-sequence registration could also be employed to compare images of the same subject, i.e. before and after treatment.

One of the first attempts to introduce temporal consistency in the intra-sequence registration algorithm is the work of Ledesma-Carbayo et al. [85]. They propose to measure cardiac deformation using spatio-temporal registration based on the B-spline transformation model. This method is applied to strain and deformation quantification from 2D ultrasound images. The approach of Ledesma-Carbayo was further extended to temporal diffeomorphic FFD (TDDFD) by DeCraene et al. [37]. To enable the recovery of a smooth displacement field where the motion at a time point depends on earlier time points, they

propose to represent the transformation using velocities instead of displacements. Synthetic data was used to compare the TDDFD algorithm to both standard 3D FFD and to [85], and the TDDFD approach proved to be more robust to noise and low temporal resolution. The method was then applied to motion and strain quantification from 3D ultrasound. Mahapatra [94] uses temporal information for joint segmentation and registration of DCE-MRI cardiac images. The novel aspect of this work is a similarity measure which ensures smoothly varying temporal curves at each voxel, which is used together with a B-spline FFD transformation model. In recent work, Shi et al. [132] propose using a sparse representation of the FFD transformation, to better capture highly localised and discontinuous deformations. To enforce sparsity, the L_1 -norm of the parametric representation is used as an additional regularisation term. This method was shown to outperform the TDDFD algorithm [37] in terms of registration accuracy.

The work of Perperidis et al. [122] falls in the second category. They address inter-sequence alignment by using a free form deformation (FFD) based on B-splines, and decouple between spatial and temporal components. The temporal transformation is a FFD using a 1D B-spline. While the global part of the transformation corrects for length differences in the cardiac cycle and for translation offsets introduced by differences in acquisition parameters, the local transformation corrects for misalignment due to different cardiac dynamic properties. Applying this approach to 4D MRI registration, they show that the spatio-temporal FFD registration gives improved results over previous spatio-temporal affine registration in terms of DICE overlap. Nevertheless, their method yields an overlap of only 80% and is likely not to capture small deformations of the heart due to the sparse control point grid. Peyrat et al. [123] solve the inter-sequence registration problem in the diffeomorphic demons framework, which yields a dense displacement field. They propose multichannel diffeomorphic demons, which is an extension of the diffeomorphic demons algorithm to vector valued images. Each time point in the sequence constitutes a separate channel. The transformation is also decoupled, with the 4D spatial component mapping trajectories of physical points and the 4D temporal component

mapping physiological states. In contrast to [122], the temporal transformation is not solely image driven, but relies on ECG data. The main advantage of this method is the 4D nature of the spatial transformation obtained by the usage of trajectory constraints. Their approach was tested on both synthetic data generated using an electromechanic cardiac model and on 4D CT images of the heart. Registration accuracy is considerably improved by using motion tracking information and multichannel registration, with respect to [122].

3.2.3 Dynamic contrast-enhanced image registration

As a DCE-MRI acquisition takes several minutes, with volumes being acquired every 5–10 seconds, the resulting time sequence is inherently affected by patient and physiological motion. This motion may introduce significant errors to the per-voxel PK model fitting, as anatomical features of interest might move to different voxel locations in subsequent volumes. To correct for this motion, image registration is needed. DCE-MRI registration is a particularly challenging problem, as observed changes throughout the time series can be either due to motion or due to contrast enhancement. Moreover, contrast arrival can give rise to image features that were not present in the baseline image. In the literature, approaches for time series motion correction broadly fall into two categories: 1. strategies that try to alleviate the effects of contrast enhancement i.e. by using a multi-modal similarity metric [161], or by restricting the applied transformation [139]; 2. strategies that implicitly derive [99] or explicitly assume [12, 22, 61] a model of contrast enhancement which is used in the registration algorithm.

Initial approaches for time series motion correction involved registering each volume of the series to the mean image volume. In the case of DCE sequences, intensity based registration approaches that use typical single-modal similarity measures are likely to fail due to the intensity changes induced by contrast enhancement across volumes. One option is using multimodal similarity measures, such as mutual information (MI) [93]. However, contrast arrival can give rise to image features that were not present in the

baseline images, causing traditional registration approaches to fail.

Among the most prominent approaches, we note the work of Buonaccorsi et al. [22], who are among the first to explore an explicit kinetic model-based registration. In that work, PK parameter estimation and registration to the model predicted sequence are performed iteratively. To be more specific, at each iteration they estimate the model parameters from the uncorrected data, generate a synthetic sequence based on the obtained parameters, and perform registration between the uncorrected and the synthetic sequences. Their work is mainly limited by allowing only 3D translations. In a more recent approach, Bhushan et al. [12] address this issue by proposing a simultaneous non-rigid motion correction and PK parameter estimation method. They use a Demons-like non-parametric transformation model, and a probabilistic cost function which enables simultaneous estimation of the transformation parameters and of the PK model parameters. However, as their approach uses a Gauss-Newton optimisation, this method is sensitive to initialisation and is likely to be trapped in local optima.

A data-driven approach, presented by Melbourne et al. [99], proposes modelling the time series data using principal component analysis (PCA). The underlying assumption is that the first few principal components will contain information about contrast enhancement trends and the remaining principal components contain noise related to motion. This assumption is valid for small peristaltic motion, but does not hold in the case of larger and periodic motion (e.g. breathing).

Hamy et al. [56] also propose a decomposition-based approach. They separate the low rank and the sparse component using robust PCA (RPCA). While the low rank component represents smooth, slowly varying changes such as breathing motion, the sparse component represents rapid intensity changes such as contrast arrival. However, this work is only able to correct for periodic motion such as breathing, and relies on the assumption that there is no peristaltic motion.

Chapter 4

Personalised Arterial Input Functions

To understand and accurately quantify tumour microvasculature changes from DCE-MRI data, an appropriate combination of tracer kinetic model and arterial input function has to be selected. This chapter analyses the effect of different design choices in the process of PK modelling. First, the standard Tofts model, which is widely trusted in clinical practice, is compared to a second generation model, the 2CTU model. Second, the choice of AIF and its impact on PK modelling are investigated. We hypothesise that population averaged models can lead to large errors in PK parameter values, as the AIF varies between patients and visits. The clinical study analysed in this chapter features both longitudinal DCE-MRI and pCT images of the same patients, enabling the extraction of patient-specific AIFs from the pCT images, as these have a much higher temporal resolution. These AIFs are then employed for PK analysis of contemporaneous DCE-MRI images acquired from the same patients, for improved PK parameter consistency. Third, microvasculature changes due to treatment are investigated by comparing pharmacokinetic parameter maps from longitudinal DCE-MRI scans acquired during the course of treatment. The main clinical question under investigation is whether Nelfinavir, an antiretroviral drug, can act as a tumour radiosensitiser by increasing tumour perfusion. The data used in this study is described in Section 4.1. Section 4.2 describes the PK parameter estimation and motion correction framework, the pCT AIF derivation method, and the way this is integrated

in DCE-MRI analysis. Experimental results are shown in Section 4.3. Part of this work was accepted for an oral presentation at SPIE Medical Imaging 2013, and published in the conference proceedings.

4.1 Materials

4.1.1 Sonatina clinical trial

The methods developed in this thesis are aimed at early stage clinical trials of novel cancer therapies. For the analyses described in this chapter, we used data from Sonatina (Study of Nelfinavir Addition to Radiotherapy Treatment in Neo-Adjuvant Rectal Cancer), a phase II clinical trial of Nelfinavir with hypofractionated radiotherapy for colorectal cancer carried out at the Oxford University Hospitals. Ethical approval was obtained from National Research Ethics Service Committee South Central (reference 10/H0604/61).

Cellular resistance to radiation is associated with the over-expression of activated oncogenes, which share molecular signalling in the phosphatidylinositol 3-kinase (PI3K)-Akt-pathway. It has been shown that the blockage of PI3K can augment response to radiation in vitro and in vivo in cells which have a constitutive activation of this pathway, however not in cells with a non-activated pathway [53]. Nelfinavir is an HIV-protease inhibitor which has been shown to inhibit Akt in standardly used clinical doses and to cause radiosensitization in vitro and vivo [52]. This trial was also motivated by pre-clinical research demonstrating that an orally administered drug, Nelfinavir, can cause significant increases in tumour blood flow that should be favourable for subsequent radiotherapy, as well as being an intrinsic radiosensitizer [127].

Ten patients with histologically proven adenocarcinoma of the rectum were recruited to this trial. They were treated with 7 days of Nelfinavir administration prior to radiotherapy (1250 mg bd) and a further 7 days of Nelfinavir during radiotherapy. For the purpose of this trial, radiotherapy was defined as a total dose of radiotherapy 25 Gy, which was delivered in 5 Gy fractions on 5 days during a 7-day period.

4.1.2 Image acquisition

Both DCE-MRI and pCT images were acquired before the administration of Nelfinavir (pre-treatment), after 7 days of drug administration (mid-treatment), and after both drug administration and radiotherapy (post-treatment), to investigate tumour microvasculature changes. Patients 1 and 5 were excluded from the analysis because they did not have a mid-treatment contrast enhancement scan, and the effect of Nelfinavir could not be observed. The post-therapy scan of patient 8 was also excluded from the analysis because of inadequate contrast enhancement. The images from the first two time points were still included in the analysis. Thus, the dataset comprised 8 colorectal cancer patients, out of which 7 had been imaged 3 times during the course of treatment, and one had been imaged pre- and mid-treatment.

The DCE-MRI exam was performed on a 3T GE Medical system. T1-weighted dynamic images of the pelvis were acquired using the LAVA protocol, which is based on a 3D spoiled gradient echo sequence. At the start of the acquisition, contrast agent was injected at a rate of 3 ml/sec, 0.1 mmol/kg body weight. The employed contrast agent was ProHance (Gadoteriol). An image volume was acquired every 9.5 seconds over a period of approximately 5 minutes. This yielded a 4D sequence of dimensions $512 \times 512 \times 52 \times 29$. The in-plane resolution of the images was 0.7813×0.7813 mm and the slice thickness was 2mm. These sequences presented a lot of intra-sequence motion due to patient breathing and peristalsis.

The pCT exam was performed on a GE Medical LightSpeed VCT device. Dynamic images of the pelvis centred over the bulkiest part of the tumour were acquired starting 10 seconds after the start of injection of an iodine based contrast agent (Omnipaque 300). This was injected at a rate of 5ml/sec, 90ml in total, using a pump injector. An image volume of the pelvis was acquired every second over a period of 45 seconds using a cine-mode acquisition (120kv, 60mA, 1s rotation speed, 36cm field of view, 17mSV effective dose). This yielded a 4D sequence of dimensions $512 \times 512 \times 8 \times 45$. The in-plane image resolution was 0.7031×0.7031 mm and the slice thickness was 5mm. Due to the short

acquisition time, these sequences presented mainly peristaltic motion.

4.2 Methods

4.2.1 Pharmacokinetic modelling

Developed in the 1990s, the Tofts model has been widely used for tracer kinetic modelling, and is regarded as the standard in the clinical community. Nevertheless, this model does not provide separate estimates for tissue perfusion and permeability, which are both contained in the parameter K^{trans} . Moreover, the model is only suitable for weakly vascularised or highly perfused tissues [134]. Featuring a higher number of parameters, second generation models can offer separate estimates for tissue perfusion and permeability. This is particularly useful for cancer imaging, as it enables more accurate assessment of tumour angiogenesis. The Tofts and Extended Tofts models were recently compared with the general 2CX model on images of cervical cancer patients [40]. The 2CX model was shown to provide a better fit to the data and to estimate physiological quantities which were assumed negligible by the Tofts model (i.e. MTT). Model comparison has also been performed for breast cancer data [87], but to date no studies have dealt with model comparison for colorectal cancer. The objective of this section is to compare the standard Tofts model with a higher order model from the 2CX family, the 2CTU model, using the data from the Sonatina clinical trial.

As mentioned before, the most widely accepted approach to pharmacokinetic modelling is the Tofts model. Besides its simplicity, the main advantage of this model is its applicability to low temporal resolution data. This is the type of data acquired in the majority of drug trials, including the one which is the focus of this work. As explained before, the Tofts model is governed by two independent parameters, the transfer constant K^{trans} , and the fractional volume of the EES, v_e , and features a third parameter which is the ratio of these two, $k_{ep} = \frac{K^{trans}}{v_e}$ (Eq. 2.16). For the AIF, $C_a(t)$, we have chosen to test the Orton bi-exponential model (Eq. 2.19). Although several other population-

averaged AIFs exist in the literature (e.g. Parker [116]), the main advantage of using a bi-exponential model is that a closed form solution for the convolution in Eq. 2.16 can be found. Furthermore, the Orton bi-exponential was previously shown to outperform the Weinmann bi-exponential [13]. Thus, by computing the convolution in Eq. 2.16 in the Laplace domain, the analytic expression for the contrast concentration in tissue, $C_t(t)$ becomes:

$$C_t(t) = \frac{A_B K^{trans}}{k_{ep} - \mu_B} \left(t e^{-\mu_B t} - \frac{e^{-\mu_B t} - e^{-k_{ep} t}}{k_{ep} - \mu_B} \right) + A_G K^{trans} \left(\frac{e^{-\mu_G t} - e^{-k_{ep} t}}{k_{ep} - \mu_G} - \frac{e^{-\mu_B t} - e^{-k_{ep} t}}{k_{ep} - \mu_B} \right) \quad (4.1)$$

For the Orton AIF, the population averaged parameters have previously been found to be: $A_B = 323 \text{mM min}^{-1}$, $A_G = 1.07 \text{mM}$, $\mu_B = 20.2 \text{min}^{-1}$, and $\mu_G = 0.172 \text{min}^{-1}$ [112].

However, as explained in Chapter 2, the Tofts model does not allow separate measurement of F_p and PS , and is further limited by the assumption that $v_p = 0$. For these reasons, we chose to compare it with a second generation model. The two compartment tissue uptake (2CTU) model is a nice compromise when the data quality required for the 2CX model cannot be attained. Assuming the ratio $\frac{PS}{v_e}$ is very small due to either slow permeability or large EES volume, the backflux of contrast from the EES can be neglected (Eq. 2.18). Following the argument from [136], the expression for the tissue response can be found:

$$F_p R(t) = F_p e^{-t/T_p} + F_p E (1 - e^{-t/T_p}) \quad (4.2)$$

Here, T_p is the plasma mean transit time, E is the extraction fraction, and the other parameters are as defined in Chapter 2. Besides the measured parameters, F_p , E and T_p , this model also enables the derivation of two more clinically relevant parameters, namely

permeability, PS , and plasma volume, v_p :

$$PS = \frac{EF_p}{1-E}, \quad v_p = (F_p + PS)T_p \quad (4.3)$$

Again, assuming the Orton AIF, the analytic expression for contrast agent concentration in tissue can be derived by computing the convolution between the tissue response function (Eq. 4.2) and the Orton bi-exponential (Eq. 2.19):

$$\begin{aligned} C_T(t) = F_p & \left[A_G \frac{1-E}{\mu_G - \frac{1}{T_p}} (e^{-\frac{t}{T_p}} - e^{-\mu_G t}) - A_G \frac{1-E}{\mu_B - \frac{1}{T_p}} (e^{-\frac{t}{T_p}} - e^{-\mu_B t}) \right. \\ & + A_G \frac{E}{\mu_G} (1 - e^{-\mu_G t}) - A_G \frac{E}{\mu_B} (1 - e^{-\mu_B t}) + A_B E \left(\frac{1}{\mu_B^2} - \frac{e^{-\mu_B t}}{\mu_B^2} - \frac{-te^{-\mu_B t}}{\mu_B} \right) \\ & \left. + A_B (1-E) \left(\frac{e^{-\frac{t}{T_p}}}{(\mu_B - \frac{1}{T_p})^2} - \frac{-e^{-\mu_B t}}{(\mu_B - \frac{1}{T_p})^2} - \frac{te^{-\mu_B t}}{\mu_B - \frac{1}{T_p}} \right) \right] \quad (4.4) \end{aligned}$$

Each of these models was fitted to the data using the Levenberg-Marquard algorithm [125], and the results were compared (Section 4.3).

To assess whether the 2CTU model, which is more complex, is more suitable for the data than the Tofts model, the corrected Akaike information criterion ($AICc$) was calculated for each voxel (Eq. 4.5) [87]. The $AICc$ takes into account both the goodness of fit of the model and the number of model parameters, and can act as a guard against overfitting. The optimal model is defined as the model with the lowest $AICc$.

$$AICc = 2k + n \ln \left(\frac{RSS}{n} \right) + \frac{2k(k+1)}{n-k-1} \quad (4.5)$$

Here, RSS is the sum of squares of the residuals, n is the number of observations (in our case, the length of the DCE-MRI sequence), and k is the number of model parameters.

4.2.2 Intra-sequence motion correction

Due to long acquisition times, DCE-MRI sequences are typically affected by patient and organ motion. At the same time, PK parameter estimation is highly sensitive to motion,

as it involves fitting the model to each voxel’s time course and assuming its spatial position does not change. Consequently, accurate pharmacokinetic parameter estimation requires intra-sequence motion correction. In this chapter, the intra-sequence motion correction was performed simultaneously with pharmacokinetic parameter estimation using software previously developed in our group [12]. A Bayesian framework enables the joint estimation of the motion parameters T and the PK model parameters $\theta = \{K^{trans}, v_e\}$:

$$P(\theta, T|X_0, Y, \sigma) \approx P(Y|X_0, T, \theta, \sigma)P(\theta)P(T) \quad (4.6)$$

Here X_0 is the baseline image, Y is the observed image sequence and σ is noise. The posterior probability, $P(\theta, T|X_0, Y, \sigma)$, can be expressed as the image prior, $P(Y|X_0, T, \theta, \sigma)$, multiplied with the prior probabilities of θ and T . The underlying assumption is that the patient motion (T) and the properties of tissue (θ) are independent. Taking the negative logarithm of Eq. 4.6, a cost function can be formulated:

$$- \ln(P(\theta, T|X_0, Y, \sigma)) = - \frac{\|T(f(\theta)) - Y\|^2}{\sigma^2} - \ln(P(T)) \quad (4.7)$$

Here, $f(\theta)$ represents the image as predicted by a set of model parameters. This cost is minimised with respect to T and θ using Gauss-Newton optimisation. In particular, the derivatives of the cost function with respect to T and θ are calculated, and the solution is iteratively updated until convergence or until the maximum number of iterations has been reached.

4.2.3 Derivation of personalised AIFs

The traditional approach involves deriving patient-specific AIFs directly from DCE-MRI intensities of in-house data, in a similar fashion to Parker et al. [115]. However, this method cannot be applied to our data due to the low temporal resolution of the sequences (9.5 seconds between volumes).

In this work we propose to derive a patient-specific AIF based on pCT data. Perfusion CT data is more suitable for AIF extraction than DCE-MRI for several reasons: it has high temporal resolution, there is a linear relationship between Hounsfield units and contrast agent concentrations, and there are no MR specific artifacts. Furthermore, there is little intra-sequence motion due to the high temporal resolution of the data (1 second between volumes) and the short total acquisition time (45 seconds). With appropriate scaling, the AIFs obtained from pCT can then be employed for PK analysis of contemporaneous DCE-MRI images of the same patients.

In order to derive the AIF from pCT data, a volume of interest (VOI) containing the femoral artery was manually selected. Signal enhancement was computed by subtracting the baseline image volume S_0 from each subsequent volume S_i , $i = 1..N$, where N is the length of the sequence. The signal enhancement was then converted to contrast agent concentration:

$$C_a(t) = k(S_i - S_0) \quad (4.8)$$

Ideally, the scaling constant, k should be measured by phantom tests using vials with varying contrast agent concentration. Since phantom data was not available for our study, we chose $k = 0.1$ based on literature [84].

This yielded the concentration time courses for each voxel, which were averaged over all the voxels in the VOI for obtaining the AIF. The obtained AIF, $C_a(t)$ is a discrete vector of 45 data points, as pCT image sequences cover only the first approximately 40 seconds after bolus arrival. This vector was padded with exponentially decaying values to extend its length up to 1.6 minutes. A pCT derived AIF obtained thus is shown in Figure 4.1, together with the corresponding AIF obtained from DCE-MRI using the approach of Parker et al. [115]. As illustrated in the figure, the DCE-MRI acquisition is too slow for the DCE-MRI AIF to be correctly sampled. To obtain an analytical form of the pCT derived AIFs and to reduce noise, Orton's bi-exponential function (Eq. 2.19) was fitted to the each of the raw AIFs and its parameters were recalculated for each individual




Figure 4.1: An example of DCE-MRI derived AIF vs. the pCT derived AIF for Patient 3. $AIF_{DCE-MRI}$ misses the peak of the uptake and has a much smaller height than AIF_{pCT} .

acquisition.

The individualised AIFs obtained from pCT were then employed for DCE-MRI pharmacokinetic analysis after appropriate normalisation. Although both MRI derived AIFs and the AIF resulting from Eq. 4.8 are typically expressed in mM , the values are still not comparable or interchangeable due to the different amounts and rates of injected contrast agent. In previous work, it has been shown that the height of the AIF is proportional to the amount of injected contrast agent [34, 84]. Thus, DCE-MRI and pCT AIFs can be made comparable by dividing each of them by the total amount of contrast agent injected in each case (given in mmol). After performing this scaling, each of the AIFs is expressed in a modality independent unit (mM/mm mol).

4.2.4 Pre-processing

A VOI of size 120×120 containing the mesorectal area was manually identified. To reduce computation time and memory requirements, each sequence was cropped down to a 4D VOI of size $120 \times 120 \times 52 \times 29$ containing the identified region, which was used for subsequent analyses.

An accurate assessment of tumour microvasculature changes requires the calcula-

tion of PK parameter values within the tumour volume. Clinical experts delineated these VOIs on the diagnostic $T2$ -weighted images because of the much better tissue contrast and resolution. In order to relate these delineations to the DCE-MRI parameter maps, the $T2$ -weighted images were registered to the baselines of the corresponding DCE-MRI sequences. Rigid registration was performed using the Image Registration Toolkit (IRTK) [137]. The obtained transformations were applied to map the masks onto the DCE-MRI image space (Fig. 4.2).




Figure 4.2: *Image registration was performed to translate the tumour VOI delineated on the diagnostic $T2$ -weighted image onto the DCE-MRI image. (Left) Slice through the $T2$ -weighted image with superimposed tumour delineation (Right) Slice through DCE-MRI image after registration, with mapped tumour delineation superimposed*

4.3 Results

4.3.1 Model comparison

The Tofts model and the 2CTU model were compared on the pre-treatment DCE-MRI images. Each of the models was fitted to the data on a voxel-by-voxel basis using Levenberg-Marquard optimisation. To compare the two models, the AIC_c was calculated based on the residual fitting error and the number of model parameters. The model with the smallest AIC_c value is considered optimal.

Figure 4.3 shows an example of the resulting parameter maps for Patient 10. For the Tofts model, the central tumour slices of the K^{trans} and v_e maps are shown, together with the corresponding AIC_c map. For the 2CTU model, the central slices of the F_p , v_p and

PS are shown, together with the corresponding AIC_c map. As illustrated in the figure, the two models provide similar fits. As the values obtained for PS are very small and F_p is very similar to K^{trans} , it can be said that the 2CTU model reduces to the Tofts model for this dataset.

