Use of Multicomponent Non-Rigid Registration to Improve Alignment of Serial Oncological PET/CT Studies

Yanni Papastavrou

UCL

A thesis submitted to University College London

for the degree of Doctor of Philosophy (Ph. D.)

in Medical Physics applied to Nuclear Medicine

January 2015

Declaration

I, Yanni Papastavrou, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Non-rigid registration of serial head and neck FDG PET/CT images from a combined scanner can be problematic. Registration techniques typically rely on similarity measures calculated from voxel intensity values; CT-CT registration is superior to PET-PET registration due to the higher quality of anatomical information present in this modality. However, when metal artefacts from dental fillings are present in a pair of CT images, a non-rigid registration will incorrectly attempt to register the two artefacts together since they are strong features compared to the features that represent the actual anatomy. This leads to localised registration errors in the deformation field in the vicinity of the artefacts. Our objective was to develop a registration technique which overcomes these limitations by using combined information from both modalities.

To study the effect of artefacts on registration, metal artefacts were simulated with one CT image rotated by a small angle in the sagittal plane. Image pairs containing these simulated artifacts were then registered to evaluate the resulting errors. To improve the registration in the vicinity where there were artefacts, intensity information from the PET images was incorporated using several techniques. A well-established B-splines based non-rigid registration code was reworked to allow multicomponent registration. A similarity measure with four possible weighted components relating to the ways in which the CT and PET information can be combined to drive the registration of a pair of these dual-valued images was employed.

Several registration methods based on using this multicomponent simi-

larity measure were implemented with the goal of effectively registering the images containing the simulated artifacts. A method was also developed to swap control point displacements from the PET-derived transformation in the vicinity of the artefact. This method yielded the best result on the simulated images and was evaluated on images where actual dental artifacts were present.

Dedication

I dedicate this work to my *Babushka*, Maria Kolomei, whose love and encouragement in life were boundless, and whose tales of Stalinist Russia were nearly impossible to fathom for us growing up in the West. Alas, she did not live to see the completion of this work.

Acknowledgements

This work would not have happened without the attentive supervision of Brian Hutton, my primary supervisor, whose guidance throughout this project was tremendously appreciated. The many discussions we had are only partially reflected in this work and I am hugely indebted to his many ideas through this research. His encouragement when most needed was hugely appreciated. I also wish to offer appreciation to my second supervisor, Dave Hawkes, for the stimulating discussions in the earlier part of this work and for inviting me to participate in the high quality and stimulating academic environment of the Centre for Medical Image Computing at UCL under his guidance. Also I'm very grateful for the valuable help and guidance from Tryphon Lambrou in the latter stages of this work, his input was most useful and much welcomed. And to the many others in the Institute for Nuclear Medicine at UCL who helped, particularly Alexander Bousse and Kjell Erlandsson. I am also thankful for the support of my co-workers, especially the guidance I received from both Alan Green and Linda Goodyear.

But most of all to my dear wife Ruthi for her support, love and advice and tremendous encouragement that helped this work into fruition, especially when the going was tough. And for putting up with much and sacrificing many weekends to patiently proof-read the draft copies of this text, as well as her help with making diagrams. And to my friends and family, especially my mother, not least for encouraging my interest in science when I was a youngster, despite her lack of technical knowledge.

Nomenclature

1-D 1 Dimensional

3-D 3 Dimensional

BGO Bismuth germanium oxide

CISG Computational Imaging Sciences Group

DICOM Digital imaging and communications in medicine

FOV Field of view

GPU Graphics processing unit

MCSM Multicomponent similarity measure

NRR Non-rigid registration

OSEM Ordered subsets expectation maximisation

PDF Probability density function

VOI Volume of interest

CPDS Control point displacement substitution

CPS Control Point Spacing

CT X-Ray computed tomography

FDG Fluourodeoxyglucose

FFD Free form deformation

HU Hounsfield Units

MRI Magnetic resonance imaging

PET Positron emission tomography

RPM Revolutions per minute

SUV Standard uptake value

VTK Visualisation toolkit (Kitware, Inc.)

$\vec{T}(\vec{r})$ Transformation acting on coordinates $\vec{r}$

δx, δy and δz Voxel spacings in the x,y and z directions

S The similarity measure

R Redundancy

$I(\vec{r}, t_0)$ The fixed image

$I(\vec{r}, t_1)$ The first floating image in the sequence of scans. The second and third floating images are dennoted by $I(\vec{r}, t_2)$ and $I(\vec{r}, t_3)$ respectively.

$I_T(\vec{r})$ Fixed image (alternative symbol for)

$I_F(\vec{r})$ Floating image (alternative symbol for)

f Intensity bin label

μ_x Linear attenuation coefficient

H Entropy

μ Step size

p Probability

N_v Number of voxels

$G(f)$ Grey level histogram

MI	Mutual information
NMI	Normalised mutual information
SUC	Symmetric uncertainty coefficient
l,m,n	Voxel indices
i,j,k	Control point indices
M	Matrix size
S	Number of slices in image

Contents

List of Figures

Chapter 1

Introduction

1.1 Introduction

Patients undergoing treatment for cancer often have serial PET/CT studies for a variety of reasons. In some cases, this is to track the response to treatment. For example, different treatment regimens may be compared by studying the anatomical and physiological response as measured using these imaging modalities. In addition, this imaging could be used for long-term follow-up.

It is useful to register these studies in order to define changes to both volume and counts. Once aligned, segmented regions can be identified and then propagated onto any of the subsequent serial images. When these regions are placed onto the PET images, the counts contained within them can be used to assess the physiological response.

Usually CT is used for registration due to its superior resolution and anatomical detail. However, registration is particularly problematic in the head and neck region because it is an area of complex and densely packed anatomy. It also has mobile structures such as the tongue and the jaw, and serial studies can involve quite complex changes due to tumour growth/regression, as well as possible changes in body habitus as a result of treatment. Accuracy of registration can also be affected by the presence

of artefacts, a common problem in the head and neck region due to dental fillings and devices.

The aim of this work was to assess registration issues in the head and neck region, and to investigate the combined use of both PET and CT data via multicomponent registration as a means of improving the registration in serial PET/CT studies. To this end, a well-established non-rigid registration code, designed for use with single modality images was re-coded to facilitate multicomponent use.

1.2 Structure of the Thesis

In chapter 2, the relevant background on PET/CT imaging in the context of this project is discussed. Other combined imaging systems currently under development are also briefly outlined. The relevant background on image registration techniques, followed by a description of the intensity-based methods used in this project, are presented in chapter 3. A discussion on information theoretic similarity measures is also presented.

Chapter 4 describes how the optimal non-rigid registration scheme was developed and evaluated for serial head and neck combined FDG PET/CT images acquired during a clinical trial, using a multi-scale B-splines based non-rigid registration algorithm. The work described in this chapter ensured the correct parameters were set to values that will achieve good results for head and neck images. A problem was identified with poor alignment for images containing dental artefacts, which then motivated the work in the following chapter.

In chapter 5 a multicomponent registration method, to help improve registration accuracy when there are dental artefacts present in the CT image, was developed and tested. Intensity information from the PET image was used, with the goal of improving the registration in the region where the artefact is present.

The multicomponent registration technique developed to mitigate the errors caused by the simulated artefacts was tested on patient data in chap-

ter 6. This was assessed and compared to a registration technique based on paired CT images.

The final chapter summarises the main methods, results and conclusions from the project. Some ideas are outlined for possible future directions in which this work could be developed.

1.3 Nature of the Contributions

The work undertaken in this thesis contributes to the large body of work in image registration in the following ways:

- Exploration of why intensity histogram transformations are useful to perform when non-rigidly registering CT images.
- Development of a framework for registering multicomponent image data (i.e. using intensity information from both PET and CT).
- Demonstration that combined use of PET and CT registration can reduce the effect of dental artefacts.

Chapter 2

Combined PET/CT Imaging

2.1 Introduction

This chapter introduces the basic principles of combined PET/CT imaging relevant to this project. A brief overview of how PET imaging works is given first, after which the principles behind CT imaging are outlined. The salient features of the combined PET/CT scanner are then described. Finally, other combined imaging systems currently under development will be briefly outlined. Throughout, the emphasis is on cancer imaging, rather than the other clinical applications that combined PET/CT is used for, such as cardiac [Flotats et al., 2011] or brain imaging [Berti, Pupi and Mosconi, 2011].

2.2 PET Imaging

2.2.1 Introduction

Positron Emission Tomography (PET) involves measuring the radiation emitted from a patient containing a radioisotope that undergoes positron decay, then reconstructing these measured data into an image. These radioisotopes are usually chemically synthesised into a compound (known as a radio-pharmaceutical) and administered to the patient, usually through

an intravenous injection. PET imaging utilises the tracer principle to explore the biochemical pathway followed by the radio- pharmaceutical through the body. There are many standard textbooks that overview the basic principles of PET imaging, such as [Smith F. W., 1998] or [Cherry, Sorenson and Phelps, 2012] which contain chapters on the fundamental aspects of PET imaging.

2.2.2 Positron Emission and Annihilation

Positron emitting isotopes undergo positron decay. The most commonly used isotope for PET imaging is ^{18}F. It is unstable and decays (with a half-life of approximately 110 minutes) to ^{18}O through positron decay by emitting a positron (β^+) with mean kinetic energy of approximately 250 keV and a neutrino (ν) [Cherry, Sorenson and Phelps, 2012]:

$$^{18}F \rightarrow {}^{18}O + \beta^+ + \nu \tag{2.1}$$

Once emitted, the energetic positron will in time reach thermal equilibrium within the local environment within the tissue following multiple interactions. The mean range of the positron in water is 0.61mm [Cal-González et al., 2011]. The range will depend on local chemistry and density of the medium (i.e. it will be larger within lung tissue which has comparatively low density) and will result in a blurring of the reconstructed PET image. Under certain imaging conditions, the effect of the positron range may become a significant limitation, for example, in small animal PET scanners, or when using isotopes that emit more energetic positrons with a large positron range, such as ^{82}Rb. Corrections for the positron range can be used [Agbeko et al., 2010; Jødal, Le Loirec and Champion, 2012; Cal-González et al., 2011; Haber, Derenzo and Uber, 1990] to de-blur the image, given knowledge of the point spread function by measuring point sources, but with the drawback of enhancing noise [Haber, Derenzo and Uber, 1990].

Once thermalised, the positron will briefly form a metastable state with an electron known as positronium [Moses, 2011] that may be described by a hydrogen-like wavefunction before annihilating. When a positron

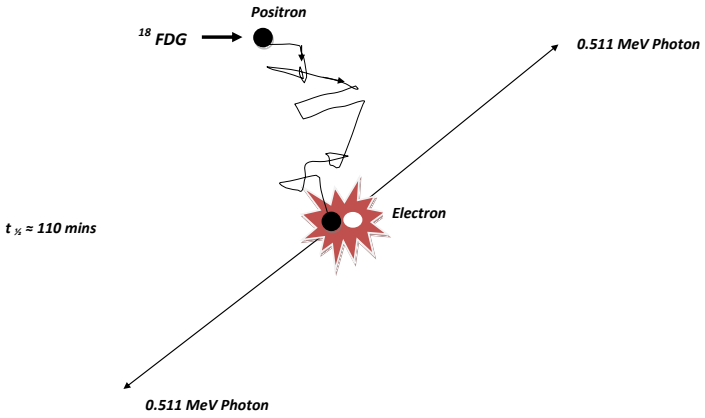

Figure 2.1: Positron annihilation

and electron annihilate, the resultant photons contain their rest-mass energy plus any kinetic energy (which is negligible by comparison in practice). The process of single photon production through annihilation is forbidden since a second body is required by the process to conserve momentum. The predominant decay mode involves the production of two 511 keV photons which are emitted at an angle of 180^{o} relative to each other to conserve momentum[1]. Departures from co-linearity, the effect of positron range and other contributions to the fundamental limits of the spatial resolution that is possible are discussed in [Moses, 2011]. Figure 2.1 illustrates the predominant processes leading up to the annihilation of the emitted positron into two co-linear 511 keV gamma photons.

2.2.3 PET Imaging Systems

PET imaging utilises the coincident detection of co-linear annihilation gamma photons to infer the line of response along which the annihila-

[1]There will be slight departures from 180^{o} that will depend on the momentum of the electron-positron pair prior to annihilation

tion event occurred. This is the principle behind dual photon coincidence imaging as compared to single photon imaging that requires the construction of a physical collimator.

The gamma photon detector consists of a volume of scintillator crystal coupled to photomultiplier devices. In the case of the GE Systems ST 16 PET/CT scanner used for this project, BGO is used for the scintillator. The gamma photon deposits energy inside the volume of the scintillator and produces visible photons that are then detected using photomultipliers that are optically coupled to the back edge of the scintillator crystal. The photomultipliers create electrical signals from which the position of interaction within the crystal and energy of the incident gamma photon may be inferred.

A PET scanner consists of a ring of detectors that are optimised to detect the 511 keV co-linear photon pairs emitted by a positron source. The annihilation event is presumed to have occurred somewhere[2] along the line of response of the positions of the two detected photons: this is done by having a coincidence logic operating in the PET pulse processing circuitry which will count an event as a true coincidence, provided detection of the photon pair falls within a small timing window. There are several kinds of coincidence events in addition to the true coincidence events, as illustrated in figure 2.2, which shows a schematic PET system of ring detectors in operation. A patient containing a positron-emitting isotope is shown inside a PET scanner. Within the patient, pairs of co-linear 511 keV annihilation gamma photons are emitted at electron-positron annihilation events at locations 1, 2, 3. Gamma interactions can occur inside the body: 1a is the Compton scattering of the photon resulting in a detection of the scattered photon and in event 3a, the scattered photon is lost to the detection system. The emergent gamma photons are detected at positions A, B, C, D and E. Event 2, a true coincidence, is correctly

[2]Time of flight based instruments do better than this by using a scintillation crystal with a substantially faster optical response than BGO and fast pulse-processing electronics to help deduce the position of the annihilation event along a segment of the line by measuring the time difference between the arrival of the two photons [Marsden, 2003; Karp et al., 2008].

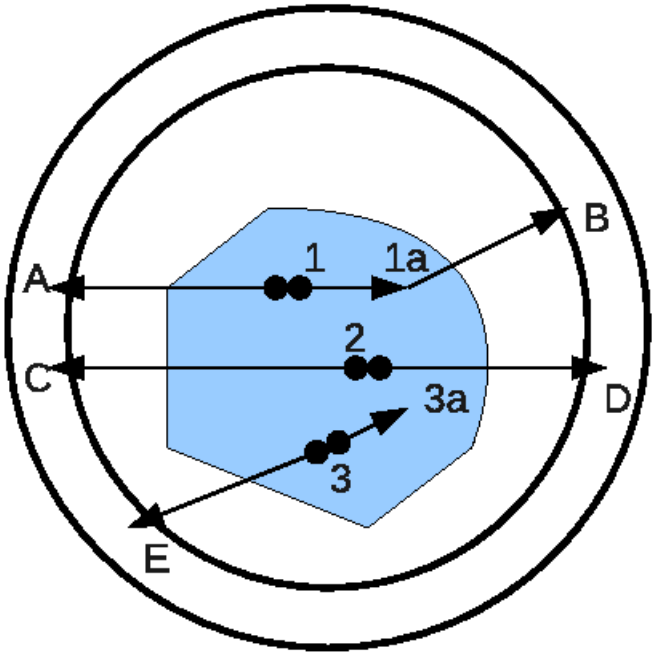

Figure 2.2: Types of detection of gamma photons arising from positron annihilation

attributed along the line of response passing through points C and D. The detected photons at A and B will contribute to noise, since the lines of response that are deduced by the instrument does not contain the actual position of the annihilation event. A further loss of signal can also occur, in addition to the cases shown in this diagram, when more than two photons are detected within the coincidence time threshold. In this case, the signal is rejected by the instrument's logic circuitry and therefore the information contained in the pairs of coincidence gammas will be lost. These are known as "multiples". The final possibility is when two photons are detected within the timing window that do not originate from the same annihilation event. These are known as "randoms" and will contribute to the noise.

When travelling through the body, the photons are attenuated by the matter they encounter. The probability of interaction per unit length is known as the linear attenuation coefficient, $\mu_x(x)$ which will vary according to the local composition. The total attenuation of a flux of photons will be given by integrating along the length l of the line through which the radiation passes:

$$I(x) = I_0 e^{-\int^l \mu_x(x) dx} \tag{2.2}$$

Correcting for this attenuation is an important step during the PET reconstruction process and presumes a knowledge of the attenuation coefficient. Prior to the advent of the combined PET/CT scanner, this correction was performed prior to the PET scan by performing a transmission scan using a positron emitting source [Zaidi and Hasegawa, 2003]. When corrected for attenuation and scatter, the counts given for a line of response are proportional to the integrated activity along that line. A set of parallel lines of response at a certain angle is known as a projection through the object at that angle. Reconstruction is performed by collecting many such projections at a set of angles through the image and then using the Filtered Backprojection Algorithm [Cherry, Sorenson and Phelps, 2012] to infer the distribution of tracer within the object by back projecting along the lines of response to construct the image. This algo-

rithm suffers from the noise characteristics inherent to the PET imaging process with its poor count statistics when compared to CT, for example.

Algorithms that can model the emission and detection process and solve the reconstruction problem iteratively [Hutton, Hudson and Beekman, 1997] through an optimisation process, are also used. These iterative methods proceed from an initial estimate of the image. From this first image estimate, forward projections are calculated. These calculated projections are then compared to the measured projections and the difference found, as evaluated using a cost function. The second image estimate is then found according to an update rule that seeks to minimise the cost function. This process is iterated and successive image estimates are generated using the update rule, until a termination criterion is satisfied - for example, when the change in the cost function is sufficiently small, the estimated image is assumed to have converged. A well-known and computationally efficient method for this is the Ordered Subsets Expectation Maximisation (OSEM) algorithm [Hudson and Larkin, 1994]. The PET scans used during this work were reconstructed using this technique. The detailed principles behind the instrumentation of a PET scanner is well described by many standard texts [Wernick M. N., 2004; Smith F. W., 1998; S., 1998; H., 2006; Cherry, Sorenson and Phelps, 2012].

2.2.4 Using ^{18}FDG to Map Metabolic Activity

^{18}Fluourodeoxyglucose (FDG) is a radio-pharmaceutical that is widely used in positron emission tomographic (PET) imaging to trace the physiological pathway of glucose within the body. The assumption being that when administered to a patient, it will behave in a biologically equivalent manner to glucose within the complex biochemical pathways that exist. It is a radio-labeled molecule; the bound ^{18}F isotope is unstable and decays (with a half-life of approximately 110 minutes). Following administration to the patient according to the relevant protocol, the patient will then be imaged in a PET scanner.

The reconstructed image forms a 3D map of activity of the tracer within the body, with a typical spatial resolution of the order of 6mm FWHM

[Teräs et al., 2007] for FDG-PET. For the clinician, the resultant image from an FDG-PET scan can be thought of as a map of glucose metabolism, which is an indicator of metabolic activity, since glucose is consumed by cells as a source of energy. The uptake of FDG is standardised by normalising the concentration of activity to the patient's weight and the injected activity: this is known as the Standard Uptake Value (SUV) [Cherry, Sorenson and Phelps, 2012].

Uptake of FDG in cells depends on how metabolically active they are, and is used as a means for detecting tumour cells, which are metabolically hyperactive. There is also uptake in other cells. The amount of uptake will depend on the intrinsic metabolic rate of the cell and on the rate of locally available FDG that gets transported to the vicinity of the cell. This will depend on the local microvasculature.

2.2.5 Some Considerations when using the SUV

The SUV is defined in terms of the activity concentration in the tissue, C, in units of MBq/ml; the injected activity I, in units of MBq; and the weight of the patient, M, in kg as follows:

$$SUV = \frac{C}{(I/M)} \tag{2.3}$$

The SUV is defined over a volume of interest (VOI) in the image, these regions may be any size or shape. Frequently, the SUV is calculated for each voxel. Depending on the clinical questions being asked, there are various ways to extract SUV parameters from the distribution present within a VOI. A typical use in oncological PET imaging is to quote the hottest pixel within a region, known as SUV_{max} since it has been shown that this is a good indicator of how likely there are to be clonogenic cells present within the region [Kidd and Grigsby, 2008]. This metric could then be used to track how well a tumour is responding to treatment, a lower value would likely indicate a positive response to treatment. Alternative measures used to track the response of a tumour to treatment are the average SUV value within an VOI SUV_{mean} [Higgins et al., 2012] where the VOI can be

defined using a variety of methods, either manually delineated using an expert observer, or using more automated segmentation techniques such as an intensity isocontour method, or using sophisticated segmentation algorithms that can give results to sub-voxel accuracy [Zeng, Shepherd and Zwiggelaar, 2012].

Suppose a VOI is defined around a tumour using an isocontour method where the threshold value is comparatively close to the SUV_{max}, and so this results in a relatively small-volumed VOI. Now, suppose the response to treatment is being quantified by propagating this VOI onto a PET scan acquired at some later date using an image registration technique to do so. If the SUV_{max} is being used as a measure to assess the response to treatment, then the accuracy of the registration method used to propagate the registration is not so important, since the likelihood is that the hottest pixel is contained within the VOI. The accuracy of the registration only becomes an issue in the case when this pixel is located outside the VOI (which is more probable for a small VOI). Therefore, a reasonable criterion to aim for, in terms of accuracy of the registration, is that it is comparable in size to the PET voxel dimensions, or slightly worse.

Consider the case when this same VOI segmentation method is used, but now the SUV_{mean} is used to track the response to treatment. Suppose there are large intensity gradients in the vicinity of the VOI. For a small VOI, there will be a large fraction of the voxels in the VOI that are on the surface. Also, when the threshold is quite close to the SUV_{max} then the surface of this VOI will likely be on a region of high SUV intensity gradient. When this volume is propagated using registration onto the subsequent PET scan, the accuracy of the registration algorithm - and also the interpolation methods - will now be more important, given the existence of large intensity gradients, and the large proportion of surface voxels. Hence, a sub-voxel registration accuracy becomes important in this case, where the intensity is varying rapidly. An accuracy significantly better than half of the voxel spacing seems appropriate. Improved registration accuracy is also important when generating groupwise statistics across many patients.

2.2.6 Limitations of Functional Imaging

The bio-distribution of FDG is therefore of a highly functional nature. It does, however, contain some anatomical structure (many tissues will also metabolise FDG, but to a much lesser extent than tumours) without which nuclear medicine physicians would have an impossible task in visualising the locations of sites of uptake. This lack of spatial context is especially true for 3-deoxy-3-^{18}fluorothymidine (FLT), a tracer that is used to map cell proliferation [Shields et al., 1998]. As a label for cell proliferation, rapidly dividing cells such as tumour cells have high uptake, but the uptake of FLT can be so low in other tissues that it is far harder to localise anatomy than for FDG-PET images.

The limited anatomical information from FDG-PET images was a motivating factor for image registration techniques developed in the late 1980's and 1990's, to place fused images acquired from PET with those from a modality that contained higher anatomical detail, such as X-ray Computed Tomography (CT) and also Magnetic Resonance Imaging (MRI), since it is important to have anatomical localisation for the FDG uptake. Image registration is described in further detail in chapter 3. This was possible mainly for brain images, since the problem is spatially constrained: the motions within the anatomy constrained within the skull is minimal compared to other parts of the body, for which image registration is more problematic due to issues such as respiratory motion and changes in posture [Beyer et al., 2010].

However, the lack of spatial information contained within PET (and other Nuclear Medicine modalities) images must be put in the context that the sensitivity of the system is extremely high, enabling extremely low quantities of tracer down to the nanomolar or picomolar levels [Weng, Ding and Volkow, 1999] to be detected.

2.3 CT Imaging

In contrast to PET imaging which involves measuring the emitted radiation to infer functional information about the patient, X-ray Computed

Tomography (CT) imaging involves sending X-rays through the patient and measuring their transmission through the patient, with the goal of inferring structural information about the patient.

A source of X-rays is generated by accelerating thermionically emitted electrons in a vacuum towards a target that is typically Tungsten. This results in a spectrum of bremsstrahlung photons, whose shape is governed by the geometry of the target and the accelerating potential of the electrons, and whose maximum energy is given by the accelerating potential, kVp, which is typically 120 kV. Characteristic X-rays are also produced and are emitted following the removal of an inner orbital electron (eg a K-shell or an L-shell electron) in the target atom by the accelerated electron. These X-rays, which are characteristic of the target material atomic number, are produced when an outer electron transitions into the missing inner orbit. These events generate characteristic lines on the spectrum. The intensity of the beam of radiation is also governed by the rate at which the electrons are produced from the heated filament. The beam of X-rays is filtered and collimated to optimise the imaging characteristics and to limit the ionising radiation dose of the beam to the patient.

On the opposite side of the patient to the X-ray source is the detector system used to measure the transmitted radiation through the patient. The detector elements consist of a scintillator that converts the X-ray photon into a visible light pulse, and photo diodes are then used to detect this signal, which is then processed by the electronics and stored.

The X-ray source and detectors are mounted on a gantry and rotated around the patient, using a slip ring to allow for continuous rotation whilst maintaining electrical connectivity.

As the source and detector rotate around the couch on which the patient is lying, the couch moves continuously through the scanning bore, with the result that the trajectory of the source-detector system is a helix. The pitch of the helix is given by the speed of couch travel compared to the speed of rotation of the source-detector assembly. Since the flux of

Material	HU Range
Air	-1000
Lung	$-600 \leq HU \leq -400$
Fat	$-100 \leq HU \leq -60$
Water	0
Muscle	$40 \leq HU \leq 80$
Bone	$400 \leq HU \leq 1000$

Table 2.1: Hounsfield Units for certain substances

photons is high, the image acquisition is quite rapid: the rate of rotation around the patient is often of the order of 120 RPM. There is also another variant for CT imaging which is the step and shoot method, where the couch is moved incrementally through the ring and each acquisition for a given fixed couch position is circular.

Within the patient, the X-ray attenuation is dependent on the locally varying attenuation coefficient, which depends on the local chemical composition and density. The goal of CT imaging is to infer these attenuation coefficients, which in turn provide structural information about the patient. This involves solving the reconstruction problem using the measured projection data as an input, using techniques such as Filtered Backprojection (or iterative reconstruction techniques). The convention is to relate these linear attenuation coefficients (μ_x) at position x to a dimensionless unit known as the Hounsfield Unit (HU) [Hounsfield, 1980] which is expressed in relative terms to the linear attenuation of water μ_{water}:

$$HU = \frac{\mu_x - \mu_{water}}{\mu_{water}}.1000 \tag{2.4}$$

The Hounsfield Units for a 120 kV CT system for various substances may be seen in table 2.1. HU values will depend on the energy of the CT scanner.

CT imaging provides good anatomical information for the radiologist about the patient's anatomy and has high spatial resolution compared to PET [Cherry, Sorenson and Phelps, 2012]. It provides good contrast between

certain structures such as the lung and bones, but is less able to provide contrast between soft tissue structures. To help improve soft tissue contrast, intravenous contrast media may be used to assist with delineating certain structures [Mamourian, 2013]. For imaging tasks where good contrast between soft tissue structures is required, Magnetic Resonance Imaging (MRI) is a better choice. Furthermore, CT imaging is limited when it comes to functional imaging, although it is possible using contrast agents to gain insights into certain physiological processes and perform bolus-type studies [van Beek and Hoffman, 2008].

2.4 Combined PET/CT Imaging

Efforts to combine brain images from PET and MRI or CT using software-based image registration techniques motivated Townsend et al [Townsend D. W.; Beyer T., 1994; Beyer et al., 2000; Townsend, Beyer and Blodgett, 2003] to develop the dual modality combined PET/CT scanner in the 1990's [Weng, Ding and Volkow, 1999]. The combined PET/CT scanner enables the functional information from the PET image to be registered to an anatomical image from the CT image by mounting both systems onto a single gantry and imaging the patient during a single session.

The PET/CT scanner has a gantry onto which is mounted a ring of PET detectors to image the emitted 511 keV photons along with a co-axially mounted multi-slice CT scanner. The requirements of the photon detection systems are too different to use the same detector types for both the transmitted X-ray photons that are of peak energy 120 kV and the emission photons of energy 511 keV (and at a much lower intensity). This therefore requires that the two different detector systems be mounted adjacently. During imaging, the patient is moved through these adjacent detection systems, which will acquire their images sequentially. In figure 2.3, a combined PET/CT scanner is shown, illustrating the adjacent detector systems. The unavoidable axial displacement between these two systems can introduce errors due to involuntary movements by the patient between successive scans, for example due to respiration. These mismatches can cause errors in registration and can induce artefacts in

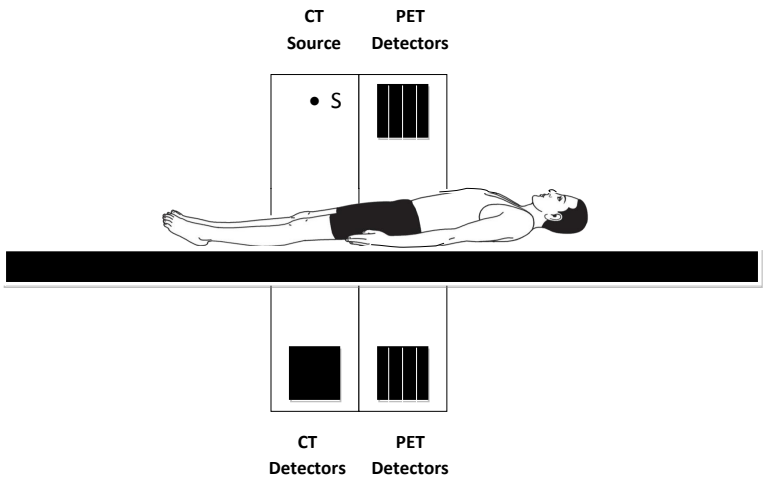

Figure 2.3: A combined PET/CT scanner

the PET image caused when performing the attenuation correction [McQuaid and Hutton, 2008].

When constructing an algorithm to register serial combined PET/CT images, the intensity information from the CT images will provide a better registration under most conditions, since CT is a more anatomical modality. But if incorporating information from the PET images for the registration is appropriate, for example, if there are metal artefacts present within the CT images, then information from the non-attenuation corrected PET images can be used instead, as mentioned in section 5.1. Otherwise, these attenuation artefacts in the PET images will have an adverse effect on the registration. So it is not likely that the kind of mismatches described by [McQuaid and Hutton, 2008] would have any knock on effects on this process of serial registration, since intensity information from non-attenuation corrected PET images would be used.

The attenuation correction for the PET imaging system is derived from

the CT scan [Burger et al., 2002; Kinahan et al., 1998]. PET/CT systems are now becoming increasingly widely available in Nuclear Medicine departments and there is now a notable decline in the supply of new PET-only systems worldwide [Townsend, 2008].

One interesting development in the design of PET/CT detector systems is a novel means of having a two layered detector, as outlined in a granted patent [Herrmann, 2012]. When operated in CT mode, the detectors of the first layer detect transmission radiation to generate the CT image, and the detectors of the second layer detect PET or SPECT radiation to generate data for reconstruction into an emission image. The patent claims that since the detectors of the first and second layers are aligned, the transmission and emission images will therefore be inherently better aligned than for a system in which the two systems are mounted adjacently.

There have been significant improvements in the diagnosis, management and planning of treatment for patients as a result of the introduction of the combined PET/CT scanner into clinical use [Vach et al., 2011]. This is in addition to the large corpus of research work that has been made possible with the advent of the combined PET/CT scanner.

2.5 Other Multi-Modality Imaging

In addition to the success of PET/CT imaging, SPECT-CT imaging also has a significant role to play in the Nuclear Medicine clinic [Schillaci, 2005]. In addition, following on from the success of PET/CT combined imaging, other hybrid combined imaging systems, such as MRI-PET, have recently entered into clinical use, following a long history of pre-clinical development work. In terms of imaging benefits, MRI offers high spatial resolution and high soft tissue contrast compared to CT, and also an ability to perform functional imaging [Catana et al., 2006; Cherry, 2009; Pichler et al., 2010]. However, the attenuation correction is not as easy to infer from the MRI image as it is from a CT image [Pichler et al., 2010]. One useful by-product of an MRI scanner's large magnetic fields is

that it will cause the emitted positrons to spiral along the magnetic field line, thereby reducing the positron range in two of the three dimensions [Catana et al., 2006]. This may be particularly advantageous for isotopes where the emitted positron has high energy and, under certain imaging conditions, for instance, for small animal scanners. In these cases, large magnetic fields can result in PET images that have improved spatial resolution. The use of magnetic fields to improve the spatial resolution for small PET animal scanners has been proposed [Hammer, Christensen and Heil, 1994] following measurements to assess the effect under different magnetic field strengths, with promising results. This effect has been long-since known about by the PET community [Iida et al., 1986].

Tri-modality imaging yields an image containing information from combined PET, CT and MRI [Veit-Haibach et al., 2013] scans: albeit from a sequential acquisition from an initial MRI scan just after the patient is injected with FDG followed by a PET/CT scan. After the MRI scan, the patient is then swiftly transferred using a shuttle onto a combined PET/CT scanner and imaged. The area of multi-modal combined imaging remains an active area of research in many areas, including the recent developments in the field of omni-tomography [Wang and Yu, 2013] which employs limited field of view reconstruction that could in principle mean that multiple detector systems [Wang et al., 2012] from differing imaging modalities could be mounted onto a single gantry. Plans to build a prototype combined CT-MRI scanner have been outlined using this paradigm [Savage, 2013]. This work relies on development of novel detector systems combined with limited field of view reconstruction algorithms that are outlined in the references just given, and some bold claims are made about how this framework will in time be a paradigm shift in the field of medical imaging. In [Wang, 2014], the following vision of the future of medical imaging is offered for our consideration:

"We envision that tomography will transcend current modality fusion towards what we call "the grand fusion" of all relevant modalities, namely, "omni-tomography", for truly simultaneous but often localized image reconstruction in terms of many contrast [sic] mechanisms such as CT, MRI,

PET, SPECT, and more. The challenge in fusing modalities for simultaneous imaging has been space conflict and other physical constraints. Traditionally, scanners of different types are longitudinally assembled but this arrangement cannot synchronize data collection. In contrast, the proposed CT-MRI project is a major step towards the holy grail of biomedical imaging – omni-tomography."

Such combined multi-modality medical imaging systems that utilise complimentary information are an area of active development. The non-rigid registration algorithm developed in this thesis that uses a multicomponent similarity measure to register combined PET/CT serial images will also be of use in future when the requirement for serial registration of multi-modality images is required.

Chapter 3

Image Registration Techniques

3.1 Introduction

This chapter provides an overview of image registration techniques, followed by a description of the specific methods used in this project. This will include introducing the relevant terminology and geometric representation of images that builds on the physics of the imaging modalities introduced in the previous chapter. The main techniques used in image registration as they relate to the work carried out in this project are surveyed.

The approach to image registration in this project was to use a voxel-based similarity measure. Images are registered by improving the similarity measure towards an optimal value by iteratively making changes to the geometric relationships in the floating image. The geometric transformation describes the geometric relationships, which may be rigid or non-rigid. The main kinds of non-rigid transformations in the literature are surveyed, including the method used in this project. The main kinds of optimisation techniques that are used in image registration literature were surveyed and some key concepts that determine the choice of optimisation scheme are outlined.

The concept of joint histograms was introduced, and it was explained how they may be used to calculate information theoretic similarity measures, and may be generalised in the $N * M$-dimensional case, where N and M are the number of intensity components in the fixed and floating images.

3.2 An Image as a Geometrically Ordered Set of Intensities

An image is a representation of an object, acquired by making measurements using an imaging system, often known informally as a scanner. The image is formed by measuring some spatially varying physical quantity, such as X-ray attenuation coefficient, activity, luminosity or temperature from the object at some acquisition time t. The image can be visualised later using a display device that renders the numerical values (known as "intensities") as an appropriate greyscale or colourscale within each element of the display, corresponding to the locations of the position of the measurement within the object. The pattern gained by the measurement process can be visualised, and clinical deductions are then made by the expert observer. These discrete elements of the image are known as voxels, for a three dimensional image, or pixels for a two dimensional image. For tomographic medical images, these inferred parameter values are derived from the reconstruction process at discrete positions in a three dimensional space, and are stored in computer memory as a representation of the object. A two dimensional section through this volume can then be displayed as a greyscale pattern. In general, an image can include multiple intensity values at each position, each representing a different physical parameter (or different frequency range for the same physical quantity), and can include a time series acquisition, known as a "dynamic sequence", that can be gated to a signal such as an ECG or a respiratory signal.

The (mono) intensity values $I(\vec{r}, t) \in \mathbb{R}$ are acquired at some particular imaging episode (time t) at discrete locations $\vec{r}$, defined within a coordinate system relative to the scanner. These intensities quantify some

typical value of the physical parameter being measured, such as an average value within the voxel at each discrete position $\vec{r} = \hat{\mathbf{i}}\mathrm{l}\delta\mathrm{x} + \hat{\mathbf{j}}\mathrm{m}\delta\mathrm{y} + \hat{\mathbf{k}}\mathrm{n}\delta\mathrm{z}$. These positions are separated by a voxel spacing of δx, δy and , δz, in each direction, and l, m, n are the voxel indices, $l = 0 \ldots M, m = 0 \ldots M, n = 0 \ldots S; l, m, n, M, S \in \mathbb{N}$ where M is known as the matrix size and S is the number of slices. The usual vector notation is used for $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$ and $\hat{\mathbf{k}}$ to denote the orthogonal unit vectors in the x, y and z directions respectively. In practice, for PET images, the voxel spacings are equal, whereas for CT images, the slice thickness is usually greater than the x-y spacings: $\delta z > \delta x = \delta y$. For CT images, the intensity values are the Hounsfield Units (HU), and the intensity values for PET images are usually the standard uptake values (SUV). The details of PET and CT imaging are covered in Chapter 2.

3.3 On Coordinate Transformations

Consider the following idealised imaging acquisitions at different acquisition times $I(\vec{r}, t_0)$ and $I(\vec{r}, t_1)$ of a test object, whose physical characteristics being imaged are time invariant. If this object is imaged at the same positional orientation with respect to the same scanner, then - for an idealised linear measurement system in the absence of noise - the images would be identical, so we can write $I(\vec{r}, t_0) = I(\vec{r}, t)$. If the spatial relationship between the object and the scanner changes between the acquisition of the first and second image because the test object was positioned differently, the images will no longer be identical and will require re-positioning as a post-imaging correction using a coordinate transformation.

This process of spatial correction is at the heart of the image registration process. However, what is known as image registration in the literature, was developed historically to align the images taken from two different imaging systems of the same patient, with the aim of combining the information gathered from the two different imaging modalities in a synergistic fashion, a far more enabling technology than a simple correction as described earlier.

A typical image registration task involves aligning images of a patient from a PET scanner acquired to visualise the FDG biodistribution that traces metabolic activity, but has limited anatomical information present with another modality that is able to distinguish anatomical features with good effect, such as CT or MRI. The registered images from each modality can then be visualised on the same display, for example using different colourscales. This is known as an image fusion process. The result is that the rather diffuse biodistribution of FDG can be visualised on top of an anatomical image that contains detailed structural information, so that the FDG can be localised accurately within the recognisable anatomical structures.

3.4 What is Image Registration?

Image registration can be defined as the process by which an optimum coordinate transformation $\vec{T}(\vec{r})$ is found that when applied to an image - known as a floating image - the salient structures within it are rendered into correspondence within a second image, known as a reference or target image. The registration process is cast as an optimisation problem whereby the floating image is iteratively transformed such that it becomes more "like" the fixed image. The "likeness" or similarity between the images is known as the similarity measure, or, in the language of optimisation theory, a cost function. The registration of the floating image, $B(\vec{r}, t_o)$, to the fixed image, $A(\vec{r})$, using a similarity measure S that is maximised[1] during optimisation may be written in very general terms as the following optimisation problem:

$$\max_{\vec{T}}\{S(A(\vec{r}), \mathrm{B}(\vec{T}(\vec{r})))\} \tag{3.1}$$

Where $\mathrm{B}(\vec{T}(\vec{r}))$ represents the transformation of the floating image's coordinates followed by an interpolation of the intensities at the voxel coordinates of the fixed image using an appropriate interpolation technique.

[1]Some similarity measures are a minimum when the images are registered, so *extremised* might be a more general word to substitute for *maximised*.

This general definition requires significant explanation, and several ingredients of the image registration process must be specified to implement a registration algorithm in practice.

3.4.1 A Generic Registration Algorithm

A typical scheme for registering a pair of images is shown in figure 3.1. The algorithm is given three inputs: the fixed image, $I_T(\vec{r})$, the floating image , $I_F(\vec{r})$, and an initial transformation, $\vec{T}_0(\vec{r})$. The output of the algorithm is the floating image transformed by the optimal transformation found by the registration algorithm, that has had the intensity values interpolated appropriately when the transformation was applied.

The registration process proceeds as follows. Given these inputs, the initial transformation is applied to the floating image. When transforming the coordinates of the floating image, this will also involve interpolating or approximating the intensity values to these new positions. There are several methods for this interpolation or approximation process, for example, nearest neighbour, linear, cubic-splines, B-splines or sinc interpolation, to list but a few. The similarity is then calculated between this transformed image and the fixed image. An optimisation algorithm is employed to update the transformation such that subsequent iterations of the algorithm will increase the similarity measure. This is repeated until a stopping condition is met, either if the similarity measure reaches a certain value, or the number of iterations is reached.

3.5 Optimisation Methods

The optimiser is the method used to modify the parameters that describe the transformation, such that the similarity is maximised and thereby driving the registration towards a solution, as shown formally in equation 3.1. The similarity measure S is an implicit function of the parameters that describe the transformation, $\vec{T}(\vec{r})$.

