Improving the Accuracy of CT-derived Attenuation Correction in Respiratory-Gated PET/CT Imaging

Robert Barnett

Thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy Faculty of Health Sciences The University of Sydney

February 2017

Abstract

The effect of respiratory motion on attenuation correction in fludeoxyglucose (^{18}F) positron emission tomography (FDG-PET) was investigated. Improvements to the accuracy of computed tomography (CT) derived attenuation correction were obtained through the alignment of the attenuation map to each emission image in a respiratory gated PET scan. Attenuation misalignment leads to artefacts in the reconstructed PET image and several methods were devised for evaluating the attenuation inaccuracies caused by this. These methods of evaluation were extended to finding the frame in the respiratory gated PET which best matched the CT. This frame was then used as a reference frame in mono-modality compensation for misalignment. Attenuation correction was found to affect the quantification of tumour volumes; thus a regional analysis was used to evaluate the impact of mismatch and the benefits of compensating for misalignment.

Deformable image registration was used to compensate for misalignment, however, there were inaccuracies caused by the poor signal-to-noise ratio (SNR) in PET images. Two models were developed that were robust to a poor SNR allowing for the estimation of deformation from very noisy images. Firstly, a cross population model was developed by statistically analysing the respiratory motion in 10 4DCT scans. Secondly, a 1D model of respiration was developed based on the physiological function of respiration. The 1D approach correctly modelled the expansion and contraction of the lungs and the differences in the compressibility of lungs and surrounding tissues. Several additional models were considered but were ruled out based on their poor goodness of fit to 4DCT scans. Approaches to evaluating the developed models were also used to assist with optimising for the most accurate attenuation correction.

It was found that the multimodality registration of the CT image to the PET image was the most accurate approach to compensating for attenuation correction mismatch. Mono-modality image registration was found to be the least accurate approach, however, incorporating a motion model improved the accuracy of image registration. The significance of these findings is twofold. Firstly, it was found that motion models are required to improve the accuracy in compensating for attenuation correction mismatch and secondly, a validation method was found for comparing approaches to compensating for attenuation mismatch.

I, Robert Barnett, hereby declare that the work contained within this thesis is my own and has not been submitted to any other university or institution as a part or a whole requirement for any higher degree.

I, Robert Barnett, hereby declare that I was the principal researcher of all work included in this thesis, including work published with multiple authors.

Name   Robert Barnett

Signed   ....

Date   19/9/2017

Acknowledgements

For their encouragement, advice, and patience during my eventful and extended time pursuing postgraduate research I acknowledge my primary and associate supervisors at the University of Sydney, Professor Roger Fulton and Professor Steven Meikle. Many aspects of this thesis would not have been possible without the feedback and support from Will Ryder, particularly in regards to data management and pipeline processing. I would also like to thank Georgios Angelis and John Gillam and other staff in our group at the BMRI for providing a forum for ideas and open discussion.

My role as a full time physicist as well as a part time student required a great deal of help and understanding from everyone I worked with. I would like to especially thank Lee Collins, David Skerrett and Rochelle McCredie from Westmead Hospital for their support. I would also like to thank both technologists and clinical staff in Nuclear Medicine and PET, especially the practical assistance in PET provided by Laura Torales.

Peculiarly, the motivation for this work came from the passing of my brother, Matthew "Greyseer" Barnett 1973-2002. It is likely that newer and more advanced PET/CT technology would have improved the staging of disease and hence management of his quality of life. It has been a privilege to use his knowledge of programming to investigate ways PET/CT technology could be advanced further. Finally this thesis would not have been possible without the support of my family, especially from my wife. I also have the utmost gratitude for the contributions by our family librarian, Pam Barnett, for the numerous proof reads of draft papers and sections of this thesis.

Publications and Presentations

The following are a list of publications and presentations arising from this thesis work

Published Papers

- Barnett R, Meikle S, Fulton R . *Cross population motion modeling applied to attenuation correction of respiratory gated F18-FDG PET.* IEEE Transactions on Nuclear Science, Volume: 63, Issue: 1, Feb. 2016

Published Abstracts and Proceedings

- R. Barnett, S. Meikle, R. Fulton. (2010). *Statistical motion modeling of the thorax applied to respiratory gated FDG PET.* In 2010 IEEE Nuclear Science Symposium & Medical Imaging Conference (pp. 2459-2464).

- R. Barnett, S. Meikle, R. Fulton. (2011). *Deformable image registration by regarding respiratory motion as 1D wave propagation in an elastic medium.* In 2011 IEEE Nuclear Science Symposium Conference Record (pp. 2956-2963).

- R. Barnett, S. Meikle, R. Fulton. (2012). *Accelerated reconstruction for identifying image regions affected by rigid body movement.* In 2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC) (pp. 3079-3082).