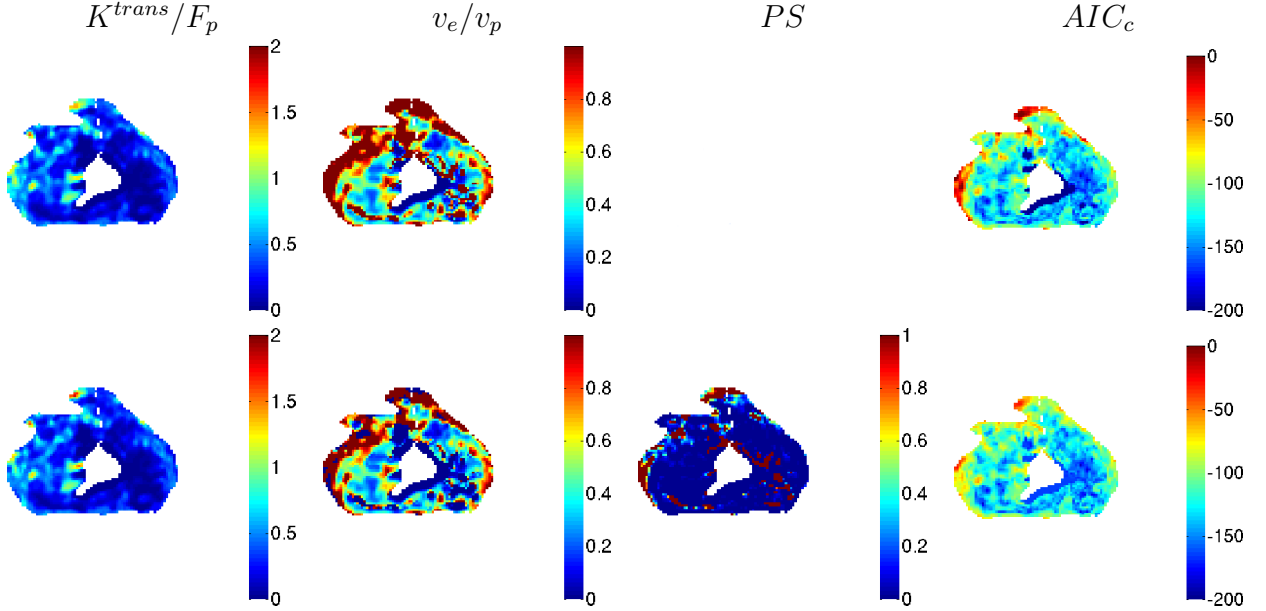


Figure 4.3: Model comparison in terms of tracer kinetic maps and the AIC_c . Top row: Tofts model parameters. Bottom row: 2CTU model parameters. Here, K^{trans} is expressed in min^{-1} , F_p and PS are expressed in $\text{mL}/\text{min}/100\text{mL}$, and v_p, v_e are fractional volumes (without unit).

The AIC_c values for the remainder of the dataset, and the median parameter values are presented in Table 4.1. As shown by the results, the AIC values of the two different models have the same order of magnitude, with the Tofts model slightly outperforming the 2CTU model (average AIC -199.33 vs. -177.22). Average F_p values are almost identical to average K^{trans} values, across the set of patients, and tissue perfusion cannot be accurately estimated (the obtained values are at the extreme boundaries of the fitting interval).

Figure 4.4 presents AIC_{min} maps based on the central tumour slice for each of the 9 patients. These maps show which is the better model for each voxel in the tumour volume, and could provide insight into tumour heterogeneity. It should be noted that the green (Tofts) and red (2CTU) regions are not randomly distributed, and that they

Table 4.1: DCE-MRI median parameter values for the tumour volume for each of the Tofts and 2CTU models. The AIC values are also listed.

Patients	AIC		K^{trans}/F_p		v_e/v_p		PS	
	Tofts	2CTU	Tofts	2CTU	Tofts	2CTU	Tofts	2CTU
P2	-128	-125	0.21	0.21	0.55	0.30	-	0.01
P3	-117	-114	0.14	0.14	0.39	0.25	-	0.01
P4	-119	-115	0.20	0.22	0.42	0.23	-	0.01
P5	-121	-119	0.44	0.49	0.51	0.39	-	0.02
P6	-146	-143	0.19	0.20	0.30	0.25	-	0.01
P7	-137	-137	0.27	0.28	0.57	0.42	-	0.02
P8	-77	-75	0.31	0.32	0.70	0.30	-	0.01
P9	-113	-112	0.32	0.33	0.66	0.55	-	0.01
P10	-116	-115	0.37	0.36	0.62	0.49	-	0.01
avg	-119.33	-117.22	0.27	0.28	0.52	0.35	-	0.01

appear to form a pattern.

4.3.2 Patient-specific AIFs

In section 4.2.3, we presented a method of calculating patient-specific AIFs from pCT data, and the normalisation employed to make these AIFs applicable to DCE-MRI images of the same patients. Intra-patient variability, inter-patient variability, and the capacity of the Orton model to fully capture the latter were further investigated in (Franklin et al.,2013). While statistical measures of these variances are presented in that paper, here we only report on two semi-quantitative measures that describe the AIF curves.

Some of the obtained AIFs are shown in Figure 4.5. On the left-hand side, the figure shows intra-patient variability of the AIFs. On the right-hand side, this figure shows the patient specific AIFs compared with the population averaged Orton curve. We note that the high inter-patient variability is not well explained by the standard Orton model, and the intra-patient variability is small.

For each of the AIF curves, peak enhancement (PE) and the area under the curve for the first 30 seconds after contrast arrival (AUC) were calculated. In Figure 4.6, intra-patient variability and variability with respect to the Orton model are represented




Figure 4.4: Central tumour slices of AIC_{min} maps for all the patients, estimated based on the pre-treatment scans. The regions where the Tofts model is better are represented in green and the regions where the 2CTU model is better are represented in red.

for all the patients using boxplots. As shown by the figures, for all the patients the intra-patient variability is smaller than the variability within the population. Moreover, the Orton model does not appear to fully capture the inter-patient variability.

In the following, the pCT derived individualised AIFs are employed for DCE-MRI PK analysis of the Sonatina trial images, and the results are compared to those obtained using the Orton model.

Figure 4.7 shows the agreement between PK parameter values obtained with the Orton AIF (AIF_{pop}) and those obtained with the pCT-derived patient specific AIF (AIF_{ind}). The high correlation coefficients for the two sets of parameters, $\rho = 0.83$ for K^{trans} , and $\rho = 0.78$ for v_e , show there is a linear relationship between the two sets of values. However, it can be seen that the usage of AIF_{pop} often leads to the over-estimation of PK parameter values.

One of the aims of the Sonatina trial is to test the clinical hypothesis that Nelfinavir alters blood flow to colorectal cancer. Preclinical results have shown that Nelfinavir induces tumour vasculature normalisation, which implies increased blood flow in that area and increased radiosensitivity [29]. We tested this hypothesis by computing K^{trans} , which is considered to be a marker of tissue perfusion, before and after treatment with

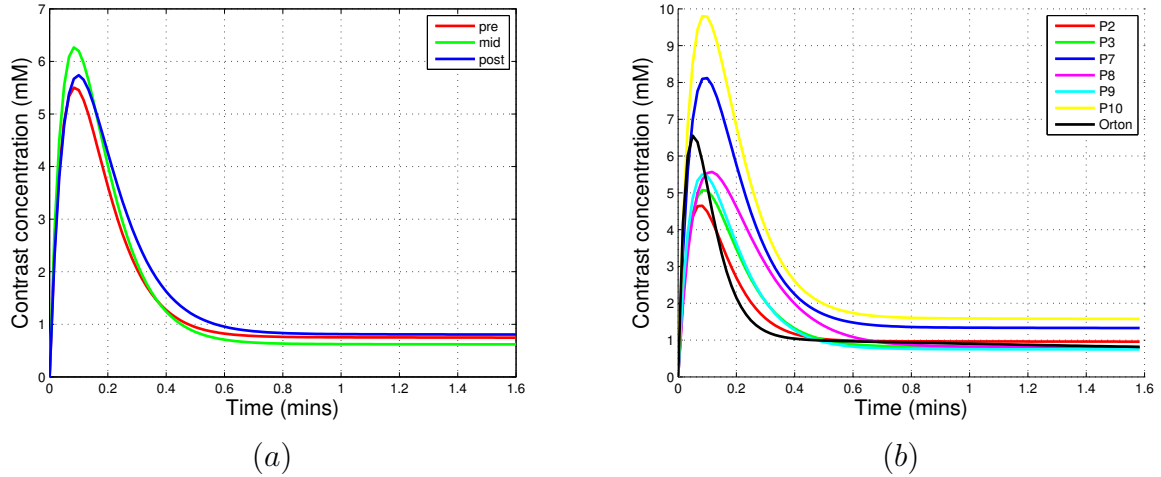


Figure 4.5: *Intra-patient variability of the AIFs vs. variability with respect to Orton: (a) Patient specific AIFs for baseline and follow-up scans of the same patient. Intra-patient variability is visibly small. (b) Patient specific AIFs for 6 different patients. The high inter-patient variability is not adequately captured by the Orton model (black)*

Nelfinavir. The images acquired after one week of Nelfinavir will be referred to as mid-treatment. K^{trans} maps were obtained using pCT-derived patient-specific AIFs and the PK estimation and motion correction strategy described in Section 4.2.

Results for two different patients are shown in Figures 4.8 (Patient 3) and 4.9 (Patient 7). One has been classified as having a good response according to MRI assessment of tumour regression grade (mrTRG) [120], the other having a poor response according to mrTRG. The mrTRG was performed on MRI images acquired 8 weeks post-therapy. In the following, we refer to patients with good response as Responders and patients with poor response as Non-responders.

Figure 4.8 shows the K^{trans} map overlapped over the anatomical image for the Responder. A slice through the VOI at pre- and mid-treatment is shown, together with K^{trans} distributions. A comparison between pre- and mid-treatment scans shows that the administration of Nelfinavir leads to an increase in K^{trans} values after treatment. This implies an increase of perfusion in the tumour volume, which is further visible from K^{trans} distributions shown in Figure 4.8 (c). Typically, tumours are characterised by a skewed distribution with a long tail towards high K^{trans} values [155]. Conventional treatment inhibits the growth of new blood vessels, resulting in shifting the histogram

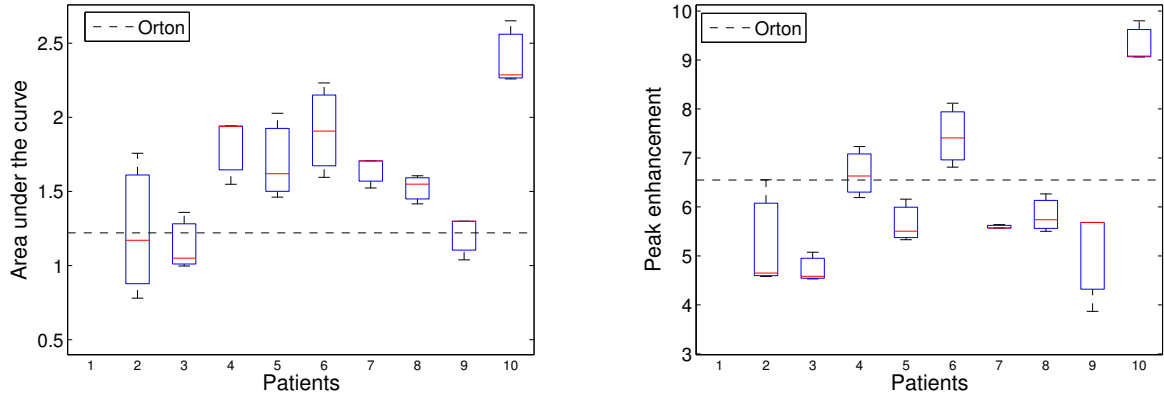


Figure 4.6: *The AIF curves are described using AUC (left) and PE (right). For all the patients, the intra-patient variability represented as a box plot is much smaller than the variability with respect to the Orton population averaged AIF (black dotted line).*

leftwards, towards low K^{trans} values. In contrast to that, Nelfinavir is expected to increase blood flow, and this treatment could be considered successful when a rightwards shift of the histogram can be observed. This can be seen in the patient illustrated in Fig. 4.8 where the mid-treatment histogram (red) has shifted towards higher K^{trans} values. This correlates with the patient being classified as a Responder.

Figure 4.9 shows the K^{trans} map overlapped over the anatomical image for a patient classified as a Non-responder according to mrTRG. Again, a slice through the VOI at pre- and mid-treatment is shown, together with K^{trans} distributions. In this case, there is no real noticeable increase in K^{trans} values when visually comparing the images; neither is there any apparent change in the pre- to mid- treatment histogram. This now correlates with the patient being classified as a Non-responder.

For the entire Sonatina dataset which consisted of 8 patients with analysable scans at mid-treatment, the shift in K^{trans} histograms correlated with mrTRG outcomes in 6 out of the 8 patients when using individualised AIFs. When using the Orton population AIF, the correlation was only observed in 5 out of the 8 patients. These 5 patients were a subset of the 6 patients for which correlation was obtained using individualised AIFs.

The post-treatment scans, which were taken after both Nelfinavir and radio-therapy, were also analysed. In this case, a mean increase of 56% in median K^{trans} was observed




Figure 4.7: The correlation between PK parameters obtained with AIF_{ind} and those obtained with AIF_{pop} . Left: Correlation of mean K^{trans} values. Right: Correlation of mean v_e values. Each data point on the graph represents the mean parameter value in the tumour for one DCE-MRI sequence. The dataset comprised 9 patients, with three scans each, and two scans were not analysable.

when using the Orton AIF, and a mean increase of 109% in median K^{trans} was observed when using the individual AIFs. As the analysed tumour volume has shrunk by 27% on average, rather than being an indicator of tumour aggressiveness, the increase in median K^{trans} is expected to be due to vascular normalisation and also inflammation after radiotherapy.

4.4 Discussion

In this chapter we have discussed several choices related to the pharmacokinetic modelling pipeline. Although the Tofts model is theoretically worse than second generation models, we cannot justify the usage of the 2CTU model for low temporal resolution data (1 volume every 9.5s), with short total acquisition time (<5 min). Individualised AIFs and their effect on PK modelling were also investigated. While the inter-patient variability is high and does not appear to be fully captured by the Orton population averaged AIF model, the intra-patient variability was found to be small for this dataset. Taking advantage of the wealth of imaging modalities present in this clinical study, we proposed deriving patient-specific AIFs from perfusion CT, and, after suitable scaling, employing

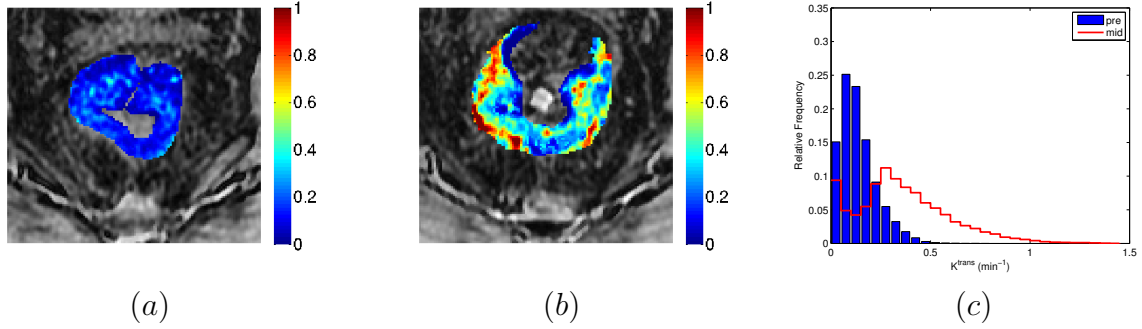


Figure 4.8: K^{trans} changes for patient 3 who was classified as a Responder according to the MRI assessment of tumour regression. (a) K^{trans} map before Nelfinavir for slice 24 through the tumour volume superimposed on the corresponding MRI slice. (b) K^{trans} map after Nelfinavir for the same slice. (c) K^{trans} histograms of the tumour volume before (blue) and after (red) treatment with Nelfinavir. The right shift of the histogram denotes an increase in perfusion as a consequence of treatment.




Figure 4.9: K^{trans} changes for patient 7 who was classified as a Non-responder according to the MRI assessment of tumour regression. (a) K^{trans} map before Nelfinavir for slice 24 through the tumour volume superimposed on the corresponding MRI slice. (b) K^{trans} map after Nelfinavir for the same slice. (c) K^{trans} histograms of the tumour volume before (blue) and after (red) treatment with Nelfinavir. Treatment does not appear to have modified much K^{trans} maps.

them for performing pharmacokinetic modelling of contemporaneous DCE-MRI scans of the same patients. Parameter estimation was performed simultaneously with intra-scan motion correction. For the first time, pCT derived AIFs were employed for DCE-MRI PK parameter estimation and compared to a population averaged model.

The analysis was performed on 9 rectal cancer patients who were imaged pre-, mid- and post-therapy. Although there was a good correlation between PK parameter values obtained with the population AIF and those obtained with individualised AIFs, using the population AIF can lead to over-estimation of PK values.

These methods were tested in the context of a novel cancer drug trial. We have at-

tempted to evaluate whether Nelfinavir alters tumour perfusion in patients with rectal cancer using DCE-MRI. Our analysis demonstrates that the administered drug, Nelfinavir, has the potential to increase blood flow in tumour tissue, acting as a radiosensitiser.

In the following, some limitations of the current study are discussed. To estimate contrast agent concentration from signal enhancement, the T_{10} relaxation time has to be known. This parameter is typically estimated based on variable flip angle sequences. However, this was not possible in this study due to sub-optimal variable flip angle sequences acquisition protocol. Consequently, we assumed a uniform T_{10} map for performing DCE-MRI analysis.

Importantly, the contrast agent (CA) injection protocols of DCE-MRI and pCT are largely different. While DCE-MRI imaging involves an injection of 0.1 mmol CA/kg body weight at a rate of 3ml/sec, pCT involves an injection of 90ml of CA at a rate of 5ml/sec. As detailed in section 4.2, we compensate for the different amounts of CA by scaling the AIFs by the total amount in each case. Since the total injection time in both cases is shorter than the blood circulation time, it is expected that the shape of the AIF will be determined by the shape of the bolus as it travels through the heart [84]. On this basis, the issue of different injection rates is currently ignored.

Furthermore, the assessment of tumour regression which is our clinical end-point, is based on MRI measurements of tumour dimension 8 weeks after the end of treatment. No histological data was available for this study, and, since it is a prospective study, the 3 year outcome is also not yet known. Moreover, due to patient recruitment problems, the study did not feature a control group, as initially intended, which makes it hard to draw a definitive conclusion about the effect of Nelfinavir.

Chapter 5

An MRF-based Discrete Optimisation Framework for Deformable Registration of Dynamic Images

Dynamic medical images are increasingly acquired in clinical practice, as they can capture both structural and functional information of the organ of interest. As these imaging techniques involve repeated scans of the volume of interest over the course of several minutes, artefacts due to patient motion are inherent. Correcting for motion in dynamic imaging scans therefore becomes an essential step in image analysis and quantification. Automated image registration techniques are increasingly used to find the optimal spatial mapping of corresponding anatomical structures. However, a robust image registration method for dynamic data has to address various challenges. Firstly, dynamic sequences can exhibit a variety of motion types, such as bulk patient motion, periodic motion caused by breathing, and random movements such as peristalsis. Secondly, contrast-enhanced modalities such as DCE-MRI and pCT pose additional challenges as image intensities change due to contrast inflow. Lastly, there is inherent acquisition noise.

In this chapter, an MRF-based discrete optimisation framework for the nonrigid registration of dynamic sequences is proposed. This attempts to address the shortcomings of

current registration methods for dynamic images, which either make restrictive assumptions about the motion pattern, or employ continuous optimisation which is sensitive to initialisation and prone to local optima. The proposed framework can capture both large and small motion, can be used in conjunction with different similarity metrics, and is generalisable across imaging modalities. Moreover, it can be applied to both contrast-enhanced and non contrast-enhanced data. The main contribution of this work is the extension of the **deeds** [64] framework for the registration of 4D dynamic sequences. Two inference strategies are proposed and compared: TRW-S message passing is used to find the optimum labelling on the full 4D graph. Loopy belief propagation is used to find the optimum labelling on the reduced 4D graph, which is obtained as a concatenation of minimum spanning trees (MST). The reduced 4D graph has the advantage of increased computational efficiency. To address the problem of intensity differences in dynamic contrast-enhanced data, a DCE-MRI tailored similarity metric which incorporates PK modelling information is formulated.

5.1 Introduction to Markov Random Fields

Graphical models provide a simple way of visualising the structure of a probabilistic model. Such a representation can offer insight into the properties of the model, is used to design and motivate new models, and can facilitate complex computations which reduce to simple graphical manipulations [15]. A graph $\mathcal{G} = (V, E)$ comprises a set of nodes V (also called vertices), and a set of edges that connect these nodes, E . Each node in the graph V_i corresponds to a random variable X_i , and the links represent probabilistic relationships between these variables. Each random variable can take values $l_i \in \mathcal{L}_i$, where \mathcal{L}_i is the set of possible labels. The assignment of label l_i to a certain node X_i is an event with the probability $p(X_i = l_i)$, or simply $p(l_i)$.

A random field X is defined as an undirected graphical model. A labelling of the random field is denoted by $\mathbf{l} = (l_1, l_2 \cdots l_n)$. The probability of assigning configuration

\mathbf{l} to the random field is denoted by $p(\mathbf{l})$. The set of all the possible configurations is denoted by \mathcal{X} . Before introducing the concept of MRFs, a couple of properties related to graph neighbourhoods have to be discussed. Two nodes are considered to be neighbours if there is an edge connecting them. A clique is a subset of nodes where every node is a neighbour of all other nodes. In other words, a clique can either be a single node or a fully connected subgraph. The order of a random field is given by size of the maximum clique minus one. Thus, a random field is said to be an MRF, provided that two conditions are fulfilled. First, the probability of any configuration $\mathbf{l} \in \mathcal{X}$ has to be positive: $p(\mathbf{l}) > 0, \mathbf{l} \in \mathcal{X}$. Second, a node depends only on its neighbours and is independent of the other nodes: $p(l_i | l_j, j \in N_i) = p(l_i | l_j, j \in V - \{i\})$. In this work, only pairwise interactions are considered, which is equivalent to a first order MRF.