Consider a 1D example where there is only one degree of freedom, θ, parameterising the transformation. If the parameter is varied and the

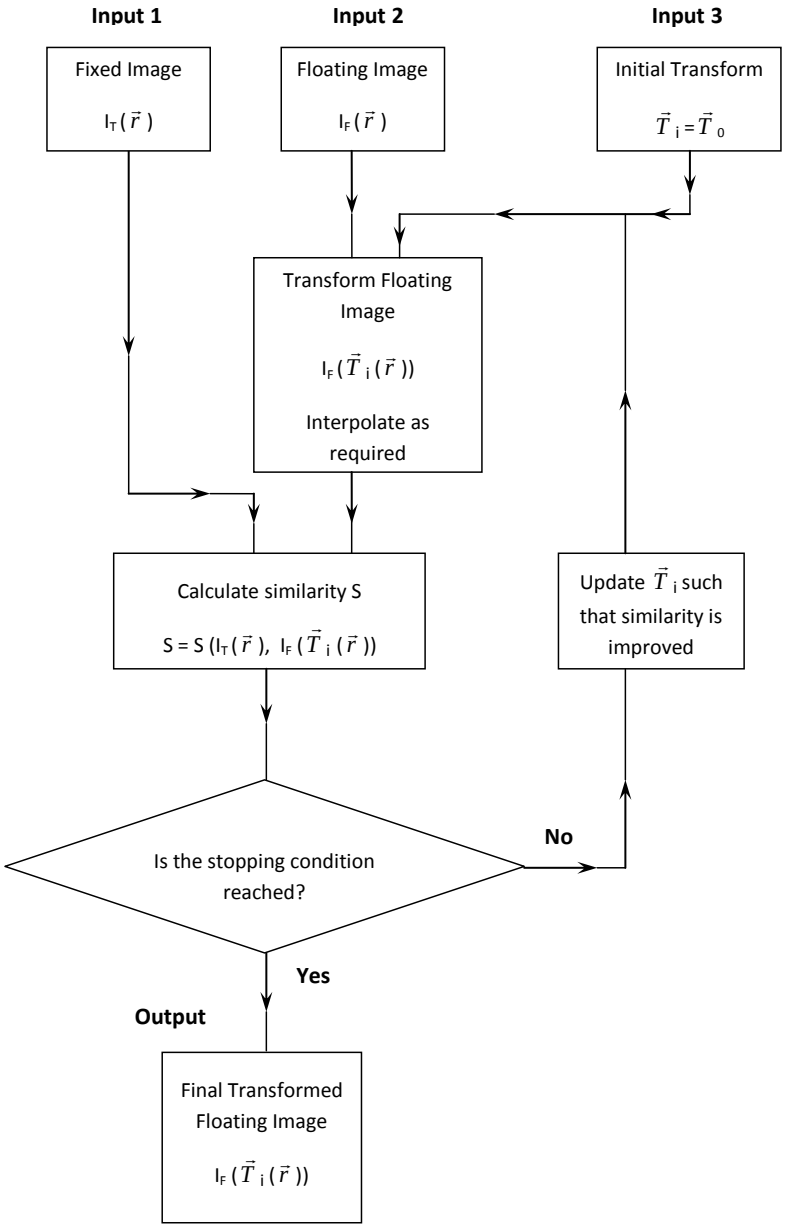

Figure 3.1: A typical registration scheme

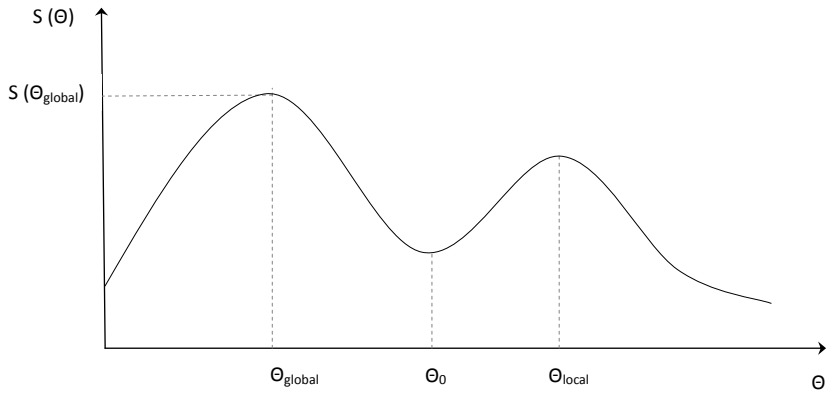

Figure 3.2: Capture range

similarity measure, $S(\theta)$, is plotted, it will be in general some kind of curve, as illustrated in figure 3.2 . The goal of an optimisation algorithm is to seek the value, θ_{global}, at which the similarity function is globally maximised. However, for most problems, the curve will have multiple local maxima for different values of θ, in addition to the global maximum value. An optimisation scheme proceeds from a starting value of $\theta_{initial}$ and then has a rule by which the next estimate is calculated, from one iteration to the next. There are many kinds of optimisation schemes used for image registration. The most often used, on account of its simplicity, is gradient ascent. This works by calculating the gradient of the similarity function and proceeding to update to the next iteration along the line of steepest ascent. It is more efficient than simple line-search methods and also, in the case of multi-dimensional problems, the gradient vector defining the update direction will point towards the steepest ascent.

Many of the simpler kinds of optimisation methods, such as steepest Gradient ascent, the Simplex algorithm, and Powell's method will, under certain conditions, find the wrong solution in the general case where there exist multiple local maxima. Such algorithms will only find the correct solution provided the initial value of the transform, $\theta_{initial}$, is within the capture range. In figure 3.2, provided $\theta_{initial}$, is within the capture range: $0 < \theta_{initial} < \theta_0$, then the solution found by the optimisation algorithm will tend towards θ_{global}. If $\theta_{initial} > \theta_0$, then the solution found by the optimisation algorithm will tend towards θ_{local} in an algorithm-dependent manner. The properties of the similarity measures in terms of the number of minima, capture range, accuracy and other ways of assessing the behaviour has been extensively studied for a range of problems, both for rigid registration [Skerl, Likar and Pernus, 2006], and for non-rigid registration [Škerl, Likar and Pernuš, 2008].

There are many possible solutions to the problem of getting stuck in an incorrect local maxima discussed in the literature. One of the most widely accepted and longest established methods is to interpolate the images to a coarse scale [Lau, Braun and Hutton, 2001; Mattes et al., 2003; Makihara, Mori and Yagi, 2011]. This has the effect of reducing the number of minima present in the similarity function when plotted as a function of the transformation parameters. The method is to interpolate the images to a coarse scale, then register the images at that scale. The solution found at the coarse scale is then an input transformation value for the next-most coarse scale. This is then repeated until the finest scale is reached. This hierarchical multi-resolution approach is an improvement and can help avoid falling into local minima, but does not guarantee it [Jenkinson and Smith, 2001].

An alternative approach is to employ an optimisation algorithm that can overcome the local maxima problem such as the simulated annealing method [Kolesov et al., 2013], that allows a non-deterministic sampling of the parameter space. A well explained introduction to optimisation algorithms can be found in Numerical Recipes in C [Press and Vetterling, 1992], although it must be noted that this text contains only a brief, some-

what dated, overview of the key techniques: the subject of optimisation theory has advanced greatly in recent years. A more recent introduction to this subject is presented in [Chong and Zak, 2013].

3.6 The Transformation Model

Image registration algorithms fall into three broad taxonomies. The first class of problems consider the registration of images from two objects that have not changed in their internal configuration between successive image acquisitions. The second and more challenging class of problems is when there have been changes within the object being imaged. These changes could include postural changes, changes due to treatment, or other physiological change. In principle, one could write down the equations that describe the time evolution of the body, based on the laws of physics, biology and so forth, and solve these equations to find out how the displacements at every point in the body relate to the points in the body at some later time, using in the first instances finite element models of the body that solve the biomechanical equations for the system [Bharatha et al., 2001; Kyriacou and Davatzikos, 1998; Samani et al., 2001]. In the second instance, biological models could be sought that model the pathological or physiological changes. For example models of tumour growth that would require knowledge of cell division rates and so on. Biophysical models of this nature whilst possible to consider, are extremely difficult in practice.

The third class of registration problem is the inter-patient case. In this case, it is not possible to build a direct, mechanistic physical model that relates the points in one object to the points in the second object, since the two images are not of the same object. There are of course common patterns present within the two images, since the objects are two examples from a group. Even so, it is possible to seek transformations to relate the geometry of one patient to another, although if there are topological differences, for example, normal variant anatomy present in one image but not the other, then this can become a very challenging problem.

This application is an extension of the original goal of image registration, whereby the information from two different imaging systems was fused. Now, the information from two different patients is being transformed into a common spatial framework, which is essentially a different class of problem.

The simple case of the rigid and affine transformations shall first be outlined and then the more complex types of non-rigid transformations will be discussed.

3.6.1 Transforming the Coordinates Within the Floating Image

Correspondence between the internal structures of the fixed and floating image[2] is achieved by modifying the coordinates of the floating image. Consider a transformation $\vec{T}$ that maps the position vectors in the floating image, $\vec{r}(x, y, z)$ to achieve this correspondence. This is written as:

$$\vec{T} : \vec{r} \rightarrow \vec{r'} \Leftrightarrow \vec{T}(\vec{r}) = \vec{r'} \tag{3.2}$$

The transformed coordinate system of the floating image is denoted by the primed coordinates: (x', y', z'). This spatial transformation $\vec{T}(\vec{r})$ is also sometimes known as a *transformation model*. Figure 3.3 illustrates how the transformation maps a point P from the floating image to the point P' in the fixed image. In this diagram, two different kinds of transformation will be required to "solve" these registration problems. In the upper problem, a rigid transformation is required. In the lower problem, in addition to the global rigid transform, locally non-rigid deformations are required to register the anatomy in the vicinity of the smiling mouth of the fixed image to the sad mouth in the fixed image.

When denoting the operation of transforming the floating image using the transformation $\vec{T}$, a remark on the notation used is required. Whenever an image is transformed, there will be an interpolation of intensity

[2]The *fixed image* is also known as the *target image* or *template image*, and the *floating image* is frequently referred to as the *source image.*

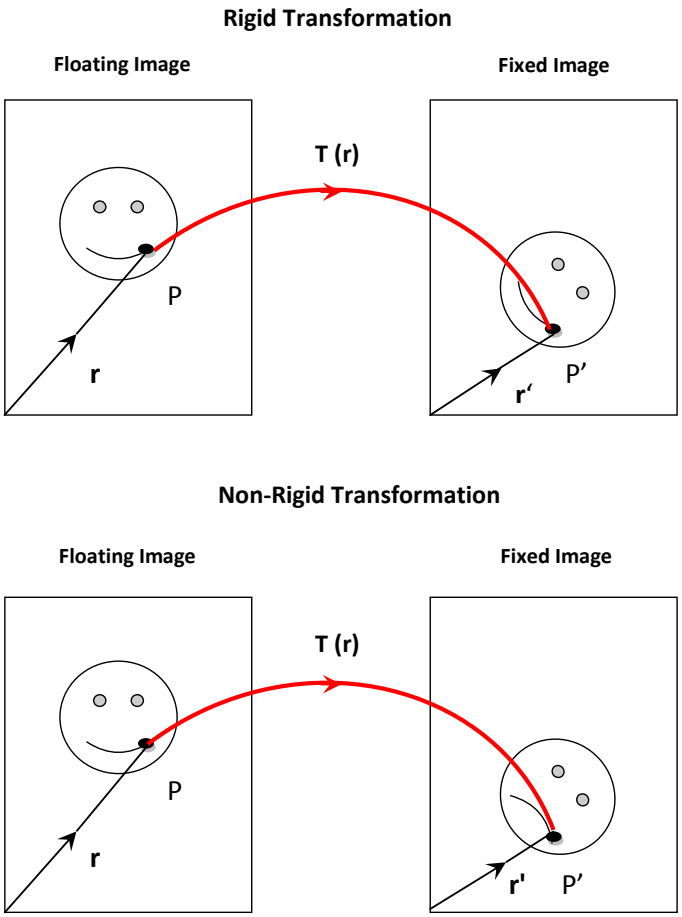

Figure 3.3: A rigid and non-rigid transform

values, apart from the unlikely case where $\vec{T}$ is a translation by an integral number of voxels. The final intensity values at each voxel in the floating image will be dependent on the kind of interpolation method used to render the intensities between the sample points at the coordinates of the voxels in the reference image.

The notation used in this work must be made explicit. Whenever a floating image is written with a transformation acting on the image coordinates like this: $I(\vec{T}(\vec{r}))$ the meaning of the notation in this context is that this represents a mapping of both the position *and the intensity* using an appropriate interpolation method to render the intensities at the voxel coordinates of the fixed image.

3.6.2 Rigid Transformations

It can be shown that a rigid transformation consists of a shift (a translation) and a rotation about the three axes. The case outlined in the first paragraph of section 3.2 would require a rigid transformation. Six degrees of freedom are required to represent a rigid body transformation: one shift parameter along each of the three orthogonal axes and one rotation parameter about each axis.

A translation by a distance of Δr_x, Δr_y and Δr_zin the x, y and z direction can be represented in matrix form as follows:

$$\begin{pmatrix} r'_x \\ r'_y \\ r'_z \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & \Delta r_x \\ 0 & 1 & 0 & \Delta r_y \\ 0 & 0 & 1 & \Delta r_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_x \\ r_y \\ r_z \\ 1 \end{pmatrix} \tag{3.3}$$

Similarly, a rotation about the x axis, through angle (θ) can be mapped using the following 4x4 matrix operation:

$$\begin{pmatrix} r'_x \\ r'_y \\ r'_z \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) & 0 \\ 0 & -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_x \\ r_y \\ r_z \\ 1 \end{pmatrix} \tag{3.4}$$

And a rotation about the y axis through ϕ is given by:

$$\begin{pmatrix} r'_x \\ r'_y \\ r'_z \\ 1 \end{pmatrix} = \begin{pmatrix} \cos(\phi) & 0 & \sin(\phi) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\phi) & 0 & \cos(\phi) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_x \\ r_y \\ r_z \\ 1 \end{pmatrix} \tag{3.5}$$

Finally, a rotation about the z axis through angle ρ is given by:

$$\begin{pmatrix} r'_x \\ r'_y \\ r'_z \\ 1 \end{pmatrix} = \begin{pmatrix} \cos(\rho) & \sin(\rho) & 0 & 0 \\ -\sin(\rho) & \cos(\rho) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_x \\ r_y \\ r_z \\ 1 \end{pmatrix} \tag{3.6}$$

By writing translations and rotations as 4x4 matrices, they may be combined via multiplication into a single 4x4 matrix that rigidly transforms the image. This is better than having multiple transformations acting on the image, since each time this happens, an interpolation of the image intensity values is required.

3.6.3 Affine Transformations

A more general geometric transformation, of which a rigid transformation is a sub-type, is the affine transformation. This is the simplest type of non-rigid transformation. In three dimensions, in addition to the six degrees of freedom required to define a rigid registration, there are an additional six degrees of freedom that also include magnification and shearing in each orthogonal direction. The transformation can be written in matrix form:

$$\begin{pmatrix} r'_x \\ r'_y \\ r'_z \\ 1 \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{23} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_x \\ r_y \\ r_z \\ 1 \end{pmatrix} \tag{3.7}$$

When a straight line is transformed according to an affine transformation, it remains a straight line. Also, ratios of distances between points are the same before and after an affine transformation.

3.6.4 Non-Rigid Transformations

There are many kinds of non-rigid transformations used in the literature covering a wide variety of applications [Holden, 2008; Sotiras, Davatzikos and Paragios, 2013]. The choice of transformation is very specific to the nature of the registration problem, and will depend on the modalities of the imaging systems; the region of anatomy; whether the registration task involved is intra-patient registration or inter-patient registration. Intra-patient registration could include the registration of serial images. Appropriate assumptions may be made about the transformation according to the requirements of the registration problem. For example, a locally rigid assumption [Fujiwara et al., 2011; Loeckx et al., 2004; Staring, Klein and Pluim, 2007] about the transformation might be valid in the intra-patient case, to penalise the deformation of rigid structures such as bones, but certainly this would be inappropriate in an inter-patient registration problem. Registering an image to an atlas is an example of an inter-patient registration problem and is useful for a wide variety

of problems, such as segmentation propagation of images of the brain in neuro-imaging [Crum, Scahill and Fox, 2001; Aljabar et al., 2009] or for radiotherapy treatment planning, for propagating segmentations for head and neck [Han et al., 2008] and for other anatomical regions of clinical interest, such as the prostate [Langerak et al., 2010]. For many registration problems, it is desirable for the deformations to be mathematically well behaved, having properties such as being continuous, having continuous derivatives, invertible and preserving the topology of the geometry (for example, it won't allow tears or folding in the geometry, nor sinks or holes). Transformations with these properties are known as diffeomorphic. This and other desirable properties of non-rigid transformations (such as volume preservation) used in registration are outlined in a review article [Holden, 2008]. However, it is quite possible to consider cases where non-diffeomorphic transformations are allowed in serial images, for example, to allow for the case of missing tissue following surgery, or to allow for tissue appearance, in the case of tumour growth [Zacharaki et al., 2009].

Some of the main types of transformation that are more widely used for image registration are outlined in the next sections. A recent review article [Sotiras, Davatzikos and Paragios, 2013] outlines the other main kinds of non-rigid transformation models used in image registration algorithms. The transformation used for our problem will be outlined in section 4.2.

3.6.5 Transformations Derived from Interpolation or Approximation

There are many kinds of geometric transformation models that can be employed. There are those [Lau, Braun and Hutton, 2001; Arsigny, Pennec and Ayache, 2003; Loktyushin et al., 2014; Seiler, Pennec and Reyes, 2011; Arsigny et al., 2006] that employ a local rigid registration at an increasingly fine scale within the image to match each block of image data in the floating and fixed images, and deriving a deformation from these blocks using interpolation. These methods only apply in the intra-patient

registration case and could be used to recover variations in posture between acquisitions, since these transformations can be well approximated by a rigid transformation at a certain choice of scale.

Purely geometric approaches have been taken in image registration that involve overlaying pseudo-landmark points of correspondence on the source and target image that form a sparse set of correspondence points, with the intermediate spatial locations found using interpolating or approximating spline functions of some kind. Several spline functions have been proposed for this task, including thin plate splines [Bookstein, 1989; Rohr et al., 2001; Rohr, Fornefett and Stiehl, 1999] and B-splines [de Boor, 1978], which have been used in a framework for deforming solid shapes in computer graphics, using a technique called free-form deformations [Sederberg and Parry, 1986]. This B-splines free-form deformation framework was incorporated into registration where a regular grid of pseudo-landmark points are placed on the floating image in the seminal work [Rueckert et al., 1999]. This approach to non-rigid registration remains popular for a wide variety of registration tasks, and has been re-coded to run in parallel on a GPU architecture [Modat et al., 2010] to speed up execution times.

One limitation that thin plate splines suffer from is that they have infinite support, so by varying the location of one correspondence point, the entire thin plate spline would need to be recomputed to define the new dense correspondence for the modified image. Since B-splines have local support, they do not suffer this limitation. B-splines based transformations will be discussed in further detail in section 4.2, since they are used for our problem.

3.6.6 Transformations Derived from Physical Models

There are approaches to non-rigid registration that involve modelling the deformation by solving the equations that model a physical process as an elastic solid. This treats the object as an elastic body [Alexander and Gee, 2000; Davatzikos, 1997; Shen and Davatzikos, 2002; Stammberger et al., 2000; Rohr et al., 2001].

Another model is the fluid model that treats the object as a compressible viscous fluid [Christensen, Rabbitt and Miller] and involves solving the Navier-Stokes equation. This has the advantage over an elastic model in that it is less constrained and highly local deformations are possible, rather analogous to the physical case, whereby a viscous fluid can flow into the corner of a cube more readily than a rubber ball could be squeezed into the same shape [Crum, Tanner and Hawkes, 2005]. This ability of capturing large changes at the local level is the reason that fluid models have been used for atlas registration problems. Another model that is appropriate for this task is the Demons algorithm [Thirion; Vercauteren et al., 2007]: the name comes from 'Maxwell's demon' thought experiment in statistical physics. It is based around effectors, or demons, which apply a force to locally transform the source image. The force was derived from an optical flow equation. The Demons algorithm is computationally inexpensive, and straightforward to parallelise [Muyan-Ozcelik et al., 2008].

A biomechanical model could, in principle, be built for certain intra-patient transformation models, to recover postural changes. Biomechanical models have also been used as a benchmark for testing non-rigid registration algorithms, for example: [Tanner et al., 2006] constructed a finite element model of a breast. More sophisticated models could also include attempts to model tumour growth [Araujo and McElwain, 2004; Alarcon, Byrne and Maini, 2003; Sherratt and Chaplain, 2001], and the resulting deformations of nearby structures that get squeezed out of their usual anatomical configuration by a dense growing mass of tumour, could be modelled using a biomechanical technique, although this would be very computationally intensive. The other difficulty in this case is that the nature of tumour growth is highly specific to local conditions [Roose, Chapman and Maini, 2007], and therefore the deformations caused by its growth will vary greatly among individuals, even with the same type of cancer. The problem of non-diffeomorphic registration was implemented [Zacharaki et al., 2009] to model tumour growth and tissue loss in the brain.

3.7 Types of Similarity Measure

There are three categories of similarity measures employed during image registration: landmark based, feature based and voxel based. Landmark based techniques involve manually identifying corresponding anatomical structures within the source and target image. Feature based techniques involve extracting features from the images such as lines, edges, surfaces or volumes. Voxel based techniques use the intensity values themselves as an automatic method of performing registration. There are advantages and disadvantages to each method.

As more information is used from the images during the registration process, the chances are that the registration process will be more robust and accurate. Voxel based methods employ a similarity measure that depends on all voxel intensity values, rather than derived information such as features, which, in addition to being tedious and time consuming to define, can inject errors due to the reliance on observers to define the points. However, voxel based methods sometimes do not take into consideration the spatial information of the intensity values.

Registration algorithms based on manually defined landmark points typically use a least squares error from each pair of points in the fixed and floating image. These techniques were typically used for rigid registration. Similarity measures can be constructed from the raw intensity values themselves, utilising the full richness of information contained in the image, with the added bonus that, for some applications, this can be an entirely automatic process. Measures that can equate the likeness of both images have been developed using ideas from information theory. This is discussed in detail in the following chapter. In general, the similarity measure, S, is a function that depends on the fixed image and the current estimate of the transformed floating image, $\mathrm{B}(\vec{r'})$ that has had its coordinates transformed according to $\vec{r'} = \vec{T}(\vec{r})$:

$$S = S(A(\vec{r}), \mathrm{B}(\vec{r'})) \tag{3.8}$$

The similarity measure is then an implicit function of the parameters

that are used to define the transformation. Voxel based similarity measures can be constructed directly from the intensity values within the image. The similarity measure will enable two rival transformations to be compared: the one with the higher similarity is deemed to be better.

Similarity measures based on absolute intensity differences can only be used where there is a direct proportionality between the intensity reading of the physical parameter being measured between the two imaging systems, for example, two images from a well calibrated CT scanner acquired using the same imaging protocol. This will not be the case if a contrast medium has been administered for one image but not the other. Either the sum of absolute differences or the sum of squared differences over intensity values is used for these problems.

For multi-modality imaging, similarity measures such as ratios of image uniformity [Woods, Cherry and Mazziotta, 1992] within regions are used, as are measures based on correlations in intensity space. Measures based on information theory were initially independently proposed in 1995 by [Viola and III, 1995] and by [Collignon et al., 1995].

The similarity measure is then used as a means of driving the optimisation, whereby the transformation is successively improved as judged by the similarity measure, such that correspondence improves.

When plotted against the N transformation parameters, the similarity measure will be an N-dimensional landscape within which the optimisation scheme can explore to find an optimal solution. The structure of this landscape can be slowly varying or rapidly varying; it can have many or few local minima that could result in the optimisation scheme selecting an incorrect local extremum. These considerations should be addressed when choosing the similarity measure and the optimisation scheme.

3.7.1 Information Theoretic Voxel-Based Similarity Measures

A similarity measure is a means of quantifying how aligned the two images are at any stage of the registration process. Voxel based similarity

measures were first proposed by Woods et al [Woods, Cherry and Mazziotta, 1992] for registering PET images without the need for predefined features. This similarity measure involved minimising the variance of intensity ratios across voxels, with the underlying assumption that the intensity values are proportional to each other. Their algorithm was also adapted [Woods, Mazziotta and Cherry, 1993] to register MRI images to PET, but this involved partitioning the MRI image. This approach was extended [Ardekani et al., 1995] to produce an algorithm using a k-means clustering technique to perform the segmentations for the partitions automatically. In both cases, the fundamental assumption was that all voxels in the MR image with a certain intensity represent the same tissue type, and that the corresponding voxels in the PET image should also be similar to one another.

Methodologies using the linear correlation coefficient assume there is a linear relationship between the voxels, which will only be applicable for the mono-modality case [Crum, Hartkens and Hill, 2004]; other correlation methods are more relaxed, but assume that the functional relationship between the intensities is monotonic, in other words, that if one value increases, so does the other $f(x) \leq g(y)$ if $x \leq y$, and the converse for a decreasing value. Information theoretic measures, such as mutual information, can take into account all types of dependence [Kojadinovic, 2005], which makes them flexible and suited for many kinds of tasks where there are complex relationships between the distributions. Information theoretic similarity measures therefore do not require this partitioning approach, since the underlying assumptions of dependency between voxel values are less specific than correlation coefficients, so as a measure of dispersion it is ideal for inter-modality comparisons as well as intra-modality comparisons. It can also be extended to a general case for comparing more than two random variables in several ways.

An overview of these information theoretic measures is now given in the next section, culminating with some generalisations to mutual information, the similarity measure that will be used for registering serial combined PET/CT images. This will require an extension of the mutual infor-

mation approach, since there is intensity information from two modalities about each object to register, requiring a kind of multi-way comparison. This is also developed further in Appendix B.

3.7.2 Entropy of an Image

Consider a random variable A in which the probabilities associated with each state i is p_i and there are N states, then the Shannon entropy for A [Shannon, 1948] is defined by:

$$H(A) = \sum_{i=1}^{N} p_i \log \frac{1}{p_i} \tag{3.9}$$

We can compute the entropy of an image $I(\vec{r})$ with intensities in the range $I_{min} \leq I \leq I_{max}$ by first finding a normalised grey level histogram for the image. Since the number of discrete intensity values can be large, the intensity values are placed into intensity bins prior to computing the histogram, and the integer variable f is the intensity bin label. f often has the range $\{f = 0, 1, .., 63\}$ when the binning is undertaken in practice, and the methods for performing the binning can be explored. A grey level histogram $G_I(f)$ calculated for image I is the number of voxels within each intensity bin f and can be calculated by summing over all N_v voxels with voxel index s at locations $\{\vec{r_s} : s = 1, 2, .., N_v\}$ according to the following summation:

$$G_I(f) = \sum_{s=1}^{N_v} \delta(f, I(\vec{r_s})) \tag{3.10}$$

where the Kronecker delta function is defined as:

$$\begin{aligned} \delta(f, I(\vec{r})) &= 1, I = I(\vec{r}) = f \\ &= 0, I(\vec{r}) \neq f \end{aligned}$$

The grey level histogram specifies the statistical distribution of intensities within the image, although it does not contain any spatial information about where these intensities come from.

If one makes the connection between frequency of occurrence and probability by normalising the grey level histogram by the total number of voxels N_v, we can get an estimate of the probability density function (PDF) for the intensities, $p_I(f)$ within the image:

$$p_I(f) = \frac{G_I(f)}{N_v} \tag{3.11}$$

Thus, we can compute the entropy of an image using equation 3.9 where we consider each intensity bin as a "state" i in the summation. We would expect a histogram that is peaked, such as the original histogram shown in figure 4.3, to have a lower entropy value than one that has a distribution that is more even, such as the transformed histogram shown in this figure. This formulation of entropy of the image based on the intensity distribution discards the spatial information present in the image. Other techniques [Ardekani et al., 1995] do take into account the spatial location of the intensities.

3.7.3 Joint Intensity Histograms and Mutual Information

A joint intensity histogram is constructed by looping over the corresponding voxels in each image pair, and calculating the resultant joint occurrence of intensities as a two dimensional plot. The construction of a joint histogram for two images is specified in more detail below.

Consider a voxel located at $\vec{r}_s$ in image one $I_1(\vec{r}_s)$, and a corresponding voxel located at this same location $\vec{r}_s$ in image two, $I_2(\vec{r}_s)$. Suppose that the intensity values in $I_1$and I_2 at this location fall into bins f_1 and f_2 respectively. A two dimensional bin in the joint histogram located at (f_1, f_2) is incremented. This process is repeated for all locations $\vec{r}_s$ and often plotted as a two dimensional plot of f_1 against f_2 and with the function values on this plot represented as an intensity. For two identical images, the result is a diagonal line through the origin that is a projection of the 1D image histogram values. The joint intensity histogram can be compactly written using the Kronecker delta function:

$$G_{I_1,I_2}(f_1,f_2) = \sum_{s=1}^{N_v} \delta(f_1, I_1(\vec{r_s})) * \delta(f_2, I_2(\vec{r_s})) \tag{3.12}$$

When normalised, this represents a joint probability density function of the two images:

$$p(f_1,f_2) = \frac{G(f_1,f_2)}{N_v} \tag{3.13}$$

where N_v is the number of voxels. The concept of a joint intensity histogram could be extended from two to many dimensions. For an ensemble of images, we could define the N-dimensional histogram to be:

$$\begin{aligned} G_{I_1,I_2...,I_N}(f_1,f_2,..,f_N) &= \sum_{s=1}^{N_v} \delta(f_1, I_1(\vec{r_s})) * \delta(f_2, I_2(\vec{r_s}))*,.., \\ &\quad *\delta(f_1, I_1(\vec{r_s})) * \delta(f_N, I_N(\vec{r_s})) \end{aligned} \tag{3.14}$$

The joint intensity histogram changes with mis-registration, becoming more dispersed, as initially observed by [Hill, 1993]. The joint entropy is computed over the joint probability density function and is defined as follows:

$$H(I_1,I_2) = \sum_{f_1,f_2} p(f_1,f_2) \log \frac{1}{p(f_1,f_2)} \tag{3.15}$$

The joint entropy measures the dispersion of the joint intensity histogram. A more peaked joint histogram, signifying a better registration, is indicated by a lower value for the joint entropy. Figure 3.4 shows a series of joint histograms computed for a single CT head and neck image that is displaced relative to itself by increasingly small distances. The amounts of displacement are (a) 17mm, (b) 10mm, (c) 9mm, (d) 3mm, (e) 2mm and (f) 1mm. The joint histogram was computed using 64 bins in each case and the occurrence in each bin is rendered as a greyscale. This figure demonstrates that the joint histogram becomes increasingly peaked as the images are more closely aligned. Therefore, registration algorithms have minimised the joint entropy, as their modus operandi, which is, it

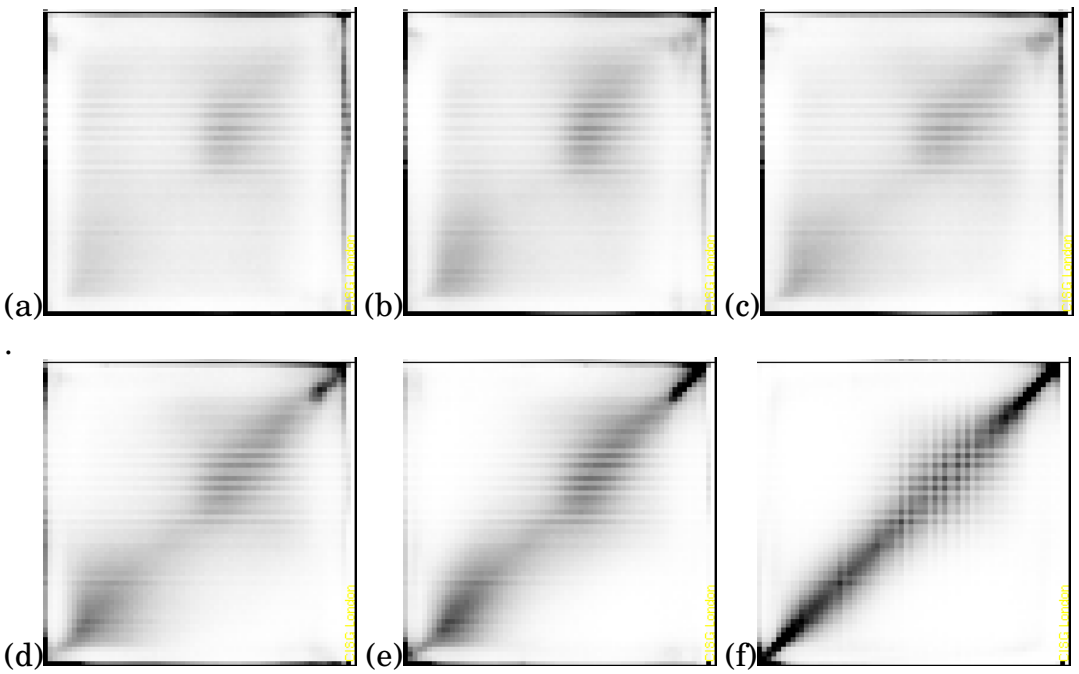

Figure 3.4: Joint intensity histograms for a CT image at progressively small translation relative to itself

must be recalled, an implicit function of the transformation. One problem with joint entropy as discussed in the literature [Pluim, Maintz and Viergever, 2003; Hajnal, Hill and Hawkes, 2001] is that a low value can be achieved when the two images are highly mis-registered, since the joint intensity histogram would have a sharp peak, caused by voxels that are in air rather than arising from the patient. This undesirable feature of joint entropy is mitigated by modifying this entropy based similarity measure to include two terms that relate to the entropy of each of the two images, but with a different sign than the joint entropy, with the result that this will counterbalance this effect. This works as a counterbalance, because as the transformation operates on the floating image to move it out of the field of view, it can also result in a decrease in entropy, since anatomy in the image can be moved out of the field of view and more air can enter it. A measure that has these attributes is mutual information and is defined as:

$$MI(I_1, I_2) = H(I_1) + H(I_2) - H(I_1, I_2) \tag{3.16}$$

Where the Shannon entropy is computed on the probability distribution of the grey values. This similarity measure was first proposed for registration by Colignon and Viola and Wells [Collignon et al., 1995; Viola and III, 1995] and has since been widely used in a variety of applications, for a wide variety of tasks. In one recent application [Rajkumar and Papastavrou, 2012], a technique involving maximising mutual information was used to register 2D histological slices. Mutual information is a symmetric measure in theory, $MI(I_1, I_2) = MI(I_2, I_1)$, although in practice during registration it is found not to be due to implementation details such as interpolation.

3.7.4 Normalised Mutual Information

Despite this counterbalancing effect that results in MI being an improvement over joint entropy, there is still an undesirable property of MI in that it depends on the volume of overlap of the two images. This was investigated by [Studholme, Hill and Hawkes] by constructing a 2D software phantom image containing a simple object: a half circle in both the fixed image and the floating image. Each image had the same width, l, the fixed image had infinite length and the floating image had a length that was varied from l to $3l$. The effect of varying the length - causing the field of view of the floating image to change (and the area of image overlap) was investigated for different values of the transformation angle, θ. The value of the similarity as a function of field of view and transformation was calculated for three different similarity measures: mutual information, joint entropy and normalised mutual information (NMI), an alternative similarity measure defined in the following way from the entropy of each image and the joint entropy:

$$NMI(I_1, I_2) = \frac{H(I_1) + H(I_2)}{H(I_1, I_2)} \tag{3.17}$$

All similarity measures showed a peak when the transformation angle was zero *and* the field of view was small, as one would expect, but the normalised mutual information also had the desirable attribute that it was constant as FOV was varied. The other similarity measures were not,

and there was a worrying demonstration for mutual information that, as the field of view increased beyond a value of approximately $l = 2$, this similarity measure had no convex property as a function of angle, rendering it useless for an optimisation scheme in this case, since the desired solution is when θ=0, yet this is a minimum rather than a maximum in similarity when $l > 2$. This is logical, because if one increases the size of the FOV, the relative importance attached to the object becomes less compared to the overlap of the background regions.

The simulation experiment proposed above [Studholme, Hill and Hawkes] was for a software phantom image that was supposed to represent a coronal view of a brain, hence the choice of geometry, but the question is open as to whether this result of overlap independence will also hold for other structures. It is also noted that the shape used for this simulation had pixel values set to a constant value inside, and zero outside, which is not representative of a head and neck PET/CT image, where the contents of the shape would have many complex distributions of pixel values in a variety of anatomical configurations, that would all contribute to the evaluation of the similarity measure. The analysis used to demonstrate the overlap invariance of normalised mutual information was critiqued [Cahill et al., 2008] and shown to be a slowly varying function of rotation. That said, normalised mutual information continues to be one of the most widely used similarity measure in the literature.

There have been other metrics for normalised mutual information in other fields within the literature [Sarndal, 1974], for example, the symmetric uncertainty coefficient SUC:

$$SUC(I_1, I_2) = \frac{MI(I_1, I_2)}{\frac{1}{2}\left[H(I_1) + H(I_2)\right]} \tag{3.18}$$

3.7.5 Multi-Dimensional Mutual Information

Despite the popularity of mutual information for a variety of tasks, a disadvantage is that it does not take into account any spatial relationships that exists among the pixels. There have been many extensions to

mutual information to attempt to include spatial information, that seek to rectify this. First spatial derivative information [Pluim, Maintz and Viergever, 2000] of the images were included in the higher dimensional mutual information, in addition to the intensity information. Mutual information was also extended to multiple dimensions by including nearest neighbours of the pixel being compared [Russakoff et al., 2004] for CT registration. This nearest pixel method was also used in [Legg et al., 2009] with the addition of gradient information at many scales of image resolution. The aim of these methods is to improve the efficacy of the mutual information by adding complementary information.

There have been general extensions to mutual information that have been proposed for application in other disciplines. Some examples are now outlined for various problems. In work published by [Li and Verma, 2011], a multivariate similarity measure was constructed which involves extracting image features as an alternate information channel to fold into the similarity measure. Recently, a multivariate similarity measure [Wachinger and Navab, 2013] was constructed to simultaneously register multiple images at a time in a groupwise manner. This simultaneous groupwise registration has the advantage over a pairwise technique (where each image in the set is registered to a chosen template image) in that it doesn't bias the problem towards the chosen template image. A novel multi-component mutual information technique termed "diffusion paired MI" for registering MRI images was developed [Studholme, 2007]. Intensities from the six different diffusion tensor components and the t1 weighted image were used to construct a form of multicomponent mutual information to register a pair of MRI images using a viscous fluid non-rigid registration algorithm.

Some mathematical definitions for mutual information when using more than two channels are now outlined. If we consider that equation 3.16 is valid for two probability distributions derived from two input images, we could extend this if we have more sources of information to work with. An N-dimensional variation on the theme of mutual information was proposed [Weinholt and Sendhoff, 1996] that extends this and is also known

as redundancy. For a set of $I_i, i = 1, ..N$ random variables, the redundancy is defined as:

$$R(I_1, I_2, .., I_N) = \sum_{i=1}^{N} H(I_i) - H(I_1, I_2, .., I_N) \tag{3.19}$$

An alternative for the mutual information for three random variables [Abramson, 1963]:

$$\begin{aligned} M(I_1, I_2, I_3) &= H(I_1) + H(I_2) + H(I_3) \\ &\quad -H(I_1, I_2) - H(I_1, I_3) - H(I_2, I_3) + H(I_1, I_2, I_3) \end{aligned} \tag{3.20}$$

This has been extended from three to N random variables [Kojadinovic, 2005]. Chapter 5 outlines an implementation of a multicomponent histogram developed in this project and used to evaluate the similarity measure within the registration algorithm described below. This enabled the construction of a multicomponent similarity measure based on vector images, where, at each voxel, instead of a scalar intensity value, the value at each point is a vector of dimension 2. One component arises from the CT image and the other from the PET image. This extension of a scalar image to a vector image could - in principle at least - be used for any number of dimensions.

Chapter 4

Registering Head and Neck Images

4.1 Introduction and Motivation

This chapter describes how the optimal non-rigid registration scheme was developed and evaluated for serial head and neck combined FDG PET/CT images acquired during a clinical trial, using a multi-scale B-splines based non-rigid registration algorithm. Regtool, a software application for image registration built using a set of C++ classes known as vtkCmic was used. This software has many parameter settings to set that control the behaviour of the registration.

The techniques used to pre-process the images are described, as are the steps involved in setting the parameters for the registration to ensure good results. The work described in this chapter ensured the correct parameters were set for our registration problem to achieve good results. Suitable methods for assessing the performance of the registration were introduced as required. This chapter finished by summarising the set-up parameters to achieve good results.

Once the optimum configuration of parameters was deduced, the registration using these settings was performed for several patients from the

clinical trial. A problem was identified with poor alignment in one of the images that contained dental artefacts.

This then motivated the work in the following chapter, which was concerned with incorporating information from the PET images to guide the registration in this region, which involved developing multicomponent registration techniques.

4.1.1 Head and Neck Image Data

A phase 1 clinical trial for patients with squamous cell head and neck cancer was designed to assess the response to treatment using radiotherapy with and without chemotherapy. Patients were imaged using combined FDG PET/CT imaging to assess the response to the treatment. The patients were treated with daily fractions of external beam radiotherapy. During treatment, they were positioned on the couch with a thermoplastic head-shell in place. This standard procedure for head and neck patients is to ensure a well defined and repeatable treatment set-up and to minimise movement, and therefore improve localisation[1] of the dose distribution. It will thereby increase the likelihood that the dose distribution delivered by the external beam radiotherapy is as close to the planned dose distribution as possible for each treatment episode (known as a fraction). The patients were also scanned on the combined PET/CT scanner in the same configuration, with a head-shell immobilisation device in place. The patients were scanned on four occasions: prior to the course of radiotherapy; on completion of the 5th fraction of radiotherapy; on completion of the 20th fraction, and on completion of the whole course of radiotherapy. These images are denoted as $I(\vec{r}, t_0)$, $I(\vec{r}, t_1)$, $I(\vec{r}, t_2)$ and $I(\vec{r}, t_3)$.

[1] The goal of advanced techniques in external beam radiotherapy treatment delivery is to optimise the localisation of the dose distribution by improved immobilisation techniques in conjunction with techniques designed to track the anatomical structures of interest through image guidance. The fluence of the external beam radiotherapy can be modified to track the target, or avoid the critical structures to maximise the dose to the tumour, whilst minimising the dose to the critical structures. This will give the highest curative rate whilst minimising the chances of side effects.

The thermoplastic immobilisation device also ensures minimal movement between successive images, and for each PET/CT acquisition performed on the combined PET/CT scanner there is minimal movement between the PET and the CT image: they are therefore 'best-case' in terms of the co-registration of the PET/CT images from the combined scanner. But there will still be several sources of movement between successive fractions for a variety of reasons, such as changes due to weight loss, changes to anatomy due to the treatment - for example, tumour shrinkage or radiation-induced effects to normal tissues such as inflammatory processes or swellings, weight loss such as a reduction in subcutaneous fat, or modifications due to posture, such as variations in the jaw position and the possible movements of the tongue.