Conference presentations - Oral

- Barnett R, Collins L, Fulton R. *Characterisation and Optimisation of Respiratory Gating Applied to a Siemens Biograph.* ANZSNM Annual Scientific Meeting 2008

- Barnett R, Fulton R. *Modification of a PET Respiratory Gating System for Amplitude Gating.* EPSM Conference 2009

- Barnett R. *A review on Respiratory motion models* ACPSEM ACT/NSW Branch Scientific Meeting 2012

- Barnett R, Meikle S, Fulton R. *Development of a wireless portable visual biofeedback device for guiding patient respiration during a medical imaging scan.* ACPSEM ACT/NSW Branch Med Phys 2012

- Barnett R, Meikle S, Fulton R. *A 1D lung deformation model including air flow and tissue elasticity.* ACPSEM ACT/NSW Branch Med Phys 2015

Conference presentations - Poster

- Barnett R, Meikle S, Fulton R. *Robust estimation of respiratory motion from PET image sequences through minimisation of voxel-wise variance.* Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2013 IEEE

Contents

Chapter 1

Introduction

A combined PET/CT scanner enables functional imaging using a PET tracer in combination with anatomical imaging using a multi-slice CT scanner. Thoracic PET/CT imaging is often used to assess lung cancer by both detecting regional metastatic disease and to stage and evaluate larger tumours in the lung. Although larger tumours can be easily seen on CT, a combined PET/CT scanner allows for the increased detection of regional metastatic disease which is not distinguishable on CT. PET also reveals additional nodes from a primary tumour, thus allowing for improved definition of organs requiring radiotherapy treatment (Greco, Rosenzweig, Cascini, & Tamburrini, 2007).
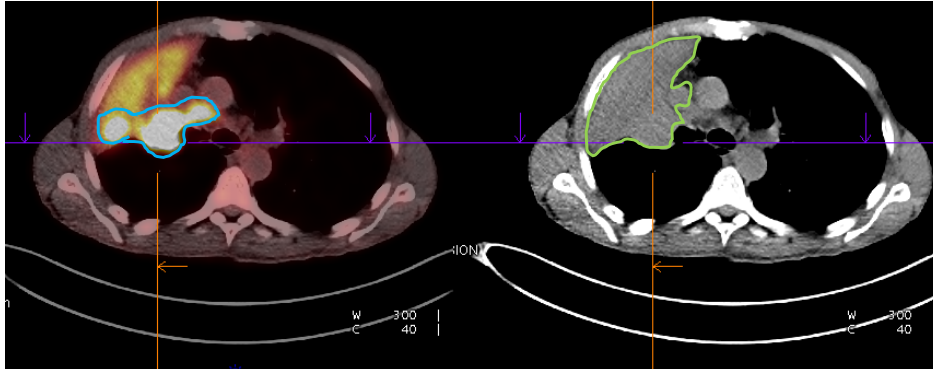
Figure 1.1: A transverse slice of a PET/CT image of the thorax using a Biograph mCT. Although the tumour mass is large, PET revealed that only a portion of the tumour mass (blue) is metabolically active in comparison to the bulk (green).

PET-CT systems provide a combination of information regarding tumour biology and precise lesion localisation (Osman, Cohade, Nakamoto, & Wahl, 2003). PET-CT imaging is especially useful in diagnosing early-stage non-small cell lung cancer (NSCLC) as there is a critical need to identify patients who are more likely to develop metastatic disease and who thus might benefit from additional therapy (Wu et al., 2016).

Accurate reporting of combined PET/CT imaging requires a thorough understanding of the normal and variant physiological distribution of tracers as well as common incidental findings and technical artefacts (Corrigan, Schleyer, & Cook, 2015). A recent study demonstrated that the visual appearance of tumours in PET/CT images can be used to predict metastatic disease (Wu et al., 2016). However, the visual appearance of image features such as small tumours can be severely affected by inaccuracies during image reconstruction, including the partial volume effect, motion and attenuation (Corrigan et al., 2015).

Respiratory gated PET/CT is a common technique used for addressing motion. This technique is particularly useful for imaging tumours in the thorax and liver which are significantly affected by motion. Although the gating technique mitigates image blur due to respiratory motion, there is a trade off between reduced motion artefacts and increased noise due to the limited counts per frame. The noise in respiratory gated scans is a significant problem, however, this can usually be overcome by either scanning for longer, increasing radiation dose, or using a more sensitive detector system.

Attenuation correction can be derived from CT by scaling the CT image to produce an attenuation map which is then incorporated into the reconstruction of the PET image. The benefit of a combined PET/CT scanner relies on the assumption that the PET and CT

images are aligned. Misalignment poses two significant problems. Firstly, the anatomical and functional images are no longer spatially aligned diminishing the clinical benefit of observing a lesion in both modalities. Secondly, the CT derived attenuation correction of PET images relies upon an aligned CT.