According to the Hammersley-Clifford theorem [55], the joint probability of a labelling in an MRF, $p(\mathbf{l})$, can be written as a product of potential functions $\Psi_C(l_C)$ over the maximal cliques in the graph:

$$p(\mathbf{l}) = \frac{1}{Z} \prod_C e^{-\Psi_C(l_C)} \quad (5.1)$$

Here, Z is a normalisation factor, $Z = \sum_{\mathbf{l}} \prod_C e^{-\Psi_C(l_C)}$, which ensures that the quantity in Eq. 5.1 is a probability. The problem of finding the configuration \mathbf{l} that maximises the probability $p(\mathbf{l})$ can be converted to an energy minimisation problem [49]. By taking the negative logarithm of Eq. 5.1, we obtain:

$$\begin{aligned} E(\mathbf{l}) &= -\ln \left(\frac{1}{Z} \prod_C e^{-\Psi_C(l_C)} \right) \\ &= \sum_C \Psi_C(l_C) + const \end{aligned} \quad (5.2)$$

Thus, the energy can be expressed as the sum of the potential function over all cliques plus a constant. The labelling $\hat{\mathbf{l}}$ that minimises the energy in Eq. 5.2 can be found as $\hat{\mathbf{l}} = \operatorname{argmin}_{\mathbf{l}}(E(\mathbf{l}))$. Applying the maximum a posteriori (MAP) principle, the configuration

which minimises the posterior energy $E(\mathbf{l}|\mathbf{y})$ can be calculated:

$$\begin{aligned}\hat{\mathbf{l}} &= \operatorname{argmin}_{\mathbf{l}}(E(\mathbf{l}|\mathbf{y})) \\ &= \operatorname{argmin}_{\mathbf{l}}(E(\mathbf{y}|\mathbf{l}) + E(\mathbf{l}))\end{aligned}\tag{5.3}$$

Here, $E(\mathbf{y}|\mathbf{l})$ represents the likelihood energy and $E(\mathbf{l})$ is the prior. In the following section, the formulation of image registration as an MRF labelling problem is briefly presented. A comprehensive description can be found in [49].

5.2 Registration as an MRF labelling problem

As mentioned in Chapter 3, continuous optimisation methods are highly dependent on initialisation and can get trapped in local optima if they are gradient based. This problem can be addressed by discretising the search space. Discrete optimisation is typically formulated as a MRF labelling problem [50]. For deformable image registration purposes, a graph is defined in which the nodes $p \in V$ represent voxels or groups of voxels, and the edges connect voxels with similar anatomical features and spatial proximity. For every node p , there is a set of labels $l_p \in \mathcal{L}$ which represent possible discrete displacements of the source image volume with respect to the target image volume. Each label l_p corresponds to a displacement vector $\mathbf{u}_p = \{u_p, v_p, w_p\}$. Finding the optimum displacement at each voxel equates to finding the labelling that minimises the MRF energy function:

$$E(l) = \sum_{p \in V} C_D(l_p) + \lambda \sum_{(p,q) \in \mathcal{N}} C_R(l_p, l_q)\tag{5.4}$$

The unary term represents the data cost C_D , which measures the similarity of a voxel in the target image to the corresponding voxel in the source image displaced with l_p . For a node p , the unary cost of assigning label l_p is independent of the labelling of its neighbours. The pairwise term represents the regularisation cost $C_R(l_p, l_q)$ and is used to smooth the displacements of directly connected voxels $(p, q) \in \mathcal{N}$. Here, \mathcal{N} represents

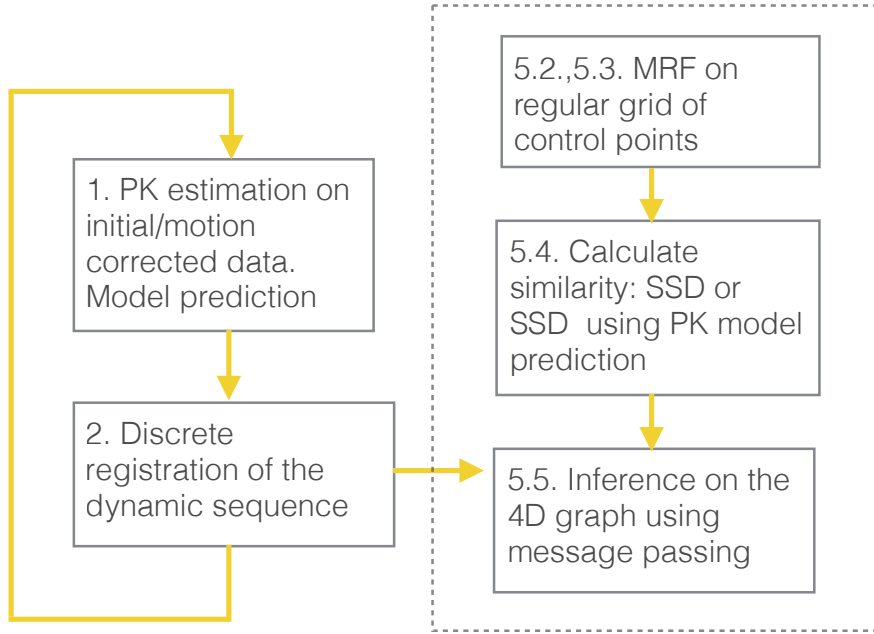


Figure 5.1: *The proposed discrete registration and pharmacokinetic estimation framework for dynamic contrast-enhanced data. For non contrast-enhanced data, Step 1 can be skipped.*

the neighbourhood of a voxel, and λ is a weighting factor that regulates the amount of regularisation.

If the chosen similarity metric is SSD, and the deformation field is regularised using the squared difference of displacements of neighbouring points, which is an approximation to the diffusion regularisation discussed in Chapter 3, Eq. 5.4 becomes:

$$E(l) = \frac{1}{N} \sum_{p \in V} (I_t(\mathbf{x}_p) - I_s(\mathbf{x}_p + \mathbf{u}_p))^2 + \lambda \sum_{(p,q) \in \mathcal{N}} \frac{\|\mathbf{u}_p - \mathbf{u}_q\|^2}{\|\mathbf{x}_p - \mathbf{x}_q\|} \quad (5.5)$$

In the following, the parameterisation of the transformation model (Section 5.3) and the choice of similarity metric, for both non-contrast enhanced and contrast enhanced data (Section 5.4) are presented. We then discuss two inference strategies that are used to solve the MRF labelling problem (Section 5.5). An overview of the proposed framework is shown in Fig. 5.1.

5.3 Transformation model: Parameterisation and properties

The image domain is subdivided into non-overlapping cubic groups of voxels using a regular grid of control points. Each control point represents a group of voxels and corresponds to a node in the MRF. The similarity cost which is associated to moving a control point in the grid is obtained by aggregating the costs of individual voxels corresponding to that control point (only a subset of K voxels are used, as explained in the next section). The regularisation cost is calculated only between control points.

A multi-level approach is adopted for control point spacing, so that the control points are sparser for coarser levels and denser for finer levels. Motion is first estimated on the coarse level with a large capture range. At finer levels, the previous solution is upsampled using B-spline interpolation and used as the prior deformation for the optimisation at that level. Using this refinement approach, different types of motion can be captured and low computational complexity is achieved. The label space is given by $\mathcal{L} = q\{0, \pm 1, \pm 2, \dots, \pm \frac{u_{max}}{2}\}^3$, where q is the quantisation step and $u_{max} * q$ is the maximum displacement on each direction. As the number of nodes increases, the maximum displacement in the label space $u_{max} * q$ is decreased correspondingly, to keep the complexity constant for all the resolution levels. The final (dense) solution is also obtained by upsampling using a first order B-spline interpolation.

Diffeomorphism is a desired property for the deformation field, as it prevents transformations that are not physically feasible, such as folding. This is particularly important for soft tissue imaging, as it agrees with the tissue incompressibility assumption. If we treat the deformation field as a velocity field, it can be transformed into a diffeomorphic mapping [64] by using the scaling and squaring method [6].

5.4 Similarity metric calculation

In the case of non-contrast enhanced dynamic images, SSD (Chapter 3) can be used as a similarity metric. To improve computational efficiency, a dense stochastic sampling was employed for calculating the similarity metric, following the approach proposed by Heinrich et al. [64]. This involves calculating the similarity term for a control point using only a random subset of all the voxels corresponding to that point.

As mentioned above, the DCE-MRI intensities of different volumes in the image sequence are not comparable using standard similarity metrics such as the sum of squared differences (SSD). To address this issue, we propose to compare the similarity between the intensities at the volume to be registered I_{t_i} , and the PK model predicted intensity at that volume $PK(I_{t_0}, K^{trans}, v_e, t_i)$:

$$C_D(l_p) = SSD(I_{t_i}(\mathbf{x}_p + \mathbf{u}_p), PK(I_{t_0}(\mathbf{x}_p), K^{trans}(\mathbf{x}_p), v_e(\mathbf{x}_p), t_i)) \quad (5.6)$$

Here, t_i is the current time point of the DCE-MRI sequence, t_0 is the first time point, and K^{trans} and v_e are parameters of the PK model. To predict the appearance of the baseline I_{t_0} at t_i , an initial estimate of the PK parameters K^{trans} , v_e is required. This initial estimate is obtained by least squares (lsq) fitting of the model to the data.

PK parameter estimation was performed using the standard Tofts model [143] with the Orton [112] bi-exponential arterial input function with population averaged parameters, which were described in Chapter 4.

5.5 Inference strategies

The optimal labelling that minimises Eq. 5.4 can be found using an inference strategy. As mentioned in Chapter 3, inference strategies fall in one of two categories: graph-cuts or message passing approaches. To date, several studies have compared the two types of approaches [74, 138], and the conclusions vary according to the application on

which the algorithms were tested. Generally, α -expansion outperforms other graph-cuts variants. Also, TRW-S, a message passing variant with sequential schedule, is sometimes able to find the (almost) global optimum when graph-cuts methods fail. An important limitation of traditional graph-cuts algorithms such as α -expansion is their applicability to submodular functions only [18]. This category of functions does not include the quadratic difference of displacements, which is often used as regulariser for registration and is not a metric. Newer approaches from the graph-cut family such as quadratic pseudo-boolean optimisation (QPBO) [17, 130] and FastPD [82] relax the metric requirements, and are applicable to energy functions with non-submodular terms as well. However, while QPBO’s applicability depends on the number of unlabelled nodes it yields, the improved performance and speed of FastPD come at the cost of high memory requirements.

This work focuses on inference strategies that fall into the message passing (belief propagation) category. Message passing strategies involve iteratively passing messages between nodes until the belief about the solution stabilises. A message m is a vector whose size is given by the number of possible labels, \mathcal{L} . Let p and q be two neighbouring nodes. Then, $m_{p \rightarrow q}$ is the message that node p sends to node q , and represents the belief at node p about the label assignment at node q . According to the *max-product*/*min-sum* algorithm [151], messages are generated using the following formula:

$$m_{p \rightarrow q}^t(l_q) = \min_{l_p} \left(C_D(l_p) + C_R(l_p, l_q) + \sum_{s \in \mathcal{N}(p) - q} m_{s \rightarrow p}^{t-1}(l_p) \right) \quad (5.7)$$

Here, $m_{p \rightarrow q}^t$ is the message sent by node p to node q at iteration t , $\mathcal{N}(p) - q$ represents the neighbourhood of p excluding node q , and $m_{s \rightarrow p}^{t-1}$ are the previously passed messages from nodes other than q . Again, $C_D(l_p)$ is the unary cost based on the image likelihood, and $C_R(l_p, l_q)$ is the pairwise prior. Practically, given label l_q at node q , node p has to decide which is the most compatible label for itself, l_p . For this, the likelihood of that label, $C_D(l_p)$, and the messages coming in from other nodes about that label are also taken into account. After a number of iterations, N , the belief at each node can be calculated as

the sum between the data cost and the incoming messages at that node:

$$b_q(l_q) = C_D(l_q) + \sum_{p \in \mathcal{N}(q)} m_{p \rightarrow q}(l_q) \quad (5.8)$$

The passing of messages is usually done according to a set schedule. On tree-like graphs, message passing is equivalent to dynamic programming. Only two passes are required to find the optimal solution, one from leaf nodes towards the root node, and the other from the root node towards the leaves. The generalisation of message passing for graphs containing loops was termed loopy belief propagation [151]. This algorithm involves iteratively passing messages from all the nodes in parallel.

In the following, two message passing schedules that are employed in this thesis are discussed. First, we describe TRW-S, which was used for optimising the full 4D graph. Then, a reduced graphical structure is proposed, which consists of a concatenation of MSTs [126]. This graph is optimised using a method akin to loopy belief propagation (LBP) [46], with the difference that a sequential message passing schedule is used.

5.5.1 Tree-reweighed message passing on the full 4D graph

An important development in the field of message passing inference strategies is the TRW-S algorithm [80, 148]. This algorithm involves LP relaxation of the initial problem by decomposing the original loopy graph into a convex combination of trees. The labelling is solved for each individual tree and the overall solution is obtained by averaging the individual results. Thus, instead of minimising the initial energy function, this algorithm involves the maximisation of a lower bound on the energy. The gap between the lower bound and the actual energy of the final solution gives an indication of how far we are from a global optimum.

The sequential message passing schedule proposed by Kolmogorov [80] is presented in Fig. 5.2 for a 2D regular grid graph. Given an ordering of the nodes, monotonically increasing chains are chosen as trees (i.e. rows in the given example). For the forward




Figure 5.2: *TRW-S message passing schedule. (a) Forward pass (b) Backward pass*

pass, the nodes are processed in increasing order and the messages from lower neighbours are also passed. For the backward pass, the same operations are performed in reverse order. This schedule reduces the memory requirements by half, as only one message per edge is passed at a time. For the complete 4D graph resulting from our optimization problem, an analogous message passing schedule can be applied. For a chosen node, p , the messages are passed on chains on the x , y , z and t dimensions.

5.5.2 Message passing on the reduced 4D graph

To reduce the computational costs which are typically associated with discrete optimisation of a full 4D graph, a graph with reduced connectivity is proposed. At each volume, the MST which best replicates the underlying anatomy is calculated using Prim’s algorithm [126], akin to the work of Heinrich et al. [64]. Starting from the six-connected neighbourhood, an MST is obtained by discarding the least relevant edges until a tree that spans all the vertices is obtained. Edges are considered to be relevant if they connect nodes with similar intensity information.

The MST topology is assumed to be identical in all the subsequent volumes, as they have the same anatomy as the baseline and differ only by contrast enhancement or motion. To enforce temporal smoothness constrains, additional links are added to connect

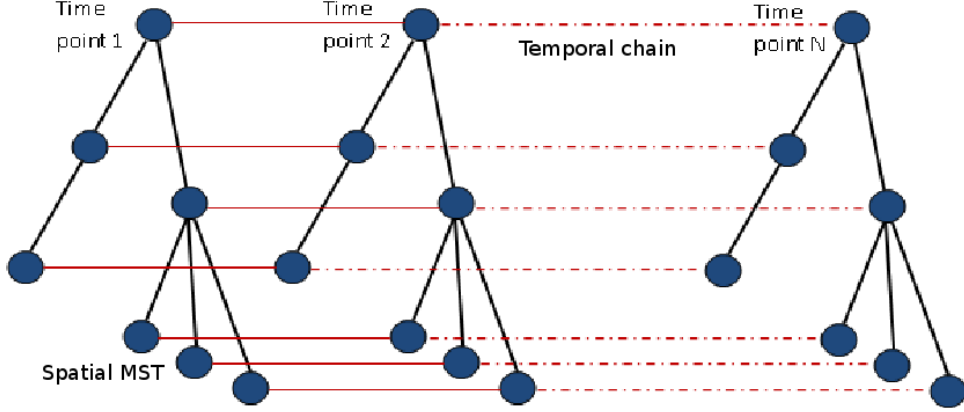


Figure 5.3: An illustration of the reduced graphical structure used in this algorithm. While the spatial connectivity is captured by minimum spanning trees (MST), the temporal continuity is captured by temporal chains.

temporally corresponding nodes across time. These nodes represent the same anatomical landmark over time, and the trajectory of this landmark is expected to be reasonably smooth. The reduced graph obtained by temporally connecting the identical MSTs representing the $n - 1$ volumes of a dynamic sequence is shown in Figure 5.3.

For each individual tree, the optimisation is possible using belief propagation on the MST [65]. At each node p , the message m_p that the node sends to its parent, q , can be found given the messages from its children c , m_c :

$$m_p^{tree}(l_q) = \min_{l_p} \left(C_D(l_p) + \lambda_{sp} C_R(l_p, l_q) + \sum_c m_c(l_p) \right) \quad (5.9)$$

For leaf nodes, Eq. 5.9 can be evaluated directly, as there are no incoming messages from children. The messages are then forward-passed from the leaves to the root node, and then backward-passed, from the root towards the leaves. For each individual temporal chain, the optimisation can be performed using belief propagation on the chain [46]. The message of the current node p , m_p , can be found given the displacement of the previous node q , l_q , and the message coming from the subsequent node r , $m_r(l_p)$:

$$m_p^{chain}(l_q) = \min_{l_p} (C_D(l_p) + \lambda_{temp} C_R(l_p, l_q) + m_r(l_p)) \quad (5.10)$$

To optimise the reduced 4D structure, for each node we perform independent message passing on each temporal chain and on each spatial tree, and average the resulting marginals. This procedure is repeated for a number of iterations akin LBP [45], where the messages towards a node in the previous iteration are added to the marginal of that node in the current iteration. The best displacement can then be found by calculating $\arg \min$ based on the marginal distribution. Although this approach is not guaranteed to converge to a global optimum, it is physically motivated, and provides a good trade-off between optimality and efficiency. The averaging of marginals from multiple graph models has previously been used in stereo processing using an approach called semi-global matching [42, 68], yielding excellent results. The MST-based proposed algorithm is detailed in Algorithm 2. The data cost is calculated using the previously estimated PK parameters, $\theta = \{K^{trans}, v_e\}$ as explained in Section 5.4. The pseudo-marginals for chains and trees are denoted by b^{chain} and b^{tree} , respectively, and the messages are denoted by m . In each optimisation step (Algorithm 2, 2.2, 2.3), the full distribution of pseudo-marginals for the space of displacements is estimated.

5.5.3 Diffusion regularisation

Naively calculating the pairwise regularisation cost in Eq 5.5 would require $|\mathcal{L}|^2$ calculations for every voxel. To reduce this cost, we employ the min-convolution technique [45], which reduces the complexity to $|\mathcal{L}|$. If we consider a 1D optimisation problem, for the diffusion regularisation term, the cost of each label l_p can be described by an upward parabola rooted at $(l_p, h(l_p))$ (Fig 5.4). Here, $h(l_p)$ represents the belief at node l_p (the data cost plus the previously passed messages). Solving the minimisation problem is then equivalent to finding the lower envelope of these parabolas. The algorithm for finding the lower envelope is detailed in [45]. The min-convolution technique can be extended to multidimensional images by applying it to each individual dimension in a fashion similar to separable convolution.

Algorithm 2 DireP: Discrete registration and Pharmacokinetic estimation

```
1: while  $n_{iterPK}$  do
2:   1. PK parameter estimation
3:   2. Registration
4:   Initialise  $\mathbf{u} \leftarrow 0$ 
5:   for  $level = 0, 1, 2$  do ▷ Three resolution levels
6:     procedure DATACOST( $I, \theta, \mathbf{u}$ )
7:     procedure REGULARISATION
8:       2.1 Initialise marginals and messages:
9:       for all nodes do  $b^{chain}[node], b^{tree}[node] \leftarrow C_D(l_p); m[node] \leftarrow 0;$ 
10:      while  $n < n_{iterMRF}$  do
11:        2.1 Re-initialise marginals: ▷ With average of previous values
12:        for all nodes do
13:           $b^{chain}[node] \leftarrow (b^{chain}[node] + b^{tree}[node])/2;$ 
14:           $b^{tree}[node] \leftarrow b^{chain}[node];$ 
15:        2.2 Pass messages on spatial MST:
16:        for  $node = leaves, root - 1$  do ▷ Forward pass
17:           $cost \leftarrow b^{tree}[node];$ 
18:           $m[node] \leftarrow \min - \text{sum}(cost);$  (see Eq. 5.9)
19:           $b^{tree}[parent] \leftarrow b^{tree}[parent] + m[node];$ 
20:        for  $node = root - 1, leaves$  do ▷ Backward pass
21:           $cost \leftarrow b^{tree}[parent] - m[node] + m[parent];$ 
22:           $m[node] \leftarrow \min - \text{sum}(cost);$ 
23:        for all nodes do  $b^{tree}[node] \leftarrow b^{tree}[node] + m[node];$ 
24:        2.3 Pass messages on temporal chain:
25:        for  $t = tdim - 1, 1$  do ▷ Forward pass
26:           $cost \leftarrow b^{chain}[node_t];$ 
27:           $m[node_t] \leftarrow \min - \text{sum}(cost);$  (see Eq. 5.10)
28:           $b^{chain}[node_{t-1}] \leftarrow b^{chain}[node_{t-1}] + m[node_t];$ 
29:        for  $t = 1, tdim - 1$  do ▷ Backward pass
30:           $cost \leftarrow b^{chain}[node_{t-1}] - m[node_t] + m[node_{t-1}];$ 
31:           $m[node_t] \leftarrow \min - \text{sum}(cost);$ 
32:        for all nodes do  $b^{chain}[node_t] \leftarrow b^{chain}[node_t] + m[node_t];$ 
33:      procedure DIFFEOMORPHIC( $\mathbf{u}$ )
```




Figure 5.4: Lower envelope computation for efficient computation of diffusion regularisation. Figure taken from Felzenszwalb and Huttenlocher [45]

5.6 Experiments on simulated data

5.6.1 Experiments

The **D**iscrete **r**egistration and **P**harmacokinetic estimation (DireP) algorithm proposed in this thesis was first tested on synthetic data, where the ground truth intra-sequence motion and PK parameters are known. To simulate a realistic dataset, synthetic images were generated as follows: 8 real DCE-MRI sequences were selected, and PK model fitting was performed on each of them. Using the resulting parameter maps and the pre-contrast image volumes, 8 model predicted sequences of size $120 \times 120 \times 52 \times 29$ each were generated. These images constituted the ground truth motion free dataset. To simulate motion, a random displacement field of size $6 \times 6 \times 4$ was generated for each image volume independently. This displacement was upsampled to the image volume size, and smoothed with a Gaussian kernel. We also applied Gaussian smoothing on the temporal dimension, as the motion of each voxel is expected to have some degree of temporal smoothness due to periodic motion patterns i.e. breathing. The parameters for the synthetic motion were chosen to obtain a diffeomorphic deformation field in the interval $\pm 7\text{mm}$ along each direction.

In order to evaluate the advantages of the proposed framework, we compared it to a state-of-the-art algorithm for joint PK estimation and non-rigid intra-sequence registration which is based on continuous optimisation (Method 1) [12]. In [12], Method 1 was evaluated against non-rigid image registration using similarity metrics which do not take into account pharmacokinetic modelling information. MI, SSD and NCC were employed for the comparison, and Method 1 was shown to outperform all of them. For this reason, in this work we focus on the comparison of discrete and continuous optimisation, and do not undertake the comparison of similarity metrics.

Three flavours of the proposed algorithm were used to register the sequences:

- **TRW-S:** Optimisation was performed on the full 4D graph using TRW-S, as detailed in Section 5.5.1.
- **DireP 3D:** Optimisation was performed on the MST-based reduced graph, as presented in Algorithm 2, with the temporal weights λ_{temp} set to 0. The registration step becomes equivalent to the **deeds** algorithm [64], with the difference that we do not ensure transformation symmetry.
- **DireP 4D:** Optimisation was performed on the 4D MST-based reduced graph, as presented in Algorithm 2.