4.1.2 Motivation for Developing Improved Registration Algorithms

A main goal of this clinical trial was to use FDG PET/CT imaging to quantify the response to radiotherapy treatment with and without chemotherapy. FDG is taken up by highly metabolising tissues, especially cancer cells. But there are also other radiation-induced mechanisms that change the FDG uptake.

This requires accurate delineation of the tumour region of interest to be sampled on the PET image, and we desire to use the transformation from the CT images gained from non-rigid registration as a vehicle to transport the segmentation from the reference geometry into the geometry in the subsequent studies. The appropriate decisions must be made regarding the nature of the non-rigid registration algorithm concerning the types of change in question. One may impose biologically appropriate constraints as required. For example, if one posits that the segmented volume being propagated ought not to change in volume - but could change in shape - then a volume-preserving deformation field would be the required choice. This is mentioned earlier in section 3.6.4. A non-rigid registration strategy for segmentation propagation, so that quantitative information within the PET regions can be studied, was the motivation

for our work.

Registration strategies for serial studies have been developed for several reasons by other workers in the literature. This includes quantifying motion [Rueckert et al., 2006], [Bissonnette et al., 2009]. Another application is to propagate a segmented structure set from a template onto a patient's image, which involves by definition a non-rigid registration, since the registration in this case is inter-patient in nature [Zhuang et al., 2008]. This is an active area of clinical research for radiotherapy treatment planning, for example, the algorithm that underpins the Atlas-Based Autosegmentation software (ABAS) system marketed by Elekta for use in clinical radiotherapy treatment planning. This system is assessed in the following work [Sims et al., 2009] in which the performance is determined for this system in the head and neck anatomy.

4.1.3 Overview of the Registration Scheme

A registration scheme was sought that will allow any of the CT images, $I(\vec{r}, t_1)$, $I(\vec{r}, t_2)$ or $I(\vec{r}, t_3)$, to be registered to the pre-treatment image, $I(\vec{r}, t_0)$. For each pair of images being registered, the pre-treatment image $I(\vec{r}, t_0)$ was taken to be the fixed image, with the other being the floating image (either $I(\vec{r}, t_1)$, $I(\vec{r}, t_2)$ or $I(\vec{r}, t_3)$). The following registration framework was then employed.

Normalised mutual information (NMI) was used as a similarity measure. As discussed in the survey article by [Pluim, Maintz and Viergever, 2003], this similarity measure has a long history of use in intensity-based registration and remains a popular choice. This survey mentions that other workers have developed CT-CT registration algorithms employing this similarity measure. For CT-CT registration, a direct intensity difference based similarity measure may also be used, such as sum of absolute differences, as reviewed by Crum et al [Crum, Hartkens and Hill, 2004].

Despite being challenging to interpret, the head and neck area of anatomy typically has better intrinsic PET/CT registration[2] than the chest and ab-

[2] *Intrinsic registration* refers to the alignment of a PET/CT study produced from a combined scanner.

domen, which will clearly suffer from respiratory motion. There are techniques to minimise motion, such as scanning the patient within a head and neck thermoplastic immobilisation device, as was the case for the previously mentioned data from the clinical trial, or by using a bite bar. When using a thermoplastic immobilisation device, the alignment of the PET to CT from a combined scanner has been reported to be less than 1mm [Thomas et al., 2014] in the head and neck region.

An example of the PET/CT data from the clinical trial is displayed in figure 4.1. The images displayed are rigidly registered[3] to the pre-treatment study - the first PET/CT study shown in the top left window - and are displayed in a common coordinate system, so the cursor positions on each pane and window represent approximately the same position in space, subject to some small error due to imperfect registration. The voxel dimensions of the CT image were $0.98mmx0.98mmx2.5mm$, and for the PET image, they were $4.69mmx4.69mmx3.27mm$.

4.2 B-splines Non-Rigid Registration

The free-form deformation non-rigid registration algorithm using approximating B-splines is a computationally efficient way of parameterising a registration that can account for large deformations at the local level. A framework for registering head and neck PET/CT images will require large changes at the local level, therefore this algorithm was chosen as a basis for a registration scheme. The vtkCmic code was used in this project and was based on the method described by [Rueckert et al., 1999] in the classic paper on this. In this section, the method outlined in that paper is discussed in more detail.

The transformation model is defined by a regular grid of control points that, when displaced, warp the underlying image at every point according to an approximating B-spline function. The control points can be thought

[3]The CT-CT derived registration transformation was applied to both CT and PET images for each of the three subsequent PET/CT studies.

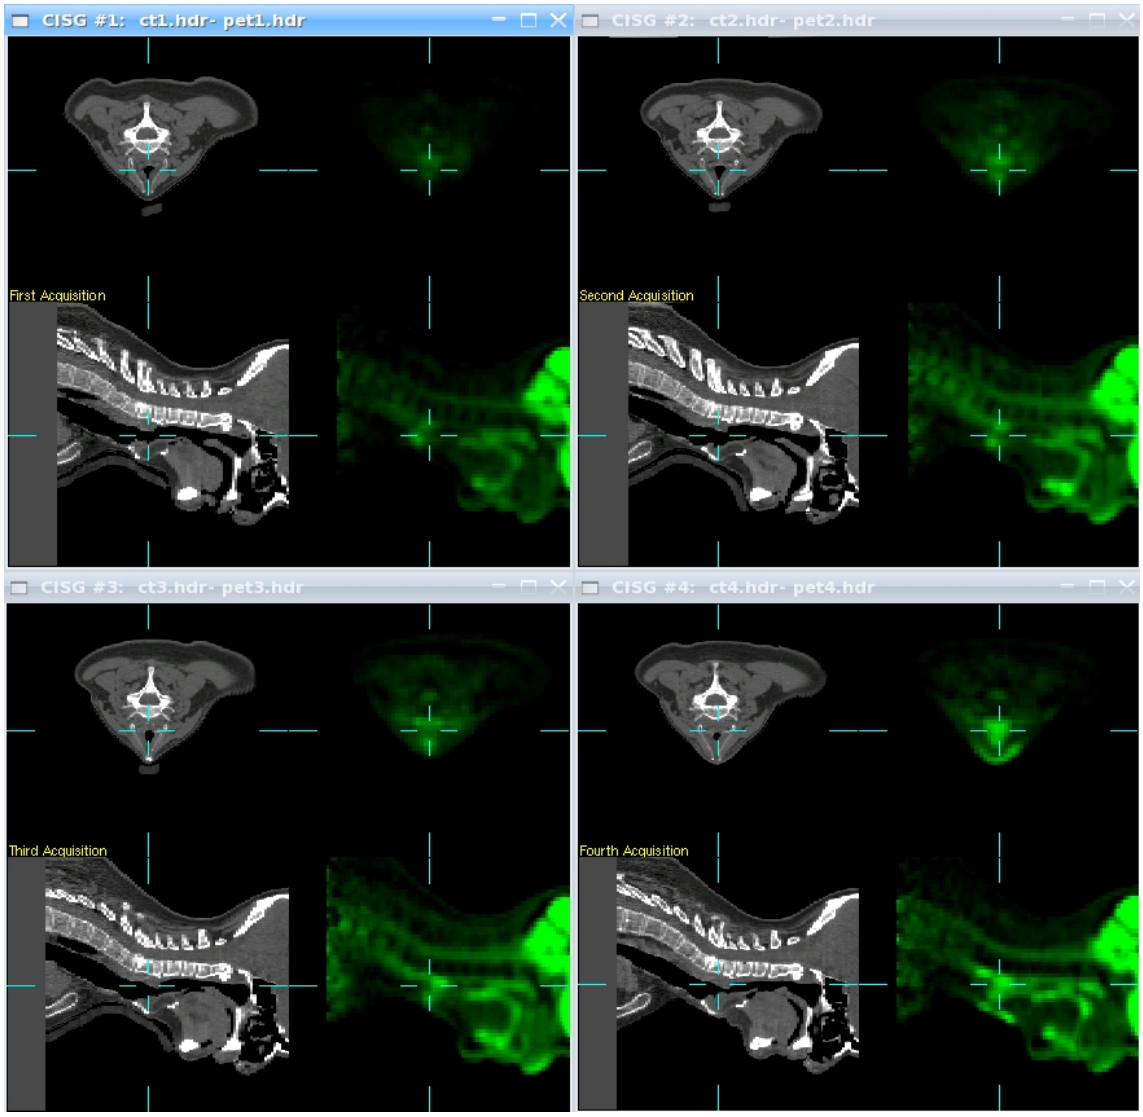

Figure 4.1: PET/CT acquisitions for patient 1. CT image in grey colourscale and PET image in green colourscale. Top left are pre-treatment images, then (clockwise) the first, third and second successive acquisitions from the study detailed in section 4.1.1.

of as “pseudo landmarks” such that when moved, the local deformation of nearby points will be described by the approximating B-spline functions.

In three dimensions, each control point on the image is identified using a triplet of integer indices $[i, j, k]$. The displacement of the control point labeled $[i, j, k]$ is written as $\vec{c}_{i,j,k} = \hat{\mathbf{i}}c_{x,i,j,k} + \hat{\mathbf{j}}c_{y,i,j,k} + \hat{\mathbf{k}}c_{z,i,j,k}$where $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$, and $\hat{\mathbf{k}}$ are the orthogonal unit vectors in the x, y and z directions, and $c_{x,i,j,k}$, $c_{y,i,j,k}$ and $c_{z,i,j,k}$ are the real-valued components of this vector, that parameterise the transformation. The control points can be thought of as attractors that, when moved, will drag the underlying image. The vector displacements of positions in the underlying image (known as warps) are given using approximating cubic B-spline functions outlined later.

If the grid spacing in the x, y and z directions respectively is δc_x, δc_y and δc_z, and there are n_x, n_y and n_z grid points in the x ,y and z directions, then the displacement vector field $\vec{v}(x, y, z)$ at a point $P(x, y, z)$ in the image is approximated according to:

$$\vec{v}(x, y, z) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \vec{c}_{\mathrm{i+l,j+m,k+n}} \tag{4.1}$$

Where the indices $[i, j, k]$ represent the first local control point in the grid that affects the displacement of the point $P(x, y, z)$ and are given by $i = \left\lfloor \frac{x}{\delta c_x} \right\rfloor - 1$, $j = \left\lfloor \frac{y}{\delta c_y} \right\rfloor$ and $k = \left\lfloor \frac{z}{\delta c_z} \right\rfloor$; also $u = \frac{x}{\delta c_x} - \left\lfloor \frac{x}{\delta c_x} \right\rfloor$, $v = \frac{y}{\delta c_y} - \left\lfloor \frac{y}{\delta c_y} \right\rfloor$ and $w = \frac{z}{\delta c_z} - \left\lfloor \frac{z}{\delta c_z} \right\rfloor$ represent the fraction of the distance from the point (x, y, z) to the next smaller control points with respect to the spacings. The symbol $\lfloor x \rfloor$ is the floor function acting on a real value x. The transformation is defined in a local way, so that for point $P(x, y, z)$ in consideration, the displacements of a local subset of the control points will uniquely determine the displacement of P. The closest two control points in each direction are selected to define this "local" grid.

The B-splines cubic functions B_i are defined as follows:

$$
\begin{aligned}
B_0(x) &= (1-x)^3/6 \\
B_1(x) &= (3x^3 - 6x^2 + 4)/6 \\
B_2(x) &= (-3x^3 + 3x^2 + 3x + 1)/6 \\
B_3(x) &= x^3/6
\end{aligned}
\tag{4.2}
$$

The components of the control point displacement vectors can be written as a set of $3N$ parameters $\{c_{x,i,j,k}, c_{y,i,j,k}, c_{z,i,j,k}; i = 0 \dots n_x, j = 0 \dots n_y, k = 0 \dots n_z\}$ that define the transformation, where $N = (n_x+1)*(n_y+1)*(n_z+1)$. This set can be written in an equivalent manner using a single vector $\vec{\Theta}$, constructed from these components to define the transformation as a single $3N$ dimensional entity according to the following scheme:

$$
\vec{\Theta} = \begin{pmatrix}
c_{x,0,0,0} \\
c_{y,0,0,0} \\
c_{z,0,0,0} \\
c_{x,1,0,0} \\
c_{y,1,0,0} \\
c_{z,1,0,0} \\
\vdots \\
c_{x,n_x,0,0} \\
c_{y,n_x,0,0} \\
c_{z,n_x,0,0} \\
\vdots \\
c_{x,n_x,n_y,0} \\
c_{y,n_x,n_y,0} \\
c_{z,n_x,n_y,0} \\
\vdots \\
c_{x,n_x,n_y,n_z} \\
c_{y,n_x,n_y,n_z} \\
c_{z,n_x,n_y,n_z}
\end{pmatrix}
$$

The cost function is an implicit function of $\vec{\Theta}$, bearing in mind that the source and template images are also required for its computation.

4.2.1 The Optimisation Scheme

The optimisation scheme for registration is steepest gradient descent [Press and Vetterling, 1992], according to the following update rule:

$$\vec{\Theta}_{s+1} = \vec{\Theta}_s + \mu \frac{\bigtriangledown \mathrm{S}(\vec{\Theta}_s)}{\left\| \bigtriangledown \mathrm{S}(\vec{\Theta}_s) \right\|} \tag{4.3}$$

Where μ is the step size, s is the iteration number and S is the similarity measure.

This is continued iteratively until there is no significant change in the gradient: $\left\| \bigtriangledown \mathrm{S}(\vec{\Theta}_s) \right\| < \varepsilon$ where ε is a small tolerance, or when the maximum number of iterations has been reached.

The step size is then halved and this process is repeated above, until the tolerance value or the maximum number of iterations is reached. This process of iteration and step size halving is repeated N_{steps} number of times.

The computation of the gradient $\bigtriangledown S(\vec{\Theta}_s)$ in equation 4.3 is a computationally expensive process, and is done using finite differences by moving each component in $\vec{\Theta}$ from $-\mu \rightarrow \mu$, noting that of course the calculation of S involves applying the transformation to the floating image at each stage.

Pseudo code representing this algorithm is adapted from [Rueckert et al., 1999] and is shown here in Algorithm 4.1 with some comments to explain how the optimisation works.

4.3 The vtkCmic Registration Code

Regtool, the software tool used during this project built using the vtkCmic C++ classes, was developed by the group at the Centre for Medical Image Computing at University College, London. A significant effort was spent refactoring the classes to support the multicomponent problem in later

Algorithm 4.1 Multi-step non-registration algorithm

// This pseudo code uses C++ style comments
// that are preceded by a double forward-slash.
// I_T and I_F are the target and fixed images, μ is the starting
// step size value and ε is a small positive real.
// Set all components of the non-rigid registration transform to zero:
Assign all components of $\vec{\Theta}$ to zero
for (i1=1,N_{steps},1) {
    for (i2=1,N_{iter},1) {
        // evaluate the similarity measure's gradient using
        //finite differences with respect to the transformation $\vec{T}_{\vec{\Theta}}(\vec{r})$:
        // (n.b. this transform is parameterised by $\vec{\Theta}$)
        Assign $\nabla S = \frac{\partial S(I_T, I_F(\vec{T}_{\vec{\Theta}}(\vec{r}))}{\partial \vec{\Theta}}$

        // if this value is below the threshold then
        // go to the next step size:
        if $\|\nabla \mathrm{S}\| < \varepsilon$ then exit loop // i.e. the inner loop over i2
        DO
            assign $\vec{\Theta}_{old} = \vec{\Theta}$
            //modify the transform parameters
            assign $\vec{\Theta} = \vec{\Theta}_{old} + \mu \frac{\nabla S}{\|\nabla S(\vec{\Theta})\|}$
            //evaluate how the similarity improves:
            assign $\delta S = S(I_T, I_F(\vec{T}_{\vec{\Theta}}(\vec{r})) - S(I_T, I_F(\vec{T}_{\vec{\Theta}_{old}}(\vec{r})))$
            until $\delta S \leq 0$ // break out of DO loop if the similarity worsens
        assign $\vec{\Theta} = \vec{\Theta}_{old}$ // because the DO loop went one step too far
        } // We have finished optimising this step so exit this i2 for loop
        // update the next step size to half of the previous step size:
        assign $\mu = \frac{\mu}{2}$
}
// Finished - the optimised parameters defining the transform are in $\vec{\Theta}$

work. It is based on the algorithm described by Rueckert et al in their classic paper [Rueckert et al., 1999] and implemented in the software developed by the CISG group at Guy's Hospital [Hartkens et al., 2002]. The previous section, 4.2, describes the basics of the algorithm. It uses the infrastructure from the VTK system developed by Kitware Inc. [Avila et al., 2010], such as the processing pipelines, by which a series of processing actions can happen whereby the output of one action on an object such as an image file reader can be fed into the input of another object, such as a method for resampling that image. Then, the output of that action can be the input for a final method that could then write that image to a file. The vtkCmic code has classes for reading and writing various medical images such as the Analyze file format, developed by the Mayo Clinic and the GIPL format, a format from Guy's Hospital that can represent many kinds of medical images. There are also classes for many registration tasks, such as computing joint histograms between two images and evaluating similarity measures. Other classes are adapted from the VTK libraries, to enable B-splines free-form transformations to be applied to images. The vtkCmic was built historically on the VTK insfrastructure used by [Hartkens et al., 2002] rather than ITK [Ibanez et al., 2005], which is more readily designed for the specific task of image registration and segmentation than visualisation. However, there were various efficiency measures that were lacking within ITK which was the original motivation for using the VTK classes as infrastructure, despite the fact that relatively little of the extensive functionality of VTK is exploited. The main focus of VTK is for computer graphics and visualisation, a subject that is beyond the scope of this thesis.

4.4 Proposed Non-Rigid Registration Transformation

A registration scheme was developed consisting of a rigid registration followed by a B-splines based free-form deformation (FFD) non-rigid registration (NRR) to recover both a global rigid transformation and non-rigid local deformations. The global rigid registration was performed first, fol-

lowed by successively local B-splines based non-rigid registrations, with the first having a control point spacing (CPS) of 40mm. The transformation generated by this stage was then used as input for a finer scale 20mm NRR after first subdividing the 40mm control point mesh to 20mm. This process was then repeated at the 10mm control point spacing. The step size for the optimisation was chosen to be 16mm, 8mm and 4mm respectively.

The registration process yields a transformation $\vec{T}(\vec{r})$ that acts on the floating image $I(\vec{r}, \mathrm{t_i})$ to map positions $\vec{r}$ onto positions $\vec{r'}$, such that the transformed[4] floating image is registered to the fixed image, $I(\vec{r}, t_0)$.

The total hierarchical transformation is written with subscripts to denote the stages of the increasingly local control point spacing used, so for example:

$$\vec{T}_{rigid,20,10}(\vec{r}) = \vec{T}_{10}(\vec{T}_{20}(\vec{T}_{40}(\vec{T}_{rigid}(\vec{r})))) \tag{4.4}$$

denotes a rigid registration followed by a non-rigid FFD B-splines registration using a control point spacing of 40mm, then 20mm, and finally 10mm. The transformation at each scale is a vector field, sometimes known as a deformation field.

4.5 Image Pre-Processing and Related Tasks

4.5.1 Overview

This section details certain practical components that were developed to get the acquired image data from the imaging workstation into a suitable form that can be read by the file reading classes within the vtkCmic image registration software. Histogram equalisation was applied to the images prior to registration. The technique used to resample the images is also described.

[4]Not to be confused with "translated", rather, modified, manipulated, or often in the NRR literature referred to as "warped".

4.5.2 Practical Aspects of Image File Conversion

The acquired image data had first to be exported from the GE Medical Systems Xeleris workstation as a DICOM series. When exporting, the data would end up in an export directory as a single slice per file. Once anonymised and copied onto the PC designated for analysis, these slices were then re-assembled into the slice correct order, and placed into a single 3D DICOM image stack. A script was written using calls to the Medcon toolkit [Nolf] to do this and then to convert this 3D image stack into the Analyze[5] file format that can be accepted as input by the vtkCmic file reader class. The images were cropped to include only the head and neck anatomy for the registration, using the CropImage utility (specified in Appendix C). The image conversion was not without difficulties, and care was taken at each stage to ensure the integrity of the intensity data, in case the file conversion process modified intensity values, for example, by applying incorrect CT windowing to rescale the intensity data. A further problem was encountered whereby the file export process from the Xeleris workstation incorrectly scaled the counts within the PET image, such that the intensity values were locally normalised (instead of globally) within each slice, generating in effect a software artefact. The cause of this was established and the problem was fixed.

4.5.3 Histogram Equalisation

Prior to registration, the CT images were histogram equalised to maintain the contrast, since within the vtkCmic code, the construction of joint intensity histograms is used to calculate the similarity measure during registration. These joint histograms have 64 bins in intensity space; currently this parameter is hard-coded within regtool. Whilst CT is an imaging modality with high spatial resolution, it can be lacking in contrast, particularly for soft tissues. PET has low spatial resolution by comparison but could be exploited in areas where CT lacks contrast or where there are artefacts on the CT image, as discussed in section 4.8.

[5]Developed by the Mayo Clinic

	Air	Lung	Fat	Water	Muscle	Bone
HU	-1000	$\{-600, -400\}$	$\{-100, -60\}$	0	$\{40, 80\}$	$\{400, 1000\}$

Table 4.1: Hounsfield Unit for various substances

CT intensity values (given in Hounsfield Units or HU) typically have a numerical range from -1000 for air to over 1000 for bone (values for metal are larger still), as shown in table 4.1, for various structures present within a medical image. The process of linearly re-binning intensities into 64 intensity bins would result in voxels relating to all the soft tissues being placed within a single, or a very small number of, intensity bins: in effect rendering those voxels indistinguishable from one another during the calculation of the similarity measure. Therefore, some form of windowing or histogram transformation is required as a prior step. Histogram equalisation was chosen as an automatic method to achieve this.

Histogram equalisation is an intensity transformation that re-assigns intensity values in the image such that the distribution in histogram space is more uniform, the effect of this can be seen in the two trans-axial CT images shown in figure 4.2, which shows a CT image before and after histogram equalisation has been applied. The intensity histogram is shown for the pre-treatment CT image from patient 1 in figure 4.3, both prior to and after applying histogram equalisation. The histogram equalisation is applied to improve the performance of the registration rather than for the purpose of aiding the interpretation of the images. When interpreting the images, the original non-histogram equalised images are used, and appropriate window levels set with the transformation applied to the floating image.

4.5.4 Setting the Histogram Equalisation Thresholding Levels

When performing the histogram equalisation, a lower and upper threshold may be set to specify the first and last bin. All intensity values less

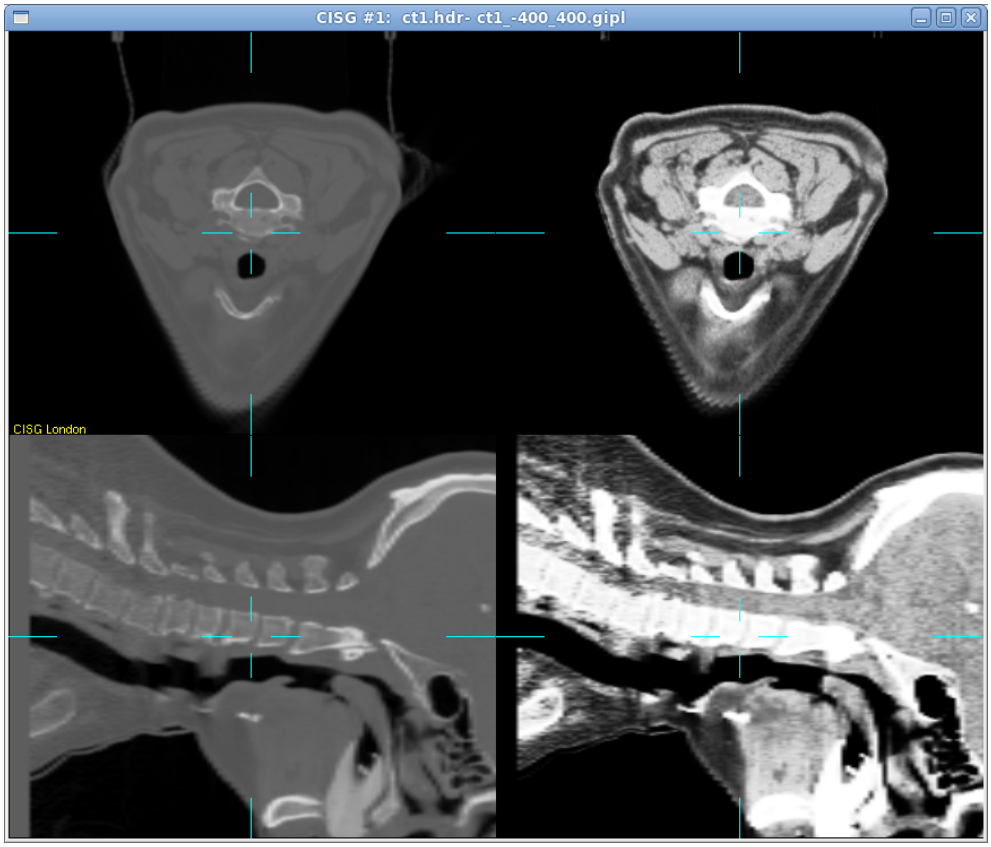

Figure 4.2: CT image before (left) and after histogram equalisation (right)

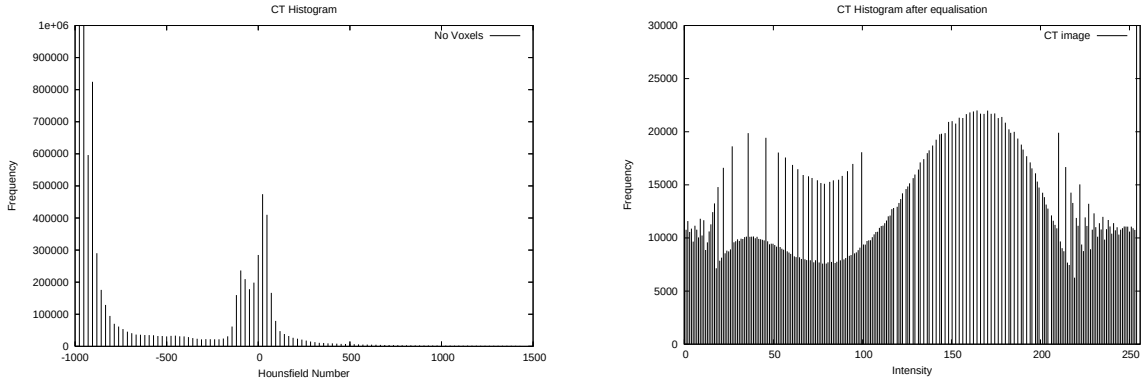

Figure 4.3: Intensity histograms before and after histogram equalisation of CT image

than the lower threshold will be counted in the first bin and similarly, all intensities greater than the upper threshold will be counted in the last bin. This section investigates how the values of the lower and upper histogram equalisation threshold levels for the CT images, as a step prior to registration, will affect the registration.

Motivation

When registering a pair of CT images, the upper and lower thresholds used for the histogram equalisation process can affect the performance of the registration. The thresholds should be chosen to reflect the Hounsfield Unit range of the tissues present within the anatomy. For example, if any significant amount of lung anatomy is present within the registration then the lower threshold should be set to approximately -600 to enable this tissue to contribute to the registration process. The aim of these tests is to find appropriate values to use for the lower and upper thresholds for the CT images such that the registration works satisfactorily well. Also, the appearance of the CT images following histogram equalisation for different thresholds will be studied for illustrative purposes.

Method

When histogram equalising an image, one may choose a lower threshold and an upper threshold value. How these thresholds may affect the performance of the registration was tested using the final CT image $I(\vec{r}, t_3)$, and the pre-treatment image, $I(\vec{r}, t_0)$. The reason that this image pair was chosen for this evaluation was that the final serial CT image will be more likely to be the most different to the pre-treatment image, and will therefore require the non-rigid registration transformation to recover larger displacements compared to the other serial images. Prior to registration, each image was histogram equalised using an upper and a lower threshold value specified in table 4.2. The number of intensity levels of the resultant histogram equalised images was set to 256, i.e. the intensities fell within the range 0 to 255. The images were then resampled to have 3mm isotropic voxels as detailed in section 4.5.6. The images were registered using a rigid, then a non-rigid FFD registration with trans-

HU L	HU U	Rigid δ	40mm FFD S	20mm FFD S	10mm FFD S
-1000	1700	2mm	1	1	1
-1000	1300	2mm	1	1	1
-400	400	2mm	2	2	2
-200	200	2mm	1	1	1
-100	100	2mm	1	1	1
-50	50	2mm	1	1	1
-25	25	>40mm	n/a	n/a	n/a

Table 4.2: Registration performance: rigid error δ and non-rigid score S for different histogram equalisation lower (L) and upper (U) thresholds

formations at 40mm, 20mm, and finally, a 10mm control point spacings performed in a hierarchical manner according to section 4.4.

Once registered, the images were assessed in two ways. Firstly, in the rigid case, by measuring the distance between anatomical landmarks[6] identified on the fixed and floating image. The anatomical landmark chosen was the dorsal tip of the C3 vertebra. For the non-rigid case, subtraction images were inspected and a scoring method was employed to globally assess the registration. A score of 1 denotes adequate global performance with a few regions containing significant errors, and a score of 2 indicates good performance throughout with no parts of the image containing significant errors.

Results

The appearance of the fixed CT image following histogram equalisation is shown for several threshold values in figure 4.4. The registration was tested for a range of thresholds as shown in table 4.2, and the performance of the registration was assessed visually by examining the images adjacently alongside each other, with a global score using the 1 to indicate good registration and 2 to indicate a very good registration performance. In the case when the lower threshold was set to -25HU and the upper

[6]It was noted that the rigid registration transform had no rotation component and was a translation only, hence a single point can be used to assess the registration

threshold was set to +25HU, the rigid registration failed completely, and therefore, the non-rigid registration could not be performed.

Conclusions

The registration performed slightly better when using the thresholds $-400HU, 400HU$ compared to when using the other Hounsfield ranges. Also, as can be seen from the histogram equalised images, these thresholds enable a sufficient distinction between the soft tissues, and will therefore provide better information for these structures to guide the registration. These thresholding values will be used in all subsequent registrations for head and neck images, a range of Hounsfield units that encompasses the structures within this area of anatomy. The rigid registration failed completely for the $-25HU, 25HU$ case. Since this failed, the non-rigid registration could not be performed as a follow-on step. This failure is due to the lack of utilisable intensity information in the image. The visual scoring method used to assess these results is a subjective test, and more quantitative methods for assessing the performance will be employed in this work in later sections when more precise metrics are required for assessing registration. This assessment must be considered to be only a useful pointer in the first instance.

4.5.5 Number of Bins Used in the Joint Histogram

Motivation

The joint intensity histogram is constructed from the input images and is used for computing various similarity measures such as NMI (as outlined in section 3.7.3) that are used during image registration. This section of work explores how the number of bins in the joint intensity histogram affects the performance of the registration. The appearance of the calculated joint intensity histograms as the number of bins changes is also explored.

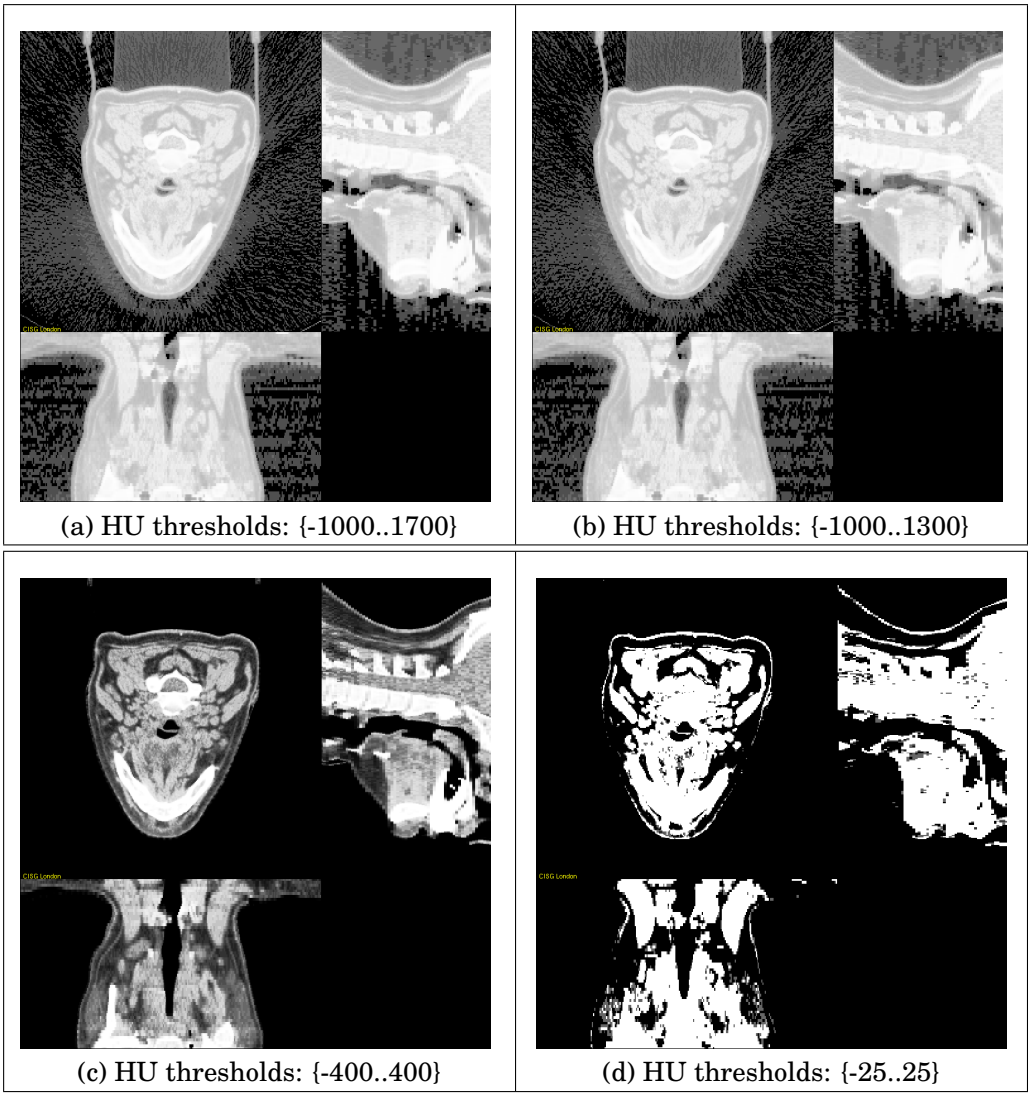

(a) HU thresholds: {-1000..1700}

(b) HU thresholds: {-1000..1300}

(c) HU thresholds: {-400..400}

(d) HU thresholds: {-25..25}

Figure 4.4: Histogram equalised images with different thresholds

Method

In the vtkCmic image registration software, the number of bins is fixed to 64. Yet it can in effect be reduced when histogram equalisation is performed as a pre-processing step. For illustrative purposes the joint intensity histograms for 128 and 256 bins were also calculated.

To investigate the appearance of the joint histograms that result from a pair of rigidly registered serial images, the following was done.

Firstly, the histogram equalised image $I(\vec{r}, t_1)$ was rigidly registered to the histogram equalised pre-treatment image, $I(\vec{r}, t_0)$ using 64 bins and upper and lower thresholds of -400HU and 400HU respectively. A starting step size of 8mm and 4 coarse-to-fine steps were used. Then, the joint histograms for a range of bin sizes were calculated for this pair of registered images using the transformation found earlier. This image pair was chosen, since the two images will be the most similar and, therefore, the appearance of their joint histograms will be best-case to maximise the chances of observing any differences when changing the number of bins.

The performance of the registration was also tested using NMI as a similarity measure for a range of different numbers of joint histogram bins. For this test, the final CT image$I(\vec{r}, t_3)$ was registered to the pre-treatment image, $I(\vec{r}, t_0)$. Again, prior to registration, the images were histogram equalised using the thresholds as previously described. Once registered, the images were assessed using the same system outlined in section 4.5.4.

Results

The joint histograms for the rigidly registered image pair for various numbers of bins can be seen be seen in figure 4.5.

The results of the tests to assess the performance of the registration when using differing numbers of bins are summarised in table 4.3 . When using 8 bins, the non-rigid registration resulted in a failure, which can be seen

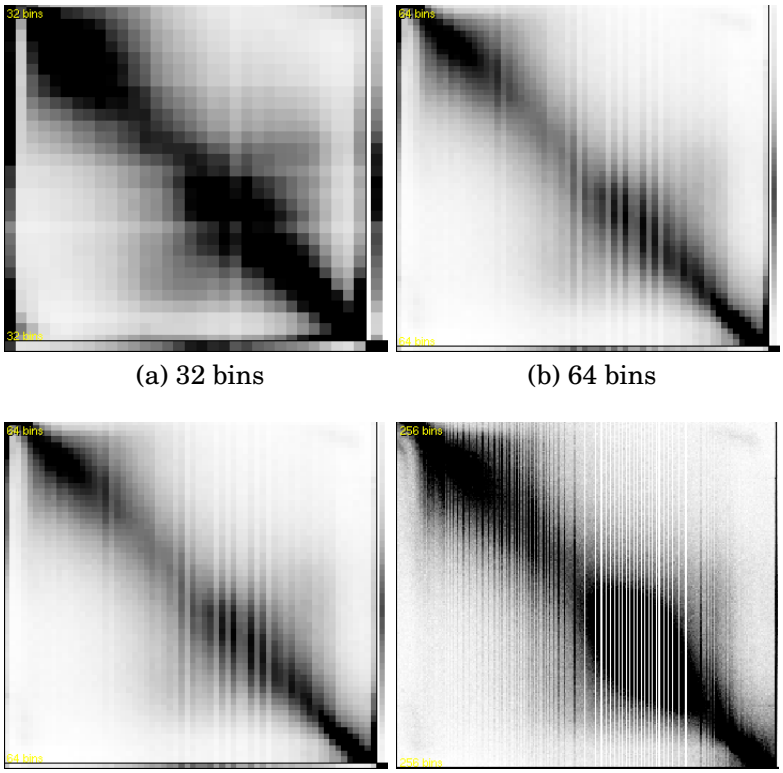

(a) 32 bins (b) 64 bins

(c) 128 bins (d) 256 bins

Figure 4.5: Joint intensity histograms for different bin sizes

nbins	Rigid δ/mm	10mm FFD S
64	2	2
32	2	2
16	2	1
8	2	fail

Table 4.3: Registration using differing bins in the joint intensity histogram

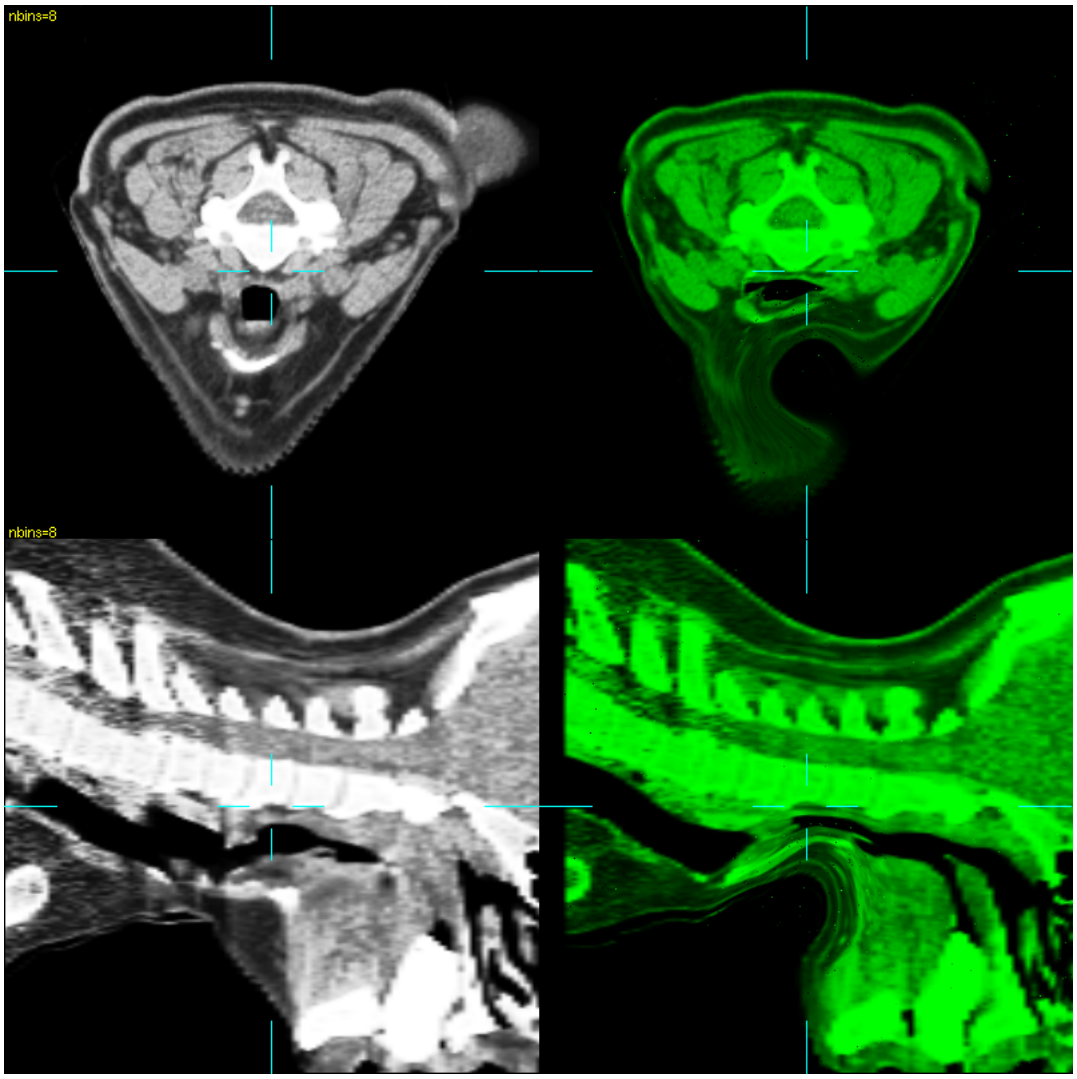

Figure 4.6: NRR failure for floating image (green) when nbins=8

in figure 4.6, which shows the floating image in green: the deformation field has created a significant artefact.