For both the synthetic and the real data (Chapter 6), the label space was chosen as $L = \{0, \pm q, \pm 2q, \dots, \pm \frac{u}{2}q\}^3$. Three resolution levels were used in the registration. Depending on the resolution level, u was 8, 6, 4, with a quantisation q of 2, 1, 0.5mm. This accounts for a maximum possible deformation of ± 8 mm, which is sufficient for the synthetic data we generated. Depending on the resolution level, groups of voxels of sizes 8^3 , 6^3 and 4^3 were used, with corresponding label spaces of size 9^3 , 7^3 and 5^3 . The experiments were run with $n_{iterPK} = 2$, $n_{iterMRF} = 5$ (see Algorithm 2), $\lambda_{sp} = 0.01$ and $\lambda_{temp} = 0.1$, which were empirically chosen.

Table 5.1: Registration results on synthetic data. The average target registration error (TRE) and residual fitting errors in PK parameters are reported, together with the corresponding standard deviations.

Measure	Before	Method1	TRW-S	DireP 3D	DireP 4D
TRE (mm)	1.86 ± 0.94	1.42 ± 0.91	0.91 ± 0.72	0.85 ± 0.56	0.87 ± 0.56
Error in k_{ep}	0.39 ± 0.03	0.15 ± 0.01	0.14 ± 0.01	0.14 ± 0.01	0.14 ± 0.01
Error in v_e	0.20 ± 0.01	0.19 ± 0.01	0.04 ± 0.003	0.03 ± 0.002	0.03 ± 0.002

5.6.2 Results

Quantitative results on the synthetic dataset are presented in Table 5.1. The target registration error (TRE), which was defined in Chapter 3, is here the average difference between the ground truth deformation field and the deformation field estimated by the algorithms. DireP 3D and DireP 4D were compared to TRW-S and Method1 [12]. The results show that all the discrete methods outperform Method1 and recover a good part of the synthetic motion. At the same time, the inference methods that use a reduced graph, DireP 3D and DireP 4D, give better results than TRW-S, which uses the full graph. DireP 3D has a slightly better performance than DireP 4D in terms of TRE. We expect this to be due to trading a solution that is globally optimal for each volume (DireP 3D) for a sub-optimal solution on the entire 4D graph (DireP 4D). However, as also noted by Metz et al. [101] who obtain the same trend in their results, temporal smoothness of the transformation is in itself desirable, even if it involves somewhat compromising on registration accuracy.

The Jacobian determinant of the recovered deformation fields was also calculated. We found that there were no negative Jacobian values for neither DireP 3D or DireP 4D. We also calculated the standard deviation of the Jacobian, $std(Jac)$, which is an indicator of transformation smoothness. Smaller standard deviations indicate smoother transformations, and $std(Jac) = 0$ indicates a perfectly smooth transformation. We obtained an average value of $std(Jac) = 0.153$ for DireP 3D and $std(Jac) = 0.143$ for DireP 4D. This result is in agreement with the fact that the 4D method enforces temporal smoothness of the deformation field.

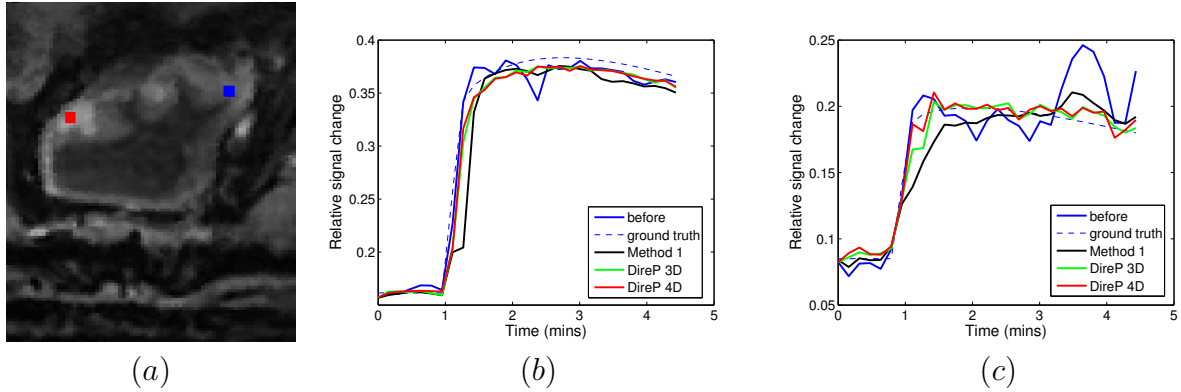


Figure 5.5: Voxel plots before and after motion correction are shown. Plot (b) corresponds to the voxel represented in red and plot (c) corresponds to the voxel represented in blue. *DireP 3D* and *DireP 4D* outperform the continuous state-of-the-art, *Method1*, yielding smoother voxel plots.

After applying each method, the PK parameters on the corrected datasets were generated by least-squares fitting. The errors in k_{ep} and v_e are defined as the average absolute difference between the ground truth PK parameters and the parameters obtained on the corrected datasets. These values were calculated on a circular region delineated around the tumour in the pre-contrast image volume (which was taken from a real dataset as explained above). The PK parameter errors decrease considerably after *DireP 3D* and *DireP 4D* registration.

Figure 5.5 shows the relative signal intensity change over time for two voxels, before registration, and after registration using *Method1*, *DireP 3D* and *DireP 4D*. As shown by the figure, *DireP 3D* and *DireP 4D* outperform the continuous state-of-the-art, *Method1*, yielding smoother voxel plots.

The computational efficiency of the different methods was also compared. On a four core Intel i7 CPU running at 1200MHz, using C++ code, runtimes are: *Method1* 14.29min, *DireP 3D* 3.15min, *DireP 4D* 7.02min. Both *Method 1* and the discrete methods were run with three resolution levels. The times reported for *DireP 3D* and *DireP 4D* are those required for one complete registration step (Algorithm 2, Step 2), and *DireP 4D* was run with $n_{iterMRF} = 5$.

The relationship between the algorithms' performance and the level of motion in the

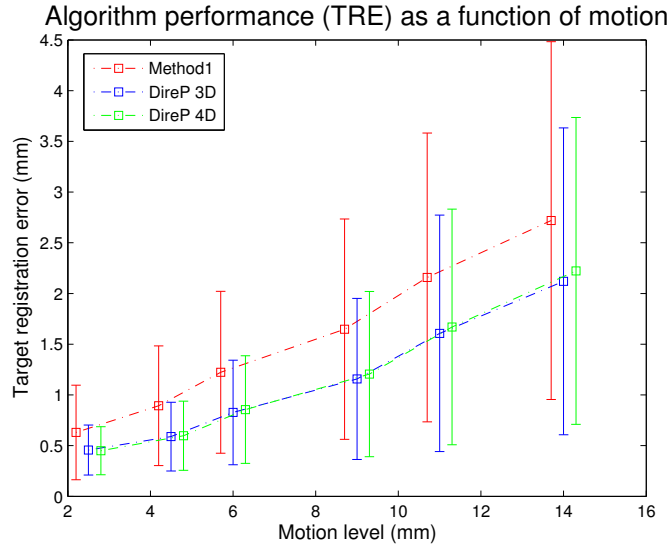


Figure 5.6: The TREs (mm) for the three different algorithms are shown as a function of the motion level. The performance of Method1 is represented with red error bars, DireP 3D is represented in blue, and DireP 4D is represented in green. The coloured error bars are slightly displaced for better visualisation, but they correspond to the same level of motion.

synthetic data was also investigated, and the results are shown in Fig. 5.6. It can be seen that DireP 3D and DireP 4D outperform Method1 regardless of the motion level, and the difference is more pronounced for bigger motion.

5.7 Discussion

We have proposed a new framework for dynamic image registration and PK parameter estimation (DireP), which is formulated on an MRF and uses discrete optimisation. This approach comes to address the shortcomings of algorithms involving continuous optimisation, which are sensitive to initialisation and prone to local optima. The PK estimation is performed iteratively with the deformable registration, and this step can also be skipped when dealing with non contrast-enhanced data. Two inference strategies were proposed: TRW-S message passing was used for optimising the full 4D graph, and loopy belief propagation was used to find the optimum labelling for the reduced 4D graph, which is built as a concatenation of MSTs. The latter strategy constitutes the 4D extension of **deeds** [64]. When tested on synthetic data, both the variants of DireP, and the TRW-S

optimisation on the full graph show an improvement over a state-of-the-art continuous optimisation [12], in terms of both TRE and computational efficiency. Although DireP $3D$ slightly outperforms DireP $4D$ in terms of TRE (0.85mm vs. 0.87mm), the key advantage of DireP $4D$ lies in its ability to impose a degree of temporal smoothness, which in itself is desirable to preserve temporal continuity or to avoid unrealistic fitting of the PK model. At the same time, high frequency motions such as peristalsis cannot be captured by a transformation model that enforces temporal smoothness. In the next chapter, we show that the positive effect of the $4D$ regularisation is much more pronounced for the registration of dynamic lung images, where low frequency motion (i.e. breathing) is predominant.

Chapter 6

Application to Cancer

In this chapter, two clinical applications of the proposed nonrigid registration framework are presented: the intra-sequence motion correction of DCE-MRI images from colorectal cancer patients and the registration of dynamic CT images of lung cancer. While the former is performed together with PK parameter estimation on data from the Sonatina clinical trial (Section 4.1), the latter is performed on a publicly available dataset (DIR-lab [25]).

6.1 DCE-MRI pharmacokinetic estimation and intra-sequence registration for colorectal cancer

As mentioned in Chapter 1, the main clinical challenge in early stage clinical trials of cancer treatment is early prediction of patient response. This would allow tailoring subsequent treatment accordingly, and avoiding unnecessary surgery or CRT cycles, which have a dramatic impact on patient quality of life. As molecular and cellular changes precede macroscopic tumour growth, these changes can act as early indicators of treatment outcome. The tumour microenvironment changes during a course of treatment can be measured non-invasively using medical imaging. DCE-MRI derived biomarkers have previously been used as predictors of radiotherapy or chemotherapy re-

sponse [159], [113], [41], [12]. More specifically, DCE-MRI signal intensity changes due to tracer in- and out-flow indicate perfusion and permeability characteristics of the underlying tumour tissue.

The DCE-MRI parameters of interest are typically derived using PK modelling as explained in Chapters 2 and 4. This is a particularly challenging problem, as observed signal changes throughout the DCE-MRI time series can be either due to motion or due to contrast enhancement. Besides choosing an appropriate model, accurate PK modelling requires a strategy to deal with patient and organ motion. Previous approaches use image registration to correct for intra-sequence motion. However, they either assume a restrictive transformation model [22], or are limited by assumptions about the type of motion present in the data [56, 99]. While Melbourne et al. [99] are able to correct for peristaltic motion, but not for breathing, Hamy et al. [56] are able to correct for periodic motion such as breathing, but do not correct for peristalsis. Recently proposed in our group, the work of Bhushan et al. [12] addresses a number of these limitations (Section 4.2.2). However, they employ continuous optimisation, making the method highly dependent on initialisation and prone to local optima.

In this section, the discrete image registration and pharmacokinetic estimation (DireP) framework proposed in Chapter 5 is applied to colorectal cancer data, with the goal of correcting for intra-sequence motion and extracting more accurate PK parameter maps.

6.1.1 Experiments

Both variants of the proposed algorithm, DireP $3D$ and DireP $4D$, were applied on the pre-, mid- and post-treatment images of the 8 analysable patients from the Sonatina clinical trial (Section 4.1). As mentioned in Section 4.1, the post-treatment scan of Patient 8 was not analysable due to inadequate contrast enhancement. Patients 1 and 5 were also excluded as they did not have a mid-treatment scan, and a further two scans were compromised by extravasation of contrast agent or inadequate enhancement. A region of interest of $120 \times 120 \times 52 \times 29$ which contained the tumour volume was




Figure 6.1: *The MST corresponding to a 2D slice from a DCE-MRI image of the mesorectum. The leaf nodes are represented in green, and the root is represented in yellow.*

selected for running the algorithms. The tumours were manually delineated by our clinical collaborators.

The experiments were run using $n_{iterPK} = 3$, $n_{iterMRF} = 5$, $\lambda_{sp} = 0.01$ and $\lambda_{temp} = 0.1$, which were empirically chosen. Four resolution levels were used for the registration. Depending on the resolution level, the search space was $u \in \{8, 6, 4, 4\}$, with a corresponding quantisation $q \in \{2, 1, 0.5, 0.5\}$ mm, which amounted to a maximum possible deformation of ± 8 mm. We used groups of voxels of sizes 8^3 , 6^3 , 4^3 , 2^3 , with corresponding label spaces of size 9^3 , 7^3 , 5^3 , 5^3 .

As explained in Chapter 5, the first step of the algorithm involves calculating the MST, which best replicates the underlying anatomy. An example of MST for the colorectal cancer application is shown in Figure 6.1. It can be seen that the obtained MSTs connect nodes with similar image intensity, resulting in a decoupling of tissue regions with different appearance.

Validating the proposed algorithms in the context of clinical DCE-MRI data is a very challenging problem, as neither the true deformation field nor the true PK parameter maps are known. For this reason, both the registration and the PK parameter estimation outcomes were assessed using surrogate measures.

To evaluate registration accuracy, the original and the motion corrected sequences

Table 6.1: *Motion scoring scale*

'1'	'2'	'3'	'4'	'5'
no motion	mild motion	moderate motion	significant motion	severe motion

were blindly graded by a clinical expert. For each sequence, the level of motion was assigned a grade on a scale from '1' to '5' (Table 6.1). Using an expert grading could be preferable in instances where the goal is to extract clinically relevant information, as clinical relevance does not always coincide with increased registration accuracy. To further visualise the results, time-cut images representing the temporal evolution of a pixel-wide line were generated before and after registration.

To assess whether the parameter maps obtained with DireP are more accurate than those estimated by simply fitting the model to the motion corrupted data, the residual error between the data and the model prediction was calculated for each of the methods. The residual error is defined as the absolute difference between the original/motion corrected image and the model predicted image, summed over all the time points of the sequence. A smaller residual error indicates a better fit of the model to the data. The results are presented in the following section.

6.1.2 Results

The effect of registration on the post treatment scan from Patient 10 is visualised using time-cut images (Fig 6.2). As shown by the figure, the proposed registration algorithm improves the temporal alignment for the presented pixel-wide lines.

Qualitative results for the entire clinical dataset are shown in Figs. 6.3, 6.4, 6.5. As shown by the figures, the observed motion level was found to decrease in 21 out of 23 images after applying our discrete registration framework in 3D and 4D, respectively. We also note that the differences between DireP 3D and DireP 4D were not detectable on a visual evaluation. The two dynamic image sequences where the proposed method did not decrease the intra-sequence motion level (Patient 8 Mid and Patient 3 Post) were




Figure 6.2: Effects of registration on an image sequence of the Sonatina dataset are shown using time-cuts. The leftmost image represents a slice through the DCE-MRI volume of interest. The middle two images represent time-cuts for line (a) in the DCE-MRI image. The rightmost two images represent time-cuts for line (b) in the DCE-MRI image. Red arrows indicate regions in the tumour where the improved alignment after registration can be observed.

difficult cases, presenting both breathing and severe peristaltic motion.

For both the 3D and the 4D algorithm, the Jacobian determinant of the deformation field was found to be positive. This indicates there is no tissue folding, which is consistent with the tissue incompressibility assumption. To evaluate the complexity of the resulting transformations, the standard deviation of the Jacobian determinant, $std(Jac)$, was also calculated. An average $std(Jac) = 0.179$ was obtained for DireP 3D, and an average $std(Jac) = 0.159$ was obtained for DireP 4D. This indicates that the 4D algorithm yields smoother transformations.

Figure 6.6 shows the residual error for Patient 3 Mid. As shown by the figure, the residual error is lower for the voxels of the tumour region when the PK parameter maps are estimated after motion correction, using the DireP framework.

The mean residual error over the tumour area is reported in Table 6.2, for all the analysable patients of the Sonatina dataset. As shown by the results, the parameters estimated using DireP yield a model with a better fit to the data than the parameters estimated on the motion corrupted sequences in all but one case (Patient 4 Mid).

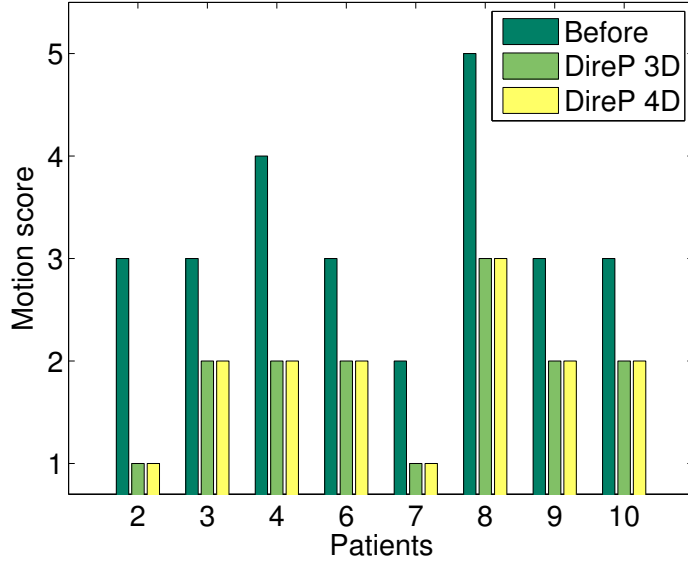


Figure 6.3: The pre-treatment DCE-MRI images were blindly graded before and after registration by a clinical expert (Table 6.1). Both DireP 3D and DireP 4D reduce the amount of motion in all the pre-treatment images of the rectal cancer dataset.

Table 6.2: Residual fitting error for the Sonatina dataset.

Residual	Pre		Mid		Post	
	Before	DireP	Before	DireP	Before	DireP
P2	0.69	0.60	0.74	0.56	0.55	0.46
P3	0.47	0.41	1.35	0.89	1.43	0.88
P4	0.78	0.62	1.22	1.23	1.08	0.77
P6	0.92	0.68	1.15	0.84	0.94	0.75
P7	1.06	0.98	0.52	0.44	0.81	0.67
P8	1.01	0.67	0.77	0.68		
P9	0.85	0.70	0.65	0.48	0.98	0.72
P10	0.73	0.62	0.74	0.61	0.68	0.57

Figure 6.7 presents the effect of motion correction on K^{trans} . Images before and after DireP 3D are shown, for a patient exhibiting severe motion (Patient 8 Pre), and for a patient exhibiting moderate motion (Patient 4 Pre). It can be seen that in both cases the K^{trans} maps become sharper after registration, with a clearer separation between individual voxels, and motion artefacts at the tumour boundary are reduced. This effect is particularly visible for the severe motion case, where the non-registered sequence yields a heavily blurred K^{trans} map.

On a four core Intel i7 CPU running at 1200MHz, using C++ code, runtimes are:

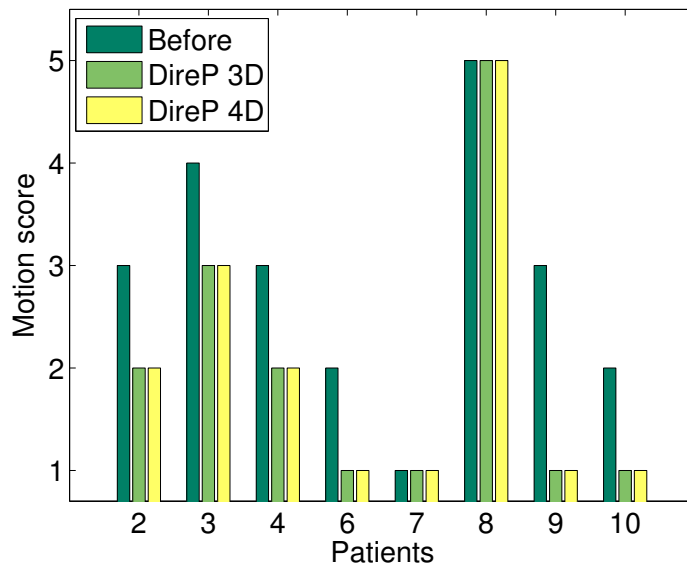


Figure 6.4: *The mid-treatment DCE-MRI images were blindly graded before and after registration by a clinical expert (Table 6.1). Both DireP 3D and DireP 4D reduce the amount of motion in six out of eight mid-treatment images.*

DireP 3D 8mins and 17sec, DireP 4D 20mins and 2sec. These times correspond to running the algorithm with four resolution levels, where groups of eight voxels (2^3) are used at the finest resolution. For three resolution levels, the runtimes are identical to those reported in Chapter 5.

6.2 Registration of lung 4D CT images

Lung cancer is the first most common cause of cancer worldwide in both men and women, with an incidence of 1.82 million new cases recorded in 2012. Moreover, lung cancer causes a fifth of all cancer related deaths worldwide [1]. One of the principal treatment options for this type of cancer is radiotherapy [2]. The main challenge in radiotherapy is target localisation, as it is desired to deliver a high radiation dose to the tumour while avoiding to irradiate proximal healthy tissue. However, target localisation is often difficult due to respiratory motion. For this reason, using a respiratory motion model for dose calculation or treatment gating has become a prerequisite of successful radiotherapy. A number of approaches employ automated image registration of dynamic lung images to derive a

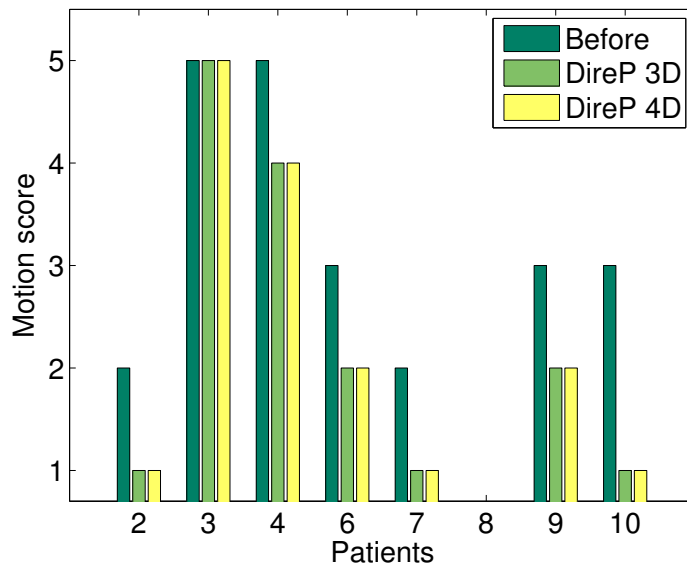


Figure 6.5: The post-treatment DCE-MRI images were blindly graded before and after registration by a clinical expert (Table 6.1). Both DireP 3D and DireP 4D reduce the amount of motion in six out of seven post-treatment images.

respiratory motion model [96]. Nevertheless, only a small number of recent registration algorithms use a spatio-temporal transformation model that enforces temporal constraints [101, 144, 153, 158].

In this section, the **discrete registration** framework (Dire) proposed in Chapter 4 is applied for respiratory motion modelling of the lung. Besides estimating respiratory motion, the goal of this application is to emphasise the effect of temporal smoothness constraints in an application where breathing motion is predominant.

6.2.1 Experiments

To evaluate the methods presented in the previous chapter in the context of an application where periodic breathing motion is predominant, we performed deformable registration on the DIR-Lab dataset [25]. This dataset comprises ten 4DCT images of the lung, which have been acquired for radiotherapy planning from patients who were treated for esophageal or lung cancer. The images had an in-plane resolution ranging from 0.97 to 1.16mm and a slice thickness of 2.5mm, and 10 image volumes had been acquired per




Figure 6.6: *Residual error before and after registration for Patient 3 Mid. (a) Axial slice through the tumour volume with the tumour delineation shown in red. (b) Residual error when the PK parameters are estimated on the motion corrupted data. (c) Residual error after PK estimation with DireP.*

breathing cycle. For each sequence, 300 corresponding anatomical landmarks have been annotated on the inhale and exhale volumes by medical experts. Of these landmarks, 75 were propagated to the intermediary volumes using image registration. For the inhale-exhale volume pairs, the average target registration error (TRE) before registration was 8.46mm and the inter-observer variability for landmark annotation was under 1mm. The most difficult cases, *C7* and *C8*, featured an average TRE of 13.01mm.

For this application, Step 1 of the proposed algorithm (Fig. 5.1) is not needed, and SSD can be used as a similarity metric. The pairwise (Dire 3D) and groupwise (Dire 4D) variants of the proposed algorithm were compared with registration performed on the complete 4D graph using TRW-S optimisation (TRWS), and a state-of-the-art groupwise registration approach [101] (Metz) using continuous optimisation. For the purpose of a complete comparison, the algorithm proposed by Metz et al. was run on Cases 6-10 of the DIR-lab dataset, on which it had not been run initially. We used the implementation provided as part of the Elastix package [78], together with the original parameter file. The registration accuracy was evaluated in terms of target registration error (Section 3.1.4). To evaluate the smoothness of the resulting transformations, the irregularity of landmark

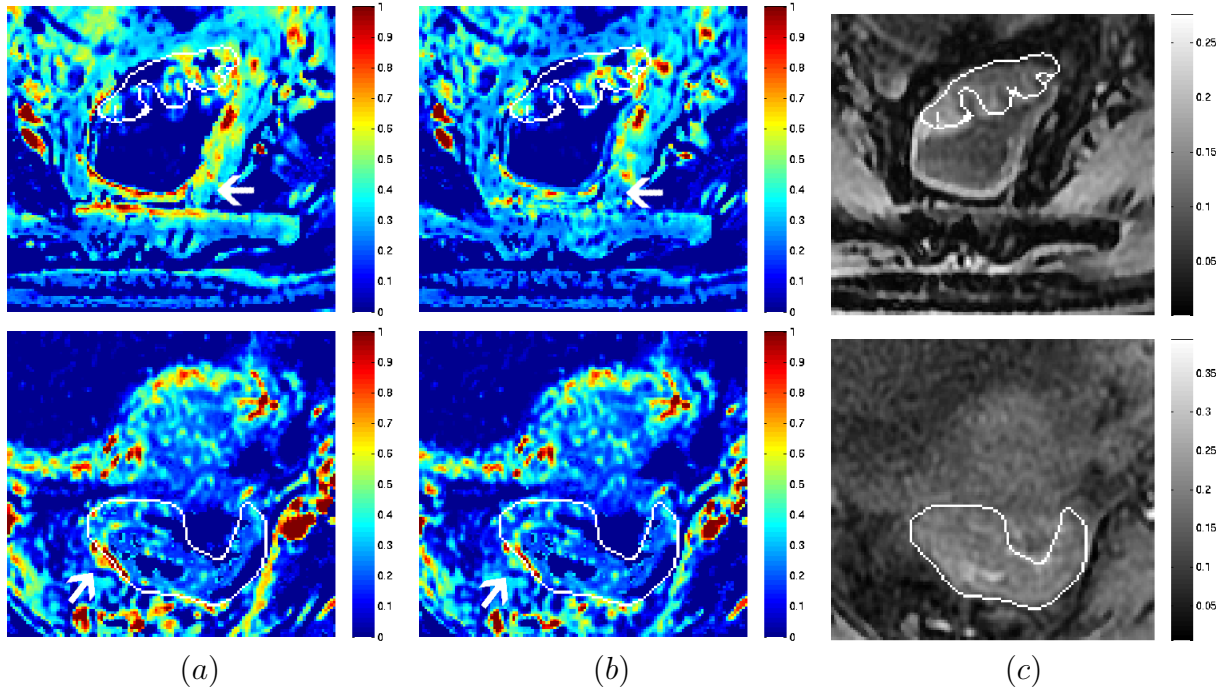


Figure 6.7: K^{trans} maps for one slice of the image volume. Images before (a) and after DireP 3D (b) motion correction are shown, together with the corresponding anatomical image (c). The top row corresponds to a patient exhibiting severe motion, and the bottom row corresponds to a patient exhibiting moderate motion.

trajectories was calculated akin Metz et al. [101]:

$$mIrr = \frac{1}{NP} \sum_{\mathbf{p}_i \in P} \sum_{t=1}^N \left\| \frac{\partial^2 T(\mathbf{p}_{i,t})}{\partial t^2} \right\|^2 \quad (6.1)$$

Here, $\mathbf{p}_{i,t}$ is landmark i in time point t , T represents the evaluated transformation, N is the total number of time points in the dynamic sequence and P is the total number of landmark points.

Three resolution levels were used for the registration. Depending on the resolution level, the search space was $u \in \{16, 12, 8\}$, with a corresponding quantisation $q \in \{2, 1, 0.5\}$ mm, which amounted to a maximum possible deformation of ± 16 mm. We used groups of voxels of sizes 8^3 , 6^3 , 4^3 , with corresponding label spaces of size 17^3 , 13^3 , 9^3 . The experiments were run with $n_{iterMRF} = 5$, $\lambda_{sp} = 0.01$ and $\lambda_{temp} = 0.125$, which were chosen empirically.

An example of MST for the lung cancer application is shown in Figure 6.8. In the




Figure 6.8: *The MST corresponding to a 2D coronal slice from a Dynamic CT image of the lung. The leaf nodes are represented in green, and the root is represented in yellow.*

case of lung images, the MST produces a decoupling between the inner and outer lung regions, which have different image intensities.

6.2.2 Results

The registration results are shown in Table 6.3 and Table 6.4 in terms of TRE. Table 6.3 reports on the registration results where the TRE is calculated only taking into account the inhale/exhale volumes. Table 6.4 reports on the registration results where the TRE is calculated on all the volumes of the sequence, based on the 75 landmarks that were propagated by registration.

As shown by the results in both Table 6.3 and Table 6.4, the proposed algorithms reduce the TRE by a factor of 4 for all of the cases. In comparison to the TRW-S registration on the full graph which only reduces the TRE by a factor of 2, this is a dramatic improvement. Although the state-of-the-art registration method (Metz) appears to outperform Dire for most of the cases, this method uses lung segmentations as prior information. Lung segmentations are obtained by thresholding and morphological closing, thus requiring additional user input for choosing an appropriate threshold. Besides involving extra computations, this approach only achieves registration inside the lungs. The Dire algorithm achieves registration for the whole image. When comparing the pro-

posed Dire 3D and Dire 4D with Metz’s algorithm without masks (w/o), the discrete algorithms are found to be much better in all the ten cases.

Table 6.3: Registration results on the DIR-lab dataset. The average target registration error (TRE) between T00-T50 for Cases 1-10 is reported.

TRE	Before	Dire 3D	Dire 4D	TRWS	Metz	Metz w/o
c1	3.89 ± 2.8	1.03 ± 0.6	1.03 ± 0.6	1.88 ± 1.5	1.02 ± 0.5	1.16 ± 1.2
c2	4.34 ± 3.9	1.00 ± 0.5	0.99 ± 0.5	1.50 ± 1.1	1.06 ± 0.6	1.45 ± 1.5
c3	6.94 ± 4.1	1.30 ± 0.7	1.32 ± 0.7	2.65 ± 2.4	1.21 ± 0.7	2.10 ± 2.0
c4	9.83 ± 4.9	1.80 ± 1.6	1.79 ± 1.6	5.59 ± 5.6	1.57 ± 1.2	2.08 ± 2.0
c5	7.48 ± 5.5	1.75 ± 1.6	1.78 ± 1.6	3.56 ± 3.4	1.70 ± 1.5	2.71 ± 2.9
c6	10.89 ± 7.0	2.33 ± 1.8	2.19 ± 1.6	5.42 ± 3.8	1.62 ± 1.3	2.52 ± 2.4
c7	11.02 ± 7.4	2.28 ± 1.9	2.05 ± 1.4	6.59 ± 5.2	1.67 ± 1.4	3.63 ± 3.7
c8	15.00 ± 9.0	4.27 ± 5.5	4.05 ± 5.3	9.78 ± 9.5	1.67 ± 1.8	6.96 ± 8.3
c9	7.92 ± 4.0	1.58 ± 1.0	1.56 ± 0.9	3.28 ± 2.4	1.43 ± 1.1	2.26 ± 1.8
c10	7.30 ± 6.3	1.75 ± 1.8	1.77 ± 1.7	3.87 ± 3.9	1.26 ± 1.2	2.47 ± 3.0

Table 6.4: Registration results on the DIR-lab dataset. The average target registration error (TRE) for the 75 landmarks in all the volumes for Cases 1-10 is reported.

TRE	Before	Dire 3D	Dire 4D	TRWS	Metz	Metz w/o
c1	2.18 ± 1.9	0.94 ± 0.7	0.92 ± 0.7	1.34 ± 1.0	0.95 ± 0.7	0.93 ± 1.1
c2	3.78 ± 3.4	0.92 ± 0.5	0.92 ± 0.6	1.34 ± 1.0	1.00 ± 0.6	1.32 ± 1.7
c3	5.05 ± 3.1	1.16 ± 0.7	1.16 ± 0.7	1.90 ± 1.5	1.15 ± 0.6	1.80 ± 1.7
c4	6.69 ± 3.8	1.57 ± 1.1	1.55 ± 1.0	3.60 ± 3.4	1.39 ± 1.0	1.71 ± 1.6
c5	5.22 ± 3.9	1.44 ± 1.3	1.47 ± 1.3	2.41 ± 2.1	1.50 ± 1.3	2.04 ± 2.2
c6	7.42 ± 5.3	1.82 ± 1.8	1.81 ± 1.7	3.49 ± 2.6	1.65 ± 1.7	2.14 ± 2.0
c7	6.66 ± 5.3	1.56 ± 1.2	1.53 ± 1.1	3.55 ± 2.6	1.53 ± 1.4	2.54 ± 2.5
c8	9.83 ± 5.8	2.65 ± 2.6	2.32 ± 2.5	5.30 ± 5.1	1.60 ± 1.8	5.05 ± 5.4
c9	5.03 ± 2.9	1.35 ± 0.9	1.35 ± 0.9	2.35 ± 1.8	1.32 ± 1.1	1.57 ± 1.3
c10	5.42 ± 4.9	1.49 ± 1.3	1.52 ± 1.3	2.65 ± 2.4	1.41 ± 1.5	2.04 ± 2.2

For each case, the algorithm with the lowest TRE that does not require any prior knowledge is shown in bold. Metz with masks was excluded from this comparison due to aforementioned limitations. As shown in Table 6.3, the temporal constraints enforced by Dire 4D improve the results in 7 out of 10 patients. This improvement is especially noticeable in difficult cases such as C6, C7, C8. The same trend can be observed in Table 6.4.

Qualitative registration results are shown in Fig. 6.9 and Fig. 6.10, for cases C2 and C8




Figure 6.9: Registration result for Case 2 of the DIR-lab dataset. Top row: differences image before registration. Second row: difference images and deformation magnitude after Metz w/o mask. Third row: results after Dire 3D. Bottom row: results after Dire 4D. Both Dire 3D and Dire 4D outperform Metz, yielding an improved alignment of both inner and outer lung structures. As shown by the deformation magnitude in the coronal slice, the discrete methods better capture the sliding motion.

of the DIR-lab dataset. While $C2$ presents an average TRE before registration of 4.34mm, $C8$ is the most challenging one of the dataset, having an average TRE before registration of 15mm. The two variants of the proposed algorithm, Dire 3D and Dire 4D are compared to Metz without masks in terms of deformation magnitude and difference images between maximum inhale and exhale volumes. As shown by the figures, the discrete methods better capture the lung sliding motion, and yield much lower difference images than the continuous method. This trend is even more pronounced for $C8$ (Fig. 6.10).

Qualitative assessment of the registration results is also performed by generating time-




Figure 6.10: Registration result for Case 8 of the DIR-lab dataset. Top row: differences image before registration. Second row: difference images and deformation magnitude after Metz w/o mask Third row: results after Dire 3D. Bottom row: results after Dire 4D. Both Dire 3D and Dire 4D outperform Metz, yielding an improved alignment of both inner and outer lung structures. As shown by the deformation magnitude in the coronal slice, the discrete methods better capture the sliding motion.




Figure 6.11: Time-cut images before (2) and after registration with Dire 4D (3) for Case 7 of the DIR-lab dataset. Subplots (a) and (b) show the temporal evolution of the pixel-wide lines marked in image (1).

cut images representing the evolution of a pixel-wide line across all time frames. Time-cut images before and after registration are shown in Fig. 6.11.

To compare the Dire 4D and the Dire 3D transformation smoothness, the landmark trajectory irregularity, $mIrr$, was calculated according to Eq. 6.1. $mIrr$ for the transformations obtained with Dire 3D and Dire 4D is presented in Figure 6.12 as a bar chart. Every group of bars corresponds to a case from the DIR-lab dataset. A smaller value for $mIrr$ corresponds to a smoother transformation. As shown by the figure, Dire 4D gives smoother landmark trajectories than Dire 3D for all the DIR-lab cases. Furthermore, these results are better than those reported by Metz et al. [101].

The quality and complexity of the resulting transformations were evaluated by calculating the Jacobian determinant and its standard deviation $std(Jac)$. There were no negative Jacobian values, and we obtained an average $std(Jac) = 0.129$ for Dire 3D and $std(Jac) = 0.120$ for Dire 4D. This confirms that the enforcement of temporal constraints yields smoother transformations.

On a four core Intel i7 CPU running at 1200MHz, using C++ code, the proposed algorithm has an average runtime of 9mins and 42sec for one sequence of the DIR-lab dataset. This represents a substantial improvement in comparison to Metz [101], who report a runtime of 1h and 30mins using an AMD Opteron processor running at 2400 MHz.




Figure 6.12: The irregularity of landmark trajectories using Dire 3D and Dire 4D is presented for all the images of the DIR-lab dataset. Dire 4D yields lower values which are equivalent to temporally smoother trajectories.

6.3 Discussion

In this chapter, two clinical applications to cancer of the proposed nonrigid registration framework were presented. We first applied the DireP framework for intra-sequence motion correction and PK parameter estimation of DCE-MRI sequences from colorectal cancer patients. The proposed framework decreases the level of motion in 21 of the 23 image sequences of the Sonatina dataset. The two cases for which the DireP algorithm fails present gross body motion, breathing motion, as well as severe peristalsis. It should be noted that previous approaches completely exclude sequences with severe motion from the analysis [22], or tailor the data acquisition protocol to reduce the level of motion. Thus, Hamy et al. [56] use a dataset where patients had received BuscopanTM to suppress peristalsis. At the same time, Melbourne et al. [99] employ breath hold protocols to reduce periodic breathing motion. The effect of motion correction on PK parameter maps was also analysed. We found that K^{trans} maps become sharper after registration, with a clearer separation between individual voxels, and motion artefacts at the tumour boundary are reduced.

To further illustrate the effect of temporal smoothness constraints, the proposed registration algorithm was also applied to data where periodic breathing motion is predominant. The Dire framework, which is equivalent to DireP without the PK estimation step, was applied for respiratory motion modelling in 4DCT images of the lung. These experiments were run on the 10 cases of the DIR-lab dataset. The proposed methods are shown to outperform the continuous optimisation state-of-the-art when no prior lung segmentations are used, in terms of both registration accuracy and computational costs. We also show that the inference strategy using MST calculation outperforms TRW-S inference on the full graph. The MST approach implicitly decouples between inner and outer lung regions, aiding the modelling of sliding motion. Moreover, results show that the Dire 4D algorithm leads to improved temporal consistency of the transformations, outperforming Dire 3D in terms of registration accuracy for 7 out of the 10 cases.

Chapter 7

Outlook

7.1 Extensions of the DireP framework

In this section, several methodological extensions of the DireP framework (Chapter 5) are proposed and another possible clinical application is described. The 4D graphical structure presented in Fig. 5.3 could be modified as to contain a MST only for the reference time point, this MST being connected with chains across the other time points (Fig. 7.1). There are two key advantages to this approach. First, the optimal labelling of graphical structure presented in Fig. 7.1 can be found using dynamic programming, as the graph does not contain any loops. Second, with this type of graph, inverse consistency of the transformations is much easier to impose.

After the registration of DCE-MRI images, we expect the intensity at each voxel to be smoothly varying in time. In addition to constraining the deformation field to be temporally smooth, further constraints could be imposed on image intensities to be smoothly changing in time. Mahapatra [94] propose having a term in the data cost that measures the temporal smoothness of the intensity signal at each voxel by a combination of its first and second temporal derivatives. Although such a data term would not be possible for our graphical structure, as the data cost would depend on more than one node, a regularisation term that takes into account intensity differences alongside displacement

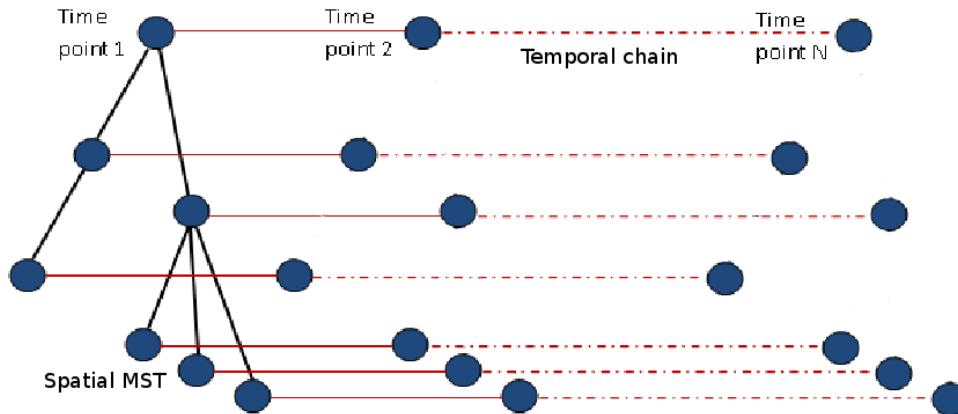


Figure 7.1: An illustration of a new possible graphical structure. While the spatial connectivity is captured by a minimum spanning tree (MST) in the baseline volume only, the temporal continuity is captured by temporal chains.

differences could be attempted for the temporal connections.

While in the case of lung data the MST offers an elegant decoupling between structures inside and outside the lung which look differently and move separately, this is not the case for colorectal images. For the latter, structures with similar intensity do not necessarily move together. To address this, a MST based on motion patterns could be calculated. Hamy et al. [56] have shown that breathing motion can be recovered using robust data decomposition. Using a similar approach, the breathing motion pattern could be found for each voxel. An MST based on vector valued data containing motion patterns would then be calculated.

For the problem of DCE-MRI intra-sequence registration and PK estimation, validation remains a key challenge, since no ground truth is available. A pseudo ground truth could be obtained akin Hamy et al. [56] by manually delineating regions of interest representing an anatomical landmark in all the time frames of the sequence. For colorectal cancer, such regions could be chosen inside the tumour volume. Nevertheless, this is a very time consuming task, and its success is dependent on operator skill.

The proposed framework could be generalised to other imaging modalities and clinical applications. For instance, DireP can be employed for the registration and PK modelling of liver pCT images. Due to its increased vasculature, the liver is the preferred site of

metastasis for other forms of cancer. Liver cancer staging and treatment monitoring have successfully been performed using pCT imaging [43, 59, 102]. To extract accurate perfusion and permeability maps from pCT images of the liver, an appropriate model has to be selected. Additional modelling challenges appear due to the dual blood supply of the liver - contrast agent is first delivered by the arterial supply and then by the venous supply. Previous approaches to modelling pCT of the liver have either used the gradient method [103] or have employed the extended Tofts model [141]. In recent work, Banerji et al. [9] have proposed combining the dual-input Materne model [95], which is used for normal tissue, with the single input extended Tofts model which is used for tumour tissue. Their work, which was initially proposed for DCE-MRI, has been extended to pCT images by our group [107].

As explained before, motion correction plays a fundamental role in the analysis of pCT [77]. Besides the already discussed motion sources such as patient repositioning and breathing, liver images present big organ shifts and displacements. Moreover, on breath-hold acquisition protocols, sudden gasps can move the entire organ outside the field of view. To our knowledge, current methods for PK modelling of the liver do not address intra-sequence motion. For this reason, we expect DireP to yield superior results as it combines PK modelling with motion correction and can capture a variety of motion types. To accommodate for the types of motion present in liver images, the DireP framework can be tailored by adjusting the search space. In particular, both the width of the search space u and the quantisation q would have to be increased, in order to capture bigger motion.

7.2 Joint modelling of pCT and DCE-MRI

As described in Chapter 2, both DCE-MRI and pCT can be used to quantify tissue microvasculature changes that occur during tumour progression or treatment. So far, we have focused on presenting the DCE-MRI modelling (Chapter 2) pipeline and the ob-

tained results (Chapter 4). Nevertheless, pCT images can also be modelled using similar methodologies (Chapter 2). Although both modalities measure the same physiological process, it is expected that there is some degree of complementarity between them. While pCT has a much higher temporal resolution for the slope of the uptake curve, and is easier to model due to the linearity between signal and contrast agent concentration, the signal enhancement is much better in DCE-MRI (up to 100% with respect to baseline, while pCT enhances 2-3%). DCE-MRI images also present better volume coverage, better signal-to-noise and do not involve ionising radiation.

Recently, several studies have shown that both DCE-MRI and pCT can be analysed using PK modelling, the analyses yielding comparable parameters [76, 106, 119, 154]. To exploit the complementarity of these two modalities, we propose to jointly model DCE-MRI and pCT image sequences of the same patient. Namely, a single set of PK parameters which is informed by both modalities is estimated. This research question can be investigated in the context of the Sonatina clinical trial (Section 4.1), where the imaging protocol involves the acquisition of both DCE-MRI and pCT images at the same time points.

The goal of this section is twofold. First, individual modelling of pCT and DCE-MRI is performed, and the PK parameter maps are compared. Second, PK parameter maps that are informed by both modalities are estimated, to investigate whether the joint modelling yields better parameter maps than the independent modelling. Answering this research question could have an important clinical impact, as one of the modalities could be excluded from the imaging protocol, provided the information it offers is redundant.

7.2.1 Preliminary results

As explained in Chapter 2, tracer kinetic modelling is a two stage process. First, contrast agent concentration is estimated from the image signal enhancement - time curve. Second, parameters of interest that indicate the underlying tissue physiology are estimated by fitting a pharmacokinetic model to the contrast agent concentration curve. While the

second stage can be identical for both DCE-MRI and pCT and contrast uptake kinetics can be modelled using one of the models described in Section 2.4, the first stage depends on the image signal. For DCE-MRI, the non-linear relationship between signal enhancement and contrast uptake can be modelled using Eq. 2.8. For pCT, the relationship between image intensities and contrast enhancement is linear (Eq. 2.10).