Conclusion

Of the different joint histograms that were computed, the ones that were the most peaked in appearance were when using either 64 bins or 128 bins. However, it is difficult to draw conclusions from the appearance of the joint histogram and the usefulness of examining them is limited. Since the similarity will be a measure of the dispersion of the joint histograms, and the optimiser works by finding the gradient of this relative to the transformation parameters $\vec{\Theta}_s$, what matters more is the magnitude of the gradient, $\left\| \bigtriangledown \mathrm{S}(\vec{\Theta}_s) \right\|$of the similarity S, rather than the absolute value of the similarity, given by the dispersion of the joint histogram. The rigid registration performed well for the values of nbins that were tested. The non-rigid registration performed best for both 64 bins and 32 bins. It seems reasonable to assume that if it was possible to use 128 bins then this would also result in a good performance of the non-rigid registration. However, this was not supported as a feature in the code at the time of processing. So, when performing a non-rigid registration of head and neck images, the default value of 64 bins for constructing the joint histogram from which normalised mutual information is then computed is an appropriate choice.

4.5.6 Effect of Image Resampling on Rigid Registration

Motivation

How the registration performance depends on the image resampling resolution is investigated in this section. To visualise the effect of image resolution on the joint histograms, they were calculated for registered image pairs that had been resampled to different resolutions. The effect of image resolution on the performance of the rigid registration was also investigated.

Method

The floating CT image $I(\vec{r}, t_1)$ was registered to the fixed CT image $I(\vec{r}, t_0)$ as per the previous section. This registration was done at their original resolution (i.e. not resampled) with voxels of dimensions $0.98mm$ x $0.98mm$ x $2.5mm$. Once registered, these images were then resampled isotropically (i. e. with cubic voxels) to different resolutions and the joint histograms were calculated for each resolution. To avoid aliasing, the images were blurred prior to resampling by convolving with a Gaussian:

$$G(x, y, z) = \frac{1}{2\pi\sigma} e^{-\frac{(x^2+y^2+z^2)}{2\sigma^2}}$$

with a width parameter σ set to half of the desired isotropic pixel spacing. The vtkImageGaussianSmooth VTK object was used to perform this 3D convolution. Then, the images were resampled using sinc interpolation using the vtkImageReslice VTK object. Both fixed CT image $I(\vec{r}, t_0)$ and floating CT image $I(\vec{r}, t_1)$ were resampled to isotropic voxels of the following dimensions for testing: $1mm^3$, $3mm^3$, $6mm^3$, $12mm^3$ and $18mm^3$.

To investigate how the registration at each image resolution varies, the floating CT image $I(\vec{r}, t_3)$ was registered to the fixed CT image $I(\vec{r}, t_0)$ using the same settings as above. The performance of the registration was assessed by comparing the transformation to a reference transformation that was a best-case transformation, which was presumed to be the one derived from registering the images that were resampled to $1mm^3$.

Results

The joint histograms are displayed in figure 4.7.

The registration errors as a function of voxel resolution may be seen in table 4.4. Figure 4.8 shows the registered images (displayed at their original resolution) when registered using $1mm^3$ resampling and the $18mm^3$ resampling. The fixed image is displayed on the left using the grey colourscale, and the floating image is shown to the right using a green colourscale.

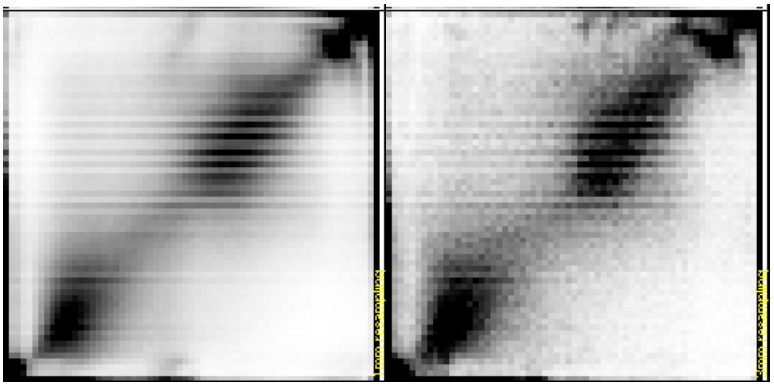

(a) 1mm resampling (b) 3mm resampling

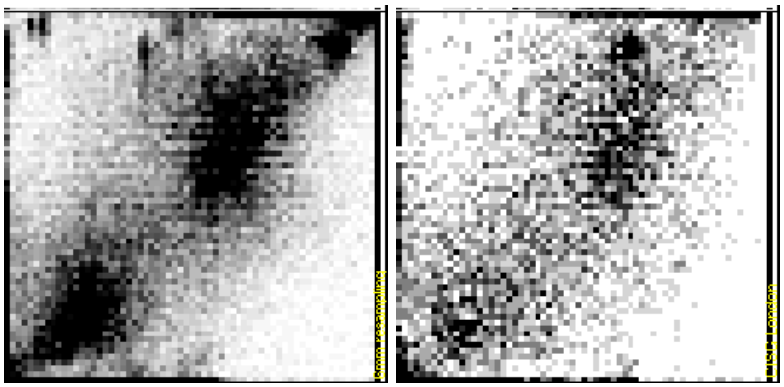

(c) 6mm resampling (d) 12mm resampling

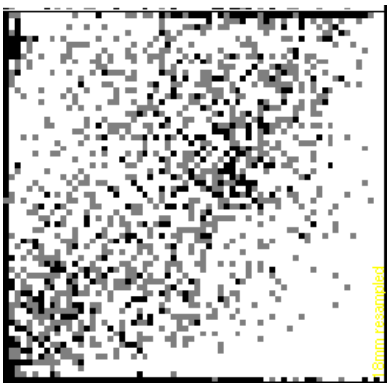

(e) 18mm resampling

Figure 4.7: Joint histograms at different image resampling

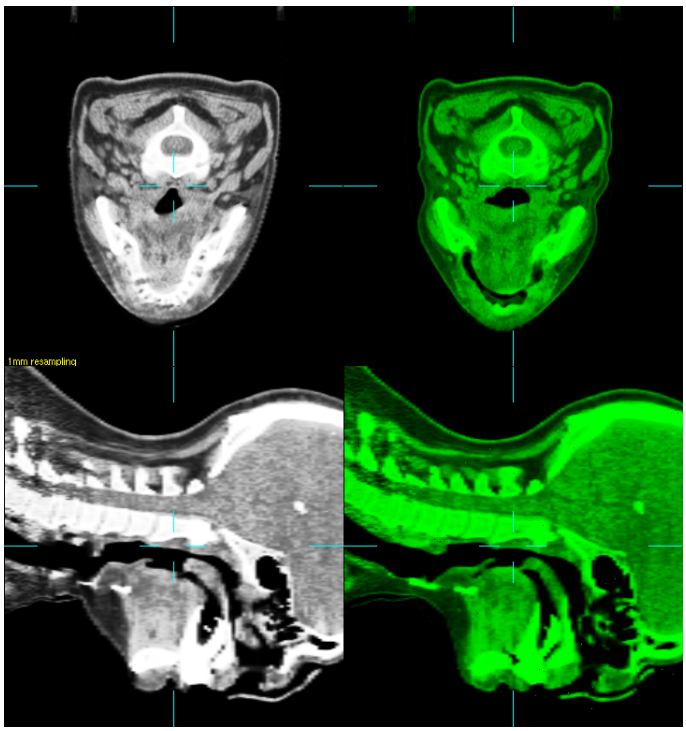

(a) 1mm resampling

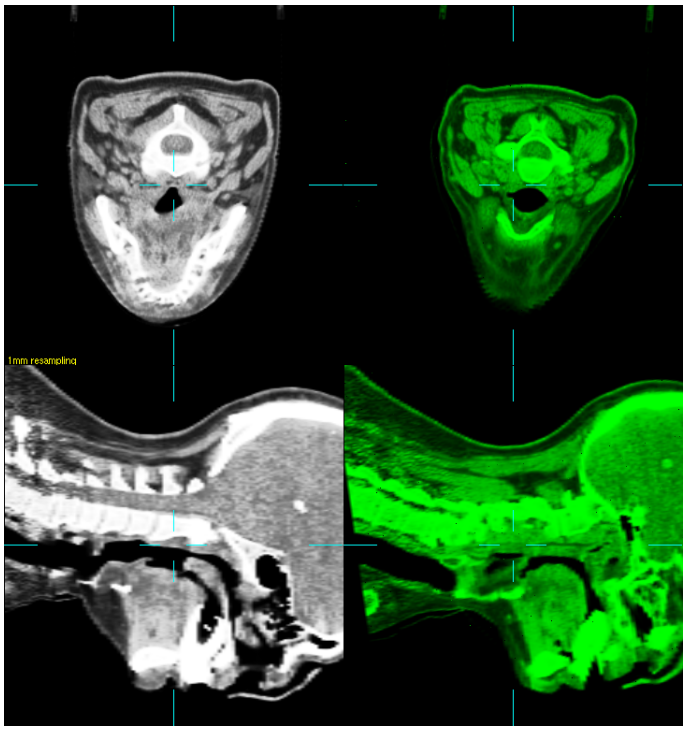

(b) 18mm resampling

Figure 4.8: Registered image pairs for two resampling levels. Floating image on right in green colourscale.

Image resolution/mm	Rigid δ/mm
1	0
3	2
4	2
6	8
8	1
12	11
18	26

Table 4.4: Image resolution and rigid registration error

Conclusion

These joint histograms calculated for a pair of registered images illustrate the appearance of the joint histograms at the end-stage of a rigid registration. The comments contained within the conclusion in the previous section on the usefulness of making deductions from the joint histograms are equally applicable here, yet it appears that there becomes a point when the resolution becomes so coarse that the joint histograms will end up tending towards uniform noise, as was the case when $18mm^3$ resampling was used. Also, for increasingly coarsely sampled images, their joint intensity histograms will reach a point at which they become too sparse to allow the similarity measure to effectively drive the registration process. Other similarity measures such as the sum of squared differences, or the correlation coefficient, could perform better under these conditions.

Estimating the registration error by comparing the transformations to a reference transformation, itself gained through registration, is a useful means of assessing the goodness of the registration, since it allows a relative comparison to be performed. There are other methods of registration assessment such as comparing landmark errors, which provide absolute comparisons, which shall be employed in the following section where the effect of image resolution on non-rigid registration will be investigated. Measurement of registration errors using fiducial markers are considered a gold standard method [West et al., 2001] in the field of

image registration assessment.

Methods for assessing the performance of point-based rigid registration for images in which there are fiducial markers visible by the imaging modalities are described by [Fitzpatrick, West and Maurer Jr, 1998]. In this work, the authors discuss the ideas of Fiducial Localisation Error (FLE), Fiducial Registration Error (FRE), and Target Registration Error (TRE). FLE is defined as the accuracy by which the coordinates of the fiducial markers are determined, usually by finding the variance of the distribution of the centre of mass of the fiducial markers. FRE is the residual error between the corresponding fiducial markers in the fixed and floating images after registration. The TRE is the registration error of the target point, which is the point of interest in the fixed and floating images, and will depend on where the point of interest resides.

Until the publication of their paper the FRE for commercial image guidance systems employing a rigid point-based registration method was quoted as a performance metric to measure the system's quality. Whilst this is a useful measure, what matters for the task is not how well the fiducial points match, but how well the target point of interest corresponds, as measured by the TRE.

For point-based rigid registration, [Fitzpatrick, West and Maurer Jr, 1998] showed that the TRE is strongly dependent on the fiducial configuration, and will improve with increasing the number of fiducial markers. A formula is given for predicting the TRE based on the distribution of fiducial points and the FLE. This is valid for rigid body registration problems in which there are fiducial markers present and could also be used, in principle, for anatomical landmarks. The methodology and analysis that describes how one may approximately model the TRE at a point is rather involved and only applicable for rigid registration, but the interested reader can refer directly to their paper for further insights into how this quantity may be estimated.

One could extend this standardisation category of registration assessment and propose that methods based on landmark errors and overlap

measures (outlined in section 6.3 later) be classed as a silver standard. Methods such as assigning a score after inspecting subtraction images, since they are strongly operator dependent, could be considered a bronze standard.

In general, the rigid registration error increases with increased image resampling. When using a very coarse resolution ($18mm^3$) the registration failed, as shown in figure 4.8. The trend of increasing registration error with increasing image resolution was not followed strictly in this case, as can be seen when the resolution is $8mm$; this results in a very small registration error of 1mm as shown in table 4.4. Sub-voxel accuracy is frequently mentioned in the literature: during the early epoch in image registration, [Collignon et al., 1995] report a registration technique using images with a voxel spacing of $1.33x1.33x4\ mm^3$ resulting in registration errors of $0.17mm$. Of course, this will depend strongly on the interpolation method used during the registration as the image is being transformed and therefore resampled: it will be reliant on high-order interpolation methods to get such low errors.

4.5.7 Effect of Image Resampling on Non-Rigid Registration

Motivation

Following the investigation of the effect of image resampling on the performance of rigid registration, the effect of image resampling on performance of non-rigid registration will be investigated. To do this, the fixed and floating image shall be resampled to different resolutions to investigate how this affects the performance of the registration.

Method

The performance of the non-rigid registration as a function of image resolution size for two non-rigid registration cases was investigated. A non-rigid registration was performed following a rigid registration. Prior to registering, we histogram-equalised both the floating CT image ($I(\vec{r}, t_3)$)

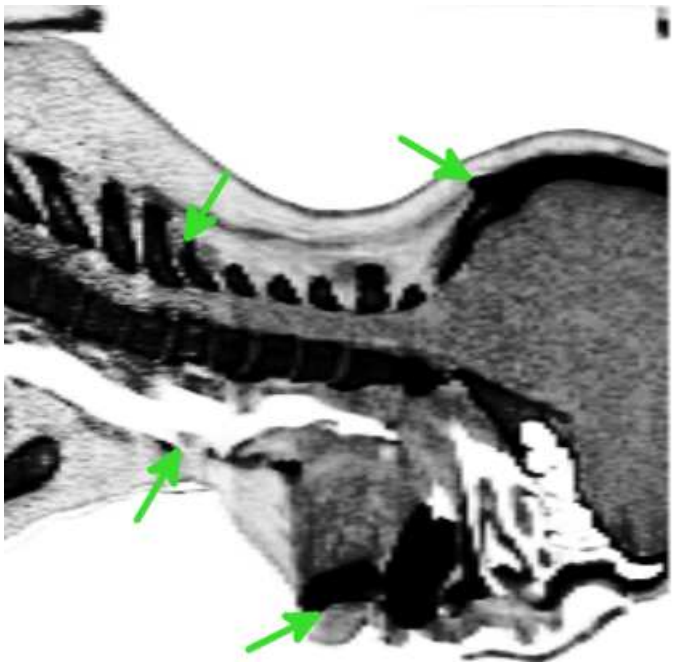

Figure 4.9: Location of anatomical landmarks

and the fixed CT image ($I(\vec{r}, t_0)$) using upper and lower thresholds of -400HU and +400 HU respectively. For the first case, the floating image was transformed with the transformation gained from the best performing rigid registration from the previous section. This was then followed by registering this rigidly transformed floating image to the fixed image using a non-rigid registration with a control point spacing of 40mm. This was performed for a range of image resolutions.

The performance was assessed by measuring anatomical landmark errors. The coordinates of the following anatomical landmark points were identified in the fixed image: the anterior edge of the mandible; the posterior tip of vertebra C6; a bony notch on the posterior edge of the skull and an anterior part of thyroid cartilage that forms part of the larynx. The location of these landmarks is shown in figure 4.9.

The errors to the landmark points were plotted against the image resolution, along with the average error computed over all landmarks.

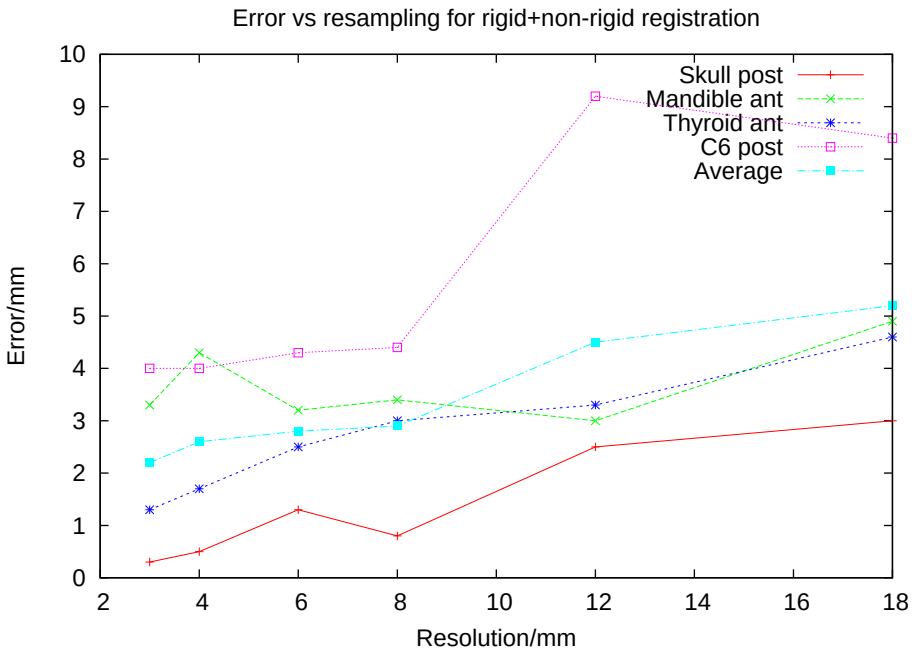

Figure 4.10: Rigid plus NRR error as a function of image resampling resolution

Results

The landmark errors were plotted as a function of the image resolution and may be seen in the graph in figure 4.10.

Conclusion

The process of identifying the landmarks with confidence was different for different points. More work must be done to gain insights on the errors associated with the landmark picking process. This was performed in section 4.7.

For all landmarks, the registration error increases with increased resampling. The average registration error when using a 40mm CPS FFD registration when resampling from 3mm voxels to 18mm voxels ranges from 2.2mm to 5.2mm, which is less than a third of a voxel. Given that this was for a very coarse-scale non-rigid registration, it is reasonable to assume that the kinds of sub-voxel registration errors reported by other authors

are not unreasonable, particularly if a more fine-scale FFD transformation is employed.

4.6 Optimisation Parameters

The preceding section describes how the images may be modified prior to registration. What follows is an assessment of the parameters that define how the registration proceeds within the optimisation process: a process whose goal is to seek an optimum spatial transformation. The registration algorithm uses gradient descent as described earlier in section 4.2.1. There are several relevant parameters that determine how the algorithm will proceed - the algorithm is detailed in the pseudo code at the end of that section. These parameters are the number of resolution levels, the number of steps and the starting step size at the finest resolution level, the maximum allowable number of iterations and, finally, the optimisation tolerance, a relative value that specifies by how much the similarity measure can change from one iteration to the next. If the similarity changes by less than this tolerance, or the maximum number of iterations is reached, then the registration proceeds to the next step; or, if the algorithm is at the final step, at the finest resolution level, then the solution is deemed to be optimum. These parameters are discussed in the following sections, and some experiments are outlined to highlight how these parameters may affect the performance of the registration to assist with setting suitable parameters values so that the registration performs well. The computational cost is also assessed in some instances.

4.6.1 Number of Multi-Resolution Levels

Introduction

The process of image registration using gradient descent involves seeking a minimum value of the image similarity measure with respect to the spatial parameters, as outlined in section 4.2.1. This method has no way of determining that the solution reached is a global minimum value or if it is a local one. So, if there is a means of modifying the problem such that

there are fewer local minima, this decreases the probability of landing in one of these sub-optimal solutions. This may be illustrated graphically in one dimension for the spatial parameter that specifies the rotation about the z axis. Firstly, the similarity at zero angle of rotation about the z axis was calculated for the rigidly registered floating CT image$I(\vec{r}, t_1)$ to the fixed CT image $I(\vec{r}, t_0)$. The floating image was then rotated about the z axis and the similarity was calculated at each angle from -180 to 180 degrees. The results of plotting this may be seen in figure 4.11. Four different similarity measures were calculated to demonstrate how they vary with rotation angle about the z axis. The SSD is normalised to the number of voxels. It can be seen that there are several local extrema[7] and demonstrates in a one-dimensional manner the existence of incorrect local minima for several similarity measures. The idea of capture range is also discussed in section 3.5.

This problem may be overcome by using several multi-resolution levels, whereby the registration operates on images resampled at a coarser scale and the resultant transformation at each resolution level is used as an initial condition for the next most fine scale resolution level. At each finer multi-resolution resolution level, the voxel spacing is halved. The resolution levels are labelled according to the reverse order in which they are are performed. Supposing one uses three multi-resolution levels, the most course level (level 3) is registered first and level 1, the final level, is run last. So for each resolution level l, the voxel resolution in each direction is a multiple of 2^{l-1}. Within each resolution level, multiple steps are used and the starting step size for each resolution level is halved for each successive resolution level.

Another feature of multi-resolution techniques is that they can greatly increase the speed of the registration process since it is much quicker to register images at the coarse-scale than at the fine-scale: the spatial parameters may arrive at an approximately optimum configuration much

[7]When registering images using SSD as a similarity measure, one seeks a minimum value of this similarity measure. For the other similarity measures, one seeks to maximise the similarity.

more rapidly than progression per iteration at a fine scale which is slower and therefore more computationally expensive. Registration at the finest scale is advisable only once the spatial parameters are already within the vicinity of an optimal solution. So, for example, suppose one chose 3 multi-resolution levels and that the starting resolution of the images is $3mm^3$ The first multi-resolution level will involve performing a registration using images with voxels with sides $2^{(3-1)}*3mm = 12mm$, and the second multi-resolution level on $2^{(3-2)} * 3mm = 6mm$, and the final resolution level on images with the starting resolution level (i.e. $2^{(3-3)} * 3mm = 3mm$).

Motivation

Previously, in section 4.5.6, the (single) image resolution was explored to investigate how this may affect the performance of the registration. We shall now try registering a serial image pair using a different number of multi-resolution levels to investigate how the registration performs for both rigid and non-rigid registration. The registration will be assessed both in terms of accuracy and also in terms of computational cost.

Method

For these tests, the final floating CT image $I(\vec{r}, t_3)$ was registered to the fixed CT image $I(\vec{r}, t_0)$. -400HU and +400HU was used for the lower and upper thresholds for the histogram equalisation of both images prior to registration. A variety of multi-resolution levels was tried. The images were resampled from their original resolution of $0.98mm$ x$0.98mm$ x$2.5mm$ to voxels of dimensions $2mm^3$.

3 steps were used where the starting step size was 8 at the highest resolution level. We used 10 iterations for the first run and then we did a second run using 6 iterations. Once registered, the images were assessed according to the method described in section 4.5.4, using an anatomical landmark.

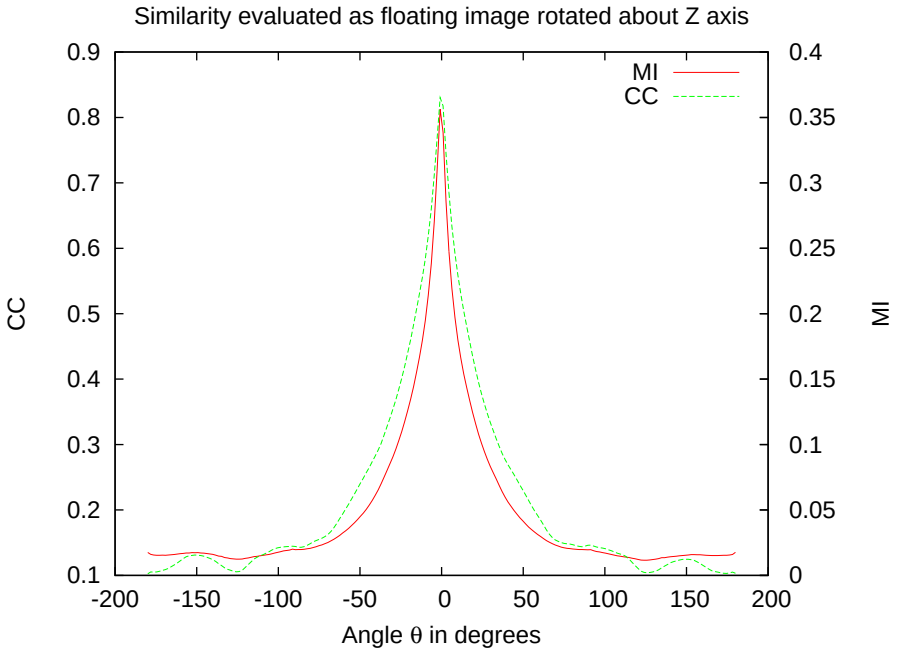

(a) CC and MI

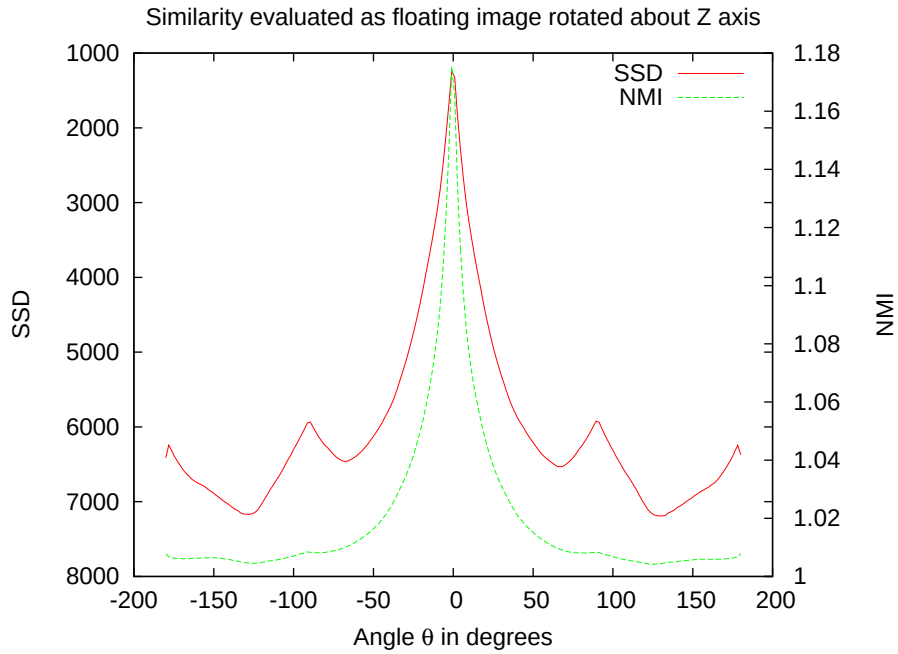

(b) NMI and SSD. The scale of the y axis is inverted for SSD.

Figure 4.11: Plots of similarity measure vs rotation about z axis

Levels	Largest Image resolution	time/s	error
1	1mm	506	1mm
2	2mm	49	1mm
3	4mm	33	1mm
4	8mm	32	1mm
5	16mm	12	fail
6	32mm	10	fail

Table 4.5: How multi-resolution levels affect rigid registration

Levels	time/s	Score
1	12240	1
2	15060	1
3	15960	1
4	16440	1

Table 4.6: How multi-resolution levels affect non-rigid registration

Results

When using between 1 and 4 resolution levels, the rigid registration transformation resulted in a solution with a very small registration error that was measured to be 1mm in each case. The parameters defining the registration did not differ by more than 0.5mm. When using 5 or 6 levels, the registration failed catastrophically. In terms of time, the performance of the registration improved considerably when increasing the number of resolution levels from 1 to 3, with a speed-up by a factor of approximately 15. Table 4.5 shows the rigid registration results, and table 4.6 shows the non-rigid results for the tests to determine the optimum number of resolution levels.

Conclusions

In the rigid case, having several resolution levels can help the registration find the solution more quickly: compared to having a single resolu-

tion level, the registration is approximately 15 times faster than when using 3 levels. However, having too many resolution levels results in an image which is too downsampled, and which, as was also confirmed previously in section 4.5.6, can result in a registration error. If the image is too coarsely resampled, there will be a point at which the registration is unable to iterate towards the correct solution, since there will be a lack of features within the image: in other words, a lack of information on which the algorithm may operate. So instead of the solution going from an initial configuration in the parameter space towards some configuration that is closer to the actual solution, the algorithm modifies the parameters in such a way that they depart rapidly to some distant configuration that the next level of resolution is then unable to correct for and bring back towards the optimal solution, since it has departed beyond the capture range.

The performance of the non-rigid registration following on from the rigid case were less clear. When each registered image pair was examined and assessed, there was no clear difference between the four cases. Also, the solution was not found more quickly when using more than one resolution level. Indeed, contrary to the rigid case, the solution was found more quickly when using a single image resolution during the registration. It is unclear why this was the case.

The difficulty in comparing one non-rigid registration solution to another, by simply assessing the images side-by-side in an image viewer to assign a score, suggests that a more quantitative means of assessing the registration is required. Assessing the performance of the non-rigid registration is a non-trivial task.

4.6.2 Number of Steps and Starting Step Size for Rigid Registration

Introduction

At a given image resolution level within the registration, the amount by which each of the spatial parameters is modified per iteration will deter-

mine the granularity by which the image may be spatially transformed during the process of registration. This amount is known as the step size and is described in section 4.2.1. A rigid transformation is specified by three shift parameters and three rotation parameters. The step size is numerically the shift in mm for the translation parameters and a rotation in degrees for the rotation parameters. It may be thought of as a granularity by which the parameter space may be explored during optimisation.

During registration, several coarse to fine steps can be utilised within each resolution level. The optimisation proceeds by iteratively updating the transformation parameters according to gradient descent until the solution for this step has been iterated to: either by the algorithm reaching the maximum number of iterations, or if the change in the similarity measure is less than the tolerance. The algorithm will then proceed to the next successively finer step, using the transformation parameters from the preceding step as an initial condition, but with half the step size for this step. In practice, one may use several of these coarse to fine steps, known simply as steps. In this section, two things that will affect the performance of the registration will be investigated: the starting step size (at the most coarse step) and the number of steps.

Aim

To see how changes in the starting step size parameter affects the performance of the rigid registration, both in terms of the registration error and also in terms of the computational cost. It shall also be explored how having several steps affects the performance of the registration. The aim is to find suitable values for the starting step length at the coarsest step and the number of steps.

Method

How the starting step size affects the performance of the rigid registration was investigated, by registering the final floating CT image $I(\mathbf{r}, t_3)$ to the fixed CT image $I(\mathbf{r}, t_0)$. The pre-processing performed prior to registration involved using -400HU and +400HU for the lower and upper

thresholds for the histogram equalisation, and resampling the images from their original resolution of $0.98mm$ x $0.98mm$ x $2.5mm$ to voxels of dimensions $2mm^3$ using the method explained in section 4.5.6. Using a single step and a maximum of 10 iterations, registering using initial step sizes of $0.5mm$, $1mm$, $2mm$, $4mm$, $8mm$, $16mm$ and $24mm$ was tried, to study how this affects the performance of the registration at a single image resolution level and using a single step. This was repeated using 2, 3 and 4 steps using these same initial starting step sizes.When assessing the registrations, one parameter set resulted in an optimal registration as judged by measuring the distance from an anatomical landmark, as previously described in section 4.5.4. This optimum transformation was designated to be the reference transformation. The registration parameters defining each transformation were recorded and it was noted that there were only translations present and no rotations. We computed the global registration error (in this case, global, since the transformation has no rotation component) by finding the Euclidean distance from each transformation to the reference transformation. For each registration, we also recorded the computation time in seconds required to compute the registration.

Results

For each registration, the error of each transform compared to the reference transform was computed and plotted along with the computational time in figure 4.12. The best-case registration - which was taken as the reference - was when using a starting step size of 2 and 4 levels. When using a starting step size of 0.5, this resulted in a poor registration regardless of the number of steps that were used. At the other extreme, when using a starting step size of 24mm, this resulted in an increasingly good registration when more steps were used, with an error of 9.8mm for 1 step to an error of 1.3mm for 4 steps.

Conclusions

The worst performance of the registration occurred when the starting step size was either much too small to allow sufficiently large distance

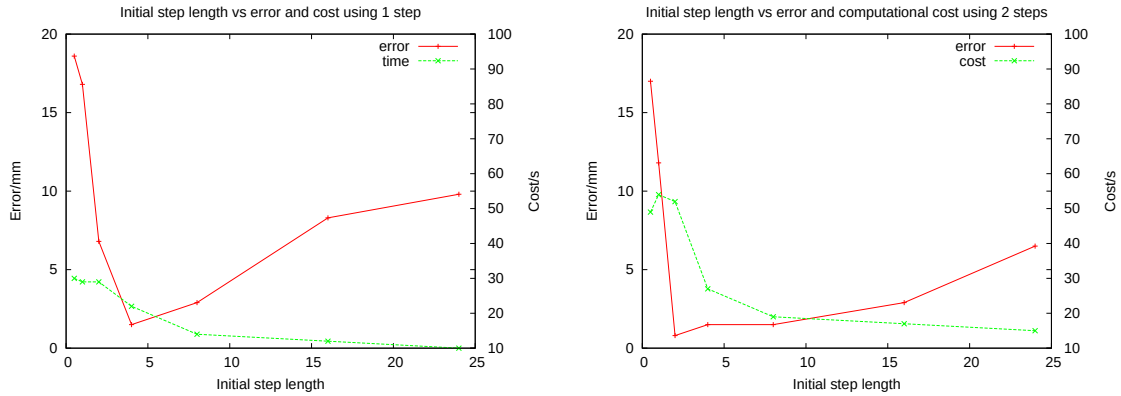

(a) A single step

(b) Two steps

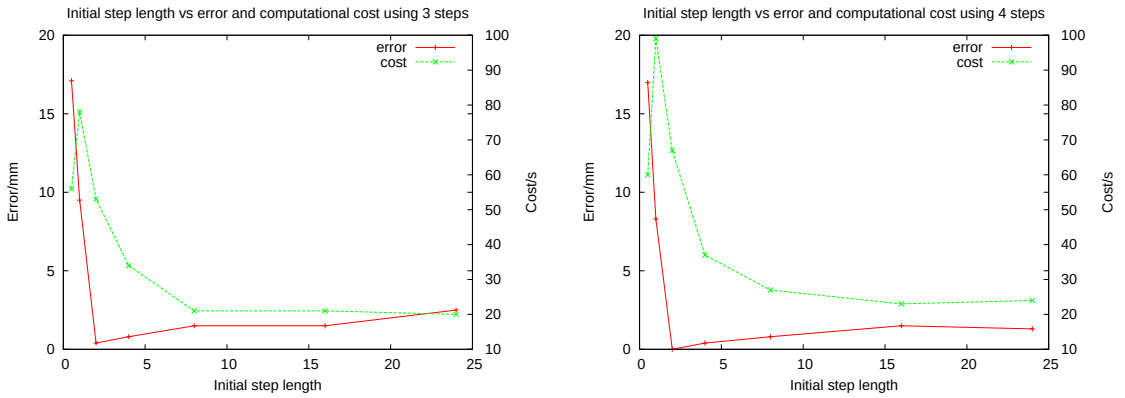

(c) Three steps

(d) Four steps

Figure 4.12: The effect of step size and number of steps on accuracy

in parameter space to bring the images into registration, or much too large. If the step size is far too large then there can be a considerable overshoot or undershoot in parameter space. If one considers this in one dimension, then the registration error will be in the range of $0 \longmapsto \frac{\mu}{2}$ for a given step size μ. Therefore it is preferable to have μ as low as possible to allow convergence to the solution. However, when the starting step size is too small, the reason that the registration will result in an error is if the product of the step length and the number of iterations is less than the capture range. For a single step, the best performance was with a starting step size of 4. When using several steps, the error becomes less dependent on the initial step length. This illustrates that a coarse to fine method with several steps is the best way to achieve a solution. In terms of computational cost, it is advisable to use a larger starting step size. This can also be a useful approach when unsure how far away from the solution you are, so there is a benefit in being able to iterate quickly but coarsely towards the solution.

Assessing the registration performance in this case was a trivial matter, since one only need compute the (global) difference between each transformation and the reference transformation, since for this problem, the required transformation did not contain any rotational components. Had that been so, then a single Euclidean distance measure would not represent the global registration error and another means of parameterising the error would be required: either some form of global translational and angular errors, or a set of local distance measures could be calculated for a set or points of interest within the images.

To summarise the findings: when registering serial CT head and neck images with voxels of dimensions $2mm^3$, the best-case choice for the number of steps is either 3 or 4, and then it is advisable to use a starting step size of either 4 or 2 respectively. This is the case when using 10 iterations. This will provide a good combination of registration accuracy without too much computational cost.

4.6.3 Starting Step Size for Rigid Plus Non-Rigid Registration

Introduction

When registering images using a B-splines non-rigid transformation, the starting step size is now the amount by which each control point that parameterises the transformation may be modified per iteration during registration. What was said previously for the rigid parameters now applies to the parameters that define the non-rigid registration in terms of the progression through each step. However, there is, in addition, another parameter that defines the granularity of how the transformation can operate, which is the control point spacing. What happens at a single, coarse control point spacing was considered in this section.

Aim

To further explore the effect of modifying the starting step size on the performance of the non-rigid registration, following a rigid registration, shall be investigated. The value used for the starting step size that was found to work in the preceding rigid case will be tested.

Method

To see how the starting step size affects the performance of the non-rigid registration, the final floating CT image $I(\vec{r}, t_3)$ was registered to the fixed CT image $I(\vec{r}, t_0)$ using a rigid registration followed by a B-splines free-form deformation with a $40mm$ control point spacing. The images were histogram equalised as per the previous section and resampled to voxels of dimensions $6mm^3$ using the method explained in section 4.5.6. Using a single step and a maximum of 10 iterations, initial step sizes of $1mm$, $4mm$, $8mm$ and $18mm$ were used, to study how this affects the performance of the registration at a single image resolution level and using a single step. This was repeated using 2 steps, using these same initial starting step sizes. For each registered image pair, the deformation field was superimposed onto the images to illustrate how the floating image has been non-locally deformed by the transformation. This aided the

assessment of the transformation in each case since it was a non-local transformation.

Results

The results of the non-rigid registration following on from a rigid registration are displayed in figures 4.13 and 4.14. The images were inspected to find where the transformations were different. This location was then displayed zoomed in on this region, as can be seen in the figure. The non-rigid deformation vectors are superimposed in white on each orthogonal image slice on the floating image for each starting step size. In each case, the fixed image is on the left with a greyscale and the floating image is on the right displayed using a green colourscale. The non-rigid registration has been applied to the floating image. The blue crosswires represent the same coordinates in each case: if the images are well registered then they should appear on the same anatomical structure in each case. If badly registered, then they will not. When using step sizes of $1mm$, the registration errors were larger across the image, when using one or two steps. If one inspects the transaxial view (the top row in figures 4.13 and 4.14) it can be seen that part of the vertebra is missing from the floating (green) image compared to the fixed (greyscale) image when the step length is $1mm$, indicating poor registration. The registration performed slightly better when using step sizes of $4mm$, $8mm$ or $18mm$, irrespective of whether one or two steps were used. On close inspection, the $18mm$ starting step size when using two steps achieved the best registration across the image.

Conclusions

The performance was not strongly dependent on the starting step size parameter when required to recover a small deformation. In all cases, the deformation was non-negligible across the whole image, as can be seen by the size of the projection of the vector deformations present in the floating images. In the region displayed in the figure, the $4mm$, $8mm$ and $18mm$ starting step sizes performed better than the $1mm$ starting step size. In the next section, we shall explore this effect when there

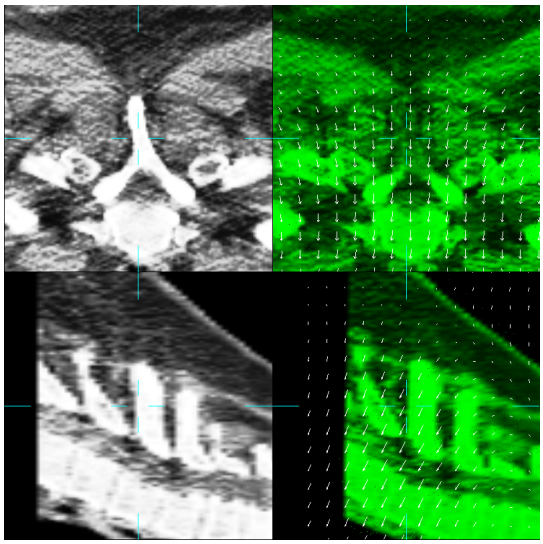

(a) 1mm step length

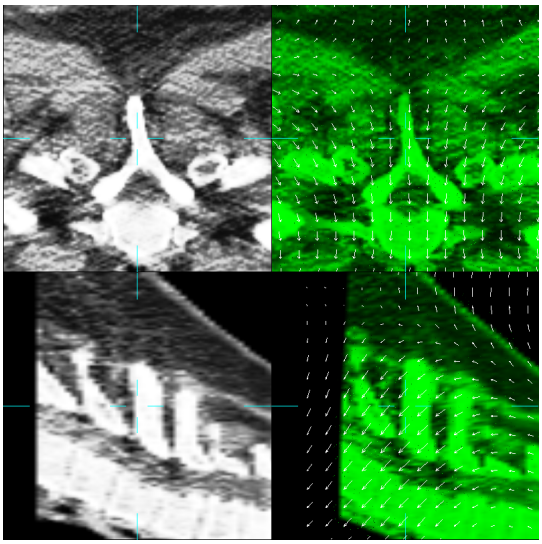

(b) 4mm step length

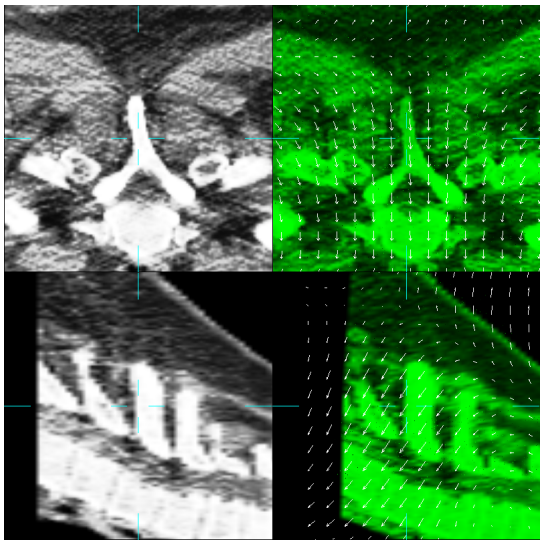

(c) 8mm step length

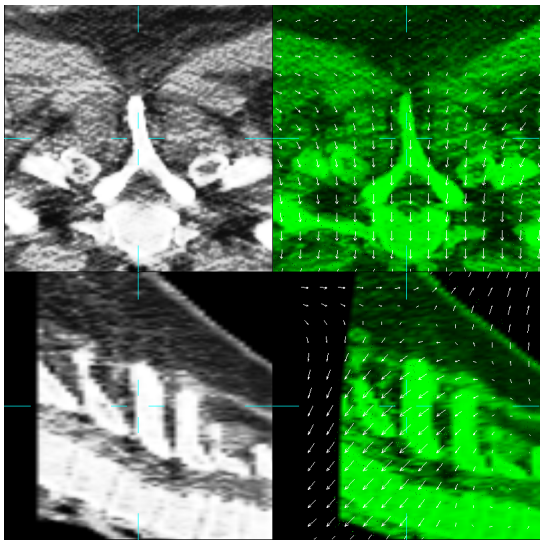

(d) 18mm step length

Figure 4.13: Rigid plus NRR results with changing step size using 1 step. Fixed CT images in grey colourscale and floating CT images in green colourscale.

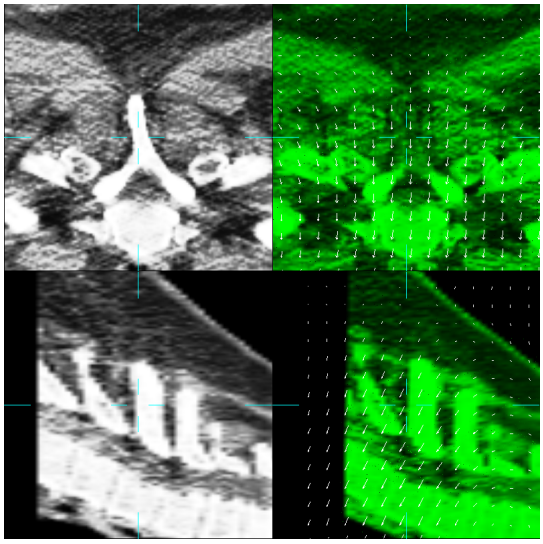

(a) 1mm step length

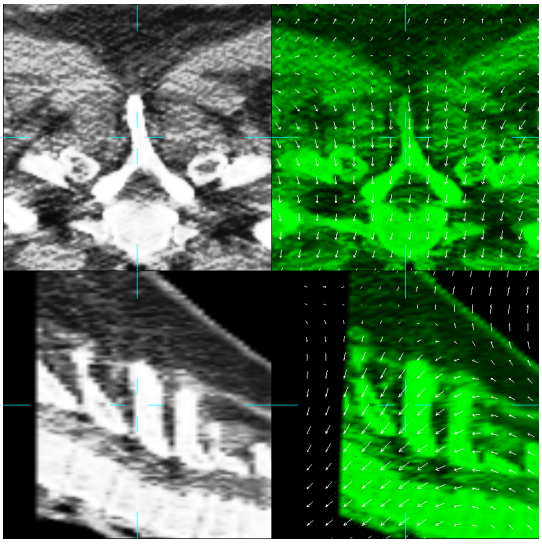

(b) 4mm step length

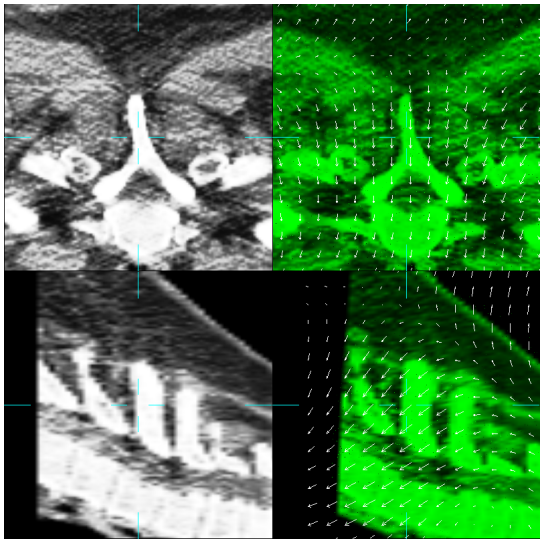

(c) 8mm step length

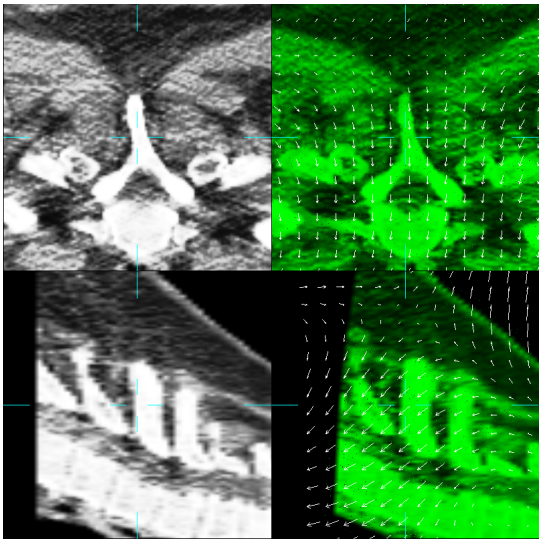

(d) 18mm step length

Figure 4.14: Rigid plus NRR with step size and 2 steps. Fixed CT images in grey colourscale and floating CT images in green colourscale.

is no initial rigid registration to study how the performance varies with starting step size when larger deformations are required to register the images. Also, a relatively small number of iterations was used during this test, which will have an adverse affect when a small starting step size is used.

4.6.4 Number of Steps and Starting Step Size for Non-Rigid-Only Registration

Aim

To see how changes in the starting step size parameter affects the performance of the non-rigid-only registration, and to explore how having either one or two steps affects the performance of the registration. The aim is to find suitable values for the starting step length in terms of the control point spacing and to investigate whether it is worth having multiple steps. In this case, the required non-rigid deformations are larger in magnitude than in the preceding case, since the deformation field must also recover the large displacement previously contained within the rigid transformation. The rigid registration had to recover a displacement of the order of $19mm$ so it will be informative to explore how the system behaves when a large motion must be recovered by the FFD with no prior rigid registration.

Method

The same process adopted in the previous section was repeated, however, the rigid rigid registration was not performed prior to performing the FFD non-rigid registration.

Results

Figures 4.15 and 4.16 shows the results of the non-rigid registration, with the non-rigid deformation vectors superimposed in white at 40mm intervals on each orthogonal image slice on the floating image for each

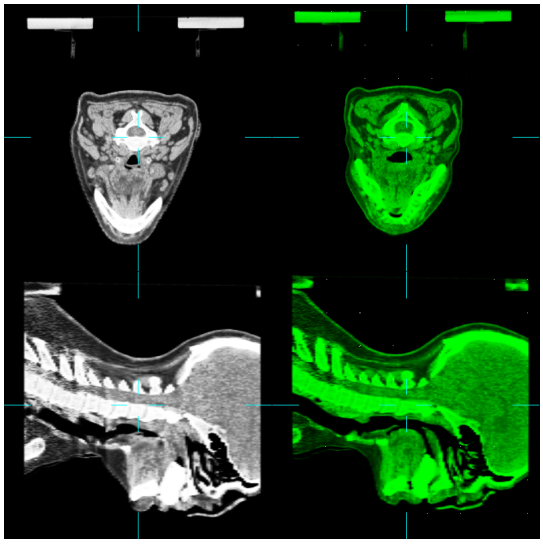

(a) 1mm step length

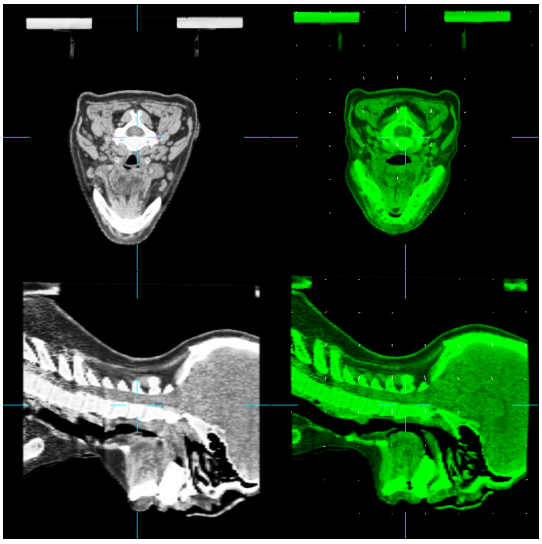

(b) 4mm step length

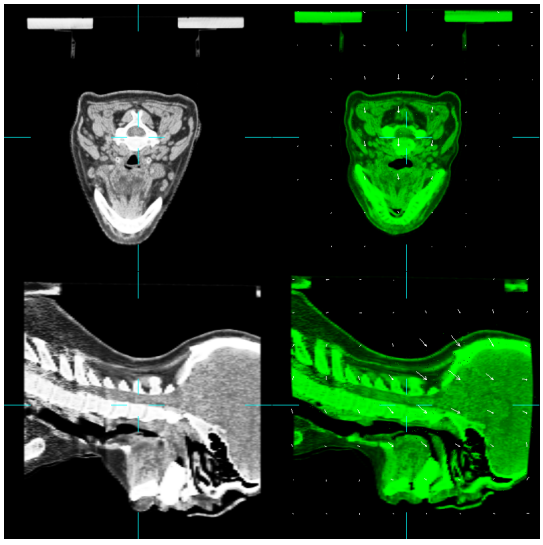

(c) 8mm step length

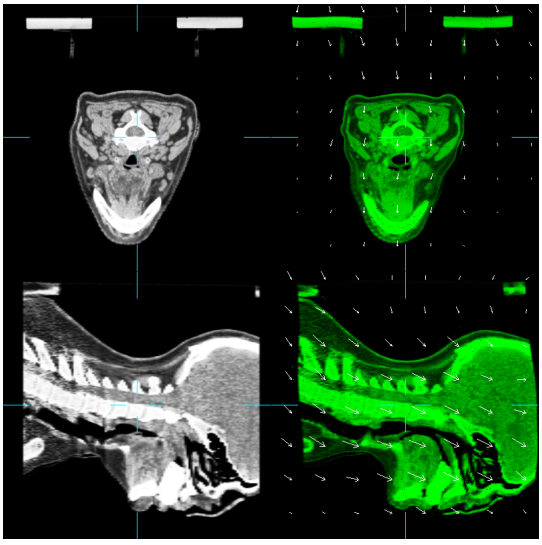

(d) 18mm step length

Figure 4.15: Non-rigid-only registration using 1 step

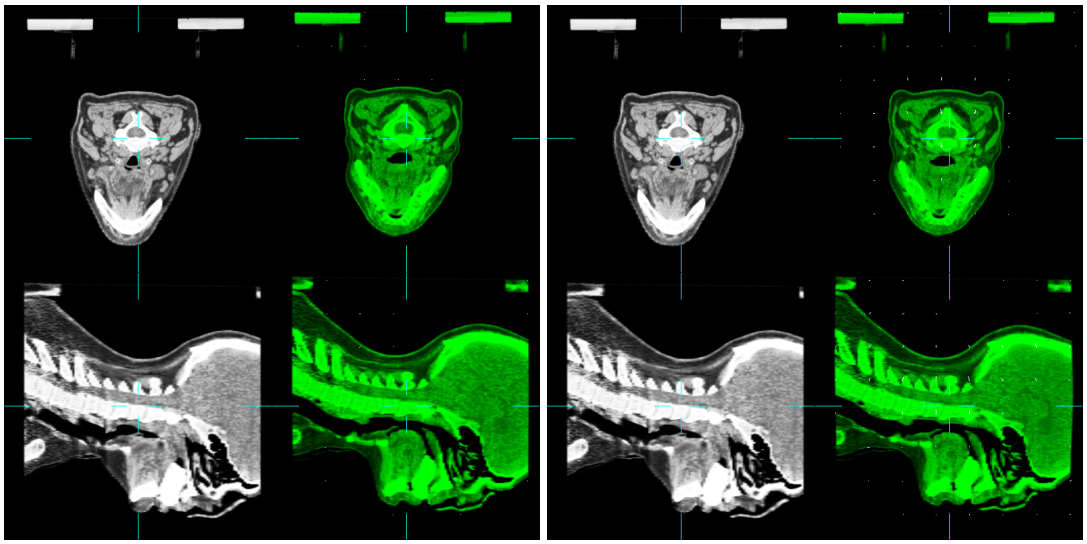

(a) 1mm step length

(b) 4mm step length

(c) 8mm step length

(d) 18mm step length

Figure 4.16: Non-rigid-only registration using 2 steps

starting step size. In each case, the fixed image is on the left with a greyscale colourscale and the floating image is on the right displayed using a green colourscale. The non-rigid registration has been applied to the floating image. The blue crosswires represent the same coordinates in each case: if the images are well registered then they should appear on the same anatomical structure in each case. If badly registered, then they will not. When using step sizes of 1mm and 4mm the registration errors were large across the image, as can be seen when using one or two steps. The size of the transformation across the image is negligible, with a maximum transformation of approximately 4mm noticed on the posterior base of the skull, but, with much of the imaging, the magnitude of the transformation was close to zero. When using a starting step size of 8mm, there is an appreciable deformation across much of the image and in certain regions, such as the posterior part of the spine, where the registration is getting close.

Conclusions

In this above test, a B-splines-based non-rigid registration was performed *without* first doing a rigid registration, which is unusual, since for a serial registration, one would normally first perform a rigid registration to make any global adjustments and to ensure the non-rigid registration is only left with small (local) deformations to recover. This was to test the algorithm's ability to traverse a large distance in the parameter space, in order to test how well it could perform under more extreme conditions than it would usually operate during serial registration. The required rigid transformation in this case is a vector consisting of a translation of magnitude 19mm. There seems little advantage to using one or two steps. What seems more important is the starting step size. When using 18mm for the starting step size, the registration is good across the image. When using a starting step size of 8mm, the registration is only good in the posterior part of the neck close to the base of the skull. The registration in other areas is poor when using this starting step size. When using starting step sizes of 4mm and 1mm the transformation is negligible in magnitude across the image. To conclude, the non-rigid registration per-

formed well when using a starting step size that is comparable in magnitude to the required deformation that must be recovered. However, it should also be considered in conjunction with the number of iterations. For example, when more iterations are used, a lower starting step size may be chosen. However, according to the literature, it should not be larger than $0.4 * C$ as previously discussed, to ensure the transformation remains diffeomorphic.

4.6.5 Maximum Number of Iterations

Aim

When the registration is at the final step, it will terminate if the similarity at the current iteration compared to the previous iteration is within the tolerance, or if the maximum number of iterations parameter is exceeded. In this section, how the maximum number of iterations parameter affects the performance of the registration is investigated.

Method

CT image $I(\vec{r}, t_3)$ was registered to CT image $I(\vec{r}, t_0)$ using a non-rigid registration with a B-splines free-form deformation without a prior rigid registration, to ensure that the size of the spatial deformation recovered by the FFD transformation would be large. Before registration, the images were histogram equalised as per the previous section and resampled to voxels of dimensions $6mm^3$ using the method detailed in section 4.5.6. A 40mm control point spacing and a single step with an initial step size of $18mm$ was used. At each iteration, the transformation was saved. The maximum number of iterations was set to 100. Four anatomical landmarks within the images were defined: the posterior tip of the vertebra closest to the skull (C1); the anterior edge of the mandible; the centre of the 4th vertebra (C4) and the 8th vertebra (T1). The coordinates of these points were noted. For each of these points, the Euclidean distance of the point in the fixed coordinate space to the floating image's coordinate space was calculated at each iteration whilst the registration occurred.

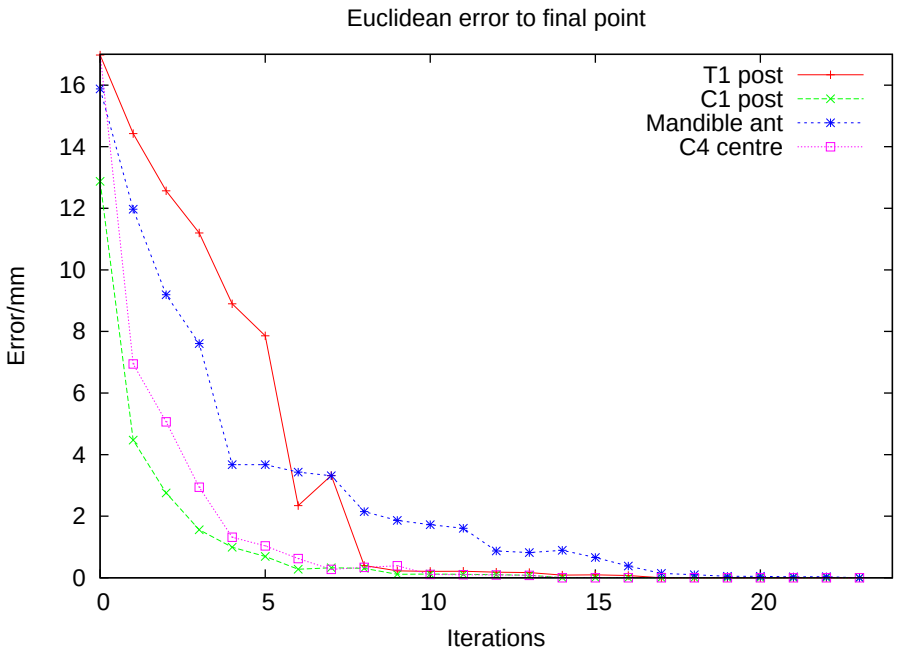

Figure 4.17: Registration error as a function of iteration

The distance error of the four corresponding anatomical landmarks was plotted as a function of iteration.

Results

A graph was plotted of the Euclidean error calculated for the corresponding anatomical points of interest within the fixed and floating images at each iteration during the registration. This can be seen in figure 4.17. The registration terminated at 24 iterations when the similarity change was lower than the threshold value, which was hard-coded within the vtkCmic software.

Conclusions

The results demonstrate that the registration error is reduced to an insignificant level for the four chosen anatomical landmark structures after about 15 iterations. This large number of iterations required for convergence was because there was no initial rigid registration and was purposely chosen so that the non-rigid registration transformation would

have to travel significantly within the parameter space to represent a worst case scenario.

For the T1 vertebra, the distance error increased from the 6th to the 7th iteration. This demonstrates that during a registration, the registration may worsen at a point. Although the registration acts locally, it still operates over a local region, and whilst it might be better on average from one iteration to the next, this need not be so for all points as has been demonstrated. This is particularly the case here since the control point spacing is 40mm so the deformation is not that local.

To conclude, the maximum number of iterations should not be set to a number less than 24, although of course this could mean that the registration could take a long time to compute if the maximum number of iterations is large. Also, the initial step size will have a strong impact on how many iterations are required: fewer will be required when the starting step size is set to a larger number.

4.7 Demonstration and Assessment of the Registration Framework

4.7.1 Introduction

In the preceding sections within this chapter, how to set up the registration parameters that determine how head and neck serial CT images can be registered was explored. Using the parameters found earlier that tailor the algorithm for our problem, registrations will be performed and assessed for several images.

When performing the histogram equalisation prior to registration, the lower and upper thresholds of -400HU and 400HU were found to give good results. It was determined that the default value of 64 bins to construct the joint intensity histograms during registrations was a suitable choice. Using several multi-resolution levels can assist the rigid registration to complete more quickly, although using more than 4 levels is ill advised.

For the rigid registration, three steps should be used and the step length should be chosen such that the step length at the final step is of a similar magnitude to the desired registration precision. For the non-rigid registration, a single step is preferable as there is little advantage in having more than one step. The step length should be chosen to be be approximately 40% of the control point spacing. The number of iterations should be set to a large enough value to ensure a compromise between computation time and accuracy. Using 25 iterations should be sufficient for the registration algorithm to converge.

Aim

To test the performance of the registration using the previously determined set-up found in the preceding sections within this chapter.

Method

Image data from four patients from the radiotherapy clinical trial was used. For each patient, the CT image $I(\vec{r}, t_3)$ was registered to CT image $I(\vec{r}, t_0)$. Prior to registration, all images were histogram equalised using the above mentioned thresholds and bins, and were smoothed and then resampled to 3mm voxels as described in section 4.5.6. The images were registered rigidly using 3 multi-resolution levels and three steps, and an initial step length of 8mm. The images were then non-rigidly registered using an increasingly fine-scale FFD registration, starting with control point spacings of 40mm, 20mm and finally 10mm, using a step size of 18mm, 9mm and 4mm, using a single step for each FFD registration and using a voxel size resampling of 6mm, 4mm and 2.5mm respectively. At each increasingly fine-scale registration, the input transformation at the higher level of control point spacing was first subdivided and then used as an input to the lower scale of transformation. For each registration, the maximum number of iterations was set at 25.

The performance of the registration was assessed by computing the difference image for the registered pair and inspecting the images. Images that are in better correspondence will result in more uniform difference images.

The assessment of the registration also involved measuring distance errors from anatomical landmarks that were identified in the fixed and floating images. These landmarks were the posterior tip of the C6 vertebra; the anterior edge of the mandible; the notch on the posterior skull surface and an anterior part of thyroid cartilage that forms part of the larynx. The anatomical landmarks were picked on the fixed and registered floating image by displaying both in the VTKView image viewer at the same time and noting the 3D location of the landmark to determine the coordinates. The registration error at each landmark point was determined as the Euclidean distance between corresponding coordinates.

The error associated with an observer identifying the corresponding landmark pair in the VTKView image viewer was estimated by simulating the landmark-picking process as follows. Firstly, an image and a transformed copy of itself were loaded into the image viewer. The transformation was an arbitrary transformation that displaces each corresponding landmark point by some small but known amount that was hidden from the observer at the time of picking the landmarks. The corresponding landmarks were identified on the fixed and floating image pair. The inverse transformation was applied to each landmark point on the floating image. Then, the Euclidean distance error was computed for each landmark point. The landmark picking was performed for each patient using the first CT image in each case. This was repeated twice, so that for each landmark point, the error was calculated 6 times. The average value for each landmark point was calculated to be the estimate for the error associated with identifying a position.

Results

The average landmark identification errors across the images were found to be 0.98mm for the mandible, 0.49mm for the skull, 1.22 mm for the thyroid cartilage and 0.62mm for the C6 vertebra. The landmark identification errors averaged across the patients are shown in figure 4.18. The estimated landmark errors are shown as error bars on the plot. As the control point spacing became increasingly fine-scale, the subtraction im-

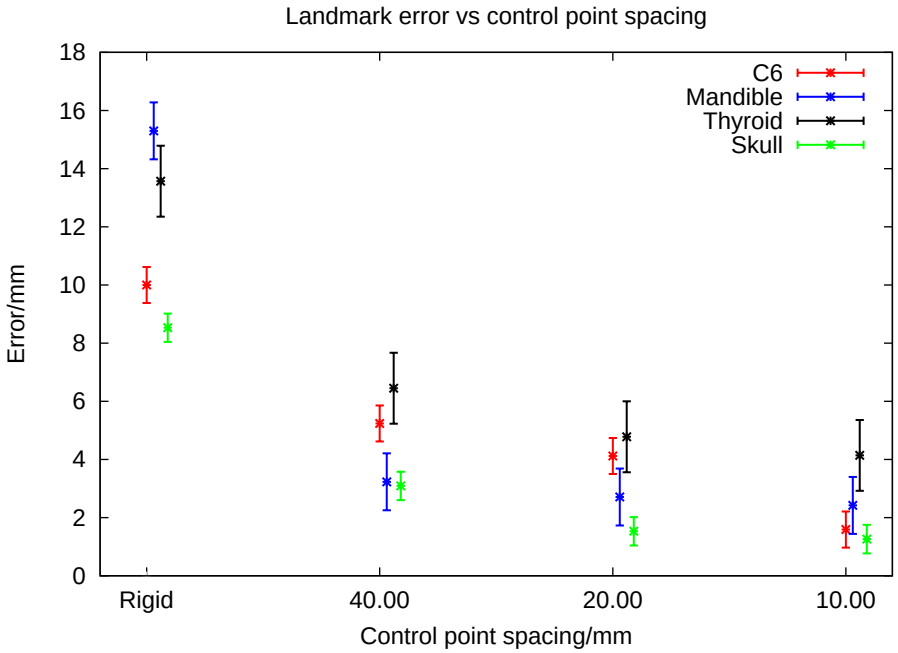

Figure 4.18: Landmark errors with decreasing control point spacing

ages were in general more uniform for all four images for increasingly fine control point spacing. The subtraction images for patient two are shown for the orthogonal planes within the image in figure 4.19. However, in the vicinity of a metal dental artefact, the subtraction images were not more uniform with increasingly fine-scale registration. The difference image was calculated in the region of the artefact and zoomed to illustrate this problem, shown in figure 4.20.

Conclusions

When using all the registration parameter settings that were investigated in the earlier parts of this chapter, this results in a successful performance of the algorithm for these head and neck images. When registering the images using an increasingly fine-scale control point spacing, the registration accuracy improves as measured using the distance errors for pairs of four corresponding anatomical markers in the fixed and floating images. The errors associated with the observer picking the landmarks were different for each landmark, since some landmarks are more

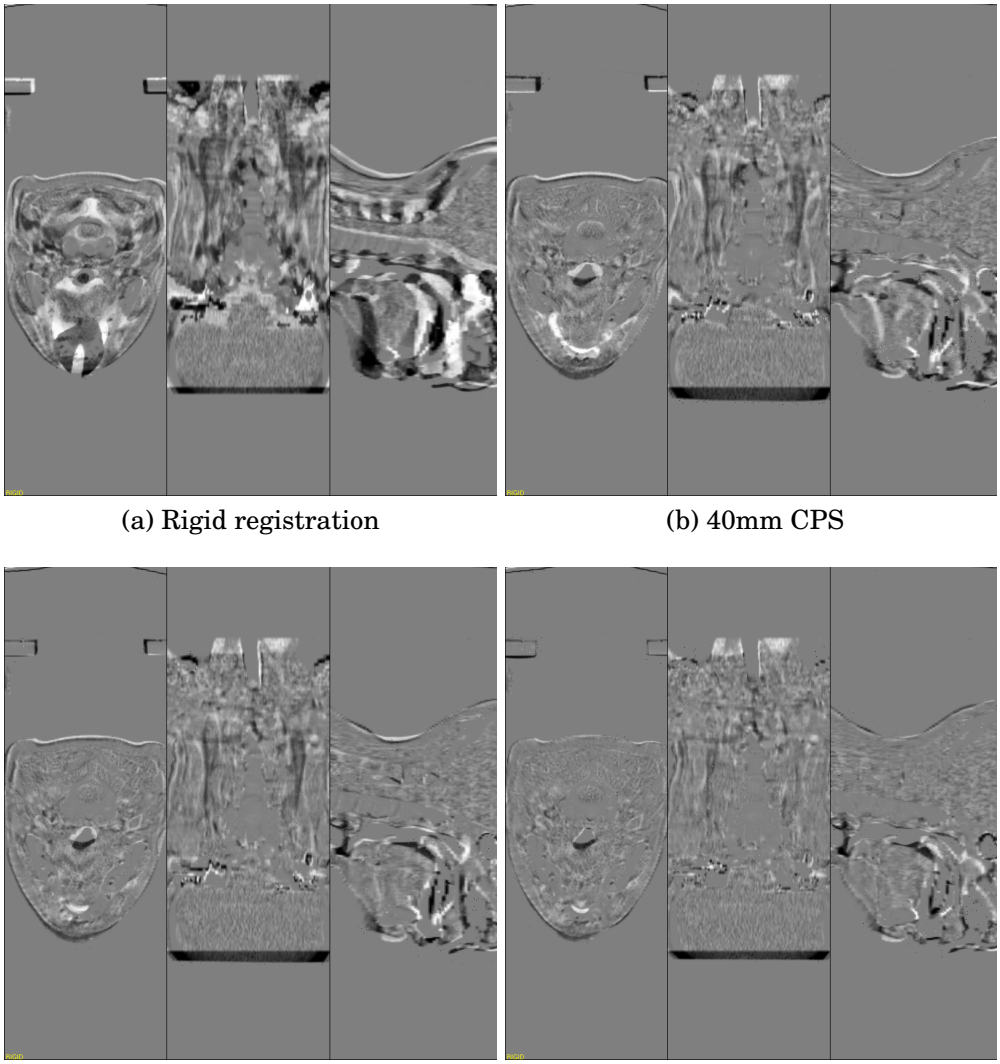

(a) Rigid registration (b) 40mm CPS

(c) 20mm CPS (d) 10mm CPS

Figure 4.19: Subtraction images at different registration levels displayed in orthogonal planes

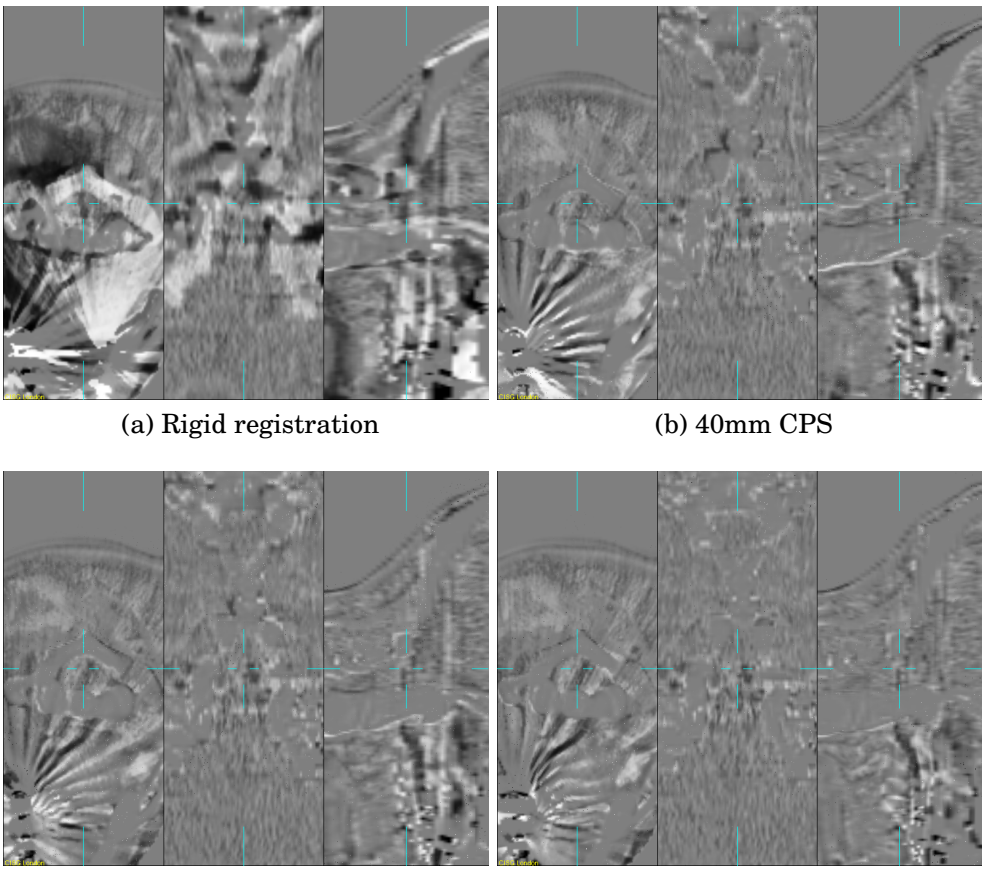

(a) Rigid registration (b) 40mm CPS

(c) 20mm CPS (d) 10mm CPS

Figure 4.20: Difference images at different registration levels zoomed in the locality of a dental artefact

easy to identify than others. That the registration framework developed in this chapter is effective for the problem is also corroborated by the fact that the difference images become more uniform in appearance as the registration becomes increasingly fine-scale.

The registration performed differently across the different landmarks, with the best accuracy found for the skull base landmark, since this is a relatively stable area of anatomy for the registration algorithm to recover the small differences: it is less likely to move in a relative sense with respect to immediate neighbourhoods within the image. Also, the landmark picking error is smaller, so it is a more precise measure. The registration performed less well for the thyroid landmark for converse reasons: that it has a higher landmark picking error, and also it is quite a mobile structure relative to its local anatomy.

On one of the patients that were registered, there were dental artefacts, which resulted in a less uniform difference image in the vicinity of the artefact. There are two reasons for this lack of uniformity in this case. Firstly, in the vicinity of the artefact, the images contain more features. Secondly, the dental artefacts will compete with the features in the image in such a manner that if the artefacts are in better alignment, then the anatomical features will be in worse correspondence (and visa versa). Since the landmark errors measured to assess the registration were not that close to the vicinity of these artefacts, there was no noticeable outlier. Since the transformation acts in a local manner, it is possible for the registration to perform well locally in some regions and less well in others. The problem caused by dental artefacts shall be studied, since they are a commonplace occurrence amongst the patient data within the clinical trial. Methods for studying this problem shall be explored in more detail in the next chapter, including a simulation technique for estimating how much effect an artefact can cause on the non-rigid registration, as well as some proposed registration strategies designed to overcome this problem.

Parameter	Setting
CT Hist. Equal. Thresholds (Min HU, Max HU)	$-400HU, 400HU$
Number of Bins Used in the Joint Histogram	64
Image Resampling for Rigid Registration	$< 3 * e$
Number of Multi-Resolution Levels (RR)	3
Number of Multi-Resolution Levels (NRR)	1
Number of Steps / Starting Step Size for RR	3 or 4, $4mm$ or $2mm$
Number of Steps / Starting Step Size for NRR	1, $0.4 * C$
Maximum number of iterations	> 24

Table 4.7: Summary of parameter values found

4.8 Effect on the Registration Performance Caused by Dental artefacts

In figure 4.20, the problem of non-rigid registration assessment in the vicinity of the artefacts when using difference images was highlighted. Dental fillings cause the CT reconstruction algorithm to create streak artefacts in the transaxial plane. This is because the attenuation of the metal is far higher than the tissues present in the body, and therefore the detection system is subject to photon starvation. When reconstructed, this results in very prominent streaks within the transaxial plane. A major problem with this is that the artefacts, will *always* be in the transaxial plane. If the floating image is rotated about the sagittal plane compared to the fixed image, then this will result in the artefacts whose features will be on an inclined plane and will introduce local errors in the non-rigid registration's transformation. Since they are strong features within the image compared to the features that represent the actual anatomy, it is quite likely that a registration algorithm will preferentially attempt to register these over the actual anatomy.

4.9 Conclusions

In this chapter, a non-rigid registration algorithm was tested to register serial CT head and neck images. This algorithm has many options

and parameters to set, to enable it to work effectively. The effect of each parameter on the performance of the registration was tested in order to determine an appropriate choice. The appropriate choice was to ensure the registration was optimum when the parameter was modified across some finite range to discover a suitable first approximation value.

Each parameter could be explored in further detail, such as the alternative methods for histogram transformations [Arici, Dikbas and Altunbasak, 2009] prior to registration for this problem, an aspect other workers have considered [Zhu, Chan and Lam, 1999]. The approach in this chapter was to ensure these parameters were set up such that an accurate result could be achieved. Further work can be done to fine-tune each of these parameters. The parameter settings found within this chapter are summarised within table 4.7. In this table, the symbol C refers to the control point spacing. The resampling should be less than $3 * e$ where e is the error that can be acceptable for the registration.

The methods employed within this chapter have implicitly assumed that it is possible to study the effect of modifying each parameter in turn. It seems reasonable to assume that there exists a causal effect between modifying one parameter *in isolation* and the impact this has on the accuracy of the registration performance to the first order. This one-at-a-time assumption with regards to choosing the appropriate settings for each parameter is also a reason why each parameter was not fine-tuned to a great degree; rather, all relevant parameters that will govern the behaviour of the registration were explored to discern approximate values for all parameters working in concert.

Our framework was then tested on a set of patients, which demonstrated that the registration gave improved results as the non-rigid registration became increasingly fine-scale. The registration was evaluated quantitatively and qualitatively: by picking landmarks and determining correspondence errors; and also by inspecting the uniformity of the calculated difference images. The performance of the registration was variable across the varying anatomical structures that were assessed, but

it improved as the registration's control point spacing was increasingly fine-scale.

A problem with metal dental artefacts was identified. The artefacts affect the image by locally degrading the image quality in the transaxial plane. In the following chapter, multicomponent registration strategies for dealing with the problems that occur when metal dental artefacts are present on the CT images shall be explored.

Chapter 5

Registration in the Presence of Simulated CT Artefacts

5.1 Introduction

This chapter firstly describes the method used for multicomponent registration in this work, which differs from the form of the multivariate similarity measures defined in equations 3.19 and 3.20 and outlined earlier in section 3.7.5. The way in which the multicomponent registration is employed for the problem is then outlined and tested. A well established non-rigid registration software was modified for the purpose of facilitating multicomponent registration. The task here involved registering a pair of dual-component images - a combined PET/CT image that is considered to be a single entity within the formalism developed here.

The goal of this development was to improve the performance of the registration when there are dental artefacts present that corrupt the CT image. The idea was to utilise intensity information from the non-attenuation corrected PET image, in the region where an artefact is present.

To guide the development of a working registration method, a means of simulating artefacts on a pair of serial CT images was developed. This enabled the registration methods developed here to be tested against a

situation in which there were no artefacts present. The initial test involved registering the serial CT images in which the simulated artefacts were present in order to assess how much the non-rigid registration was adversely affected by their presence. Then several methods were developed for incorporating PET information to drive the registration process. The first method involved a multicomponent similarity measure that was a weighted sum of terms. The second method involved masking intensity information from the CT image in the vicinity of the artefact and using PET intensity information only within that region. The final method also involved acting in a spatially-varying manner across the image. This time, instead of acting on the domain of the intensity information within the image, the domain of action were the control points that define the transformation. Control points were substituted from the transformation derived from the PET-only registration into the CT-only registration in a region within the vicinity of the artefact.

For each method tested, the performance of the non-rigid registration was assessed in the region close to the artefact and also in regions of the image far from the artefact.

5.2 Implementing a Multicomponent Similarity Measure

5.2.1 Introduction

This section outlines the ideas that are used to enable a multicomponent similarity measure to be defined from a pair of multicomponent images. This includes a brief description of how a well-known registration code was adapted to facilitate this algorithm. How this relates in practice to our problem of registering a pair of serial PET/CT images is then discussed.

5.2.2 Refactoring a B-Splines Based FFD Registration Algorithm to Support Multicomponent Images

During this project, substantial effort was invested in modifying the vtkCmic classes to allow Regtool to enable a pair of multicomponent images to be registered. The mathematical details for this formalism are discussed in Appendix B, which describes the general case that was implemented in the vtkCmic classes, and then the registration application, regtoolMC was implemented using these objects to register pairs of dual-component images. This significant piece of work involved performing some extensive tests to ensure the code was operating correctly and that the similarity measures were correctly calculated. Aspects of this work included writing a file reader method to load the CT and PET images to define the dual component fixed and floating images. This included combining the fixed and floating CT and PET image into a single multicompoment image. The classes that were used to compute the similarity measure were updated to facilitate the evaluation of a multicomponent histogram to be used during registration. This then involved calculating a multicomponent similarity measure based on a weighted sum of components.

During this project, several tools were written by the author using the vtkCmic classes, to perform various tasks, such as MaskPETImage, MaskCTImage, DofCompare, JointHist, CropImage and finally RegToolMC, the re-coded registration tool that supports as input a pair of dual-valued images. These are specified in Appendix C.

5.2.3 Representing Combined PET/CT Images as a Single Dual-Component Image

For this problem, the combined head and neck PET/CT images are represented as a single dual component image. The goal was to develop a scheme to register a pair of serial PET/CT scans where there are artefacts present. Since the patient had a head-shell immobilisation device whilst being scanned, the intrinsic registration from the PET image to the CT image was extremely good, particularly within this anatomy in

which, unlike the thorax, the motions caused by breathing are minimal. The aim was to register images where there are dental artefacts on the CT image, so the PET image that has not been attenuation corrected was used as the additional source of information. Otherwise, the artefacts present within the attenuation artefact would adversely affect the PET image. The assumption is that, for corresponding points from the fixed to the floating PET image, there will be a comparable scaling factor due to a comparable lack of attenuation correction in each case.

To construct the single dual-valued image, the CT and the PET image were first histogram equalised and then resampled to the same voxel dimensions to have isotropic voxels of size $3mm$ as discussed in the previous chapter.

5.2.4 Registering Dual-Component Images

A multicomponent similarity measure outlined can be constructed from a dual-valued PET/CT image pair.

The following multicomponent similarity measure for a pair of these 2D images was used:

$$\begin{aligned} S(\vec{I}_T, \vec{I}_F) &= w_0 NMI(I_{T,C}, I_{F,C}) + w_1 NMI(I_{T,C}, I_{F,P}) + \\ &\quad w_2 NMI(I_{T,P}, I_{F,C}) + w_3 NMI(I_{T,P}, I_{F,P}) \end{aligned} \tag{5.1}$$

Where $NMI(I_T, I_F)$ is the NMI and the relative factors $w_0..w_3$ weight the different components of similarity measure. The subscripts T and F denote the target image and the floating image estimate and the subscripts C and P denote CT and PET components respectively.

The work in later sections in this chapter presumed that statistical dependence (as measured using normalised mutual information) for the voxel-based comparison for the CT-CT and the PET-PET will be far stronger than the PET-CT or the CT-PET components so the weights for these

cross-terms in the similarity measure were set to zero. The validity of this assumption was not tested, nevertheless it appears a reasonable one.

5.3 Simulated Dental Artefacts on CT Images

In order to explore the potential registration problems associated with CT dental artefacts, CT simulation was used used[1]. This was achieved using code written in MATLAB (Mathworks Inc.) that is listed in Appendix A. The code was used as in the following way.

Firstly, a set of 4 slices (extending $10mm$ in the superior-inferior direction, the slice thickness of $2.5mm$) were identified in which the simulated artefact was to be generated. On each slice, a set of small areas of interest were defined, in which the high intensity values were inserted to simulate the dental artefacts during reconstruction, as described below.

Then, a set of forward projections from this slice at multiple angles ranging from $0^o \leq \theta \leq 179^o$ were constructed using radon transforms. For each projection line at angle θ, the shadow cast along the projection line by the small regions of interest was determined. This geometric ray-tracing was done empirically by forward projecting the points inside each region onto the projection line and then noting the locus of points where high values appeared onto each line.