The joint estimation of PK parameters can be formulated mathematically using Bayes' theorem:

$$P(\theta|X_{01}, X_{02}, Y, \sigma) \approx P(Y|\theta, X_{01}, X_{02}, \sigma)P(\theta) \quad (7.1)$$

Here θ are the PK parameters which are informed by both modalities and Y represents the data from both modalities. X_{01} and X_{02} are the known pre-contrast images from DCE-MRI and pCT and σ is an unknown noise parameter.

We are interested in calculating the full posterior probability $P(\theta|X_{01}, X_{02}, Y, \sigma)$ from which we can extract the best set of parameters $\theta = \{k_{ep}, v_e, ot\}$ that model the noisy image data Y , together with their variances. In particular, the Tofts model parameters $\theta = \{k_{ep}, v_e\}$ can be estimated jointly from both modalities, as the underlying tissue properties are expected to be the same. The time of contrast arrival, ot , is estimated separately for each modality, as it depends on the imaging protocol.

This problem can be solved using the Variational Bayes (VB) framework, by taking a mean field approximation to the full posterior [27]. For the non-linear model fitting problem considered here, this reduces to an iterative procedure that takes a Taylor expansion to the model akin Gauss-Newton. The optimisation can be performed using the FABBER software with a custom-made model [26].

The algorithm was tested on one dataset from the Sonatina clinical trial, comprising a DCE-MRI and a pCT sequence which were acquired at the same time point. To be able to perform joint modelling of corresponding voxels, the pCT and DCE-MRI baseline image volumes were non-rigidly aligned using a B-splines transformation with NMI as a similarity metric. The pCT data modelling was a challenging task, due to the high level




Figure 7.2: *Ktrans* maps for (a) Individual dceMRI estimation (b) Individual pCT estimation (c) Joint modelling

of spatio-temporal noise which covered the useful signal enhancement. Several strategies were attempted to improve the signal-to-noise ration. First, the images were downsampled by averaging over 2×2 voxels in plane. Second, a bilateral noise reduction method proposed by Mendrick et al. [100] especially for pCT was implemented and applied to the pCT sequence. On the one hand, each pixel value in the image was replaced with a weighted average of similar and nearby pixels. On the other hand, a time-intensity profile similarity (TIPS) term is introduced, which preserves time-intensity profiles while still reducing noise. For the colorectal data in this experiment, a kernel of size $3 \times 3 \times 3$ was used, and σ_d , σ_s were chosen to be 1.4 and 4, respectively.

After preprocessing, the pCT and DCE-MRI images were modelled both independently and jointly in the Variational Bayes framework, and the results are shown in Figure 7.2. As seen in the figure, the parameter map obtained from pCT is quite smudged and the underlying anatomy is not clearly visible, which indicates that the employed pre-processing methods were not effective for separating the useful data from the noise. Moreover, it appears that the DCE-MRI data dominates the model parameter maps. This could be explained by the fact that the signal enhancement in DCE-MRI is about 30 times higher than the enhancement in pCT.

7.2.2 Future work

Future work should first address the problem of noise in pCT data, to better separate the useful signal enhancement from the noise component. The optimal combination of parameters for TIPS filtering can be investigated. Furthermore, because of the much higher signal enhancement in DCE-MRI, it appears that this modality dominates the joint modelling. To balance the contributions of DCE-MRI and pCT, a scaling factor could be employed. The value of this scaling factor needs to be optimised. The joint modelling framework can also be extended to incorporate intra-sequence motion correction, which would yield improved PK parameter maps. After the challenges related to modelling are addressed, this work can be further tested on clinical data. Of particular interest is Rhythm [3], a new clinical trial with the aim of characterising tumour hypoxia, which is currently ongoing at the Oxford University Hospitals. In this trial, pCT image acquisition is performed using an amended protocol, which should yield better quality data. The acquisition time is extended to several minutes without increasing the radiation dose. In particular, an image volume is acquired every 1.5 seconds for the first 40 seconds, followed by acquiring an image volume every 30 seconds for an additional 2 minutes. This new protocol enables improved PK modelling of the contrast enhancement curves.

7.3 Tumour growth modelling from longitudinal DCE-MRI scans

Quantifying tumour progression and/or subject response to treatment are key challenges in cancer imaging. This quantification is performed by comparing tumour changes based on longitudinal imaging data, i.e. longitudinal DCE-MRI. Currently, the comparison can be performed by looking at statistics across the tumour area (mean, median, histogram). However, such an approach discards all spatial information about tumour growth and heterogeneity, and is therefore limited. A way of addressing this issue is to perform

longitudinal registration of the DCE-MRI images by incorporating a model of tumour growth/ tumour response to treatment. The first part of this work was directed towards modelling tumour growth using image registration, and the results are presented in the following section. This work was published in ISMRM 2014. In Section 7.3.2, several avenues towards incorporating a tumour growth model or a response to treatment model are discussed.

7.3.1 Preliminary results

Tumour growth measurement from longitudinal data was performed using image registration. We have adapted a previously developed multichannel non-rigid registration algorithm [32] to align subsequent time points over the tumour growth cycle to the first time point (Tp). The novelty of this approach resides in using the PK parameter maps k_{ep} and v_e , alongside the image intensity as channels in the registration (Fig. 7.3). The Jacobian of these transformations is used as an indicator of tumour volume change. We also investigated whether K^{trans} at the first time point has predictive value for tumour growth in subsequent images.

These analyses were performed on pre-clinical data, where uninhibited tumour growth can be observed. Three mice were injected with MC38 (colon carcinoma) tumour cells to develop subcutaneous tumours and imaged with DCE-MRI at three time points over the tumour growth cycle (days 21, 23 and 24 post injection). This enabled the investigation of tumour growth and heterogeneity. The animals were immobilised for imaging using gaseous anaesthesia. DCE-MRI was performed at 4.7T (Agilent VNMRs) using a respiratory-gated 3D gradient echo sequence (TE 0.6ms, TR 1.1ms and a nominal flip angle of 5°). $128 \times 64 \times 64$ voxels covered a field of view of $54 \times 27 \times 27$ mm encompassing the entire volume of the birdcage coil. 50-100 repeats of this scan, each lasting ca. 8-12s were acquired with Gadolinium contrast agent (Omniscan, 30 ul in 5s). To quantify the T1 change, T1(0) mapping was performed using an array of flip angles, and correction of the B1 inhomogeneities was performed using a respiration gated modification of the

actual flip angle method [157].

Pharmacokinetic modelling of each time-point was performed using the Tofts model [143] and the Weimann bi-exponential arterial input function with population averaged parameters as derived by Heilmann et al. [62]. To enable a spatial comparison of the PK parameter maps, the differences in mouse positioning were eliminated using rigid registration. We assume the remaining differences are due to tumour growth. To capture these differences, we use the Hybrid Feature-based Diffeomorphic [32] (HFD) registration method. In this multichannel approach we use both image intensity and the PK parameter maps to drive the registration. HFD consists of the estimation of two types of forces for each input channel: (1) Demons-like forces estimated locally, and (2) forces computed with the block-matching strategy, which encode spatial correspondences. This amounts to the computation of a family of $n = 6$ forces F_n . These are then smoothed with Gaussian kernels K_{σ_n} , $\sigma_n = 2, 2, 2, 4, 4, 4$ mm within a diffeomorphic framework. The final diffeomorphic deformation encodes the contribution of each smooth update field given by an individual force. We refer the reader to Cifor et al. [32] for more details.

Fig. 7.3 shows the resulting PK parameter maps for Mouse 2, for all the three Tps. These images were rigidly registered to ensure spatial correspondence. On a visual evaluation, the PK maps do not radically change over time, the tumour maintaining a necrotic core and an enhancing rim. Non-rigid multichannel image registration was then performed between Tp2 (day 23) - Tp1 (day 21) and Tp3 (day 24) - Tp1 for all the mice. The success of HFD multichannel is reported in Table 7.1 in terms of Dice overlap (%). To illustrate the importance of including the PK maps as channels, the method was compared with HFD based only on image intensity, and is shown to outperform it. 'Rigid' represents the results after rigid registration only.

The Jacobian matrices of the deformation were calculated to evaluate tumour growth. These results are shown in Fig. 7.4. The Jacobian maps (in grey) generally indicate volume expansion (bright grey) in the tumour core and volume preservation ($Jac = 100$) at the enhancing rim. We also computed the expected volume change by integrating over




Figure 7.3: Pharmacokinetic parameter maps together with corresponding DCE-MRI baseline images: (a) DCE-MRI (b) k_{ep} (c) v_e . Each row corresponds to a longitudinal time point: Row 1 (day 21), Row 2 (day 23), Row 3 (day 24)

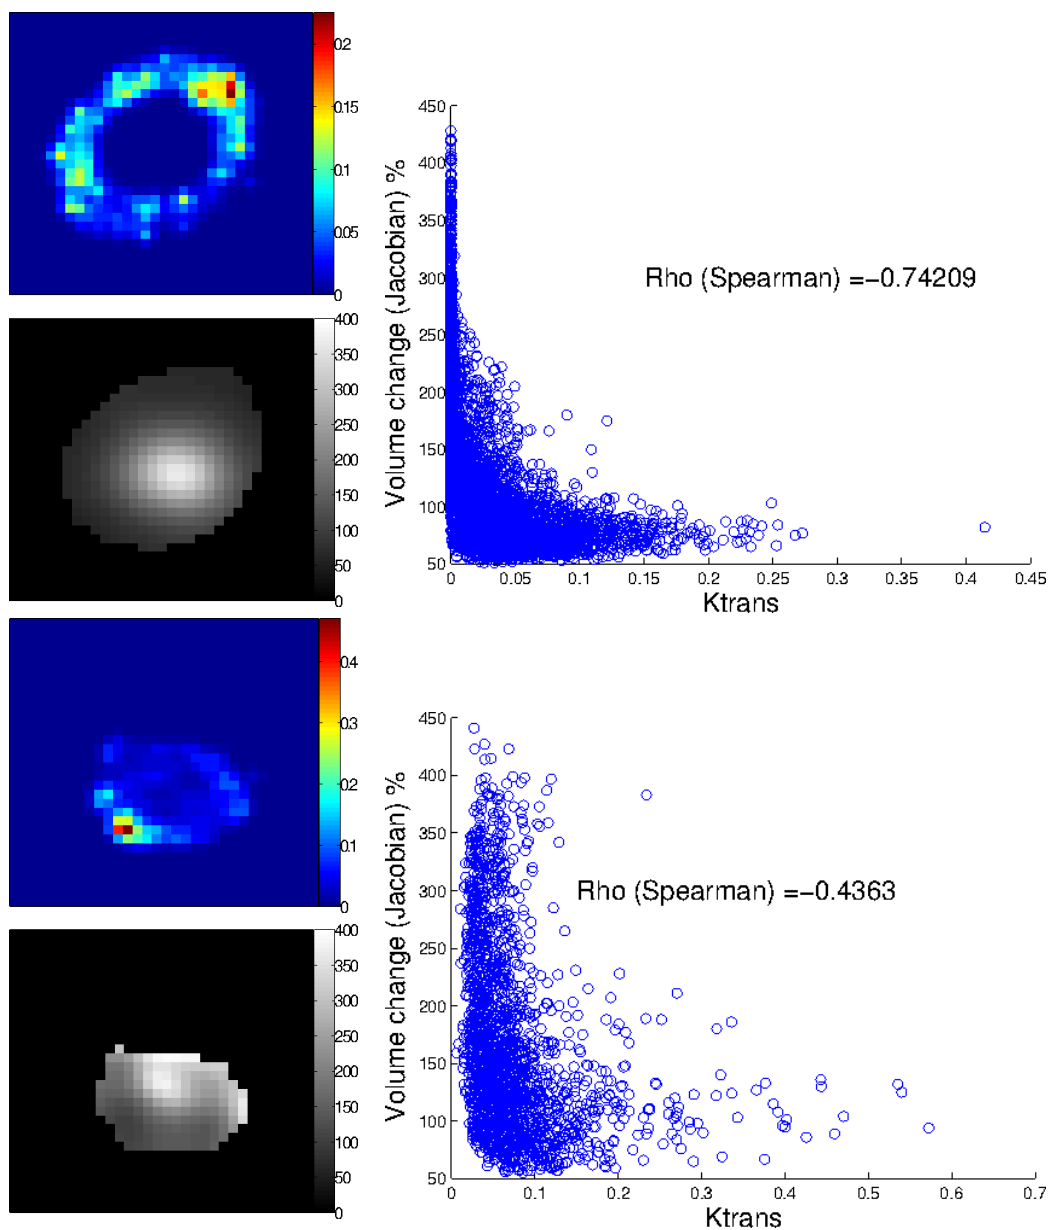


Figure 7.4: The correlation between the K^{trans} map (Tp1) (in colour) and the Volume change indicated by the Jacobian (in grey). Jacobian=100 indicates volume preservation. The top plot represents the volume change between Tp1-Tp2 for Mouse 1 (more homogenous PK maps). The bottom plot represents the volume change between Tp1-Tp2 for Mouse 3 (more heterogeneous PK maps).

Table 7.1: *DICE overlap % for the tumour area*

DICE(%)	Mouse 1		Mouse 2		Mouse 3	
	Tp1-Tp2	Tp1-Tp3	Tp1-Tp2	Tp1-Tp3	Tp1-Tp2	Tp1-Tp3
Rigid	78.56	84.42	81.20	63.65	75.07	64.57
HFD intensity	79.57	79.41	66.83	72.23	61.94	79.24
HFD multichannel	80.47	90.55	84.19	86.96	87.22	87.69

the Jacobian. This yielded an average volume increase (over all the mice) of $55.7 \pm 10\%$ for Tp2 and $116.7 \pm 27\%$ for Tp3. This is an overestimation compared to the volume calculation from ground truth tumour delineations ($36.8 \pm 13\%$ for Tp2 and $84.3 \pm 19\%$ for Tp3). To investigate whether PK parameters at the first Tp can be a predictor of volume change, we have investigated the relationship between Ktrans maps at Tp1 and the Jacobians, which indicate volume change between Tp1-Tp2 and Tp1-Tp3. The scatter plots of Volume change- K^{trans} for two mice, Tp1-Tp2 are shown in Fig. 7.4. The measures appear to be correlated for more homogeneous tumours (row 1). In the case of more complex, heterogeneous tumours (row 2), the Jacobian appears to be a too simplistic model for tumour growth.

In this section, we have proposed a simple, registration based approach to quantify tumour growth from longitudinal DCE-MRI images. Subsequent Tps were registered to the first Tp using a multichannel registration method based on PK parameter maps. This method was shown to successfully recover tumour growth, on data where the tumour change is small and gradual. However, the registration is based on non-linear deformation, which is not able to account for the deposition of new tumour cells. Moreover, the Jacobian of the deformation field also predicts deformation of the necrotic core. In other words, the Jacobian could indicate growth in regions where the tissue is not growing, but it is expanding by getting more necrotic cells.

In the following section, several avenues for including a more complex tumour model are discussed. We have also investigated whether K^{trans} at the first time point can act as a predictor of localised tumour growth. For more homogeneous tumours presenting a well defined necrotic core and an enhancing rim, there appears to be a relationship between

K^{trans} and tumour growth as indicated by the Jacobian. A more complex model will be required for modelling more heterogeneous tumour growth.

7.3.2 Future work

In this section, we propose a model of tumour growth which is personalised based on noninvasive imaging data, and which could be used to predict tumour evolution across time. Although vast amounts of research have been conducted on mathematical tumour growth modelling, only recently has image-based tumour growth modelling become an active research topic [7, 28, 33, 83, 88, 152, 156].

Most tumour growth models that have been used to date rely on the reaction-diffusion equation, which describes the rate of tumour cell density change:

$$\frac{\partial c}{\partial t} = \nabla(D\nabla c(\mathbf{x}, t)) + R(c(\mathbf{x}, t)) \quad (7.2)$$

Here, $c(\mathbf{x}, t)$ is the tumour cell density at position \mathbf{x} and time t . The first term on the right-hand side, also known as the diffusion term, represents tumour cell diffusion into surrounding tissue and is characterised by D , the diffusion tensor. The second term $R(c(\mathbf{x}, t))$, also known as the reaction term, represents tumour cell proliferation. Tumour cell proliferation has been described using a number of models: the exponential model $R(c(\mathbf{x}, t)) = \rho c(\mathbf{x}, t)$, the Gompertz model $R(c(\mathbf{x}, t)) = \rho c(\mathbf{x}, t) \ln(1/c(\mathbf{x}, t))$, the logistic model $R(c(\mathbf{x}, t)) = \rho c(\mathbf{x}, t)(1 - c(\mathbf{x}, t))$, etc. Here ρ denotes the proliferation rate.

There have been some recent attempts to personalise models based on imaging. For instance, Clatz et al. [33], who have worked on modelling glioblastoma, use the model presented in Eq. 7.2, where proliferation is described using the logistic model, and the anisotropic diffusion coefficient is derived based on diffusion tensor MR images (DT-MRI). The mechanical interaction of the tumour with the invaded tissue (mass effect) is captured using a mechanical model of tumour growth.

Atuegwu et al. [7] model breast cancer response to treatment using diffusion-weighted

MRI (DW-MRI) and DCE-MRI. DW-MRI is employed to map the apparent diffusion coefficient (ADC), which indicates the rate of water diffusion in tissue and was shown to inversely correlate with tumour cellularity. DCE-MRI is employed to derive the extravascular extracellular (EES) fractional volume v_e and the plasma volume v_p , using the extended Tofts model. These measures are then used to estimate the number of tumour cells:

$$N(t) = \theta \frac{ADC_w - ADC(t)}{ADC_w - ADC_{min}} v_{TC}(t) \quad (7.3)$$

Here, $N(t)$ is the number of tumour cells, θ is the carrying capacity, ADC_{min} is the ADC for the voxel with the maximum number of cells, and ADC_w is the ADC of voxels that contain only water. The fractional volume occupied by tumour cells within a voxel is estimated using either $v_{TC}(t) = 1$ or $v_{TC}(t) = 1 - v_e(t) - v_p(t)$. The estimated numbers of cells at two longitudinal time points are used as inputs in the logistic tumour growth model, and the proliferation rate ρ is estimated. The model can then be used to predict cellularity at further time points.

Liu et al. [88] extend the model in Eq. 7.2 by incorporating an advection term, which accounts for tumour cell displacement due to mechanical deformation (Eq 7.4). Their model is personalised using dual phase CT and FDG-PET. In particular, the intracellular volume fraction is estimated using dual phase CT, and this is used as a surrogate for the cell number, assuming all cells have similar volumes. The amount of nutrients required for tumour growth is calculated based on the glucose metabolic rate extracted from FDG-PET. This method was applied for pancreatic tumour modelling. However, the main limitation of their work is the attempt to estimate a high number of parameters from limited data.

$$\begin{aligned}
\frac{\partial c}{\partial t} &= D\Delta c - \nabla(c\mathbf{v}) + \rho c(1 - c) \\
0 &= \nabla\sigma + \mathbf{f}_{\text{ext}} \\
\sigma &= \lambda\nabla\mathbf{u} + \mu(\nabla\mathbf{u} + \nabla\mathbf{u}^{\mathbf{T}}) \\
\mathbf{v} &= \frac{\partial\mathbf{u}}{\partial t}
\end{aligned} \tag{7.4}$$

Here, $\nabla(c\mathbf{v})$ represents the advection term, where \mathbf{v} is the tumour cell drift velocity. \mathbf{v} depends on the displacement \mathbf{u} induced by the balance between the stress tensor σ and the external force f_{ext} applied on the tissue. The external force is proportional to the local gradient of the tumour cell density: $f_{\text{ext}} = -\gamma\nabla c$. λ and μ are unknown Lamé coefficients.

In a recent extension of the work presented by Liu et al., Wong et al. [152] propose employing a hyperelastic biomechanical model instead of the previously used linear model. They also reduce the number of parameters to be estimated, and use a gradient-free nonlinear optimisation to perform the inference.

Several other imaging modalities can be used to derive physiological quantities that could inform tumour growth models. Not only does DCE-MRI offer estimates of plasma and EES volume, it also provides estimates of tissue perfusion, which an indicator of tumour vasculature and could be correlated with tumour growth. Furthermore, FMISO-PET images can be used as an indicator of hypoxia.

For the case of colorectal tumours which was investigated in this work, we propose to begin with a simple model that accounts for tumour proliferation only (Eq. 7.5). This model is governed by two unknowns, the tumour cell density c and the proliferation rate ρ . The logistic tumour growth model is chosen, as it can account for both proliferation and necrosis, giving a net proliferation rate.

$$\frac{\partial c}{\partial t} = \rho(\mathbf{x}, t)c(\mathbf{x}, t)(1 - c(\mathbf{x}, t)) \tag{7.5}$$

As the tumour progresses, the parts with sufficient supplies of nutrients and oxygen are expected to proliferate, and those which are poorly supplied with nutrients are expected to undergo necrosis. For this reason, the proliferation rate $\rho(\mathbf{x}, t)$ in Eq. 7.5 is made spatially varying. As the nutrient supply is dependent on the underlying tissue vasculature, $\rho(\mathbf{x}, t)$ could be informed by DCE-MRI images:

$$\rho(\mathbf{x}, t) = AF_p(\mathbf{x}) \quad (7.6)$$

Here, F_p represents tissue perfusion, which is related to the amount of nutrients being delivered to a specific voxel \mathbf{x} . A is a scaling constant. It would be desirable to also use DCE-MRI to estimate tumour cellularity, as this does not involve any extra costs related to the acquisition of other modalities. An initial attempt towards estimating tumour cellularity could use the intracellular volume given by $1 - v_e(t) - v_p(t)$.

It would then be desirable to also account for tumour cell displacement due to the mass effect. This could be performed with the aid of a biomechanical model such as the one presented in Eq. 7.4. It is expected that tumour cell diffusion would be much lower in the case of colorectal cancer than in the case of glioma. For this reason, we chose to neglect the effect of diffusion in the proposed model. Further extensions of this model should attempt to incorporate multi-modal information for informing the model. For instance, glucose metabolism information can be extracted from FDG-PET and used alongside tissue perfusion for calculating the proliferation rate in Eq. 7.6.

Chapter 8

Conclusions

This thesis has presented a comprehensive framework for PK modelling and intra-sequence motion correction of dynamic imaging data. This framework combines a number of methodological contributions which are at the interface between image registration and pharmacokinetic modelling. In Chapter 4, we investigated the effect of different design choices in the process of PK modelling. The main contribution presented in this chapter is the derivation of patient-specific AIFs from pCT data, which were used for the analysis of contemporaneous DCE-MRI images of the same patients. For the first time, pCT derived AIFs were employed for DCE-MRI PK parameter estimation and compared to a population averaged model. We found that the intra-patient variability was much smaller than the variability with respect to the population model for this dataset. Moreover, using the population AIF can lead to over-estimation of the PK values.