Finally, for each sector of each projection line created by the radon transform, a high numerical value was then inserted, whose magnitude was determined empirically. The image was then re-created from this modified set of projections by using an inverse radon transform for back-projection. This re-created image contains the simulated artefacts.

The results of this process on CT image data, when two regions of high-intensity values were created, are shown alongside the original CT image

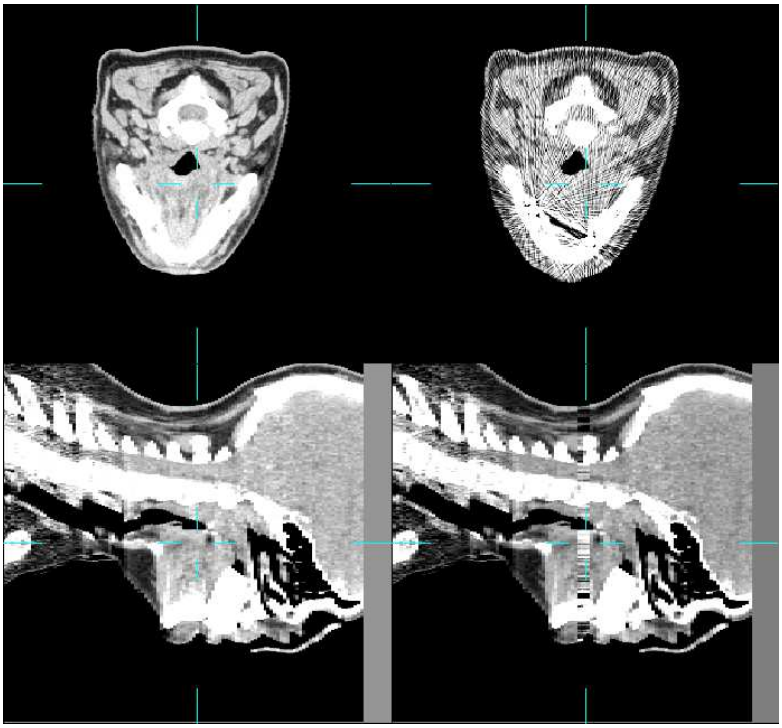

Figure 5.1: The original CT data next to the simulated dental artefact

data in figure 5.1.

5.4 Assessing How the Simulated Artefacts Affect Registration

In this section, the effect that the simulated artefacts have on the performance of the image registration of serial images will be demonstrated. In section 4.8, the phenomenon of dental artefacts was discussed, specifically that the artefacts are a strong feature in the transaxial plane. Given that the artefact is a predominant feature within the transaxial plane, the effect that a small rotation in the sagittal plane has on the registration performance when there are dental artefacts present shall be modelled. The registration performance under these conditions will then be compared to the case where there were no simulated artefacts present.

[1]I acknowledge advice and assistance from Dr Tryphon Lambrou in producing the simulated CT artefacts

5.4.1 Simulating Artefacts on a Pair of Serial Images

From the clinical trial outlined in section 4.1.1, image data were selected in which there were no dental fillings present. To study the effect of artefacts on registration, metal artefacts were simulated in the pre-treatment image, $I(\vec{r}, t_0)$ - the fixed image - using the method described in section 5.3. The floating image $I(\vec{r}, t_1)$ was rotated by a small angle θ^o in the sagittal plane. This angle θ^o shall be known as the *artefact angle* in subsequent discussions. Simulated artefacts were then created on this image in the same location. Both fixed and floating images were histogram equalised using the upper and lower thresholds of -400HU and 400HU respectively. Figure 5.2 shows the appearance of these simulated artefacts on the fixed and floating images, using a rotation angle of $+5^o$ prior to the creation of the artefact on the floating image. The figure shows the floating image following the rigid registration (a rotation of -5^o). The performance of the rigid registration is unaffected by the presence of the simulated artefact: see section 5.4.2 and the following discussion.

These operations on the floating image may be represented by the following mapping:

$$I(\vec{r}, t_1) \Longrightarrow I'(\vec{r}, t_1) = A\{I(\vec{T}_\theta(\vec{r}), t_1))\} \tag{5.2}$$

Where $I(\vec{T}_\theta(\vec{r}), t_1)$ represents the rotation of image $I(\vec{r}, t)$ through an angle of θ^o in the sagittal plane and $A\{I(\vec{r}, t)\}$ represents the process of creating the simulated artefact onto $I(\vec{r}, t)$. The operation applied to the fixed image was simply to create the artefact: $I(\vec{r}, t_0) \Longrightarrow I'(\vec{r}, t_0) = A\{I(\vec{r}, t_0)\}$.

5.4.2 Registering the Serial Images Containing Simulated Artefacts

Aim

The fixed and floating images onto which the artefacts had been created according to the description above were registered using a B-splines FFD

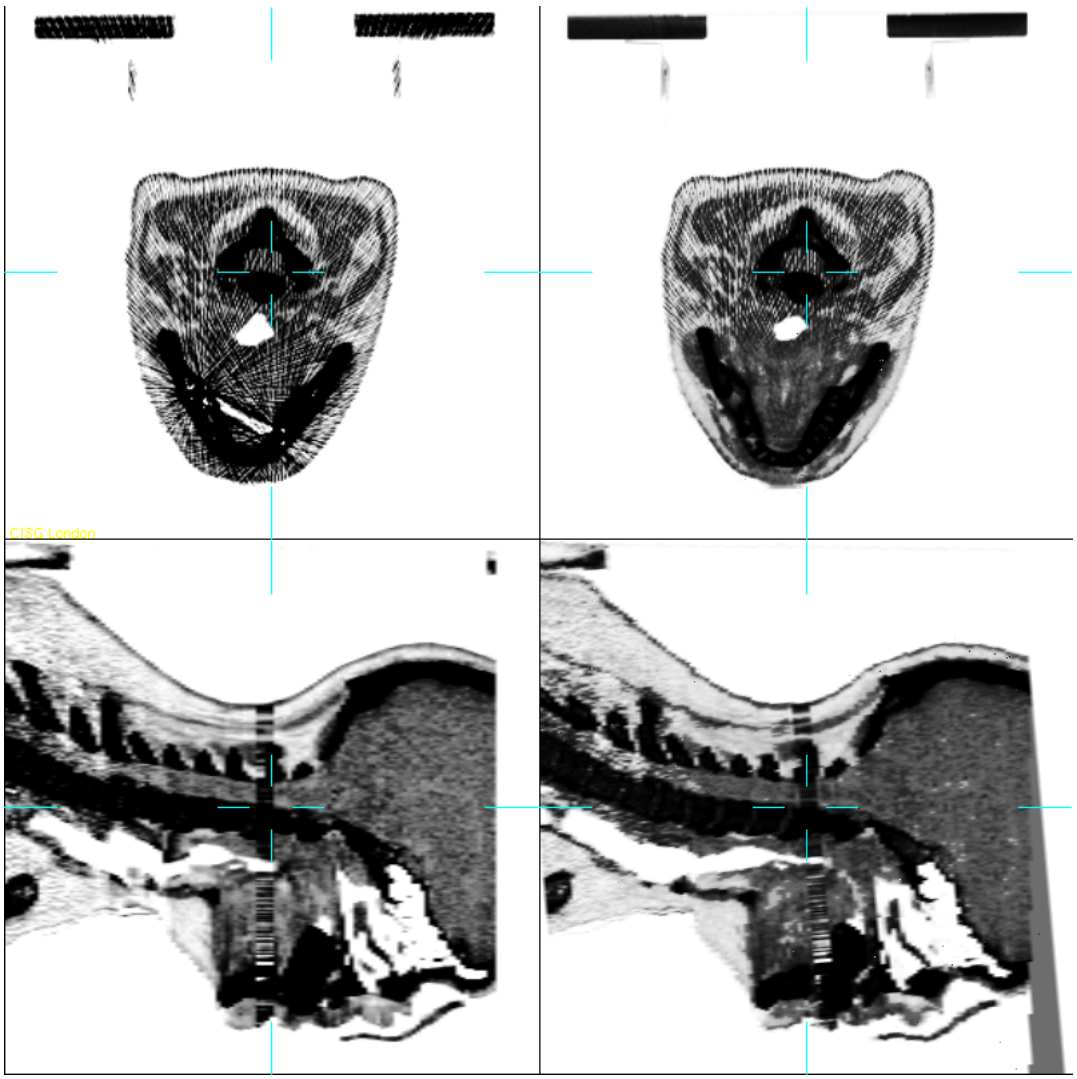

Figure 5.2: The simulated artefacts: fixed image on the left and floating image on the right after rigid registration

registration. This was to estimate how the registration is adversely affected by the presence of these simulated artefacts.

Method

The histogram equalised fixed and floating images, $I'(\vec{r}, t_0)$ and $I'(\vec{r}, t_1)$, which have had artefacts created on them, were registered using a rigid registration, followed by a multiscale non-rigid registration using control point spacings of 40mm, 20mm and finishing with a 10mm FFD control point spacing to give the following combined transformation:

$$\vec{T}(\vec{r}) = \vec{T}_{rigid,40,20,10}(\vec{r}) = \vec{T}_{10}(\vec{T}_{20}(\vec{T}_{40}(\vec{T}_{rigid}(\vec{r})))) \tag{5.3}$$

A single registration step was used in each case such that the final step size was 4mm at the 10mm control point spacing. The results of the registration were visualised using the VTKView viewer to assess how the simulated artefact affected the registration in the vicinity of the artefact. The deformation vectors were projected onto the slices to assist with understanding the behavior of the registration in this region.

A more quantitative assessment was performed as follows. A registration was performed using the original images $I(\vec{r}, t_0)$ and $I(\vec{r}, t_1)$ (i. e. without simulated artefacts) to yield a reference base-line deformation field, since, despite the patient being within an immobilisation device, there will still be slight non-rigid motions that will require local changes. Then, the deformation field captured from the registration of the images in which there were simulated artefacts could be compared to this base-line transformation. The Euclidean distance between these two deformation fields can then be evaluated at any desired number of points within the image.

Suppose $\vec{T}$ is the (non-artefact) deformation field that was the result of registering image $I(\vec{r}, t_1)$ to $I(\vec{r}, t_0)$ and also that $\vec{T}'$ is the (artefact) deformation field that was the result of registering $I'(\vec{r}, t_1)$ to $I'(\vec{r}, t_0)$. The Euclidean distance between these two deformation fields at any desired point defined by the position vector $\vec{r}$ can be calculated using:

$$\Delta r = \| \vec{T}(\vec{r}) - \vec{T}'(\vec{r}) \| \tag{5.4}$$

Where $\| \vec{v} \|$ denotes the scalar length (magnitude) of the vector $\vec{v}$.

A non-rigid registration, using a rigid transformation followed by a multi-level FFD transformation to give the transform $\vec{T}_{rigid,40,20,10}(\vec{r})$, was performed for two cases: when the artefact angle was 5^o and -3^o. At each control point spacing level, the performance of the registration was initially assessed by examining the registered images within the VTKView image viewer to help identify sub-volumes where the registration was most adversely affected. Having identified these sub-volumes, further quantitative analysis was performed using the Euclidean distances within these regions. For the 5^o artefact angle case, a sub-volume was defined that was more posteriorly located and centred around the spine. For the second case, when the artefact angle was -3^o, a sub-volume was defined that was more anteriorly located and centred around the jaw. For each of these volumes, the artefact fell approximately mid-way along the volume in the superior-inferior axis (z).

For each sub-volume, the following three sets of points were chosen to calculate these distances. Firstly, a set of points was chosen to be distributed within the volume in which the artefact was present. The second set of points was chosen within a volume located superiorly to the artefact. The third set of points was chosen within a volume located inferiorly. Within each set, the coordinates were chosen to be evenly distributed within the volume. Five points were chosen within each volume, and the mean, maximum (an indication of outlier behaviour) and minimum Euclidean distance Δr, calculated using equation 5.4. The sub-volumes used in the two cases for the calculations can be seen in figure 5.3.

Results

In all cases, the rigid registration managed to recover the inverse applied rotation exactly. For the first two non-rigid registration control point spacing levels that yielded transformations $\vec{T}_{rigid,40}(\vec{r})$ and $\vec{T}_{rigid,40,20}(\vec{r})$,

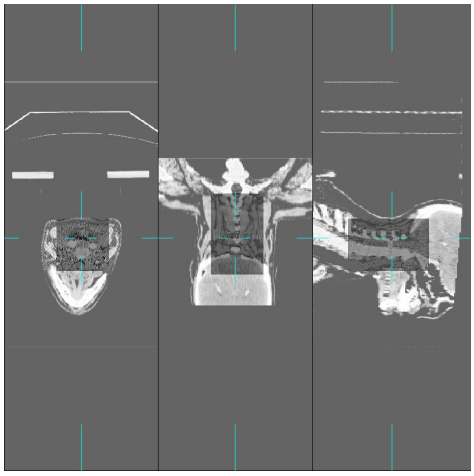

(a) The spine VOI

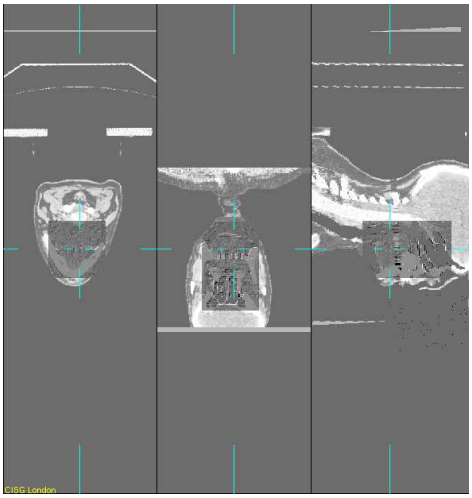

(b) The jaw VOI

Figure 5.3: The volumes of interest for assessment

there was no detectable adverse effect caused by the artefact when inspecting the deformation fields within the VTKViewer. The effects caused by the artefacts were observed at the fine scale registration, which resulted in the transformation $\vec{T}_{rigid,40,20,10}(\vec{r})$.

The results of the registration using a 10mm CPS FFD in the vicinity of the artefact (using a zoom factor) can be seen in figure 5.4. The transaxial and coronal CT slices are displayed with inverted colourscales in the VTKView viewer, with the cross-hairs indicating common geometrical reference. In figure 5.4a the vicinity of the artefact is shown after the rigid registration but prior to the non-rigid registration. The fixed image is on the left and the floating image on the right. The upper pane is a transaxial view showing the streaking effect caused by the simulated artefact. In figure 5.4b, the results of the registration are shown with the projection of the deformation vectors also shown visualised. The overall upward[2] deformation field can be seen in the vicinity of the artefact.

The average distances, along with the maximum and minimum distances within each region, are plotted for each case. These results are shown in figure 5.5.

Discussion and Conclusion

The rigid and the more coarse-scale non-rigid registrations were unaffected by the presence of the simulated artefacts. Since the majority of the image was unchanged by the presence of the artefact, there are many voxels unaffected by the artefact to drive the registration at the global level to ensure an accurate registration result. As the registration becomes increasingly local, the ratio of artefact voxels to non-artefact voxels becomes increasingly large in the vicinity of a control point's region. This effect becomes increasingly significant at a lower control point spacing, until at some level, the registration at a particular control point will become entirely governed by artefact voxels. It was found that this effect became dominant: the artefact caused problems when the control-point

[2]The upward direction being with respect to the page, since this is a coronal view. The local movement caused by the simulated artefact is in the superior direction.

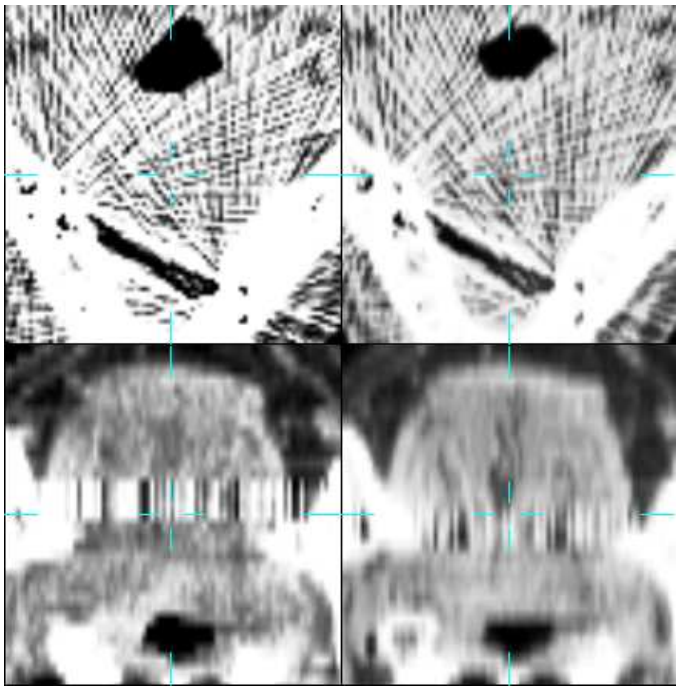

(a) Vicinity of artefact before non-rigid registration

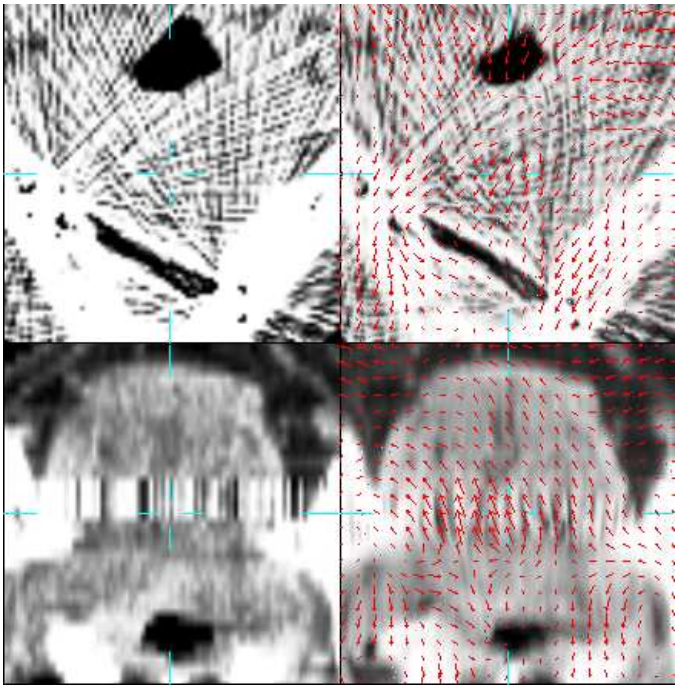

(b) Vicinity of artefact after non-rigid registration

Figure 5.4: Effect of artefact on non-rigid registration

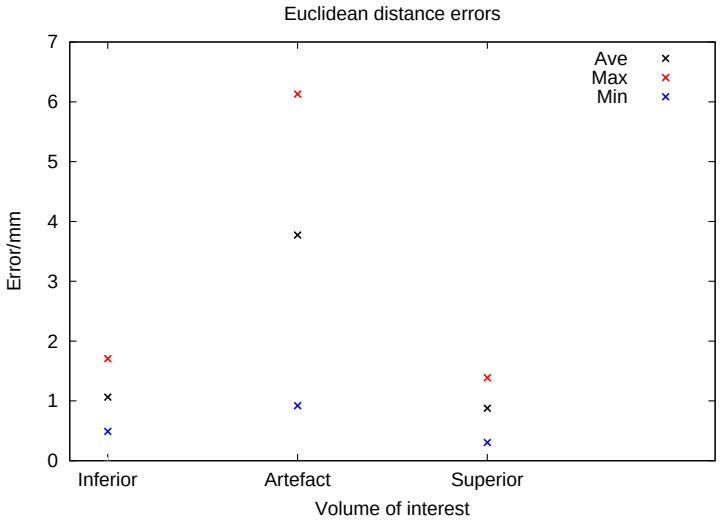

(a) The spine VOI

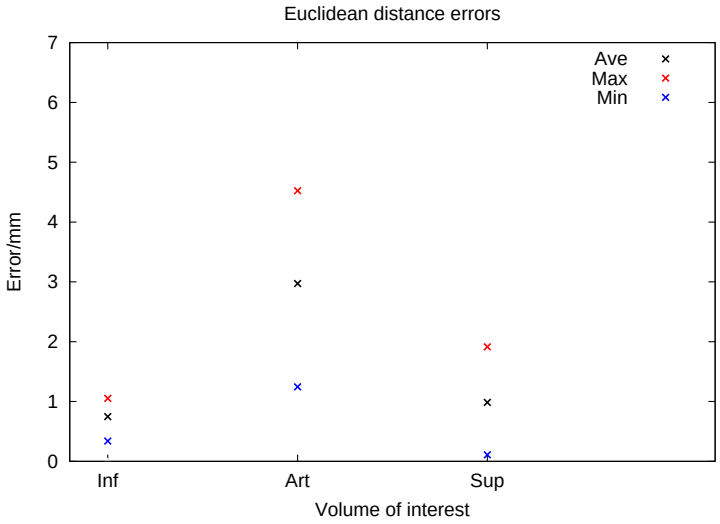

(b) The jaw VOI

Figure 5.5: Distance error to reference transform for different volumes of interest

spacing is 10mm. One presumes it would be even more pronounced at the 5mm scale, yet since this will increase the number of degrees of freedom of the problem by a factor of 8 it will also greatly increase the computational cost.

The adverse affect on the registration caused by the simulated artefact can be seen in the deformations displayed in figure 5.4b, showing a transaxial and coronal view, in which to some degree, the deformation field tends to align the two artefacts together. That the two artefacts do not more strongly register is because there are also competing physical structures within the images - the underlying anatomy - that act as a counterbalance. The reason behind using the first post-treatment image $I(\vec{r}, t_1)$ instead of the subsequent ones at t_2 or t_3, was to ensure there will be minimal underlying changes for the non-rigid registration to recover so that the effect caused by the artefact might be easier to detect. Any disturbances caused by the simulated artefact will be more easily detectable if the underlying deformation field is smaller.

This effect is localised to the vicinity of the artefact, as can be seen from the results shown in the two graphs in figure 5.5. Within the artefact region, the average registration difference in each case was significantly larger compared to the other regions. The minimum and maximum difference in the artefact regions were also larger.

Since B-splines FFD transformations have local support, it is quite possible for the transformation to correspond well within one particular region within the image, and for corruptions to a distant part of the image to have no adverse effect on that region during registration. Indeed, one might naively expect that the distance would be exactly zero in the regions that were far from the artefact, due to this property of local support. However, since the landscape of the similarity measure will be different in the presence of artefacts, the journey through the parameter space as the optimiser seeks the registration solution will in practice be a different one and the configuration of the control point spacing will be close, but not exactly in perfect agreement.

Having quantified how the registration is adversely affected in the presence of simulated artefacts, multicomponent registration strategies are explored with the aim of improving the registration within the vicinity of the artefact, whilst also seeking to ensure that the performance of the registration isn't compromised within the other regions.

5.5 Multicomponent Registration Strategies

5.5.1 Introduction

Strategies for minimising the adverse effects on the registration caused by the simulated artefacts will be investigated in this section. The aim was to get results that are as close as possible to the case in which there were no artefacts present, across the entire image. By combining the extra information contained in the PET scans, this additional intensity information can help where there is a degradation caused by the simulated artefacts on the CT images. In each case, the method was compared to the performance of the CT-only method described in section 5.4.2.

5.5.2 Registration Using a Multicomponent Similarity Measure

Aim

A multicomponent similarity measure was used to drive the registration, to investigate how this affects the performance of the registration where there are simulated artefacts. The goal is to investigate the performance of the registration when different weights are used in the similarity measure.

Method

Prior to the registration, the PET image was also histogram equalised to have the same intensity range in each case as the histogram equalised CT image. Both images were resampled to the same 3mm voxel spacing prior

to creating the single dual valued vector image. The registration was performed according to the description in section 5.4.2. The multicomponent similarity measure defined in equation 5.1 was used, and the cross-term weights w_1 and w_2 were set to zero. Using $w_0 = 1$, a range of values for w_3 were tested to increasingly include more of the information from the PET component of the similarity measure to drive the registration. Registrations were performed for the following values $w_3 = \{0.25, 0.5, 1, 2, 10\}$. The form of the similarity measure was therefore simply:

$$S(\vec{I}_T, \vec{I}_F) = NMI(I_{T,C}, I_{F,C}) + w_3 NMI(I_{T,P}, I_{F,P}) \tag{5.5}$$

The assessment was performed on the spine region, since this was the region that was most adversely affected by the simulated artefact. The regional errors in the vicinity of the artefact - as well as in distal regions - were assessed using the method described in section 5.4.2, by comparing each transformation to a reference transformation that was the result of a registration when there were no artefacts present. The average, minimum and maximum errors within each region were plotted for the range of weights w_3.

Results

Graphs of the results are shown in figure 5.6 for the three regions.

Compared to the CT-only method described in section 5.4.2, the average error within the artefact region had improved from $3.8mm$ to between $2.7mm$ when $w_3 = 2$ to $3.1mm$ when $w_3 = 0.25$. There was little difference for the minimum error for all values of $w_3$within the artefact region: all values were below $1mm$. The maximum error was similar to the CT-only method described in section 5.4.2, with little improvement apart from when $w_3 = 2$, where it had improved from $6.1mm$ to $4.8mm$. For the artefact region, $w_3 = 2$ is the value for which the artefact had the least effect on the registration.

Within the other image regions far from the artefact, the performance of the multicomponent registration is consistently worse than the CT-only

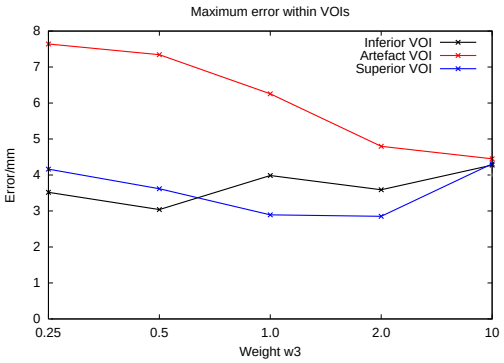

(a) Maximum error

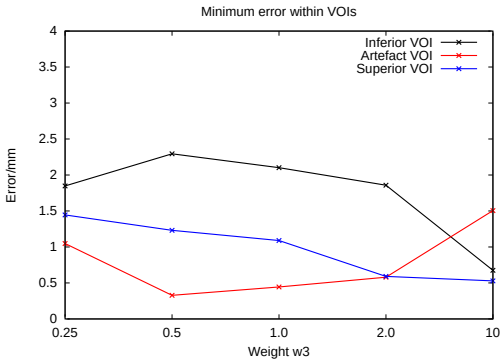

(b) Minimum error

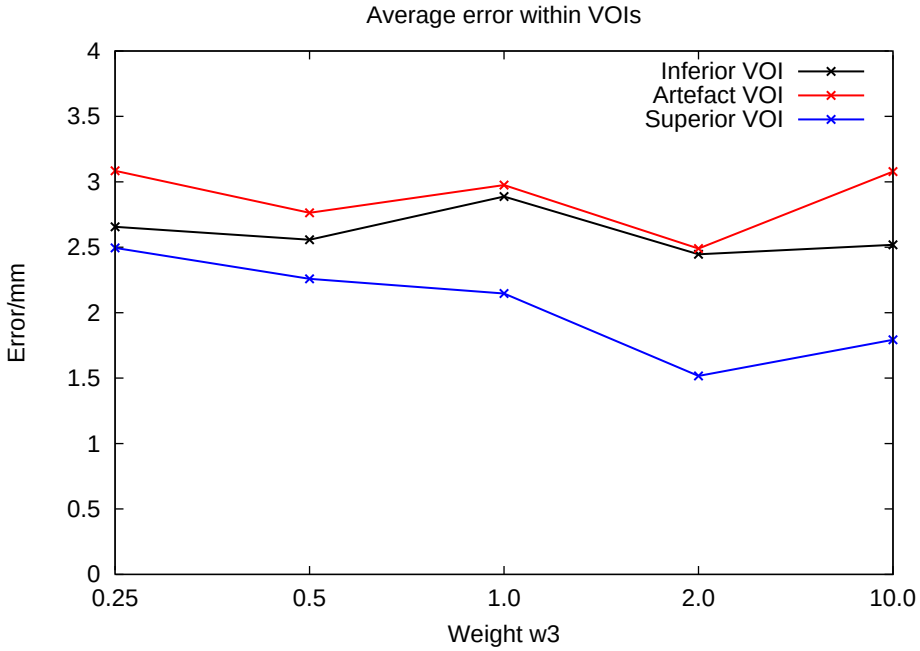

(c) Average errors

Figure 5.6: Errors for different weights within regions

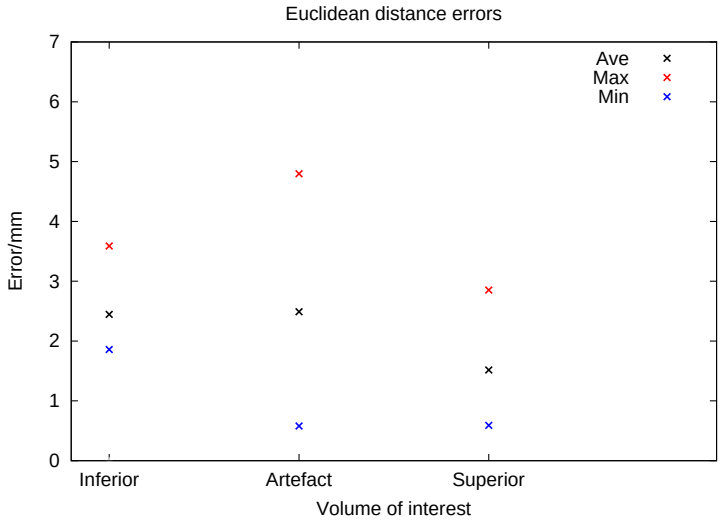

Figure 5.7: Distance errors for regions for a multicomponent NRR with $w_3 = 2$

case. In the CT-only case, the average error for both regions was $1.0mm$ with maximum and minimum errors of $1.7mm$ and $0.3mm$ for both regions combined. For the multicomponent case, the lowest average error for both regions combined was $2.0mm$, again when $w_3 = 2$, with the largest combined average being 2.6mm when $w_3 = 0.25$. The range of maximum error values for the different weights was from $3.6mm$ when $w_3 = 0.5$, to 4.3mm when $w_3 = 10$, and the range of minimum error values was from 0.5 when $w_3 = 10$ to 1.4 when $w_3 = 0.25$. The distance errors for the regions were plotted for the value of $w_3 = 2$, the value that yielded the best overall registration result, all things considered. This graph can be seen in figure 5.7.

Discussion and Conclusion

The spine volume case was studied in further detail, since it was found previously that the registration was more adversely affected by the simulated artifacts. The graphs in figure 5.6 demonstrate that, within the artefact region, the average error is marginally improved for all values of w_3.

The best result by the multicomponent similarity measure in the vicinity

of the artefact was achieved when $w_3 = 2$. Comparing the distance errors within the three regions plotted in figure 5.7 to the distance error when CT-only registration was performed - shown in figure 5.5 - shows a reduction of the average error in the artefact region by approximately $1mm$ at the expense of greater errors elsewhere.

An improvement on this compromise would be to selectively include information in a spatially varying manner into the registration process. The idea would be to have the registration process dominated by intensity information from the PET component within the vicinity of the artefact, and, conversely, to have the registration process dominated by information from the CT component in locations within the image that are not close to the artefact.

Methods for incorporating PET information into the registration process in a spatially varying manner shall now be proposed and tested with the goal of improving the multicomponent non-rigid registration.

5.5.3 Registering Masked Dual-Component Images

Aim

The previous multicomponent similarity measure using information from both PET and CT reduced the registration error in the vicinity of the artefact but increased the errors elsewhere. A method of image masking was therefore proposed, with the aim of maintaining this reduced registration error in the artefact region but without the increased errors for the rest of the image. The method involves a graduated masking of the CT intensities in the vicinity of the artefact and masking the PET intensities elsewhere so that only PET intensities are present in the vicinity of the artefact.

Method

The CT image and the PET images were masked using a graduated mask function $0 \leq B(z) \leq 1$ that varied in each case in the superior-inferior (z) axis. The graduated mask function was constructed to mask out the CT

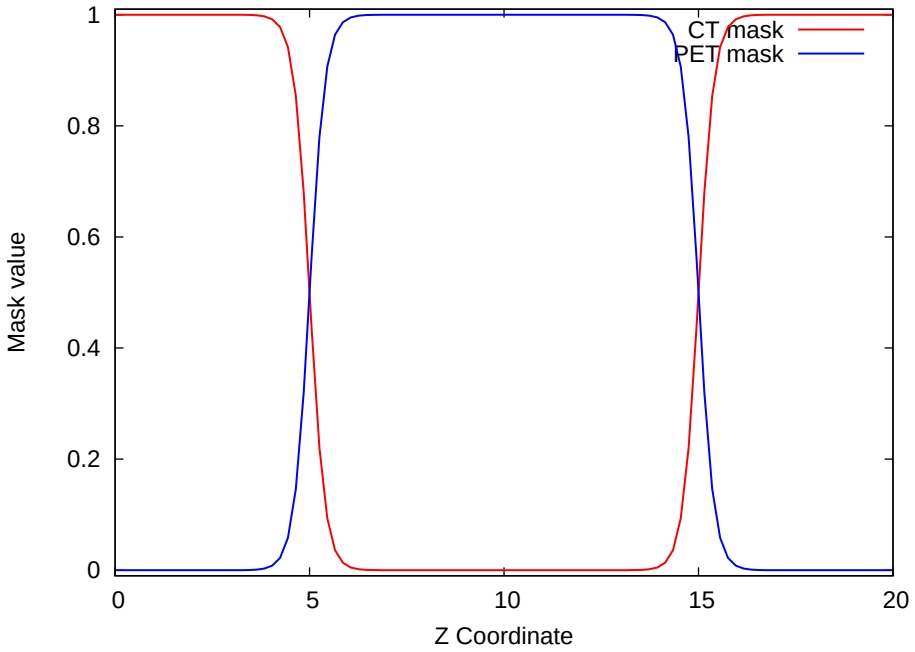

Figure 5.8: Graduated masks for PET and CT intensities

intensities in the z coordinate from z_{min} to z_{max} using a width factor of σ defined using a sum of sigmoid functions. The CT mask is defined by the following equation:

$$B_{CT}(z) = \frac{1}{1 + e^{\frac{Z - Z_{min}}{\sigma}}} + \frac{1}{1 + e^{\frac{Z_{max} - Z}{\sigma}}} \tag{5.6}$$

And the PET mask is conversely defined. It acts to mask intensities that are outside $z_{min}$to z_{max} ; so it is set to one in this region and is zero elsewhere, again subject to using a smooth roll-off factor σ:

$$B_{PET}(z) = \frac{1}{1 + e^{\frac{Z_{min} - Z}{\sigma}}} + \frac{1}{1 + e^{\frac{Z - Z_{max}}{\sigma}}} - 1 \tag{5.7}$$

Figure 5.8 illustrates the graduated mask functions defined in equations 5.6 and 5.7 with $\sigma = 0.2$, $z_{min} = 5$ and $z_{max} = 15$.

To ensure the artefact region was fully masked out from the CT images, the masking was performed on the CT and PET images by setting $z_{min} = 22$, $z_{max} = 35$ and a width of $\sigma = 1.5$ slices to ensure a graduated intensity

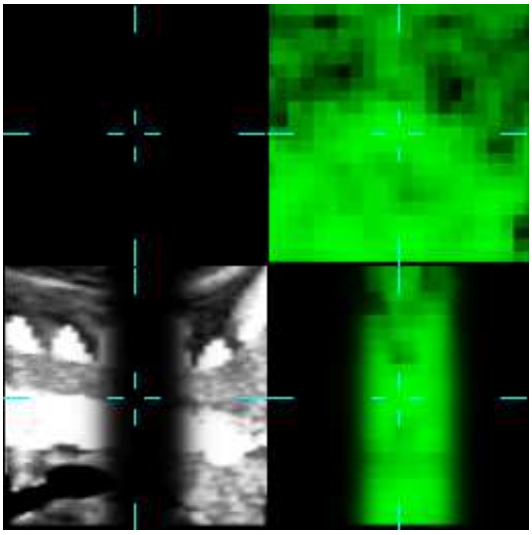

Figure 5.9: Masking near artefact

roll-off. Figure 5.9 shows the CT mask applied to the CT intensities and the PET mask applied to the PET intensities. The software tools that were written for this purpose using the vtkCmic classes, MaskPETImage and MaskCTImage are specified in Appendix C.

A non-rigid registration was performed in the same way as described in section 5.2.4. It was previously concluded that a multicomponent similarity measure with $w_3 = 2$ yielded the best results, so this was used for the registration.

The performance of the registration was assessed using the VTKView image viewer. The distance errors were determined from the transformation to the reference transformation within the three regions as described earlier.

Results

In figure 5.10, the orthogonal views of the registered CT image pair are displayed using the VTKView viewer. For comparison, the result from the reference transformation (gained from registering the CT images with no simulated artifacts) is shown, along with the results from the multicom-

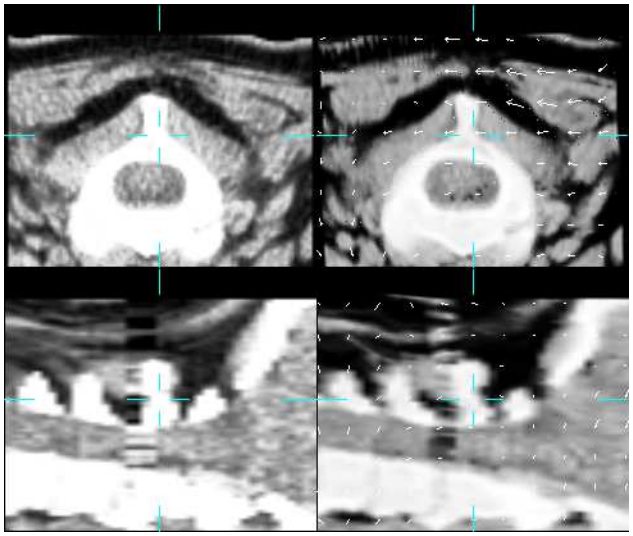

(a) The reference transform

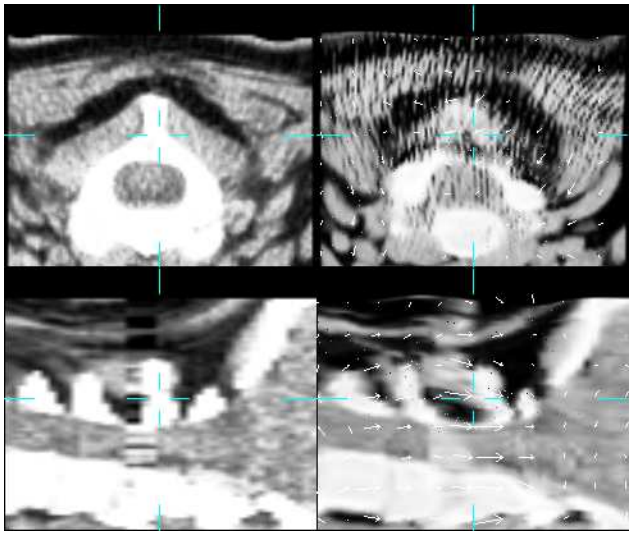

(b) Transform from the masked multicomponent registration

Figure 5.10: Non-rigid registration results

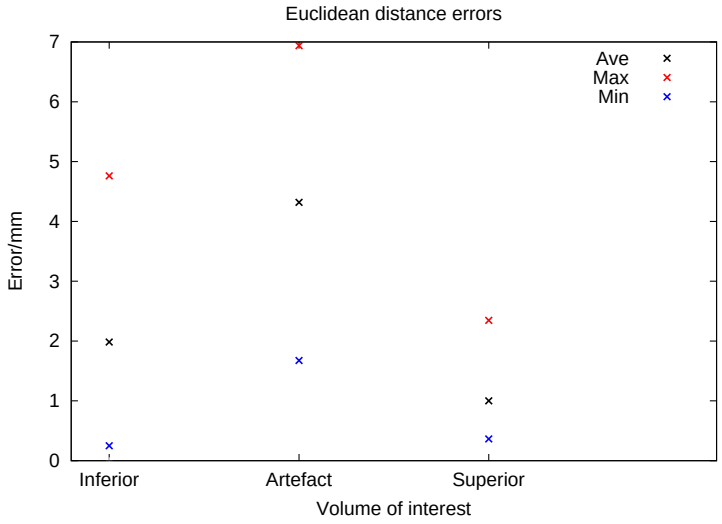

Figure 5.11: Multicomponent masked registration errors

ponent registration. The errors are shown in figure 5.11.

Discussion and Conclusion

Introducing sharp edges into the images will clearly be problematic for a FFD non-rigid registration technique, since they are strong features. So instead of using a sharp mask defined using a combination of shifted Heavyside step functions (or, equivalently, letting $\sigma \rightarrow 0$ in equations 5.6 and 5.7), a graduated mask was employed with the aim of not generating strong features within the image. It was hoped that this would enable a gradual masking out of CT intensity values within the artefact region in the z axis, and a gradual masking in of the PET intensities to enable a multicomponent registration to be performed. Comparing the results plotted in figure 5.11 to those when CT-only registration was performed using the images containing the simulated artifacts, in figure 5.5, it is clear that not only is the registration error slightly greater within the artefact region using this masking technique, but that it is significantly worse for the rest of the images, as shown by the outlying large values for the maximum errors in the superior and inferior regions. The modification of the intensity values by masking, despite using a graduated mask, has introduced a significant feature within the images, which is a likely

reason for the poor performance of this technique. An alternative way of combining PET and CT information in a spatially varying manner shall be considered in the next section.

5.5.4 Control-point Displacement Substitution from PET to CT Transform

Aim

A method for registering the images by selectively combining the transformation yielded from the CT-CT registration with that from the PET-PET in a spatially varying manner will be investigated. For each registration at each control point spacing level, the registration will be performed independently and then the transformation will be combined and then input to the next level of transformation.

Method

At each control point spacing level of dimension s in mm, the registration is performed independently for the CT-CT I_c components and for the I_P components of the floating and target images, to yield transformations $\vec{T}_{CT-CT,s}$ and $\vec{T}_{PET-PET,s}$ respectively. The results of the transformation from each independent registration will then be combined by blending the transformation into a single transformation, $\vec{T}_{CT-PET,s}$. This is done by identifying $z_{min} \leq z \leq z_{max}$ that defines a slab of image in the superior-inferior dimension within which control point displacements will be substituted from transformation $\vec{T}_{PET-PET}$ into $\vec{T}_{CT-CT}$. Figure 5.12 illustrates how $\vec{T}_{CT-CT,s}$ and $\vec{T}_{PET-PET,s}$ are blended: the z coordinate runs along the page, showing how the green-coloured control points will have their displacements substituted in from the PET transformation within the artefact zone. The values z_{min} and z_{max} were identified to encompass the entirety of the artefact region.

The rigid registration is performed first using the CT components, to give the rigid transformation $\vec{T}_{rigid}(\underline{r})$. The blended transformation at the control point spacing s in mm forms the input transformation to each of the

next two independent transformations at the lower scale. This can be illustrated as follows for a sequence of non-rigid registrations that are performed for control point spacings of $40mm$, $20mm$, and $10mm$:

$$\vec{T}_{CT-PET,40,20,10}(\vec{r}) = \vec{T}_{CT-PET,10}(\vec{T}_{CT-PET,20}(\vec{T}_{CT-PET,40}(\vec{T}_{rigid}(\vec{r})))) \quad (5.8)$$

As in previous sections, the registration was assessed by comparing the resultant transformation to the reference transformation and computing the errors in the three regions. The maximum, minimum and average errors were plotted for each region.

Results

When inspecting the results of the image registration within the VTKView image viewer, there were no discernible parts of the image that appeared to be adversely affected by the artefact. The errors within the regions of interest may be seen in the graph in figure 5.13.

Discussion and Conclusion

This registration technique of blending the transformation by swapping in the control point displacements from the PET-derived transformation within the vicinity of the artefact has proved to be a useful method. The average error within the majority of the image was within $0.2mm$ of the error that was found from the CT-only registration. The average error within the artefact region was improved from $3.8mm$ to $2.7mm$, which is a significant improvement. Also, it is less prone to outlier errors as the maximum error is less, at $4mm$ compared to $6mm$. This is the same average error as was obtained by the multicomponent registration when $w_3 = 2$.

There are several permutations for constructing schemes for this control point substitution approach to multicomponent registration. Since it was previously deduced that the artefact only had an adverse effect at the lowest scale, then the following improvement can remove some of the computation steps and is therefore more computationally efficient:

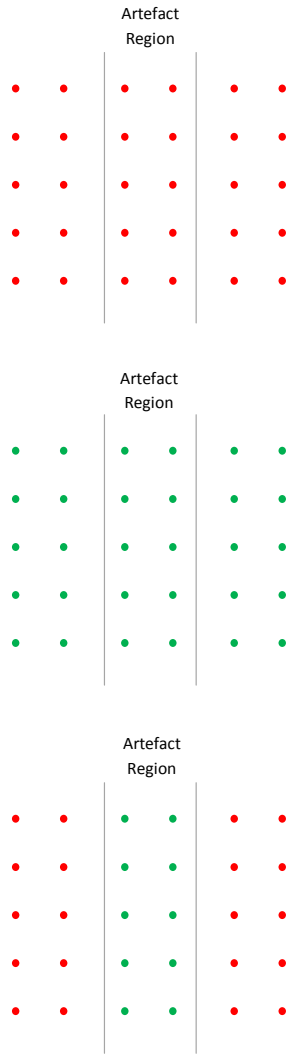

Figure 5.12: The red control points are from the CT transformation and the green control points are from the PET transformation. The blended transformation is shown in the control-point grid below

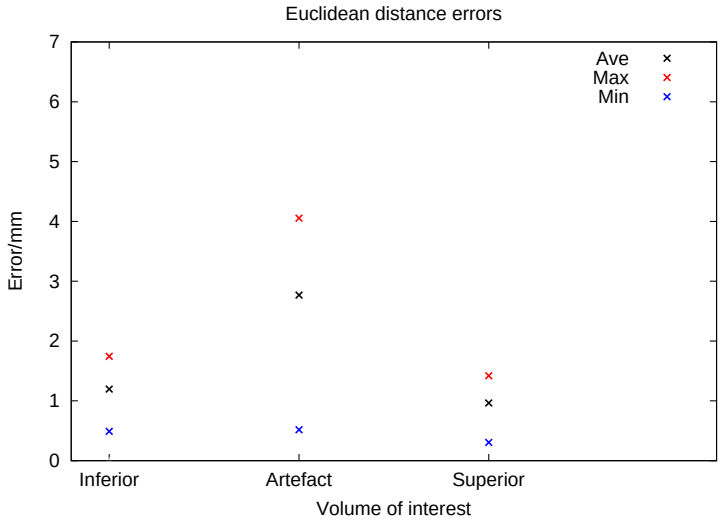

Figure 5.13: Errors within regions when using a blended transformation method

$$\vec{T}_{CT-PET,40,20,10}(\vec{r}) = \vec{T}_{CT-PET,10}(\vec{T}_{CT-CT,20}(\vec{T}_{CT-CT,40}(\vec{T}_{rigid}(\vec{r})))) \tag{5.9}$$

This method of multicomponent registration that involves the blended transformations from the PET and CT-based transformations during a multi-scale registration is a significant improvement when there are simulated artifacts present. This technique of control point displacement substitution registration shall be known as the CPDS registration technique.

5.6 Discussion

In summary, the work in this chapter has modelled the effect that an artefact causes on the performance of the registration. Various registration approaches were tested against a reference registration in which there were no artifacts present, using the information from the CT images to perform the registration. This CT-only registration approach was then tested on the images that contained the simulated artifacts, and it was found that the artefact had no effect on the rigid registration. The simulated artifacts gave rise to registration errors when doing a fine-scale

FFD registration using a control point spacing of $10mm$.

Given that the registration when using information from the CT images alone is affected by the presence of artifacts, a multicomponent registration method was developed. This involved implementing this within a well-known registration code to support multi-valued intensity vector images, from which a similarity measure may be calculated that has many terms relating to the possible permutations of the intensity components. This substantial piece of work to implement within the vtkCmic software was then put to use in this section of work, whereby the images were considered to be dual-valued at each spatial point.

The multicomponent registration was then tested to see how it performed in the case where there were simulated artifacts present in the CT component. A similarity measure consisting of two terms was tested in two ways: firstly, by including all intensity values from both the CT and PET intensities into the image; and secondly, by spatially masking the intensities in such a way as to selectively include PET intensity information in the vicinity of the artefact. Whilst this second method attempted to solve a problem discovered in the first method, it actually made the results worse. It was considered that this was because the masking process in itself was introducing a large spatial feature into the image that was causing the registration to erroneously attempt to register to it. Ironically, it was perhaps a larger artefact than the original simulated artefact that this method was seeking to mitigate. The final method involved control point displacement substitution from the PET transformation into the CT transformation in the region of the image affected by the artefact. This managed to maintain a registration performance within the rest of the image that was comparable to the CT-only registration whilst lowering the registration error within the artefact region. Table 5.1 summarises the registration errors for each of the methods. For each method, the minimum, the average (in bold) and the maximum error in mm to the reference transformation are listed for each region.

So to conclude, the ranking of the methods tested in this chapter are,

VOI	CT-only	MCSM	MCSM mask	Blended CPDS
Superior	0.5, **1.1**, 1.7	1.9, **2.4**, 3.6	0.2, **2.0**, 4.8	0.5, **1.2**, 1.7
Artefact	0.9, **3.8**, 6.1	0.6, **2.7**, 4.8	1.7, **4.3**, 6.9	0.5, **2.7**, 4.0
Inferior	0.3, **0.9**, 1.4	0.6, **1.5**, 2.9	0.4, **1.0**, 2.3	0.3, **1.0**, 1.4

Table 5.1: Summary of registration performance

from worst to best:

1. The multicomponent similarity measure using the graduated masking method.
2. A CT-only registration.
3. A multicomponent similarity measure with $w_3 = 2$.
4. The control point substitution method.

Now that a registration method has been determined that can reduce the errors caused by the simulated artifacts, this shall be tested and evaluated on actual clinical data in which there are dental artifacts present. In the next chapter, the registration methods shall be tested on clinical data.

Chapter 6

Registration in the Presence of CT Artefacts: Clinical Assessment

6.1 Introduction

The registration techniques developed in the previous chapters were tested on real patient data sets in which there are dental artefacts present. This involved introducing another technique for assessing the registration performance, the Jaccard coefficient. This was used to help assess how the registration performs when there are artefacts present. Two sets of patient data were compared: one set containing dental artefacts and the second set without artefacts

6.2 Assessment Using Anatomical Landmarks

6.2.1 Background

In the previous chapters, registration performance was assessed using several methods: visual scoring, anatomical landmarks and finally, a

"point-cloud" technique that involved comparing a registration to a reference. The reference transformation was the one obtained when there were no simulated dental artefacts present, so any local deviations from this in the vicinity of the simulated artefact were assumed to be caused by this artefact. This seems a reasonable assumption. However, when assessing the registration techniques for images in which there are dental artefacts present, it is not possible to use the point-cloud method for comparing multiple points within the artefact region, since there is no reference transform for comparison.

The performance of the registration was compared for three patients with artefacts present and three patients with no artefacts present. The initial assessment was based around the landmark picking methods. The landmarks used were the same as those used in section 4.5.7 and shown in figure 4.9. The only landmark point that resided in the vicinity of the artefact was the point located on the posterior skull region. The results of the landmark picking error demonstrated that the performance of the registration was slightly worse for the skull notch point, which was located in the vicinity of the artefact. When using the two methods for registration improvement - the MCSM and the CPDS registration methods, the latter technique showed an improvement.

6.2.2 Assessing the Performance of the Registration

Aim

To assess how the registration performance is affected by the presence of dental artefacts by using anatomical landmarks.

Method

From the data acquired as part of the clinical trial described in section 4.1.1, three patients were selected with image data that had no dental artefacts. A second group of three patients were chosen that had dental artefacts present. For each patient, the pre-treatment PET/CT image was chosen to be the reference image and the final PET/CT image in the series was chosen to be the floating image.

The image pairs from the first group were registered using the coarse-to-fine FFD scheme with the optimal parameters as described in section 4.7 and using a similarity measure derived from information only from the CT images. For the group of patients in which there were dental artefacts present, the images were registered in two ways: firstly, using the method that employs a similarity measure derived from information only from the CT images; and secondly, using the CPDS method described in section 5.5.4 and using the PET information at the finest scale, according to the transformation described in equation 5.9. The landmarks chosen were the same as those used in section 4.5.7 and illustrated in figure 4.9.

The landmark picking errors were determined for the new set of patients containing the artefacts.

Results

The landmark picking errors for the new set of patients containing the artefacts were found to be comparable to the values determined previously, however, a difference was noted for the skull landmark, which was determined to be 0.99mm compared to 0.49mm previously (see section 4.7 for comparison).

The average landmark errors at each level of registration were calculated and plotted for the three patients in which there were no artefacts, and registered using the registration scheme based on using information from the CT images alone. The graph may be seen in figure 6.1. The landmark errors for the three patients in which there were artefacts present are shown in the graphs in figure 6.2, which shows the results for the errors based on information from CT alone, and on the results after performing a registration using the CPDS method.

Conclusion

The performance of the two groups of three patients improved with increasingly fine control point spacing, as one would expect. The results shown in the graph in figure 6.2 show only a very slight difference in registration performance at the $10mm$ CPS level for the skull landmark,

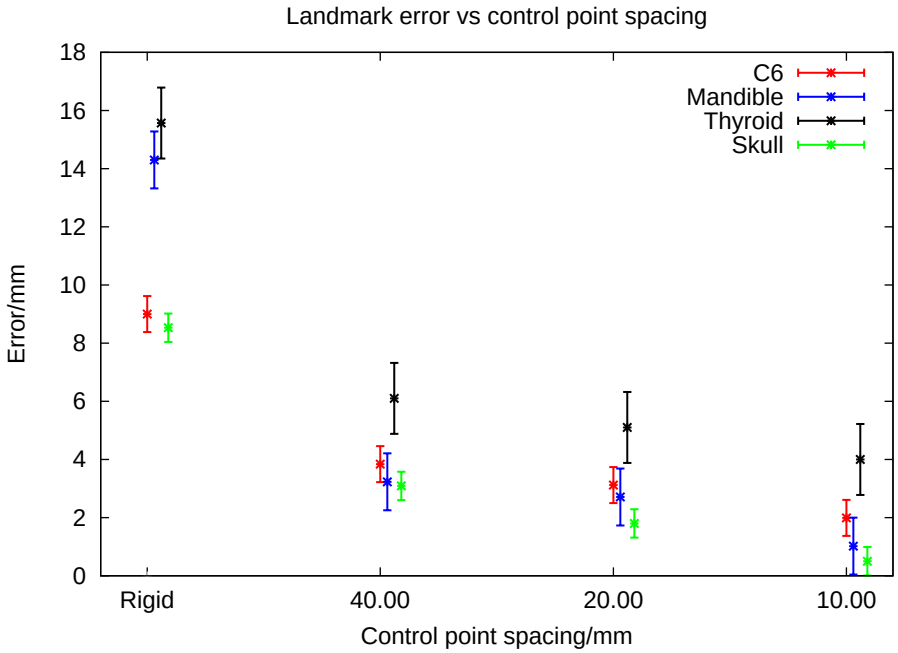

Figure 6.1: Landmark error with decreasing control point spacing for images with no dental artefacts

which was located at the edge of the artefact region. An improvement in the registration of the thyroid landmark was also measured, although this was in a region of anatomy far from the region adversely affected by the artefacts.

The problem with using anatomical landmark points was that in the location of the image that is adversely affected by the artefact, there were no readily available landmark points in this region of anatomy - the closest landmark point was the skull landmark, which was located at the superior edge of the region. This is a simple fact of anatomy rather than being due to the imaging process.

It is also important to determine how well the following additional soft tissue structures are registered: the many muscles, lymph nodes, and adipose tissues. These structures do not have an easily identifiable point that may be reliably identified for the purpose of registration assessment, therefore, an alternative means of assessing the registration may be used

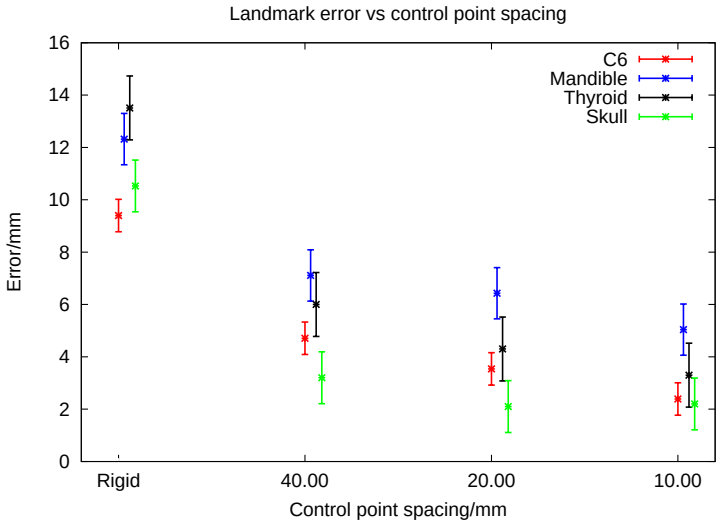

(a) Errors when registered using CT-CT based similarity measure

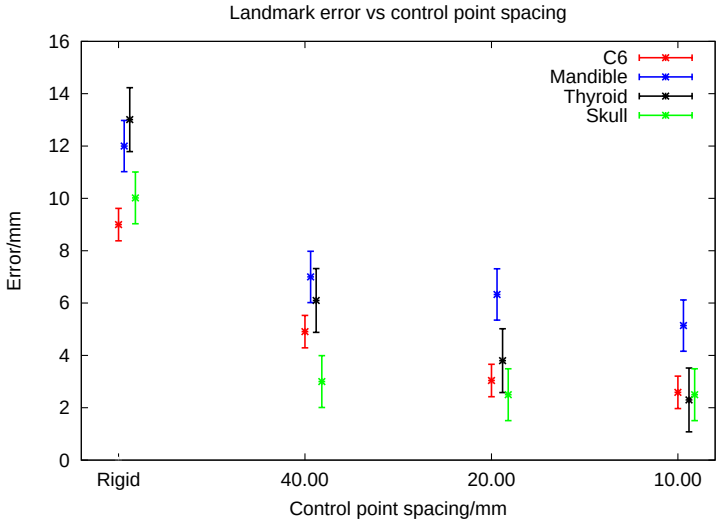

(b) Errors when registered using the CPDS technique

Figure 6.2: Average landmark error with decreasing control point spacing for images with dental artefact present

based on on a relative volume metric to determine how well a corresponding segmented structure is registered from the source to the target image. In addition to these soft tissues, the superior vertebra do not have an easily identifiable dorsal tip, as do the lower vertebra, so they too could benefit from a measure that is volume-based.

No firm conclusions may be drawn from this initial assessment of the registration, other than the error in identifying the landmark on the skull was greater, since it was harder to visually identify when the images contained the streak artefacts.

6.3 Assessment Using the Jaccard Index

6.3.1 Background

Overlap metrics based on quantifying how well two corresponding segmented volumes are aligned on a registered floating and fixed image have been used extensively in the field of image registration assessment. They enable a figure of merit to compare registration methods. A whole set of segmented regions [Crum, Camara and Hill, 2006] on the images may be used to calculate a single figure of merit. For a pair of corresponding segmented volumes, $V_{reference}$and $V_{floating}$ that represent some kind of anatomical structure within the reference and floating images, the Jaccard Index, O, can quantify how well these two are aligned:

$$O = \frac{N(V_{reference} \cap V_{floating})}{N(V_{reference} \cup V_{floating})} \tag{6.1}$$

where N denotes the total number of voxels in the set. This was implemented in practice in code, by setting all voxels inside the set $V_{reference} = 1$ and also $V_{floating} = 1$ and all voxels not contained within these segmentations were set to zero. The overlap was then calculated by looping over all voxels of index $i = 1, \ldots, N_{voxels}$ where $N_{voxels}$is the number of voxels in the images (which both have the same dimensions) and using the following logic to determine whether a voxel is contained within the intersection of the segmentations, the union, or both:

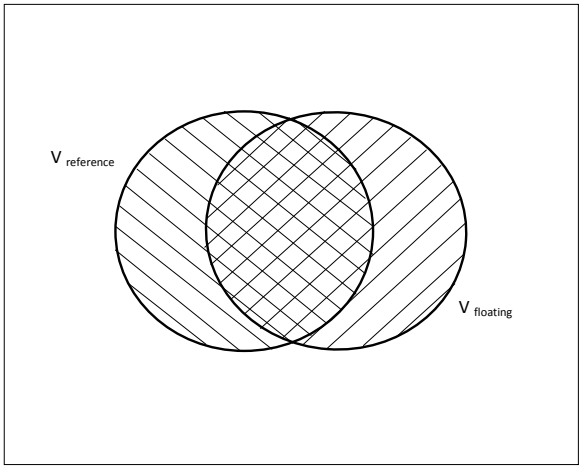

Figure 6.3: Overlap coefficient for a pair of corresponding segmented structures on the fixed and floating image, defined as the ratio of the intersection volume to the union volume

$$O = \frac{\sum_i I_{reference,i}.AND.I_{floating,i}}{\sum_i I_{reference,i}.OR.I_{floating,i}} \tag{6.2}$$

Where $.OR.$ and $.AND.$ are the logical operators that act on the binary values that are inside the fixed and floating images $I_{reference}$and $I_{floating}$.

Figure 6.3 illustrates how these segmented volumes are used to define the Jaccard coefficient. This coefficient will have the range $0 \leq O \leq 1$ with the value 1 when the two segmented regions are identically aligned and 0 when no part of either volume intersects. As a fractional volume estimate, it will depend on both the size and shape of the two volumes. It will tend to vary more rapidly for shapes with a higher tortuosity as a function of mis-registration [Crum, Camara and Hill, 2006].

6.3.2 Clinical Assessment

Aim

To assess how the registration performance is affected by the presence of dental artefacts using the overlap metric for corresponding segmented regions on the fixed and registered floating images.

Method

The registrations of the two groups of patients described in the preceding section were assessed using segmented regions on the fixed and floating images and calculating the overlap index for several tissues of interest. The segmentation was performed in a manner that could realise tissue types based on the Hounsfield Unit. A thresholding technique was performed to segment the tissues of interest within cuboid sub-volumes, chosen empirically to include the relevant structures of interest.

For the images that contained artefacts, the segmented structures were chosen within the vicinity of the artefact. A duplicate set of structures was also chosen for each tissue type in an inferior region, some distance from the artefact.

For the images that did not contain artefacts, the regions defined were also chosen in a similar anatomical region to where there were artefacts on the other images, to compare like with like.

Three different tissue types were chosen that could be readily distinguished using Hounsfield Units by performing a thresholding segmentation on the original CT images, prior to the histogram equalisation preprocessing step that is performed as part of the registration. The tissue types chosen were adipose tissue, muscles and bone.

Adipose tissue is present in the head and neck anatomy between other anatomical structures and also as a subcutaneous layer. Segmentation of the subcutaneous fat and other adipose tissue was performed using a Hounsfield threshold range $-150 < H < -10$. Voxels with intensities

lying within this range, within a cuboid sub-volume, were labelled as being a member of the segmented volume. This resulted in a segmentation within the sub-volume that can be seen in figure 6.4. This figure shows the segmentations performed on the fixed image in the region of the artefact and in an inferior region in which there were no artefacts.

Muscles are also present within this region of anatomy, especially on the posterior side of the neck, which is where the sub-regions for the segmentations were chosen. The Hounsfield threshold range of $5 < H < 80$ was used to perform the segmentations, which can be seen in figure 6.5.

Bone tissue in the form of vertebra was the final structure chosen to segment. To segment this structure, voxels with a Hounsfield unit $H > 200$ were labelled within the sub-volume, as shown in figure 6.6.

Once the segmentations were performed on the registered fixed and floating image for each registration level, the overlap was calculated for the corresponding regions.

Results

For the set of registered images in which there were no dental artefacts present, the performance of the registration was plotted. The overlap measures for each segmented region can be seen as a function of the registration's control point spacing, in figure 6.7.

The performance of the registration for the set of images in which there were dental artefacts present was also plotted. For the segmented regions, the overlap as a function of control point spacing may be seen in figure 6.8. The overlap coefficients were calculated for corresponding regions within a segmented part of the spine, for adipose tissue and for the muscle tissue. For each registration technique used, the graphs were plotted for the overlaps calculated from structures within the artefact region and in an inferior region of anatomy. The results are shown in table 6.1 for $10mm$ control point spacing, which is the control point spacing at which the differences between the two registration techniques will become apparent.

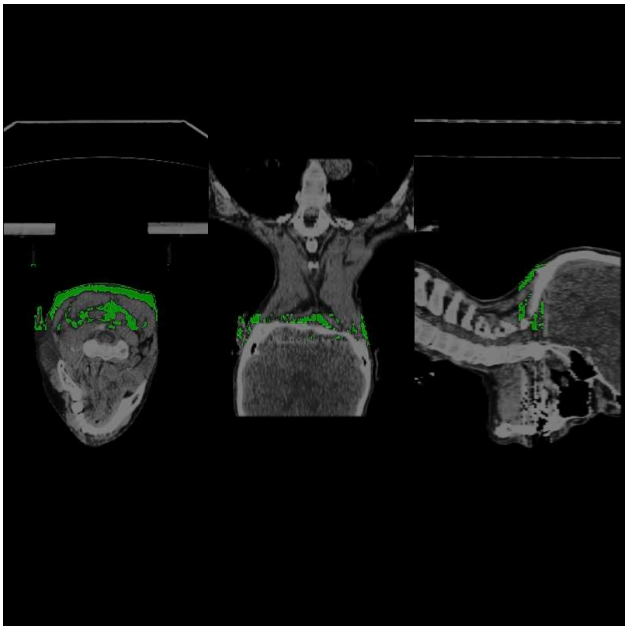

(a) Artefact region segmentation

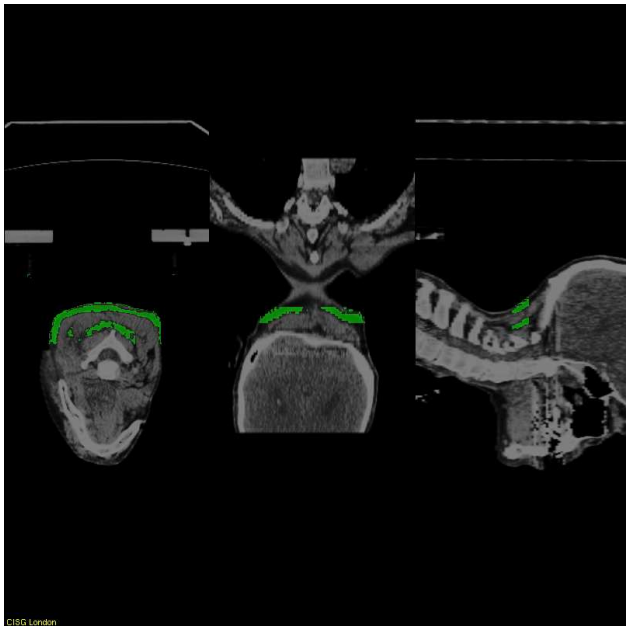

(b) Inferior region segmentation

Figure 6.4: Segmentations used for adipose tissue

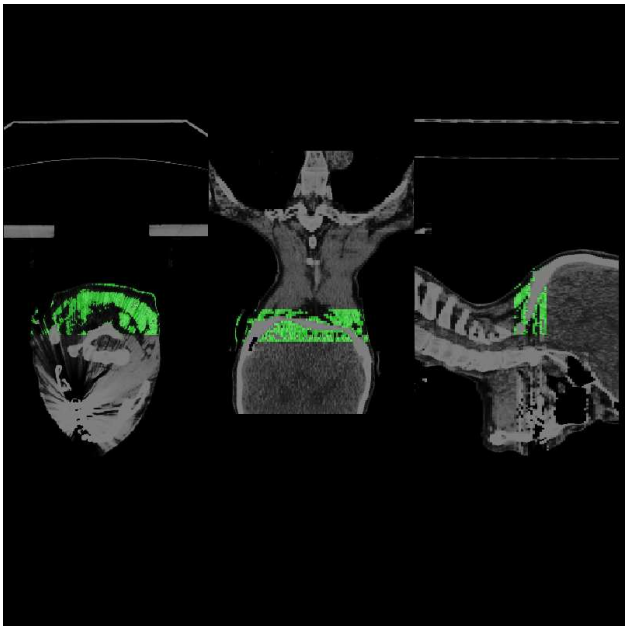

(a) Artefact region segmentation

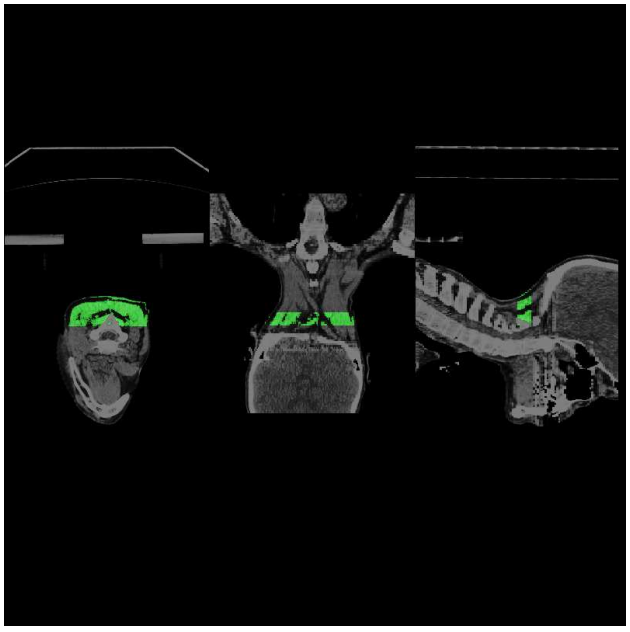

(b) Inferior region segmentation

Figure 6.5: Segmentations used for muscle tissue

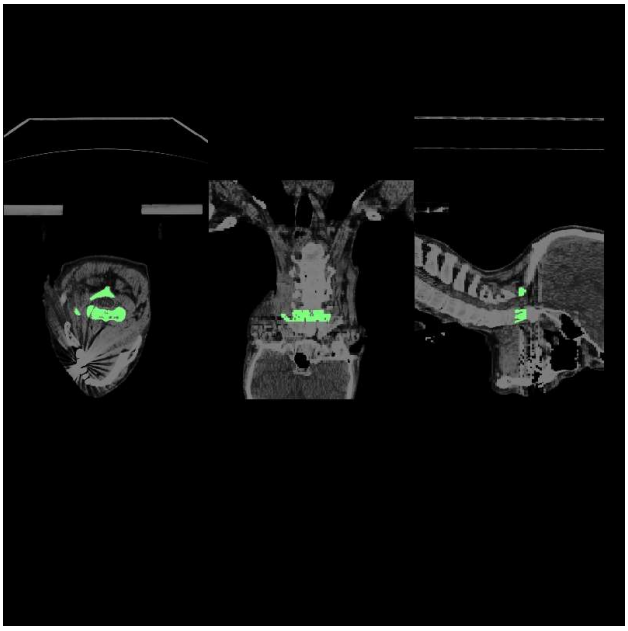

(a) Artefact region segmentation

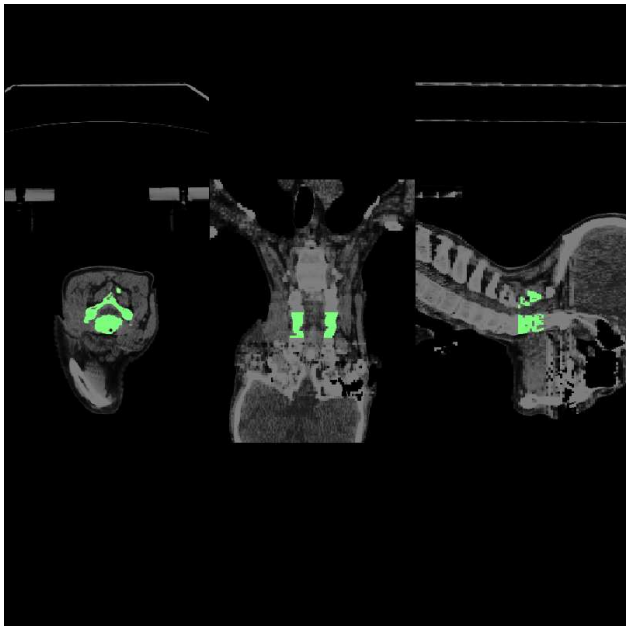

(b) Inferior region segmentation

Figure 6.6: Segmentations used for bone tissue

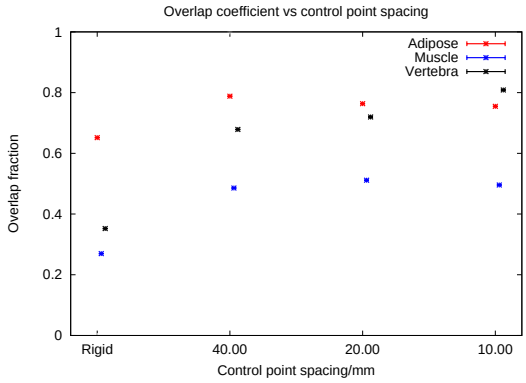

Figure 6.7: Registration performance for images with no dental artefacts

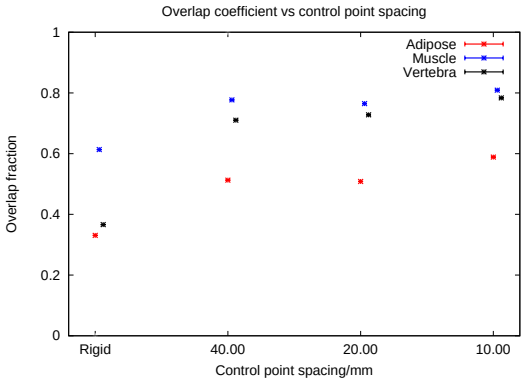

(a) Inferior region CT-CT registration

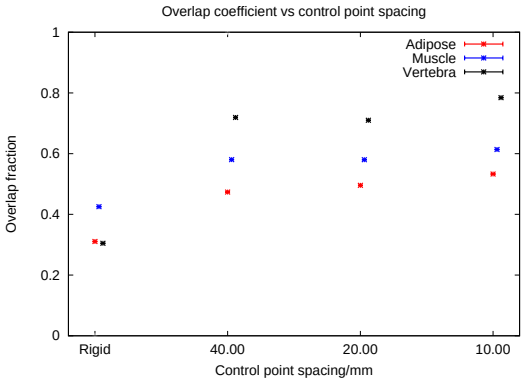

(b) Artefact region CT-CT registration

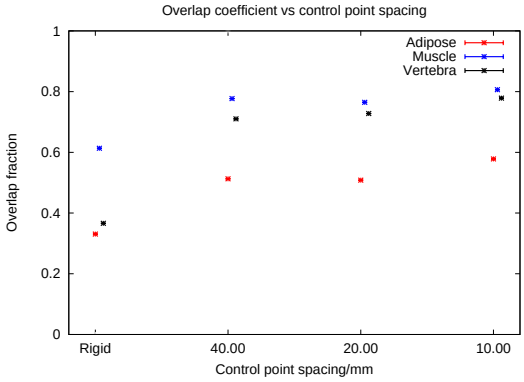

(c) Inferior region CPDS registration

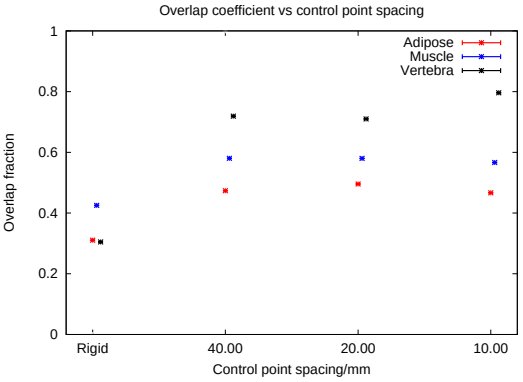

(d) Artefact region CPDS registration

Figure 6.8: Registration performance for different registration techniques in different image regions for images containing artefacts

	Adipose	Muscle	Vertebra
Inferior CPDS	0.5782	0.8062	0.7785
Inferior CT-CT	0.5886	0.8090	0.7837
Artefact zone CT-CT	0.4667	0.5663	0.7962
Artefact zone CPDS	0.5328	0.6137	0.7844

Table 6.1: Overlap coefficients for the structures at $10mm$ CPS registration in the different zones for each registration technique

Discussion and Conclusion

The results plotted in figures 6.8 and 6.7 show a general trend whereby the overlap index improves when the registration transformation becomes increasingly fine-scale. At the $10mm$ CPS level, the results shown in table 6.1 of the overlap coefficients for the group of patients in which there were artefacts present, show a very small difference between the two registration techniques in the region of anatomy that was inferior to the artefact region. This is because the deformation field within this anatomy will become increasingly dominated by the CT-derived transformation, with increasing distance in the superior-inferior direction. Beyond a certain point, there should be no difference between the overlap metrics for the two techniques, since for the CPDS technique, the deformation will consist purely of the CT-derived transformation.

Within the artefact region, the departures in the overlap measures for the two registration techniques become more apparent. The overlap coefficient for the vertebra is higher in this region for the CPDS technique but is lower for the other two techniques.

6.4 Discussion and Conclusion

In this chapter, the performance of the CPDS registration technique was assessed in practice using real clinical data in which there are dental artefacts present. In the previous section, the registration performance was assessed using the overlap techniques for segmented structures. The aim was to compare how well in practice the CPDS technique performed,

compared to a registration technique where the similarity measure was defined using information from the CT images only. The results of this analysis were suggestive for this technique, as assessed using this method, although the improvements were small. Several comments should be made concerning this practical comparison.

The first reason why the improvements were small is that in the limited number of patients there was minimal misregistration present; the problem with the artefacts only occurs if there is misregistration including specific rotations. Ideally a controlled experiment is needed in patients where artefacts are already identified; this is not logistically easy to achieve and it would be difficult to justify additional CT exposure for this research.

There were limitations here in the choice of anatomy available on which to perform the segmentation. Ideally, when using the overlap method, a shape that is smooth is preferable over one that is tortuous, for example, a shape that is petal-shaped, as discussed in [Crum, Camara and Hill, 2006] will result in the overlap coefficient being much more sensitive to small angular motions than a shape that is circular in cross section, when rotated about its axis. Within the head and neck anatomy, the structures present there are all very tortuous and therefore overly sensitive to small variations in the registration's deformation field.

There were limitations inherent to the segmentation technique used for the overlap method: since the segmentation was based on using CT thresholding techniques, the segmented volumes themselves were prone to being adversely affected by the artefacts. The benefit of a segmentation based around image intensity values is that it is a more objective method and it is operator independent; and reliant only on the inherent information that is contained within the images. Having expert observers segmenting structures will inject certain prejudicial bias into the segmentations, which will depend on the observer's level of experience, type of education, and even where they were educated might affect their interpretation of the images. Inter-observer variability has been described for vari-

ous circumstances in the literature [Muenzel et al., 2012; Lundqvist et al., 2012].

Further work could be done to investigate alternative methods for segmentation that are less prone to being affected by the presence of dental artefacts for the overlap technique to be more robustly implemented. An alternative would be to employ different techniques for assessment, such as a distance measure between corresponding surfaces, defined on the fixed and floating images, yet it seems likely that an automated technique employed to define these surfaces will also become corrupted by the presence of the artefacts. This is a non-trivial problem to address and will require a substantial effort that is beyond the scope of this work, which was primarily on the development and testing of a multicomponent similarity measure for use in non-rigid registration.

Finally, there was a limited amount of clinical data available from the clinical trial data. Of the original cohort of patients recruited onto this trial, the data from only 6 patients was available at the time this work was carried out. Ideally, clinical data from more patients is required to perform analysis to gain more statistically meaningful results.

Chapter 7

Summary, Conclusions and Further Work

7.1 Introduction

The work in this thesis was concerned with developing, implementing and then assessing a non-rigid registration technique to improve the alignment of serial oncological FDG PET/CT studies. These images were acquired during a phase 1 clinical trial to assess the response to radiotherapy treatment by acquiring serial images on four different occasions. Once in alignment, segmented regions can be identified on the reference image and then propagated onto any of the subsequent serial images using the transformation gained using the registration. The segmented regions would define the tumour(s), and the changes in SUV within these region(s) could parameterise the response to treatment. Other regions could also be defined if desired. For example: an acute inflammatory response can result from radiotherapy treatment, so a segmentation could be defined to assess the time-development of this effect.

Whilst these serial images were acquired with the patient wearing an immobilisation mask to minimise misalignments, there will still be non-rigid inter-scan motions. This can be due to slight postural changes, variations in tongue or jaw position, or physiological changes which may be

caused by the treatment or the disease. These physiological changes can be from inflammatory processes, which could result in localised swellings, or shrinkages to other structures (such as the tumour), or caused by the loss of adipose tissue due to loss of appetite. So a registration technique was developed using a non-rigid transformation to generate the required localised correspondence.

The non-rigid registration technique initially developed relied on similarity measures calculated from voxel intensity values from the CT image pairs, since there is a higher degree of anatomical information than in the PET images. However, it was noticed that when metal artefacts from dental fillings are present in a pair of CT images, a non-rigid registration will incorrectly attempt to register the two artefacts together since they are strong features within the image compared to the features that represent the actual anatomy. This resulted in localised registration errors in the deformation field in the vicinity of the artefacts. Several registration techniques were developed and implemented with the goal of improving the localised correspondence in the vicinity of the artefacts using combined information from both modalities.

To help develop and compare these registration techniques, artefacts were artificially simulated on images. Of the techniques developed, the CPDS method was found to perform best. It was then compared to the CT-CT registration technique for images with actual dental artifacts present.

7.2 Summary of Work in Each Chapter

The work done in each chapter - apart from the introduction and this chapter - will be summarised in the following subsections, with an outline of the primary focus of the chapter, the main techniques developed and implemented, the results acquired, and the conclusions drawn from these results. Any novel developments will be mentioned as appropriate: whether they were existing techniques used in a novel manner, or novel techniques developed and deployed for this work.

7.2.1 Summary of Chapter 2: Combined PET/CT Imaging

The primary focus of this chapter was to outline the relevant physics and instrumentation techniques behind combined PET/CT imaging to provide the relevant background information. Other kinds of multi-modality imaging systems were mentioned. This has relevance for the techniques developed within this project as there are advances within combined imaging that have moved on from PET and SPECT/CT. PET/MRI is now in clinical use and other techniques such as omnitomography are on the horizon. Therefore, image registration techniques reliant on intensity information from multiple sources that may be considered to have been acquired by the same scanner, are timely and an important area of development and research.

7.2.2 Summary of Chapter 3: Image Registration Techniques

In this chapter, the subject of image registration was introduced and surveyed, with a special focus on those techniques developed in the later sections within this thesis. This included a conceptual overview of what image registration is. The coordinates that go into defining an image - referred to as an "ordered set of intensities within a space" - were discussed, and it was explained why image registration is often required to modify the underlying geometric relationships that help to specify what an image is. Once this concept was introduced, the different kinds of image registration techniques were surveyed and in particular, those automatic iterative techniques reliant on similarity measures constructed from voxel intensities were discussed in more detail.

The general approach to image registration, in which a similarity measure is optimised by iteratively making changes to the geometric relationships in the floating image, was discussed. The transformation describes the ways in which these geometric relationships are modified. The main types of transformations were described: rigid, affine and, finally, non-linear transformations, also known as non-rigid transformations. The

main kinds of non-rigid transformations in the literature were surveyed, and the type used in this project was mentioned, with the details given in the following chapter. The main kinds of optimisation techniques that are used in image registration literature were surveyed, along with some other key concepts such as capture range, and how the nature of the similarity measure may affect the choice of the optimisation technique, as well as how coarse-to-fine techniques can be employed to simplify the similarity measure to prevent the solution becoming trapped in an incorrect local minima.

Voxel based similarity measures were outlined and the concept of joint histograms was introduced, and it was explained how they may be used to calculate mutual information using the joint histogram. A related similarity measure, normalised mutual information was also discussed, since it is used in later work. An N-dimensional joint intensity histogram was defined, along with a definition of multidimensional mutual information that can be calculated when comparing N floating images to a target image. This is an equivalent idea to the concept developed and implemented later, in which the floating and fixed image pairs each have N intensity components at each point. However, the details of the multidimensional joint histogram and the multidimensional mutual information differs from the form of these that are developed in chapter 5.