We also investigated the issue of optimal PK model selection for capturing the biological processes of interest. The Tofts model is the preferred approach in clinical practice. However, it does not provide separate estimates for perfusion and permeability, and is only suitable for weakly vascularised or highly perfused tissue. For this reason, we chose to compare it to a second generation model, the 2CTU, which addresses these limitations. The comparison was performed in terms of the Akaike information criterion. Although the Tofts model is theoretically worse than second generation models, we found that the

usage of the 2CTU model cannot be justified for low temporal resolution data (1 volume every 9.5 seconds), with short total acquisition time (<5min).

In Chapter 5, we addressed the problem of motion correction for dynamic image sequences. To this end, a novel framework for nonrigid registration of 4D dynamic sequences was proposed. This framework is formulated on an MRF, and can be applied across imaging modalities, being suitable for both contrast-enhanced and non contrast-enhanced data. To address the change in image intensities due to contrast inflow, a similarity metric which incorporates PK modelling information was formulated. A graph with reduced connectivity was proposed. The MST which best replicates the underlying anatomy was calculated for the baseline volume. This topology was assumed to be identical in all subsequent volumes, and temporal smoothness was enforced by adding temporal links that connect corresponding nodes across time. The labelling problem for the reduced graph was solved using loopy belief propagation. This inference method was compared with optimisation of the full 4D graph using TRW-S. When tested on simulated data where the ground truth is known, the proposed framework (DireP) was shown to outperform a state-of-the-art method [12] using continuous optimisation. However, we also found that the positive effect of the temporal constraints was not visible on this dataset that mimics real colorectal DCE-MRI, as high frequency motion such as peristalsis was predominant over low frequency motion (i.e. breathing).

The methods proposed in Chapters 4 and 5 were tested on colorectal images from a novel cancer drug trial (Sonatina). The main clinical question is whether the administered drug, Nelfinavir, has the potential to increase blood flow acting as a radiosensitiser. To this end, the trial involved acquiring both DCE-MRI and pCT images at three time points over the course of treatment.

To assess the effect of Nelfinavir, K^{trans} distributions at pre- and mid-treatment (after 7 days of Nelfinavir) were compared. A right shift in the K^{trans} distribution indicates increased perfusion. Treatment is considered to have had effect, when such a trend is observed. This trend was observed in half of the analysed patients, indicating that

Nelfinavir has the potential to increase tumour perfusion. Nevertheless, this clinical study presents some limitations: First, the T_{10} relaxation time could not be estimated due to sub-optimal variable flip angle sequences acquisition protocol. Consequently, an uniform T_{10} map was assumed. Second, the clinical end point was tumour regression grading based on MRI measurements of tumour dimension 8 weeks after the end of treatment. No histological data was available for this study. Moreover, due to patient recruitment problems, the study did not feature a control group as initially intended.

When applying the framework proposed in Chapter 5 to the Sonatina dataset, we found it to decrease the level of motion in 21 out of the 23 image sequences. These sequences correspond to the 8 analysable patients at pre- mid- and post-treatment. The effect of motion correction on the PK parameter maps was also analysed. We found that K^{trans} maps become sharper after registration, with a clearer separation between individual voxels. Moreover, motion artefacts at the tumour boundary are reduced. In other words, some voxels that were classified as highly perfused before registration have lower perfusion values after registration, as the initial values appear to have been motion artefacts. The residual fitting error before and after motion correction was calculated. We found the residual error to be lower when PK parameter maps are estimated after motion correction in 22 out of the 23 sequences.

The registration algorithm part of the framework proposed in Chapter 5 was also applied for respiratory motion estimation from dynamic CT images of lung cancer. In this application, periodic breathing motion is predominant, and we could illustrate the positive effect of temporal smoothness constraints. The data used in these experiments is from DIR-lab [25], a publicly available dataset comprising 10 4DCT images that have been acquired for radiotherapy planning. The proposed method was shown to outperform the continuous optimisation state-of-the-art (Metz et al. [101]) when lung segmentations are not used as prior knowledge. The improvement is both in terms of registration accuracy and computational costs. The results also highlight the effect of temporal constraints, which lead to an improvement in registration accuracy for 7 out of 10 cases.

Chapter 7 presented three different projects that stem from this work, which are potential areas of subsequent research. First, several improvements to the proposed framework were proposed, and another possible clinical application was discussed. Second, initial work towards the joint modelling of DCE-MRI and pCT was presented. Third, we proposed modelling tumour growth from longitudinal DCE-MRI data. To summarise, this thesis has presented a comprehensive framework for PK modelling and intra-sequence motion correction of dynamic imaging data and its applications to cancer therapy.

Appendix A

List of publications

- M Enescu, MP Heinrich, EJ Hill, J Franklin, RA Sharma, MA Chappell and JA Schnabel, An MRF-Based Discrete Optimization Framework for Nonrigid Registration of Dynamic Medical Images, *Medical Image Analysis*, *In preparation*.
- M Enescu, MP Heinrich, EJ Hill, RA Sharma, MA Chappell and JA Schnabel, An MRF-Based Discrete Optimization Framework for Combined DCE-MRI Motion Correction and Pharmacokinetic Parameter Estimation, *MICCAI: Bayesian and Graphical Models for Biomedical Imaging*, 2014.
- M Enescu, A Cifor, V Kersemans, S Gilchrist, J Beech, D Allen, S Smart, MA Chappell and JA Schnabel, Pharmacokinetic modelling of longitudinal dceMRI scans for assessment of tumour growth, *ISMRM*, 2014.
- M Enescu, M Bhushan, EJ Hill, J Franklin, EM Anderson, RA Sharma, JA Schnabel, pCT derived arterial input function for improved pharmacokinetic analysis of longitudinal dceMRI for colorectal cancer, *Proceedings of SPIE Medical Imaging 2013 - Image Processing*, Orlando, Florida, USA, 9-14 February 2013.
- EJ Hill, C Roberts, JM Franklin, M Enescu, N West, TP MacGregor, KY Chu, L Boyle, C Blesing, LM Wang, S Mukherjee, EM Anderson, G Brown, S Dutton, SB Love, JA Schnabel, P Quirke, R. Muschel, WG McKenna, M Partridge, RA

Sharma, Clinical Trial of Oral Nelfinavir Before and During Radiation Therapy for Advanced Rectal Cancer, *Clinical Cancer Research*, *Submitted*.

- EJ Hill, M Enescu, N West, JM Franklin, K-Y Chu, J Li , D Whittington, T P MacGregor, S Azraf, C Blesing, S Mukherjee, M Betts, A Slater, EM Anderson, G Brown, JA Schnabel, M Partridge, P Quirke, RA Sharma, Oral nelfinavir before and during radiation therapy for rectal cancer: Changes in tumor perfusion and correlation between tissue and radiological markers of response, GI ASCO, 2014.
- B Irving, L Tanner, M Enescu, M Bhushan, EJ Hill, J Franklin, EM Anderson, RA Sharma, JA Schnabel, and M Brady, Personalised Estimation of the Arterial Input Function for Improved Pharmacokinetic Modelling of Colorectal Cancer Using dceMRI, MICCAI: Abdominal imaging, 2013.
- M Bhushan, M Enescu, C Galerneau, M Anderson, F Gleeson, M Brady, JA Schnabel, M Jenkinson, "Quantifying response to CRT in colorectal cancer patients from dynamic imaging, Novel technology that shapes Radiology: EIBIR presents IMAGINE", European Congress of Radiology 2013.
- J Franklin, M Enescu, EJ Hill, RA Sharma, JA Schnabel, FV Gleeson, EM Anderson, "A perfusion CT-derived patient-specific arterial input function for pharmacokinetic modeling of dynamic contrast enhanced MRI", European Congress of Radiology 2013.
- EJ Hill, J Franklin, M Enescu, T Pwint, JA Schnabel, EM Anderson, G Brown, WG McKenna, T. Maughan and R.A. Sharma, "Phase 0-I clinical trial of nelfinavir with hypofractionated radiotherapy for rectal cancer: Clinical safety and changes in tumour perfusion observed", The National Cancer Research Institute (NCRI) 2012.

Appendix B

List of abbreviations

2CTU : two compartment tissue uptake

2CX : two compartment exchange model

AATH : adiabatic approximation to tissue homogeneity

AIC : Akaike information criterion

AIF : arterial input function

ADC : apparent diffusion coefficient

ATP : adenosine triphosphate

AUC : area under the curve

CRC : colorectal cancer

CRT : chemo-radiotherapy

DCE-MRI : dynamic contrast-enhanced magnetic resonance imaging

DP : distributed parameter

DW-MRI : diffusion-weighted magnetic resonance imaging

dPET : dynamic positron emission tomography

EES : extravascular extracellular space

FFD : free form deformation

FDG-PET : fludeoxyglucose positron emission tomography

FMISO-PET : fluoromisonidazole positron emission tomography

LBP : loopy belief propagation

MAP : maximum a posteriori

MRF : Markov Random Field

MST : minimum spanning tree

MTT : mean transit time

NCC : normalised cross-correlation

NMI : normalised mutual information

PCA : principal component analysis

pCT : perfusion computed tomography

PE : peak enhancement

PK : pharmacokinetic

PS : permeability surface area product

RF : radio frequency

SPGR : spoiled gradient echo sequence

SSD : sum of squared differences

TE : echo time

TR : repetition time

TRW-S : tree-reweighted message passing

VOI : volume of interest

Bibliography

- [1] Cancer Research UK. <http://www.cancerresearchuk.org/cancer-info/>.
- [2] National Health Service UK. www.nhs.uk.
- [3] Rhythm: modulation of radiotherapy according to hypoxia. www.octoxford.org.uk/alltrials/trials/Rhythm1.html.
- [4] M. Aklilu and C. Eng. The current landscape of locally advanced rectal cancer. *Nature Reviews Clinical Oncology*, 8(11):649–659, 2011.
- [5] P. Armitage, C. Behrenbruch, M. Brady, and N. Moore. Extracting and visualizing physiological parameters using dynamic contrast-enhanced magnetic resonance imaging of the breast. *Medical Image Analysis*, 9(4):315–329, 2005.
- [6] V. Arsigny, P. Commonwick, X. Pennec, and N. Ayache. A log-euclidean framework for statistics on diffeomorphisms. In *MICCAI Proc.*, 2006.
- [7] N. Atuegwu, L. Arlinghaus, X. Li, B. Chakravarthy, V. Abramson, M. Sanders, and T. Yankeelov. Parameterizing the logistic model of tumour growth by DW-MRI and DCE-MRI data to predict treatment response and changes in breast cancer cellularity during neoadjuvant chemotherapy. *Translational Oncology*, 6(3):256–264, 2013.
- [8] L. J. Bains, D. M. McGrath, J. H. Naish, S. Cheung, Y. Watson, M. B. Taylor, J. P. Logue, G. J. M. Parker, J. C. Waterton, and D. L. Buckley. Tracer kinetic analysis

- of dynamic contrast-enhanced MRI and CT bladder cancer data: A preliminary comparison to assess the magnitude of water exchange effects. *Magnetic Resonance in Medicine*, 64(2):595–603, 2010.
- [9] A. Banerji, J. H. Naish, Y. Watson, G.C. Jayson, G. A. Buonaccorsi, and G. J. M. Parker. DCE-MRI model selection for investigating disruption of microvascular function in livers with metastatic disease. *Journal of Magnetic Resonance Imaging*, 35(1):196–203, 2012.
- [10] R. G. H. Beets-Tan and G. L. Beets. Rectal cancer: Review with emphasis on MR imaging. *Radiology*, 232(2):335–346, 2004.
- [11] F.M. Beg, M.I. Miller, A. Trounev, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(139-157), 2005.
- [12] M. Bhushan, J. A. Schnabel, L. Risser, M. P. Heinrich, J. M. Brady, and M. Jenkinson. Motion correction and parameter estimation in dceMRI sequences: Application to colorectal cancer. In *MICCAI Proc.*, 2011.
- [13] M. Bhushan, J. A. Schnabel, L.N. Tanner, F. V. Gleeson, J. M. Brady, and M. Jenkinson. Estimation of motion and pharmacokinetic parameters in DCE-MRI sequences for discrimination of responders in colorectal cancer. In *Proc. ISMRM*, 2012.
- [14] G. Binefa, F. Rodriguez-Moranta, A. Teule, and M. Medina-Hayas. Colorectal cancer: From prevention to personalized medicine. *World Journal of Gastroenterology*, 20(22):6786–6808, 2014.
- [15] C.M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [16] F. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions PAMI*, 11:567–585, 1989.

- [17] E. Boros, P.L. Hammer, and X. Sun. Network flows and the minimization of quadratic pseudo-boolean functions. Technical report, RUTCOR, 1991.
- [18] V. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions PAMI*, 23:1222–1239, 2001.
- [19] G. Brix, J. Griebel, F. Kiessling, and F. Wenz. Tracer kinetic modelling of tumour angiogenesis based on dynamic contrast-enhanced CT and MRI measurements. *European Journal of Nuclear Medicine and Molecular Imaging*, 37(1):30–51, 2010.
- [20] C. Broit. *Optimal Registration of Deformed Images*. PhD thesis, University of Pennsylvania, 1981.
- [21] M. BroNielsen and C. Gramkow. Fast fluid registration of medical images. *Visualization in Biomedical Computing*, 1131(1):265–276, 1996.
- [22] G. A. Buonaccorsi, J. P. B. O’Connor, A. Caunce, C. Roberts, S. Cheung, Y. Watson, K. Davies, L. Hope, A. Jackson, G. C. Jayson, and G. J. M. Parker. Tracer kinetic model-driven registration for dynamic contrast-enhanced MRI time-series data. *Magnetic Resonance in Medicine*, 58(5):1010–19, 2007.
- [23] P. Cachier, E. Bardinet, D. Dormont, X. Pennec, and N. Ayache. Iconic feature based nonrigid registration: the pasha algorithm. *Computer Vision and Image Understanding*, 89(1):272–298, 2003.
- [24] E. Castillo, R. Castillo, J. Martinez, M. Shenoy, and T. Guerrero. Four-dimensional deformable image registration using trajectory modelling. *Physics in Medicine and Biology*, 55(1):305–327, 2010.
- [25] R. Castillo, E. Castillo, R. Guerra, V.E. Johnson, T. McPhail, A.K. Garg, and T. Guerrero. A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. *Physics in Medicine and Biology*, 54(1):1849–1870, 2009.

- [26] M. Chappell. Fsl toolbox. <http://www.fmrib.ox.ac.uk/fsl/fabber>.
- [27] M. Chappell, A. Groves, and M. Woolrich. Variational bayesian inference for a non-linear forward model. *IEEE Transactions on Signal Processing*, 57(1):223–236, 2009.
- [28] X. Chen, R. Summers, and J. Yao. Kidney tumour growth prediction by coupling reaction-diffusion and biomechanical model. *IEEE Transactions on Biomedical Engineering*, 60(1):169–173, 2013.
- [29] W. Chow, C. Jiang, and M. Guan. Anti-HIV drugs for cancer therapeutics: back to the future? *Lancet*, 10(1):61–71, 2009.
- [30] G.E. Christensen, S.C. Joshi, and M.I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Medical Imaging*, 16(1):864–877, 1997.
- [31] A.C. Chung, W.M. Wells III, A. Norbash, and W.E.L. Grimson. Multi-modal image registration by minimising kullback-leibler distance. In *MICCAI*, 2005.
- [32] A. Cifor, L. Risser, D. Chung, E.M. Anderson, and J. A. Schnabel. Hybrid feature-based diffeomorphic registration for tumour tracking in 2d liver ultrasound images. *IEEE Transactions on Medical Imaging*, 32(9):1647–56, 2013.
- [33] O. Clatz, M. Sermesant, P. Bondiau, H. Delingette, S.K. Warfield, G. Malandain, and N. Ayache. Realistic simulation of the 3d growth of brain tumours in mr images coupling diffusion with biomechanical deformation. *IEEE Transactions on Medical Imaging*, 24(1):1334–46, 2005.
- [34] T. E. Conturo, E. Akbudak, M. S. Kotys, M. L. Chen, S. J. Chun, R. M. Hsu, C. C. Sweeney, and J. Markham. Arterial input functions for dynamic susceptibility contrast MRI: Requirements and signal options. *Journal of Magnetic Resonance Imaging*, 22(6):697–703, 2005.

- [35] W.R. Crum, T. Hartkens, and D. L. G. Hill. Non-rigid image registration: theory and practice. *The British Journal of Radiology*, 77(1):140–153, 2004.
- [36] E. D’Agostino, F. Maes, D. Vandermeulen, and P. Suetens. A viscous fluid model for multimodal non-rigid image registration using mutual information. *IEEE Transactions on Medical Imaging*, 27(1):442–456, 2003.
- [37] M. DeCraene, G. Piella, O. Camara, N. Duchateau, E. Silva, A. Doltra, J. D’hooge, J. Brugada, M. Sitges, and A. Frangi. Temporal diffeomorphic free-form deformation: Application to motion and strain estimation from 3d echocardiography. *Medical Image Analysis*, 16(1):427–450, 2012.
- [38] S. C. L. Deoni, B. K. Rutt, and T. M. Peters. Rapid combined T-1 and T-2 mapping using gradient recalled acquisition in the steady state. *Magnetic Resonance in Medicine*, 49(3):515–526, 2003.
- [39] S. B. Donaldson, G. Betts, S. C. Bonington, J. J. Homer, N. J. Slevin, L. E. Kershaw, H. Valentine, C. M. L. West, and D. L. Buckley. Perfusion estimated with rapid dynamic contrast-enhanced magnetic resonance imaging correlates inversely with vascular endothelial growth factor expression and pimonidazole staining in head-and-neck cancer: A pilot study. *International Journal of Radiation Oncology Biology Physics*, 81(4):1176–1183, 2011.
- [40] S. B. Donaldson, C. M. L. West, S. E. Davidson, B. M. Carrington, G. Hutchinson, A. P. Jones, S. P. Sourbron, and D. L. Buckley. A comparison of tracer kinetic models for T1-weighted dynamic contrast-enhanced MRI: Application in carcinoma of the cervix. *Magnetic Resonance in Medicine*, 63(1):691–700, 2010.
- [41] H. Dongfeng, M Daqing, and J. Erhu. Dynamic breast magnetic resonance imaging: pretreatment prediction of tumour response to neoadjuvant chemotherapy. *Clinical Breast Cancer*, 12(2):94–101, 2012.

- [42] A. Drory, C. Haubold, S. Avidan, and Hamprecht F.A. Semi-global matching: a principled derivation in terms of message passing. *Pattern Recognition*, 8753(1):43–53, 2014.
- [43] P.E. Dugdale, K. A. Miles, I. Bunce, B. Kelley, and A.C. Leggett. CT measurement of perfusion and permeability within lymphoma masses and its ability to assess grade, activity, and chemotherapeutic response. *Journal of Computer Assisted Tomography*, 23(4):540–547, 1999.
- [44] S.M. Eschmann, F. Paulsen, M. Reimold, H. Dittmann, S. Welz, G. Reischl, H.J. Machulla, and R. Bares. Prognostic impact of hypoxia imaging with 18f-misonidazole PET in non-small cell lung cancer and head and neck cancer before radiotherapy. *Journal of Nuclear Medicine*, 46(2):253–60, 2005.
- [45] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient belief propagation for early vision. *Intl. J. Computer Vision*, 70(1):41–54, 2006.
- [46] P. F. Felzenszwalb and R. Zabih. Dynamic programming and graph algorithms in computer vision. *IEEE Transactions PAMI*, 33(4):721–740, 2011.
- [47] R. Garcia Figueiras, V. Goh, A. R. Padhani, A. Bermudez Naveira, A. Gomez Caa-
mano, and C. Villalba Martin. The role of functional imaging in colorectal cancer. *American Journal of Roentgenology*, 195(1):54–66, 2010.
- [48] J.C. Gee and R. Bajcsy. Elastic matching: Continuum mechanical and probabilistic analysis. *Brain warping*, 1999.
- [49] B. Glocker. *Random fields for image registration*. PhD thesis, Technische Universität München, 2011.
- [50] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios. Dense image registration through MRFs and efficient linear programming. *Medical Image Analysis*, 12(6):731–741, 2008.

- [51] L.H. Gray, A.D. Conger, M. Ebert, S Hornsey, and O.C. Scott. The concentration of oxygen dissolved in tissues at the time of irradiation as a factor in radiotherapy. *British Journal of Radiology*, 26(1):638–48, 1953.
- [52] A.K. Gupta, G.J. Cerniglia, and R. Mick. HIV protease inhibitors block akt signalling and radiosensitize tumour cells both in vitro and in vivo. *Cancer Research*, 65(1):8256–8265, 2005.
- [53] A.K. Gupta, W. G. McKenna, and C.N. Weber. Local recurrence in head and neck cancer: relationship to radiation resistance and signal transduction. *Clinical Cancer Research*, 8(885-892), 2002.
- [54] J.V. Hajnal, D.L. Hill, and D.J. Hawkes. *Medical image registration*. CRC, 2001.
- [55] J. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. Unpublished manuscript, 1971.
- [56] V. Hamy, N. Dikaios, S. Punwani, A. Melbourne, A. Latifoltojar, J. Makanyanga, M. Chouhan, E. Helbren, A. Menys, S. Taylor, and D. Atkinson. Respiratory motion correction in dynamic MRI using robust data decomposition registration - application to DCE-MRI. *Medical Image Analysis*, 18(1):301–313, 2014.
- [57] D. Hanahan and R.A. Weinberg. The hallmarks of cancer. *Cell*, 100(1):57–70, 2000.
- [58] D. Hanahan and R.A. Weinberg. Hallmarks of cancer: The next generation. *Cell*, 144(5):646–674, 2011.
- [59] C. Harvey, A. Dooher, J. Morgan, M. Blomley, and P. Dawson. Imaging of tumour therapy responses by dynamic CT. *European Journal of Radiology*, 30(3):221–226, 1999.
- [60] R. Hashemi, W. Bradley, and C. Lisanti. *MRI: The Basics*. Lippincott and Williams and Wilkins, 2010.

- [61] P. Hayton, M. Brady, L. Tarassenko, and N. Moore. Analysis of dynamic MR breast images using a model of contrast enhancement. *Medical image analysis*, 1(3):207–224, 1997.
- [62] M. Heilmann, C. Walczak, J. Vautier, J.L. Dimicoli, C.D. Thomas, M. Lupu, J. Mispelter, and A. Volk. Simultaneous dynamic T1 and T2* measurement for AIF assessment combined with DCE-MRI in a mouse tumour model. *MAGMA*, 20(4):193–203, 2007.
- [63] M. P. Heinrich, M. Jenkinson, M. Bhushan, T. Matin, F. V. Gleeson, J. M. Brady, and J. A. Schnabel. Mind: Modality independent neighbourhood descriptor for multi-modal deformable registration. *Medical Image Analysis*, 16(7):1423–35, 2012.
- [64] M. P. Heinrich, M. Jenkinson, M. Brady, and J. A. Schnabel. MRF-based deformable registration and ventilation estimation of lung CT. *IEEE Transactions on Medical Imaging*, 32(7):1239–48, 2013.
- [65] M. P. Heinrich, I. Simpson, M. Jenkinson, M. Brady, and J. A. Schnabel. Uncertainty estimates for improved accuracy of registration-based segmentation propagation using discrete optimization. In *MICCAI SATA Workshop*, 2013.
- [66] R. Hermans, M. Meijerink, W. Van den Bogaert, A. Rijnders, C. Weltens, and P. Lambin. Tumour perfusion rate determined noninvasively by dynamic computed tomography predicts outcome in head-and-neck cancer after radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 57(5):1351–6, 2003.
- [67] E.J Hill, M. Enescu, N. West, J.M. Franklin, K.Y. Chu, J. Li, D. Whittington, T.P. MacGregor, C. Azraf, C. Blessing, S. Mukherjee, M. Betts, A. Slater, E.M. Anderson, G. Brown, J. A. Schnabel, M. Partridge, P. Quirke, and R.A. Sharma. Oral nelfinavir before and during radiation therapy for rectal cancer: Changes in tumour perfusion and correlation between tissue and radiological markers of response. In *GI ASCO*, 2014.