7.2.3 Summary of Chapter 4: Registering Head and Neck Images

In this chapter, an optimal non-rigid registration framework was developed to register pairs of serial combined head and neck FDG PET/CT images, using the CT images for the registration. The vtkCmic image registration software was used. This code can perform rigid and non-rigid registration using a B-splines non-rigid registration transformation, using a gradient descent optimisation technique. The salient features of the algorithm were described.

The relevant parameters were investigated that determine how the image registration scheme performed for the head and neck image data. So

the main focus of this chapter was concerned with developing an optimal framework for non-rigid registration. This was broken down into a series of tests for each parameter. Each parameter was modified in turn to explore how the modifications changed the performance of the registration. Appropriate methods of testing the performance of the registration were employed.

The relevant pre-processing parameters investigated were the thresholds used for the histogram equalisation, the number of bins used in the joint intensity histogram and the image resampling resolution. The relevant optimisation parameters investigated were the number of resolution levels, the number of steps and the starting step length. The registration was assessed using visual scoring techniques, either using subtraction images and inspecting how uniform they are, or by inspecting the registered image pair side by side in a viewing application. The other technique used was to identify a set of anatomical landmarks in the fixed and floated images. This also involved getting an estimate of the error associated with identifying these anatomical markers.

Once the optimum configuration of parameters was deduced, the registration using these settings was performed for several patients from the clinical trial. A problem was identified with poor alignment in one of the images that contained dental artefacts.

This then motivated the work in the following chapter, which was concerned with incorporating information from the PET images to guide the registration in this region, which involved developing a multicomponent similarity measure.

The work in this chapter included a novel application of an existing technique. A histogram equalisation method was employed and the number of bins used in the joint histogram was investigated, to understand how this improves the registration accuracy. This has been done for rigid and affine registration in earlier work [Rajwade, Banerjee and Rangarajan, 2006], but not for non-rigid registration. Histogram equalisation was used to improve the performance of the CT-CT registration by ensuring

that the soft tissues can be distinguished from one another. This is a more important consideration for non-rigid registration (than for rigid registration) because local changes in soft tissue positioning must be recovered and, in doing so, it is important for the algorithm to be able to distinguish between the soft tissues. This could not happen if all soft tissues were placed within a single histogram bin - those tissues would be indistinguishable. A large improvement was noticed when using a histogram equalisation compared to when not using it.

7.2.4 Summary of Chapter 5: Registration in the Presence of Simulated CT Artefacts

In this chapter, a formalism was developed for multicomponent registration, with the aim of implementing a system that could accurately register images in which there are dental artefacts present. The idea was to include intensity information from the PET image to help improve accuracy within the regions affected by the artefact. This involved refactoring a well-known image registration code to facilitate this formalism developed for registering a pair of multicomponent images. For this application, the images were dual-valued but, in principle, there could be any number of intensity components present in the image pair. This new registration technique that was developed and coded was tested on images in which dental artefacts were simulated.

The focus of the work in this chapter was two-fold. Firstly, to develop a formalism for defining a multicomponent joint histogram from a pair of multicomponent images from which a similarity measure may be calculated. Secondly, to test how this MCSM could be used to register images in which there are artificially simulated dental artefacts on the CT images. The question was investigated for our problem: how do CT artefacts adversely affect the image registration? A method was implemented for simulating dental artefacts on the CT images by numerically inserting large values into pre-determined locations within the forward-projected sinogram. This enabled the study of how the registration was affected by the presence of artefacts, which will *always* be in the transaxial plane.

If the floating image is rotated about the sagittal plane relative to the fixed image, then this will result in artefacts whose features will be on an inclined plane and will introduce local errors in the non-rigid registration's transformation. How the registration was affected by the presence of these simulated artefacts which were at an angle with respect to one another, was explored. This was studied in detail for the registration techniques developed within this chapter.

A reference non-rigid registration transformation was found by registering the two serial head and neck images prior to generating the CT artefacts on them, but otherwise in the same geometric configuration. Then, the registration techniques that were developed were compared to this reference transformation. When assessing the registration, three regions were defined: a region that was superior to the artefact, a region inferior to the artefact, and finally a region in the vicinity of the artefact. Within each region, a set of points was defined, and the difference from the reference transform to the transform under consideration was calculated for each point. For each of the three zones, the average, minimum and maximum errors were calculated. The following registration techniques were developed and assessed:

- A weighted MCSM containing terms from both CT-CT comparison and PET-PET comparison, with the cross terms having weights set to zero. The registration was performed for different values of the weight in front of the PET-PET term.
- The above similarity measure was used for images that were masked as follows: the PET images had their intensity values masked to zero outside the artefact region and the CT images had their intensity values masked to zero within the artefact region.
- Independent CT-CT and PET-PET registrations were performed and the control-point displacements were substituted from the PET transform to the CT transform within the vicinity of the artefact to produce a piecewise continuous and smooth deformation field across

the boundary zone.

The first method resulted in a lower error in the vicinity of the artefact, but a higher error elsewhere. The second method resulted in an error that was larger than the error induced by the artefact. The final method was found to be the most accurate. This method was then tested on images in which there were artefacts present, in the final chapter.

The idea of having a similarity measure in which there is more than one term is a well-established idea in the literature. Usually, the additional term is from some property of the transformation. For example, to ensure that the transformation is well behaved in some manner, there are methods to have a term calculated from the bending energy of the transformation to ensure it does not become too mobile, a smoothing term dependent on the transformation is introduced [Rueckert et al., 1999]. Or, to ensure that some other feature of the transformation is preferred, such as volume preservation [Rohlfing et al., 2003].

This work was the first time that a registration method has been developed to employ intensity information from both CT and PET images from a single combined scanner to improve the registration accuracy. It is also the first time that an intensity based registration technique has been proposed to overcome image artefacts. It was also the first systematic study of how dental artefacts adversely impact non-rigid registration, using the method employed to generate those artefacts artificially. This is better than the alternative method of performing imaging experiments using phantoms, since it would be impractical to construct a phantom that is sufficiently anthropomorphic. Most phantoms in clinical use are too regular and immobile, and therefore would not generate a sufficiently realistic image that would exercise the registration in a realistic manner. The patterns of intensities would be too well-ordered compared to an image of a real human.

The final novel aspect of this work was the control point substitution

method, whereby the displacements in different regions of the image are given from the results of two independent registration solutions.

7.2.5 Summary of Chapter 6: Registration in the Presence of CT Artefacts: Clinical Assessment

The registration techniques developed in the previous chapters were tested on patient data containing dental artefacts. The registration techniques were tested on two sets of patient data. One set of patient data contained dental artefacts and the other did not.

The registration was assessed using anatomical landmarks in the first instance. Since there are only very limited landmark points available within this region of anatomy, a volume overlap technique, involving calculating the Jaccard index, was employed to assess the performance of the registration of some soft tissue structures. A thresholding technique was employed to segment the structures of interest. The registration performance at each registration level was assessed using these methods.

The registration performance did not change noticeably when the CPDS method was used compared to the CT-CT registration alone. Some reasons for this were discussed in the concluding section of this chapter.

7.3 Further Work

This section presents some ideas for further work. During any body of research work, there are often occasions when a particular method is developed to a certain degree and not improved, since it would distract from the main thrust of the work. There are three areas in which this work may be taken further.

In this work, several techniques were used that were, to a first order, a good solution for the problem at hand. An example of this is the overlap

method used to assess the performance of the registration. This could be improved in two ways. More sophisticated segmentation algorithms could be employed to segment the anatomical structures of interest. There is a large body of work in this sister subject to image registration, which could be explored for the purpose of improving the segmentation.

Alternative methods for constructing the joint intensity histograms such that the information from the CT images will be used in an optimal manner could be explored. The technique employed in this thesis involved performing a histogram equalisation between two thresholds for each CT image, and then constructing the joint intensity histogram from this during registration, to calculate the desired similarity measure. NMI was used in preference to SSD since it has better accuracy. Alternative histogram transformation techniques could be investigated, and there are techniques that involve adaptive bin sizes that can be optimised.

When the CPDS method was employed, the location of the swapping region was determined manually. This could be performed using an automated method. Also, the location of the grid of control points relative to the image could be modified to study how this affects the performance of the registration. Whilst this technique remains popular today, it is surprising that such a well-established technique has not yet been tested in this manner.

When using the MCSM similarity measure, it was assumed that the cross terms in the similarity measure would not help the registration converge. Whilst there is a good conceptual basis for this assumption, it could be worth questioning this assumption and making some tests to determine the effect of changing the weights for these inner terms.

This work, in employing voxel information from multiple image modalities when registering a pair of images is a timely development since imaging is becoming increasingly multi-modal. Combined scanners have moved on from PET/CT and SPECT/CT to PET/MRI which is now entering clinical use. Multicomponent registration techniques could prove to

be useful for problems in other combined imaging modalities. By incorporating information from multiple sources, other more complex registration problems may be addressed, such as atlas-based techniques that involve inter-patient registration. Improved performance for the registration of inter-patient studies by using a multicomponent voxel similarity has been previously described [Papastavrou et al., 2009], but it is envisioned that this will be particularly useful if the concept of omnitomography is realised in practice, as discussed in section 2.5. Of course, any long-term predictions concerning future[1] work in a particular field are considered speculative.

[1]There is a saying - falsely attributed to Grouch Marx - that is worth mentioning here: *'The only thing you cannot predict is the future*'

Bibliography

Abramson, N.: Information theory and coding. McGraw-Hill, 1963

Agbeko, NN et al.: Positron range correction in PET using an alternating EM algorithm. In Nuclear Science Symposium Conference Record (NSS/MIC), 2010 IEEE. IEEE 2010, 2875–2878

Alarcon, T., Byrne, HM and Maini, PK: A cellular automaton model for tumour growth in inhomogeneous environment. Journal of Theoretical Biology, 225 2003, Nr. 2, 257–74

Alexander, DC and Gee, JC: Elastic Matching of Diffusion Tensor Images. Computer Vision and Image Understanding, 77 2000, Nr. 2, 233–250

Aljabar, Paul et al.: Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy. Neuroimage, 46 2009, Nr. 3, 726–738

Araujo, R.P. and McElwain, D.L.S.: A history of the study of solid tumour growth: The contribution of mathematical modelling. Bulletin of Mathematical Biology, 66 2004, Nr. 5, 1039–1091

Ardekani, B A et al.: A fully automatic multimodality image registration algorithm. Journal of Computer Assisted Tomography, 19 Jul 1995, Nr. 4, 615–623

Arici, Tarik, Dikbas, Salih and Altunbasak, Yucel: A histogram modification framework and its application for image contrast en-

hancement. IEEE Transactions on Image Processing, 18 2009, Nr. 9, 1921–1935

Arsigny, Vincent et al.: A Log-Euclidean Polyaffine Framework for Locally Rigid or Affine Registration. In Biomedical Image Registration: Third International Workshop, WBIR 2006, Utrecht, The Netherlands, July 9-11, 2006, Proceedings. Volume 4057, Springer 2006, 120

Arsigny, Vincent, Pennec, Xavier and Ayache, Nicholas: A novel family of geometrical transformations: Polyrigid transformations. Application to the registration of histological slices. INRIA, 2003 (4837). – Research report

Avila, Lisa S et al.: The VTK User's Guide. Kitware New York, 2010

Beek, Edwin JR van and Hoffman, Eric A: Functional imaging: CT and MRI. Clinics in chest medicine, 29 2008, Nr. 1, 195–216

Berti, Valentina, Pupi, Alberto and Mosconi, Lisa: PET/CT in diagnosis of dementia. Annals of the New York Academy of Sciences, 1228 2011, Nr. 1, 81–92

Beyer, Thomas et al.: MR/PET–Hybrid Imaging for the Next Decade. MAGNETOM Flash.(Online)(Cited 2013 Nov 13). Available from URL: www. siemens. com/magnetom-world 2010

Beyer, Thomas et al.: A combined PET/CT scanner for clinical oncology. Journal of Nuclear Medicine, 41 2000, Nr. 8, 1369–1379

Beyer T., Kinahan P. E., Townsend D. W. et al: The use of x-ray CT for attenuation correction of PET data". IEEE Nuclear Science Symposium 1557 1994

Bharatha, A. et al.: Evaluation of three-dimensional finite element-based deformable registration of pre-and intraoperative prostate imaging. Medical Physics, 28 2001, 2551

Bissonnette, J. P. et al.: Quantifying interfraction and intrafraction tumor motion in lung stereotactic body radiotherapy using respiration-correlated cone beam computed tomography. International Journal of Radiation Oncology Biology Physics, 75 Nov 2009, Nr. 3, 688–695

Bookstein, Fred L.: Principal warps: Thin-plate splines and the decomposition of deformations. IEEE Transactions on pattern analysis and machine intelligence, 11 1989, Nr. 6, 567–585

Boor, Carl de: A Practical Guide to Splines. Springer-Verlag, 1978

Burger, C et al.: PET attenuation coefficients from CT images: experimental evaluation of the transformation of CT into PET 511-keV attenuation coefficients. European Journal of Nuclear Medicine and Molecular Imaging, 29 2002, Nr. 7, 922–927

Cahill, Nathan D et al.: Revisiting overlap invariance in medical image alignment. In IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshop. CVPRW'08. IEEE 2008, 1–8

Cal-González, Jacobo et al.: Study of CT-based positron range correction in high resolution 3D PET imaging. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 648 2011, S172–S175

Catana, Ciprian et al.: Simultaneous acquisition of multislice PET and MR images: initial results with a MR-compatible PET scanner. Journal of Nuclear Medicine, 47 2006, Nr. 12, 1968–1976

Cherry, Simon R: Multimodality imaging: Beyond PET/CT and SPECT/CT. In Seminars in nuclear medicine. Volume 39, Elsevier 2009, 348–353

Cherry, Simon R, Sorenson, James A and Phelps, Michael E: Physics in Nuclear Medicine. Elsevier Health Sciences, 2012

Chong, Edwin KP and Zak, Stanislaw H: An introduction to optimization. Volume 76, John Wiley & Sons, 2013

Christensen, G. E., Rabbitt, R. D. and Miller, M. I.: Deformable templates using large deformation kinematics. IEEE Transactions on Image Processing,, Nr. 10, 1435–1447

Collignon, A. et al.: Automated multi-modality image registration based on information theory. Information Processing in Medical Imaging, 1995, 263–274

Crum, W R, Hartkens, T and Hill, D L G: Non-rigid image registration: theory and practice. British Journal of Radiology, 77 Spec No 2 2004, 140–153

Crum, William R, Camara, Oscar and Hill, Derek L G: Generalized overlap measures for evaluation and validation in medical image analysis. IEEE Transactions on Medical Imaging, 25 Nov 2006, Nr. 11, 1451–1461, Evaluation Studies

Crum, William R, Scahill, Rachael I and Fox, Nick C: Automated hippocampal segmentation by regional fluid registration of serial MRI: validation and application in Alzheimer's disease. Neuroimage, 13 2001, Nr. 5, 847–855

Crum, WR, Tanner, C and Hawkes, DJ: Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging. Physics in Medicine and Biology, 50 2005, Nr. 21, 5153

Davatzikos, C.: Spatial transformation and registration of brain images using elastically deformable models. Computer Vision and Image Understanding, 66 1997, Nr. 2, 207–222

Fitzpatrick, J Michael, West, Jay B and Maurer Jr, Calvin R: Predicting error in rigid-body point-based registration. Medical Imaging, IEEE Transactions on, 17 1998, Nr. 5, 694–702

Flotats, Albert et al.: Hybrid cardiac imaging: SPECT/CT and PET/CT. A joint position statement by the European Association of Nuclear Medicine (EANM), the European Society of Cardiac Radiology (ESCR) and the European Council of Nuclear Cardiology (ECNC). European Journal of Nuclear Medicine and Molecular Imaging, 38 2011, Nr. 1, 201–212

Fujiwara, Kent et al.: Locally rigid globally non-rigid surface registration. In Computer Vision (ICCV), 2011 IEEE International Conference on. IEEE 2011, 1527–1534

H., Zaidi; H, Zaidi, editor: Quantitative Analysis in Nuclear Medicine Imaging. Springer, 2006

Haber, SF, Derenzo, Stephen E and Uber, D: Application of mathematical removal of positron range blurring in positron emission tomography. IEEE Transactions on Nuclear Science, 37 1990, Nr. 3, 1293–1299

Hajnal, J.V., Hill, D.L.G. and Hawkes, D.J.: Medical Image Registration. CRC Press, 2001

Hammer, Bruce E, Christensen, Nelson L and Heil, Brian G: Use of a magnetic field to increase the spatial resolution of positron emission tomography. Medical Physics, 21 1994, Nr. 12, 1917–1920

Han, Xiao et al.: Atlas-based auto-segmentation of head and neck CT images. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2008. Springer, 2008, 434–441

Hartkens, Thomas et al.: VTK CISG Registration Toolkit An Open Source Software Package for Affine and Non-rigid Registration of Single-and Multimodal 3D Images. In Bildverarbeitung für die Medizin 2002. Springer, 2002, 409–412

Herrmann, C.: Pet-CT system with single detector. Patent App. PCT/IB2011/055,081. 2012

Higgins, Kristin A et al.: Analysis of Pretreatment FDG-PET SUV Parameters in Head-and-Neck Cancer: Tumor SUV< sub> mean</sub> Has Superior Prognostic Value. International Journal of Radiation Oncology* Biology* Physics, 82 2012, Nr. 2, 548–553

Hill, Derek L. G.: Combination of 3D medical images from multiple modalities. Ph. D thesis, Image Processing Group, Radiological Sciences, UMDS, Guys Campus, St Thomas's Street, University of London, 1993

Holden, Mark: A review of geometric transformations for nonrigid body registration. IEEE Transactions on Medical Imaging, 27 2008, Nr. 1, 111–128

Hounsfield, GN: Nobel lecture, 8 December 1979. Computed medical imaging. Journal de Radiologie, 61 1980, Nr. 6-7, 459

Hudson, H Malcolm and Larkin, Richard S: Accelerated image reconstruction using ordered subsets of projection data. IEEE Transactions on Medical Imaging, 13 1994, Nr. 4, 601–609

Hutton, Brian F, Hudson, H Malcolm and Beekman, Freek J: A clinical perspective of accelerated statistical reconstruction. European Journal of Nuclear Medicine, 24 1997, Nr. 7, 797–808

Ibanez, Luis et al.: The ITK software guide second edition updated for ITK version 2.4. 2005

Iida, H et al.: A simulation study of a method to reduce positron annihilation spread distributions using a strong magnetic field in positron emission tomography. IEEE Transactions on Nuclear Science, 33 1986, Nr. 1, 597–600

Jenkinson, Mark and Smith, Stephen: A global optimisation method for robust affine registration of brain images. Medical Image Analysis, 5 2001, Nr. 2, 143–156

Jødal, Lars, Le Loirec, Cindy and Champion, Christophe: Positron range in PET imaging: an alternative approach for assessing and

correcting the blurring. Physics in Medicine and Biology, 57 2012, Nr. 12, 3931

Karp, Joel S et al.: Benefit of time-of-flight in PET: experimental and clinical results. Journal of Nuclear Medicine, 49 2008, Nr. 3, 462–470

Kidd, Elizabeth A and Grigsby, Perry W: Intratumoral metabolic heterogeneity of cervical cancer. Clinical Cancer Research, 14 2008, Nr. 16, 5236–5241

Kinahan, PE et al.: Attenuation correction for a combined 3D PET/CT scanner. Medical Physics, 25 1998, Nr. 10, 2046–2053

Kojadinovic, Ivan: On the use of mutual information in data analysis: an overview. In 11th international symposium on Applied Stochastic Models and Data Analysis (ASMDA'05). Brest, France, May 2005, 738–747

Kolesov, Ivan et al.: A stochastic approach for non-rigid image registration. In IS&T/SPIE Electronic Imaging. International Society for Optics and Photonics 2013, 86550U–86550U

Kyriacou, S.K. and Davatzikos, C.: A biomechanical model of soft tissue deformation, with applications to non-rigid registration of brain images with tumor pathology. MICCAI, 1998, 531–538

Langerak, Thomas Robin et al.: Label fusion in atlas-based segmentation using a selective and iterative method for performance level estimation (SIMPLE). IEEE Transactions on Medical Imaging, 29 2010, Nr. 12, 2000–2008

Lau, Y.H., Braun, M. and Hutton, B.F.: Non-rigid image registration using a median-filtered coarse-to-fine displacement field and a symmetric correlation ratio. Physics in Medicine and Biology, 46 2001, Nr. 4, 1297–1319

Legg, Philip A et al.: A robust solution to multi-modal image registration by combining mutual information with multi-scale derivatives.

In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009. Springer, 2009, 616–623

Li, Yang and Verma, Ragini: Multichannel image registration by feature-based information fusion. IEEE Transactions on Medical Imaging, 30 2011, Nr. 3, 707–720

Loeckx, Dirk et al.: Nonrigid image registration using free-form deformations with a local rigidity constraint. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2004. Springer, 2004, 639–646

Loktyushin, Alexander et al.: Blind multirigid retrospective motion correction of MR images. Magnetic Resonance in Medicine 2014

Lundqvist, Eva et al.: Measurement of transplanted pancreatic volume using computed tomography: reliability by intra-and inter-observer variability. Acta Radiologica, 53 2012, Nr. 9, 966–972

Makihara, Yasushi, Mori, Atsushi and Yagi, Yasushi: Temporal super resolution from a single quasi-periodic image sequence based on phase registration. In Computer Vision–ACCV 2010. Springer, 2011, 107–120

Mamourian, Alexander C: CT Imaging: Practical Physics, Artifacts, and Pitfalls. Oxford University Press, 2013, ISBN 978–0199782604

Marsden, Paul K: Detector technology challenges for nuclear medicine and PET. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 513 2003, Nr. 1, 1–7

Mattes, David et al.: PET-CT image registration in the chest using free-form deformations. IEEE Transactions on Medical Imaging, 22 2003, Nr. 1, 120–128

McQuaid, Sarah J and Hutton, Brian F: Sources of attenuation-correction artefacts in cardiac PET/CT and SPECT/CT. European

Journal of Nuclear Medicine and Molecular Imaging, 35 2008, Nr. 6, 1117–1123

Modat, Marc et al.: Fast free-form deformation using graphics processing units. Computer Methods and Programs in Biomedicine, 98 2010, Nr. 3, 278–284

Moses, William W: Fundamental limits of spatial resolution in PET. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 648 2011, S236–S240

Muenzel, Daniela et al.: Intra-and inter-observer variability in measurement of target lesions: implication on response evaluation according to RECIST 1.1. Radiology and oncology, 46 2012, Nr. 1, 8–18

Muyan-Ozcelik, Pinar et al.: Fast deformable registration on the GPU: A CUDA implementation of demons. In Computational Sciences and Its Applications, 2008. ICCSA'08. International Conference on. IEEE 2008, 223–233

Nolf, E.: XMedCon an open-source medical image conversion toolkit. European Journal of Nuclear Medicine; Vol. 30 (suppl. 2); 2003; pp S246; TP39.

Papastavrou, Y et al.: A Multi-component similarity measure for improved robustness of non-rigid registration of combined FDG PET-CT head and neck images. In 4th European Conference of the International Federation for Medical and Biological Engineering. Springer 2009, 433–435

Pichler, Bernd J et al.: PET/MRI: paving the way for the next generation of clinical multimodality imaging applications. Journal of Nuclear Medicine, 51 2010, Nr. 3, 333–336

Pluim, Josien P W, Maintz, J B Antoine and Viergever, Max A: Mutual-information-based registration of medical images: a survey.

IEEE Transactions on Medical Imaging, 22 Aug 2003, Nr. 8, 986–1004

Pluim, Josien PW, Maintz, JB Antoine and Viergever, Max A: Image registration by maximization of combined mutual information and gradient information. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2000. Springer 2000, 452–461

Pohl, Cle and Van Genderen, JL: Review article multisensor image fusion in remote sensing: concepts, methods and applications. International Journal of Remote Sensing, 19 1998, Nr. 5, 823–854

Press, B. P. Flannering, S. A. Teukolsky and Vetterling, W. T.: Numerical Recipes in C. Cambridge University Press Cambridge, 1992, ISBN 0521431085

Rajkumar, Vineeth and Papastavrou, Yanni et al: A comparative study of PDGFR inhibition with imatinib on radiolabeled antibody targeting and clearance in two pathologically distinct models of colon adenocarcinoma. Tumor Biology, 33 2012, Nr. 6, 2019–2029

Rajwade, Ajit, Banerjee, Arunava and Rangarajan, Anand: Continuous image representations avoid the histogram binning problem in mutual information based image registration. In Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on. IEEE 2006, 840–843

Rohlfing, Torsten et al.: Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. IEEE Transactions on Medical Imaging, 22 2003, Nr. 6, 730–741

Rohr, Karl, Fornefett, Mike and Stiehl, H Siegfried: Approximating thin-plate splines for elastic registration: Integration of landmark errors and orientation attributes. In Information Processing in Medical Imaging. Springer 1999, 252–265

Rohr, Karl et al.: Landmark-based elastic registration using approximating thin-plate splines. IEEE Transactions on Medical Imaging, 20 2001, Nr. 6, 526–534

Roose, Tiina, Chapman, S Jonathan and Maini, Philip K: Mathematical models of avascular tumor growth. Siam Review, 49 2007, Nr. 2, 179–208

Rueckert, Daniel et al.: Quantification of growth and motion using non-rigid registration. In Computer Vision Approaches to Medical Image Analysis. Springer, 2006, 49–60

Rueckert, Daniel et al.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE Transactions on Medical Imaging, 18 1999, Nr. 8, 712–721

Russakoff, Daniel B et al.: Image similarity using mutual information of regions. In Computer Vision-ECCV 2004. Springer, 2004, 596–607

S., Webb; S., Webb, editor: The physics of Medical Imaging. Oxford University Press, 1998, ISBN 9780852743492

Samani, A. et al.: Biomechanical 3-D finite element modeling of the human breast using MRI data. IEEE Transactions on Medical Imaging, 20 2001, Nr. 4, 271–279

Sarndal, C.E.: A comparative study of association measures. Psychometrika, 39 1974, Nr. 2, 165–187

Savage, Neil: Path found to combined MRI and CT Scanner [News]. Spectrum, IEEE, 50 2013, Nr. 4, 16–18

Schillaci, Orazio: Hybrid SPECT/CT: a new era for SPECT imaging? European Journal of Nuclear Medicine and Molecular Imaging, 32 2005, Nr. 5, 521–524

Sederberg, Thomas W and Parry, Scott R: Free-form deformation of solid geometric models. In ACM Siggraph Computer Graphics. Volume 20, ACM 1986, 151–160

Seiler, Christof, Pennec, Xavier and Reyes, Mauricio: Geometry-aware multiscale image registration via OBBTree-based polyaffine log-demons. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2011. Springer, 2011, 631–638

Shannon, Claude E.: A Mathematical Theory of Communication. The Bell System Technical Journal, 27 july 1948, 379–423

Shen, D. and Davatzikos, C.: HAMMER: hierarchical attribute matching mechanism for elastic registration. IEEE Transactions on Medical Imaging, 21 2002, Nr. 11, 1421–1439

Sherratt, J.A. and Chaplain, M.A.J.: A new mathematical model for avascular tumour growth. Journal of Mathematical Biology, 43 2001, Nr. 4, 291–312

Shields, Anthony F et al.: Imaging proliferation in vivo with [F-18] FLT and positron emission tomography. Nature medicine, 4 1998, Nr. 11, 1334–1336

Sims, Richard et al.: A pre-clinical assessment of an atlas-based automatic segmentation tool for the head and neck. Radiotherapy and Oncology, 93 2009, Nr. 3, 474–478

Skerl, Darko, Likar, Bostjan and Pernus, Franjo: A protocol for evaluation of similarity measures for rigid registration. IEEE Transactions on Medical Imaging, 25 2006, Nr. 6, 779–791

Škerl, Darko, Likar, Boštjan and Pernuš, Franjo: A protocol for evaluation of similarity measures for non-rigid registration. Medical Image Analysis, 12 2008, Nr. 1, 42–54

Smith F. W., Gemmell H. G., Sharp P.F: Practical Nuclear Medicine. Oxford University Press, 1998, ISBN 0199630321

Sotiras, Aristeidis, Davatzikos, Christos and Paragios, Nikos: Deformable medical image registration: A survey. IEEE Transactions on Medical Imaging, 32 2013, Nr. 7, 1153–1190

Stammberger, T. et al.: Elastic registration of 3 D cartilage surfaces from MR image data for detecting local changes in cartilage thickness. Magnetic Resonance in Medicine, 44 2000, Nr. 4, 592–601

Staring, Marius, Klein, Stefan and Pluim, Josien PW: A rigidity penalty term for nonrigid registration. Medical Physics, 34 2007, Nr. 11, 4098–4108

Studholme, C., Hill, D. L. G. and Hawkes, D. J.: An overlap invariant entropy measure of 3D medical image alignment. Pattern Recognition, January, Nr. 1, 71–86

Studholme, Colin: Incorporating DTI data as a constraint in deformation tensor morphometry between T1 MR images. In Information Processing in Medical Imaging. Springer 2007, 223–232

Tanner, Christine et al.: Factors influencing the accuracy of biomechanical breast models. Medical Physics, 33 Jun 2006, Nr. 6, 1758–1769

Teräs, M et al.: Performance of the new generation of whole-body PET/CT scanners: Discovery STE and Discovery VCT. European Journal of Nuclear Medicine and Molecular Imaging, 34 2007, Nr. 10, 1683–1692

Thirion, J. P.: Image matching as a diffusion process: an analogy with Maxwell's demons. Medical Image Analysis, September, Nr. 3, 243–260, ISSN 1361–8415

Thomas, CM et al.: Specific recommendations for accurate and direct use of PET-CT in PET guided radiotherapy for head and neck sites. Medical physics, 41 2014, Nr. 4, 041710

Townsend, David W: Combined positron emission tomography–computed tomography: the historical perspective. In Seminars in Ultrasound, CT and MRI. Volume 29, Elsevier 2008, 232–235

Townsend, David W, Beyer, Thomas and Blodgett, Todd M: PET/CT scanners: a hardware approach to image fusion. In Seminars in Nuclear Medicine. Volume 33, Elsevier 2003, 193–204

Townsend D. W., Beyer T., Kinahan P. E. et al: The SMART scanner: a combined PET/CT tomograph for clinical oncology. In IEEE Nuclear Scinece Symposium and Medical Imaging Conference.

Vach, Werner et al.: Generating evidence for clinical benefit of PET/CT in diagnosing cancer patients. Journal of Nuclear Medicine, 52 2011, Nr. Supplement 2, 77S–85S

Veit-Haibach, Patrick et al.: PET–MR imaging using a tri-modality PET/CT–MR system with a dedicated shuttle in clinical routine. Magnetic Resonance Materials in Physics, Biology and Medicine, 26 2013, Nr. 1, 25–35

Vercauteren, Tom et al.: Non-parametric diffeomorphic image registration with the demons algorithm. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2007. Springer, 2007, 319–326

Viola, P. and III, W. M. Wells: Alignment by maximization of mutual information. International Conference on Computer Vision, 1995, 16–23

Wachinger, Christian and Navab, Nassir: Simultaneous registration of multiple images: similarity metrics and efficient optimization. IEEE Transactions on Pattern Analysis and Machine Intelligence, 35 2013, Nr. 5, 1221–1233

Wang, Ge: First CT-MRI Scanner for Multi-dimensional Synchrony and Multi-physical Coupling. arXiv preprint arXiv:1402.2647 2014

Wang, Ge and Yu, Hengyong: The meaning of interior tomography. Physics in Medicine and Biology, 58 2013, Nr. 16, R161

Wang, Ge et al.: Towards omnitomography grand fusion of multiple modalities for simultaneous interior tomography. PloS one, 7 2012, Nr. 6, e39700

Weinholt, W. and Sendhoff, B.: How to determine the redundance of noisy chaotic time series. International Journal of Bifurcation and Chaos [in Applied Sciences and Engineering], 6 1996, Nr. 1, 101–117

Weng, Xuchu, Ding, Yu-Shin and Volkow, Nora D: Imaging the functioning human brain. Proceedings of the National Academy of Sciences, 96 1999, Nr. 20, 11073–11074

Wernick M. N., John N. Aarsvold J. N.: Emission Tomography: The Fundamentals of PET and SPECT Hardcover. Academic Press, 2004

West, Jay B et al.: Fiducial point placement and the accuracy of point-based, rigid body registration. Neurosurgery, 48 2001, Nr. 4, 810–817

Woods, R P, Cherry, S R and Mazziotta, J C: Rapid automated algorithm for aligning and reslicing PET images. Journal of Computer Assisted Tomography, 16 Jul 1992, Nr. 4, 620–633

Woods, R P, Mazziotta, J C and Cherry, S R: MRI-PET registration with automated algorithm. Journal of Computer Assisted Tomography, 17 Jul 1993, Nr. 4, 536–546

Zacharaki, Evangelia I et al.: Non-diffeomorphic registration of brain tumor images by simulating tissue loss and tumor growth. Neuroimage, 46 2009, Nr. 3, 762–774

Zaidi, Habib and Hasegawa, Bruce: Determination of the attenuation map in emission tomography. Journal of Nuclear Medicine, 44 2003, Nr. 2, 291–315

Zeng, Ziming, Shepherd, Tony and Zwiggelaar, Reyer: Unsupervised tumour segmentation in PET based on local and global intensity fitting active surface and alpha matting. In Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE. IEEE 2012, 2339–2342

Zhu, Hui, Chan, Francis HY and Lam, Francis K: Image contrast enhancement by constrained local histogram equalization. Computer Vision and Image Understanding, 73 1999, Nr. 2, 281–290

Zhuang, X. et al.: An atlas-based segmentation propagation framework locally affine registration–application to automatic whole heart segmentation. Medical Image Computing and Computer-Assisted Intervention–MICCAI, 11 2008, Nr. Pt 2, 425–433

Appendix A

Artefact Simulation using MATLAB

The MATLAB code is listed below that was used to create the artefacts. It is described in section 5.3.

```
%Matlab Code for artefact generation by Tryphon Lambrou

function I1 = Test_Reconstruction_v4(I1, x1, y1, x2, y2, w)

[X, Y, Z] = size(I1);

theta = 0:1:179;

Rim = radon(I1, theta);

P2 = zeros(X,Y,Z);

rx = randi([-1 1], 180, 1);

for x = 1:w

  for y = 1:w

   I1(x1+x,y1+y) = 6000.0;

    I1(x2+x,y2+y) = 6000.0;
```

```
  P2(x1+x,y1+y) = 6000.0;

   P2(x2+x,y2+y) = 6000.0;

end

end

P = I1;

clear I1;

R = radon(P,theta);

% I1 with the artifacts inserted R2 = radon(P2,theta); % Zeros with the artifacts inserted [~, c] = size(R2);

% The number of projections

for x = 1:c % loop over the projections

rr1 = R(:,x); % I1 projection at angle theta with the artifacts inserted

   rr2 = R2(:,x); % Zeros projection at angle theta with the artifacts inserted

   kk = rr2 > 0.0; % A set of indices along the line for which the artefact exists on the zeroed image

    kkk = circshift(kk, [0 rx(x)]); %

  rr1(kk) =  Rim(kkk,x) + 25000.0; %The original image, forward projected – add a large intensity value to the range of indices

    R(:,x) = rr1; % let's make R(theta) be equal to the projection end

I1 = iradon(R,theta);
```

Appendix B

Multicomponent Images and Joint Histograms

This appendix details how multicomponent images are represented in the general case and how a joint histogram may be defined.

Images that have more than one intensity value at each point are known as multispectral images in remote sensing, [Pohl and Van Genderen, 1998]. In this field, different bands of the radiation are imaged by satellite and the measured radiation in each band is represented as one intensity value in each pixel. A more commonplace example of multicomponent image is a colour image, where the three primary colours: red, green and blue are assigned a different intensity value. At each point in the image, there will be three numbers that describe the intensity of each primary light component. One could generalise this to the case where there are N intensity components in an image. Extending the notation as introduced in section 3.2, the multicomponent intensity values $\vec{I}(\vec{r})$ at locations $\vec{r}$ are represented as a column vector:

$$\vec{I}(\vec{r}) = \begin{pmatrix} I_1(\vec{r}) \\ I_2(\vec{r}) \\ \vdots \\ I_N(\vec{r}) \end{pmatrix} \tag{B.1}$$

From a pair of multicomponent images, a multicomponent joint intensity histogram can be calculated:

The multicomponent joint intensity histogram for two multicomponent images: the target and floating images that have already been histogram binned (so the intensity values here range from $0-63$) ; $\vec{I}(\vec{r},t_0)$ and $\vec{I}(\vec{r},t_1)$, having M and N intensity components respectively, is represented by a joint histogram $\vec{G}(\vec{I}_T(\vec{r}),\vec{I}_F(\vec{r}))$ with $M*N$ components that are defined by the following vector relationship:

$$\vec{G}(\vec{I}_T(\vec{r},t_0),\vec{I}_F(\vec{r},t_1)) = \begin{pmatrix} G(I_1(\vec{r},t_0),I_1(\vec{r},t_1)) \\ G(I_1(\vec{r},t_0),I_2(\vec{r},t_1)) \\ \vdots \\ G(I_1(\vec{r},t_0),I_N(\vec{r},t_1)) \\ G(I_2(\vec{r},t_0),I_1(\vec{r},t_1)) \\ G(I_2(\vec{r},t_0),I_2(\vec{r},t_1)) \\ \vdots \\ G(I_2(\vec{r},t_0),I_N(\vec{r},t_1)) \\ \vdots \\ G(I_M(\vec{r},t_0),I_1(\vec{r},t_1)) \\ G(I_M(\vec{r},t_0),I_2(\vec{r},t_1)) \\ \vdots \\ G(I_M(\vec{r},t_0),I_N(\vec{r},t_1)) \end{pmatrix} \tag{B.2}$$

Where the joint intensity histogram for the $p^{th}$and q^{th} component of the fixed and floating images is given by:

$$G(I_{T,p},I_{F,q}) = \sum_{\vec{r}} \delta(I_{T,p}(\vec{r},t_0),I_{F,q}(\vec{r},t_1)) \tag{B.3}$$

An alternative notation is to decompose the multicomponent image pair that are of dimensions M and N respectively then to evaluate $M*N$ joint histograms. For the purposes of registration, given that the floating image is being transformed, it is, from the point of view of implementation

of image registration, a convenient implementation method, which is why this notation was developed.

This multicomponent joint histogram can be appropriately normalised to the number of voxels to yield a multicomponent probability distribution, from which various similarity measures may be defined. For the work in chapter 5, a combined PET/CT image can be considered as a multicomponent image with two components. In that case, the joint histogram will therefore have 4 components as described within that chapter.

Appendix C

Synopsis of vtkCmic Applications Developed During This Work

MaskPETImage

Usage: maskctimage <image_in> <masked_image> <zmin> <zmax> <width>

Masks-in a PET image in the z dimension from <zmin> to <zmax>, around where the dental metal artefact resides. The masked image is created and written into an image with filename <masked_image>. The roll-off factor is specified by setting <width>

MaskCTImage

Usage: maskpetimage <image_in> <masked_image> <zmin> <zmax> <width> from <zmin> to <zmax>, around where the dental metal artefact resides. The masked image is created and written into an image with filename <masked_image>. The roll-off factor is specified by setting <width>

Masks-out a ct image in the z dimension -

DofCompare

Usage: dofcompare <dof1> <dof2> <points_filenams>

Given two FFD transformations, calculate the average distance (ie the Euclidean error) between the displacement vectors at the points specified within the file called points_filename.

JointHist

Usage: jointhist <image1> <image2> <transform> <joint_hist_image> <nbins> nbins is the number of bins

Computes the joint historogram of two images using nbins number of bins. The joint histogram created and written into a file with filename <joint_hist_image>

CropImage

Usage: cropimage <image_in> <cropped_image> <xmin> <xmax> <ymin> <ymax> <zmin> <zmax>

A convenience utility that crops an image.

Overlap

Usage: overlap <image1> <image2> <transform> <H_Lower> <H_Upper> <x1> <x2> <y1> <y2> <z1> <z2>

The transformation in <transform> is applied to image2 and the overlap is calulated using CT thresholding using <H_Lower> <H_Upper> within the sub-region specified by the coordinates <x1> <x2> <y1> <y2> <z1> <z2>

RegToolMC

An updated version of regool that supports four images and built using vtkCmic objects that were modified to support multicomponent images. The arguements are the same as for regool and are listed below; however, there are the additional arguements that allow the associated PET image to be loaded with each source (floating) and target (fixed) image. The weights for the terms in the multicomponent similarity measure must also be specified.

Usage: regtoolmc [options] The default parameter values are listed in square brackets []

-si Source image

-ti Target image

-nd Degrees of freedom (6, 7, 9, 15, or 0=FFD) [6]

-dd Select 3D or 2D registration classes (3=3D, 2=2D) [3]

-dc Control point spacing [20.0,20.0,20.0]

-cw Weights for the multicomponent similarity measure [1.0,0.0,0.0,1.0]

-pt PET target image [NULL]

-ps PET source image [NULL]

-sm Source mask [NULL]

-tm Target mask [NULL]

-xi Initial DOF file [NULL]

-xo Output DOF file [NULL]

-id Invert input displacements

- for old CISG ASCII files (0=no, 1=yes) [0]

-xc Write output dof file in CISG format (0=no, 1=yes) [0]

-io Write out intermediate registration results (0=no, 1=yes) [0]

-tp Target padding background [-1.0] (conf: pad_tgt)

-ns Number of steps [4] (conf: steps) -ni Iterations per step [20]

-nl Multiresolution levels [1] (conf: levels) -ds Initial step length [4.0]

-st Subdivide the initial transform (0=no, 1=yes) [0]

-sl Subdivide the transform at each resolution level (0=no, 1=yes) [0]

-dl (unimplemented) register using a new FFD field level (0=no, 1=yes) [0]

bs Source blurring (<0 for none) [-1.0,-1.0,-1.0]

-bt Target blurring (<0 for none) [-1.0,-1.0,-1.0]

-rs Source resample (-1 for orig) [-1.0,-1.0,-1.0]

-rt Target blurring (-1 for orig) [-1.0,-1.0,-1.0]

-nt Threshold for bitmap masks [0]

-mg Switch on/off modified gradient [1]

-ep Tolerance for improvement of optimization [0.0]

-ge Improvement Tolerance for each gradient step in optimization [0.0]

-bp Control points in background areas will not be optimised (0=no, 1=yes) [1]

-rc Read configuration file [NULL]

-help Print this message -version Identify vtkCmic version of this reg-toolmc