- [68] H. Hirschmueller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions PAMI*, 30(2):328–341, 2008.
- [69] M. Holden. A review of geometric transformations for nonrigid body registration. *IEEE Transactions on Medical Imaging*, 27(1):111–128, 2008.
- [70] M. A. Horsfield, J. S. Thornton, A. Gill, H. R. Jager, A. N. Priest, and B. Morgan. A functional form for injected MRI Gd-chelate contrast agent concentration incorporating recirculation, extravasation and excretion. *Physics in Medicine and Biology*, 54(9):2933–49, 2009.
- [71] M. Ingrisch and S. Sourbron. Tracer kinetic modeling of dynamic contrast-enhanced MRI and CT: a primer. *J. Pharmacokinetics and Pharmacodynamics*, (40):281–300, 2013.
- [72] S. Kabus, T. Klinder, K. Murphy, B. van Ginneken, C. Lorenz, and J.P.W. Pluim. Evaluation of 4D-CT lung registration. In *MICCAI Proc.*, 2009.
- [73] Willi A. Kalender. *Computed Tomography*. Publicis, 2011.
- [74] J.H. Kappes, A. Bjoern, Hamprecht F.A., C. Schnorr, S. Nowozin, D. Batra, S. Kim, Kausler B.H., T. Kroger, J. Lellmann, N. Komodakis, B. Savchynskyy, and C. Rother. A comparative study of modern inference techniques for structured discrete energy minimization problems. *preprint*, 2014.
- [75] L. V. Karsa, T. A. Lignini, J. Patnick, R. Lambert, and C. Sauvaget. The dimensions of the CRC problem. *Best Practice and Research in Clinical Gastroenterology*, 24(4):381–396, 2010.
- [76] R. G. J. Kierkels, W. H. Backes, M. H. M. Janssen, J. Buijsen, R. G. H. Beets-Tan, P. Lambin, G. Lammering, M. C. Oellers, and H. J. W. L. Aerts. Comparison between perfusion computed tomography and dynamic contrast-enhanced magnetic

- resonance imaging in rectal cancer. *International Journal of Radiation Oncology Biology Physics*, 77(2):400–408, 2010.
- [77] S.H. Kim, A. Kamaya, and J.K. Willmann. CT perfusion of the liver: Principles and applications in oncology. *Radiology*, 272(2):322–344, 2014.
- [78] S. Klein, M. Staring, K. Murphy, M.A. Viergever, and J.P.W. Pluim. Elastix: a toolbox for intensity based medical image registration. *IEEE Transactions on Medical Imaging*, 29(1):196–205, 2010.
- [79] S. Klein, M. Staring, and J.P.W. Pluim. Evaluation of optimization methods for nonrigid medical image registration using mutual information and b-splines. *IEEE Transactions on Image Processing*, 16(1):2879–90, 2007.
- [80] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Transactions PAMI*, 28(10):1568–83, 2006.
- [81] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions PAMI*, 26(1):147–59, 2004.
- [82] N. Komodakis and G. Tziritas. Approximate labelling via graph cuts based on linear programming. *IEEE Transactions PAMI*, 29(1):1436–53, 2007.
- [83] E. Konukoglu, O. Clatz, B. Menze, B. Stieltjes, M.A. Weber, E. Mandonnet, H. Delingette, and N. Ayache. Image guided personalization of reaction-diffusion type tumour growth models using modified anisotropic eikonal equations. *IEEE Transactions on Medical Imaging*, 29(1):77–95, 2010.
- [84] J. G. Korpelaar, C. A. T. van den Berg, M. J. P. van Osch, G. Groenendaal, M. van Vulpen, and U. A. van der Heide. Phase-based arterial input function measurements in the femoral arteries for quantification of dynamic contrast-enhanced (DCE) MRI and comparison with DCE-CT. *Magnetic Resonance in Medicine*, 66(5):1267–1274, 2011.

- [85] M. J. Ledesma-Carbayo, J. Kybic, M. Desco, A. Santos, M. Suhling, P. Hunziker, and M. Unser. Spatio-temporal nonrigid registration of ultrasound cardiac motion estimation. *IEEE Transactions on Medical Imaging*, 24(9):1113–1126, 2005.
- [86] X. Li, L. Arlinghaus, B. Chakravarthy, J. Farley, I. Mayer, V. Abramson, M. Kelley, I. Meszoely, J. Means-Powell, and T. Yankeelov. Early DCE-MRI changes after longitudinal registration may predict breast cancer response to neoadjuvant chemotherapy. In *WBIR*, 2012.
- [87] X. Li, B. Welch, B. Chakravarthy, L. Xu, L. Arlinghaus, J. Farley, I. Mayer, M. Kelley, I. Meszoely, J. Means-Powell, V. Abramson, A. Grau, J. Gore, and T. Yankeelov. Statistical comparison of dynamic contrast-enhanced MRI pharmacokinetic models in human breast cancer. *Magnetic Resonance in Medicine*, 68(1):261–271, 2012.
- [88] Y. Liu, S. Sadowski, A. Weisbrod, E. Kebebew, R. Summers, and J. Yao. Multi-modal image driven patient specific tumour growth modelling. In *MICCAI Proc.*, 2013.
- [89] D. Loeckx, P. Slagmolen, F. Maes, D. Vandermeulen, and P. Suetens. Nonrigid image registration using conditional mutual information. *IEEE Transactions on Medical Imaging*, 29(1):19–29, 2010.
- [90] X. Lu and Y. Kang. Hypoxia and hypoxia-inducible factors: Master regulators of metastasis. *Clinical Cancer Research*, 16(24):5928–35, 2010.
- [91] R. Luypaert, S. P. Sourbron, and J. de Mey. Validity of perfusion parameters obtained using the modified tofts model: a simulation study. *Magnetic Resonance in Medicine*, 65(5):1491–1497, 2011.
- [92] H. Lyng, O. Haraldseth, and E.K. Rofstad. Measurement of cell density and necrotic fraction in human melanoma xenografts by diffusion weighted magnetic resonance imaging. *Magnetic Resonance in Medicine*, 43(6):828–36, 2000.

- [93] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16(2):187–198, 1997.
- [94] D. Mahapatra. Joint segmentation and groupwise registration of cardiac perfusion images using temporal information. *Journal of Digital Imaging*, 26(1):173–182, 2013.
- [95] R. Materne, B.E. Van Beers, A.M. Smith, I. Leconte, J. Jamart, J.P. Dehoux, A. Keyeux, and Y. Horsmans. Non-invasive quantification of liver perfusion with dynamic computed tomography and a dual-input one compartmental model. *Clinical Science*, 99(6):517–525, 2000.
- [96] J.R. McClelland, D. J. Hawkes, T. Schaeffter, and A.P. King. Respiratory motion models: A review. *Medical Image Analysis*, 17(1):19–42, 2013.
- [97] D. M. McDonald and P. L. Choyke. Imaging of angiogenesis: from microscope to clinic. *Nature Medicine*, 9(6):713–25, 2003.
- [98] M. Medved, G. Karczmar, C. Yang, J. Dignam, T. F. Gajewski, H. Kindler, E. Vokes, P. MacEneany, M. T. Mitchell, and W. M. Stadler. Semiquantitative analysis of dynamic contrast enhanced MRI in cancer patients: Variability and changes in tumor tissue over time. *Journal of Magnetic Resonance Imaging*, 20(1):122–128, 2004.
- [99] A. Melbourne, D. Atkinson, M. J. White, D. Collins, M. Leach, and D. Hawkes. Registration of dynamic contrast-enhanced MRI using a progressive principal component registration (PPCR). *Physics in Medicine and Biology*, 52(17):5147–56, 2007.
- [100] A. M Mendrik, E. Vonken, B. van Ginneken, H.W. de Jong, A. Riordan, T. van Seeters, E.J. Smit, M.A. Viergever, and M. Prokop. Tips bilateral noise reduction

- in 4D CT perfusion scans produces high-quality cerebral blood flow maps. *Physics in Medicine and Biology*, 56(13):3857–3872, 2011.
- [101] C.T. Metz, S. Klein, M. Schaap, T. van Walsum, and W. Niessen. Nonrigid registration of dynamic medical imaging data using nd+t b-splines and a groupwise optimization approach. *Medical Image Analysis*, 15(1):238–249, 2011.
- [102] K. A. Miles and M. R. Griffiths. Perfusion CT: a worthwhile enhancement? *British Journal of Radiology*, 76(904):220–31, 2003.
- [103] K. A. Miles, M.P. Hayball, and A.K. Dixon. Functional images of hepatic perfusion obtained with dynamic CT. *Radiology*, 188(2):405–411, 1993.
- [104] K.A. Miles, M.R. Griffiths, and M.A. Fuentes. Standardized perfusion value: universal CT contrast enhancement scale that correlates with FDG PET in lung nodules. *Radiology*, 220(2):548–553, 2001.
- [105] M. Modat, T. Vercauteren, G. R. Ridgway, D. J. Hawkes, N. C. Fox, and S. Ourselin. Diffeomorphic demons using normalised mutual information, evaluation on multi-modal brain MR images. In *SPIE Medical Imaging*, 2010.
- [106] J. H. Naish, D. M. McGrath, L. J. Bains, K. Passera, C. Roberts, Y. Watson, S. Cheung, M. B. Taylor, J. P. Logue, D. L. Buckley, J. Tessier, H. Young, J. C. Waterton, and G. J. M. Parker. Comparison of dynamic contrast-enhanced mri and dynamic contrast-enhanced ct biomarkers in bladder cancer. *Magnetic Resonance in Medicine*, 66(1):219–26, 2011.
- [107] E. Naydenova, J. A. Schnabel, and R.A. Sharma. Tumour segmentation and characterisation based on spatial heterogeneity. Technical report, University of Oxford, 2013.
- [108] K. Newbold, I. Castellano, E. Charles-Edwards, D. Mears, A. Sohaib, M. Leach, P. Rhys-Evans, P. Clarke, C. Fisher, K. Harrington, and C. Nutting. An exploratory

- study into the role of dynamic contrast-enhanced magnetic resonance imaging or perfusion computed tomography for detection of intratumoral hypoxia in head-and-neck cancer. *International Journal of Radiation Oncology Biology Physics*, 74(1):29–37, 2009.
- [109] J. P. B. O’Connor, A. Jackson, G. J. M. Parker, C. Roberts, and G. C. Jayson. Dynamic contrast-enhanced MRI in clinical trials of antivascular therapies. *Nature Reviews Clinical Oncology*, 9(1):167–177, 2013.
- [110] J. P. B. O’Connor, P. S. Tofts, K. A. Miles, L. M. Parkes, G. Thompson, and A. Jackson. Dynamic contrast-enhanced imaging techniques: CT and MRI. *British Journal of Radiology*, 84(S2):112–120, 2011.
- [111] B.D.P. O’Neill, G. Brown, R. Heald, D. Cunningham, and D.M. Tait. Non-operative treatment after neo-adjuvant chemoradiotherapy for rectal cancer. *Lancet Oncology*, 8(1):625–633, 2007.
- [112] M. R. Orton, J. A. D’Arcy, S. Walker-Samuel, D. J. Hawkes, D. Atkinson, D. J. Collins, and M. O. Leach. Computationally efficient vascular input function models for quantitative kinetic modelling using DCE-MRI. *Physics in Medicine and Biology*, 53(5):1225–1239, 2008.
- [113] A.R. Padhani, C. Hayes, L Assersohn, T. Powles, A. Makris, J. Suckling, M.O. Leach, and J.E. Husband. Prediction of clinicopathologic response of breast cancer to primary chemotherapy at contrast-enhanced MR imaging: initial clinical results. *Radiology*, 239(2):361–371, 2006.
- [114] A.R. Padhani and J.E. Husband. Dynamic contrast-enhanced MRI studies in oncology with an emphasis on quantification, validation and human studies. *Clinical Radiology*, 56(8):607–620, 2001.

- [115] G.J.M. Parker, C. Roberts, A. Macdonald, G.A. Buonaccorsi, S. Cheung, D.L. Buckley, A. Jackson, Y. Watson, K. Davies, and G.C. Jayson. Automated arterial input function extraction for T1-weighted dce-MRI. In *Proc. ISMRM*, 2003.
- [116] G.J.M. Parker, C. Roberts, A. Macdonald, G.A. Buonaccorsi, S. Cheung, D.L. Buckley, A. Jackson, Y. Watson, K. Davies, and G.C. Jayson. Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced MRI. *Magnetic Resonance in Medicine*, 56(5):993–1000, 2006.
- [117] G.J.M. Parker, J. Suckling, S.F. Tanner, A.R. Padhani, P.B. Revell, J.E. Husband, and M.O. Leach. Probing tumor microvasculature by measurement, analysis and display of contrast agent uptake kinetics. *Jmri-Journal of Magnetic Resonance Imaging*, 7(3):564–74, 1997.
- [118] G.J.M. Parker, S.F. Tanner, and M.O. Leach. Pitfalls in the measurement of tissue permeability over short time-scales using multi-compartment models with a low temporal resolution blood input function. In *Proc. ISMRM*, 1996.
- [119] K. Passera, L. Mainardi, D. McGrath, J. Naish, D. L. Buckley, S. Cheung, Y. Watson, A. Counce, G. Buonaccorsi, J. P. Logue, M. B. Taylor, C. Taylor, J. C. Waterton, H. Young, and G. J. M. Parker. A non-linear registration method for dce-mri and dce-ct comparison in bladder tumors. In *ISBI*, 2008.
- [120] U.B. Patel, F. Taylor, L. Blomqvist, C. George, H. Evans, P. Tekkis, P. Quirke, D. Sebag-Montefiore, B. Moran, R. Heald, A. Guthrie, N. Bees, I. Swift, K. Pennert, and G. Brown. Magnetic resonance imaging-detected tumour response for locally advanced rectal cancer predicts survival outcomes: Mercury experience. *Journal of Clinical Oncology*, 29(28):3753–60, 2011.
- [121] X. Pennec, P. Cachier, and N. Ayache. Understanding the "Demon's Algorithm": 3D non-rigid registration by gradient descent. In *MICCAI*, 1999.

- [122] D. Perperidis, R.H. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *Medical Image Analysis*, 9(5):441–456, 2005.
- [123] J. Peyrat, H. Delingette, M. Sermesant, C. Xu, and N. Ayache. Registration of 4D cardiac CT sequences under trajectory constraints with multichannel diffeomorphic demons. *IEEE Transactions on Medical Imaging*, 29(7):1351–1368, 2010.
- [124] R. E. Port, M. V. Knopp, and G. Brix. Dynamic contrast-enhanced MRI using Gd-DTPA: Interindividual variability of the arterial input function and consequences for the assessment of kinetics in tumors. *Magnetic Resonance in Medicine*, 45(6):1030–38, 2001.
- [125] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.
- [126] R. Prim. Shortest connection networks and some generalizations. *Bell System Technical Journal*, (36), 1957.
- [127] N. Qayum, R. J. Muschel, J. H. Im, L. Balathasan, C. J. Koch, S. Patel, W. G. McKenna, and E. J. Bernhard. Tumor vascular changes mediated by inhibition of oncogenic signaling. *Cancer Research*, 69(15):6347–54, 2009.
- [128] J.G. Rajendran, D.L. Schwartz, J. O’Sullivan, L.M. Peterson, P. Ng, J. Scharnhorst, J.R. Grierson, and K.A. Krohn. Tumour hypoxia imaging with [F-18] fluoromisonidazole positron emission tomography in head and neck cancer. *Clinical Cancer Research*, 12(18):5435–41, 2006.
- [129] T. Rohlfing. Image similarity and tissue overlaps as surrogates for image registration accuracy: Widely used but unreliable. *IEEE Transactions on Medical Imaging*, 31(2):153–163, 2012.

- [130] C. Rother, V. Kolmogorov, V.S. Lempitsky, and M. Szummer. Optimizing binary mrfs via extended roof duality. *CVPR*, 2007.
- [131] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.
- [132] W. Shi, M. Jantsch, P. Aljabar, L. Pizzaro, W. Bai, H. Wang, D. O’Regan, X. Zhuang, and D. Rueckert. Temporal sparse free-form deformations. *Medical Image Analysis*, 17(1):779–789, 2013.
- [133] A. Sotiras, C. Davatzikos, and N. Paragios. Deformable medical image registration: A survey. *IEEE Transactions on Medical Imaging*, 32(7):1153–1190, 2013.
- [134] S. P. Sourbron and D. L. Buckley. On the scope and interpretation of the tofts models for DCE-MRI. *Magnetic Resonance in Medicine*, 66(3):735–745, 2011.
- [135] S. P. Sourbron and D. L. Buckley. Tracer kinetic modelling in MRI: estimating perfusion and capillary permeability. *Physics in Medicine and Biology*, 57(2):1–33, 2012.
- [136] S. P. Sourbron and D. L. Buckley. Classic models for DCE-MRI. *NMR in Biomedicine*, 26(1):1004–1027, 2013.
- [137] C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition*, 32(1):71–86, 1999.
- [138] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother. A comparative study of energy minimization methods for markov random fields. In *ECCV*, 2006.
- [139] C. Tanner, J. A. Schnabel, D. Chung, M. J. Clarkson, D. Rueckert, D. L. G. Hill, and D. J. Hawkes. Volume and shape preservation of enhancing lesions when

- applying non-rigid registration to a time series of contrast enhancing MR breast images. *MICCAI Proc.*, 2000.
- [140] J. P. Thirion. Image matching as a diffusion process: an analogy with maxwell's demons. *Medical image analysis*, 2(3):243–60, 1998.
- [141] P. S. Tofts. Modeling tracer kinetics in dynamic gd-dtpa mr imaging. *Jmri-Journal of Magnetic Resonance Imaging*, 7(1):91–101, 1997.
- [142] P. S. Tofts, G. Brix, D. L. Buckley, J.L. Evelhoch, E. Henderson, M. V. Knopp, H.B.W. Larsson, T.Y. Lee, N.A. Mayr, G. J. M. Parker, R. E. Port, J. Taylor, and R. M. Weisskoff. Estimating kinetic parameters from dynamic contrast-enhanced T1 weighted MRI of a diffusable tracer: Standardized quantities and symbols. *Journal of Magnetic Resonance Imaging*, 10(1):223–232, 1999.
- [143] P. S. Tofts and A. G. Kermode. MEASUREMENT OF THE BLOOD-BRAIN-BARRIER PERMEABILITY AND LEAKAGE SPACE USING DYNAMIC MR IMAGING .1. FUNDAMENTAL-CONCEPTS. *Magnetic Resonance in Medicine*, 17(2):357–367, 1991.
- [144] J. Vandemeulebroucke, S. Rit, J. Kybic, P. Clarysse, and D. Sarrut. Spatiotemporal motion estimation for respiratory-correlated imaging of the lungs. *Medical Physics*, 38(1):166–78, 2011.
- [145] J. Vandemeulebroucke, D.Sarrut, and P. Clarysse. The POPI-model, a point-validated pixel-based breathing thorax model. In *Intl. Conf. Comput. Radiation Therapy*, 2007.
- [146] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Non-parametric diffeomorphic image registration with the demons algorithm. In *MICCAI Proc.*, 2007.
- [147] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Diffeomorphic demons: Efficient non-parametric image registration. *Neuroimage*, 45(1):S61–S72, 2009.

- [148] M. Wainwright, T. Jaakkola, and A.S. Willsky. Map estimation via agreement on trees: message-passing and linear programming. *IEEE Transactions on Information Theory*, 51(11):3697–3717, 2005.
- [149] O. Warburg. On the origin of cancer cells. *Science*, 123:309–14, 1956.
- [150] H. J. Weinmann, M. Laniado, and W. Mutzel. Pharmacokinetics of gddtpa dimeglumine after intravenous-injection into healthy-volunteers. *Physiological Chemistry and Physics and Medical Nmr*, 16(2), 1984.
- [151] Y. Weiss and W.T. Freeman. On the optimality of solutions of the max-product belief propagation algorithm in arbitrary graphs. *IEEE Transactions on Information theory*, 47(2):723–735, 2001.
- [152] K. Wong, R. Summers, E. Kebebew, and J. Yao. Tumour growth prediction with hyperelastic biomechanical model, physiological data fusion and nonlinear optimization. *MICCAI Proc.*, 2014.
- [153] G. Wu, Q. Wang, J. Lian, and D. Shen. Estimating the 4d respiratory lung motion by spatiotemporal registration and super-resolution image reconstruction. *Medical Physics*, 40(3), 2013.
- [154] C. Yang, W. M. Stadler, G. Karczmar, M. Milosevic, I. Yeung, and M.A. Haider. Comparison of quantitative parameters in cervix cancer measured by dynamic contrast-enhanced MRI and CT. *Magnetic Resonance in Medicine*, 63(6):1601–1609, 2010.
- [155] X. Yang and M. V. Knopp. Quantifying tumour vascular heterogeneity with dynamic contrast-enhanced magnetic resonance imaging: A review. *Journal of Biomedicine and Biotechnology*, 2011.
- [156] T. Yankeelov. Integrating imaging data into predictive biomathematical and biophysical models of cancer. *ISRN Biomathematics*, 2012.

- [157] V.L. Yarnykh. Actual flip-angle imaging in the pulsed steady state: a method for rapid three-dimensional mapping of the transmitted radiofrequency field. *Magnetic Resonance in Medicine*, 57(1):192–200, 2007.
- [158] M. Yigitsoy, C. Wachtinger, and N. Navab. Temporal groupwise registration for motion modelling. *Information Processing in Medical Imaging*, 22(1):648–59, 2011.
- [159] M.A. Zahra, K.G. Hollingsworth, E. Sala, D.J. Lomas, and L.T. Tan. Dynamic contrast-enhanced MRI as predictor of tumour response to radiotherapy. *Lancet Oncology*, 8(1):63–74, 2007.
- [160] D. Zikic. *Contributions to Medical Image Registration*. PhD thesis, Technische Universitat Munchen, 2011.
- [161] F. G. Zoellner, R. Sancee, P. Rogelj, M. J. Ledesma-Carbayo, J. Rorvik, A. Santos, and A. Lundervold. Assessment of 3D DCE-MRI of the kidneys using non-rigid image registration and segmentation of voxel time courses. *Comput. Med. Imag. Graphics*, 33(3):171–181, 2009.