CT derived attenuation correction is often inaccurate due to misalignment caused by patient motion. Respiratory motion is the largest contributor to misalignment, and image artefacts can be clearly seen in cases where the misalignment due to motion has altered the appearance of the PET image. Attenuation mismatch artefacts are commonly visible near the upper liver due to under-attenuation correction owing to larger lung volumes on the CT than on the emission PET (Osman et al., 2003).
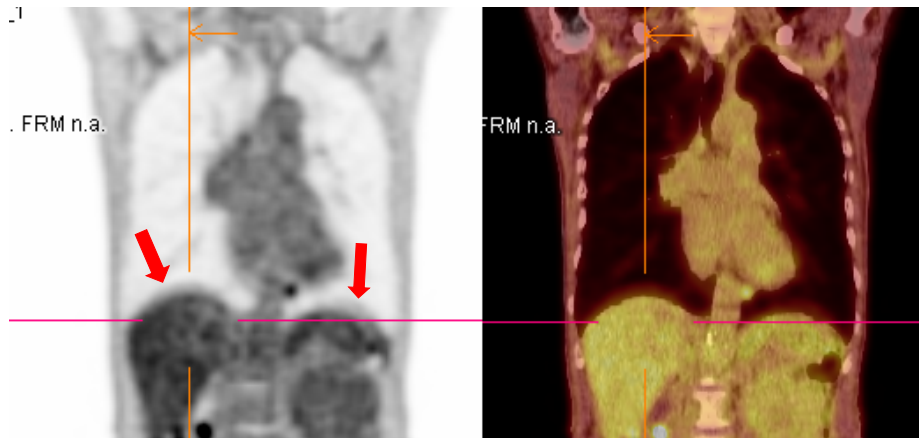
Figure 1.2: Respiratory motion typically produces a band artefact resulting from under correction of attenuation around the diaphragm in AC PET images.

There have been numerous investigations into improving respiratory gated PET images using various approaches to motion compensation. Given the increased sensitivity of PET systems, it has been argued that motion compensation only provides incremental benefit since motion corrupted data from PET systems could be discarded without jeopardising the signal to noise ratio (SNR) in PET images (Liu et al., 2010). Despite numerous developments in compensating for respiratory motion in PET, not all of these methods have been extended to address attenuation misalignment even though it remains a clinically challenging issue (Corrigan et al., 2015)

There has been extensive work on the task of modelling respiratory motion for the purposes of estimating, tracking and predicting motion (McClelland, Hawkes, Schaeffter, & King, 2013) with the most common method of tracking achieved through the use of image registration applied to respiratory gated frames. Some methods for compensating for motion in PET/CT incorporate motion estimation into PET image reconstruction. Although these approaches are intended to improve the SNR, it has been shown that noise increases for some reconstruction algorithms (Tsoumpas et al., 2013). This thesis proposes that a more valuable goal is to track internal movement primarily for the purposes of attenuation correction. It focuses on developing and comparing approaches to address mismatch between PET and CT in an effort to address the clinical issues of diagnosing disease from a PET scan.

1.1 Contributions

An initial investigation into comparing existing methods of image registration was completed for the purposes of tracking the internal deformation of organs due to respiratory motion.

Existing registration packages often feature a variety of methods for modelling tissue deformation, however, the performance of image registration packages is dependant on many factors. Most packages use the tuning parameters as the main way to improve registration. Registration can also be improved by finding optimal methods of tissue deformation, algorithm regularisation and estimating image similarity. Because existing registration packages feature a variety of methods a systematic approach was taken to validation whereby direct comparisons between packages was possible.

The poor image quality of PET images can be addressed by simplifying image registration. This was achieved by reducing the number of control points or through extensive regularisation. Although these approaches were known to increase the robustness of registration, the investigations in this thesis showed that these were not adequately complex to accurately estimate respiratory motion (Barnett, Meikle, & Fulton, 2011) .

Some methods of image registration have been evaluated using digital anthropomorphic phantoms to determine suitability for PET/CT alignment (Bai & Brady, 2011). Although this method was a useful first step, a limitation of this approach is that image registration may appear accurate for simulations, yet still encounter difficulties in circumstances of abnormal anatomy (Barnett et al. (2011) - in appendix).

The shortcomings of image registration were addressed by considering a variety of respiratory motion models. It was thought that this was an area lacking development because, although many motion estimation algorithms were developed for CT imaging, few had been successfully used in PET.

Existing models of respiratory motion have been used to track the motion of a tumour (Liu, Alessio, & Kinahan, 2011), however, alternative models were sought which could also track the movement of organs surrounding the tumour. This was especially challenging because these organs contributed to attenuation mismatch, yet were not clearly visible on the PET image.

An initial investigation highlighted difficulties in modelling motion in regions of low contrast and high noise on the PET image. A number of existing models were investigated which were based on the physics of respiratory motion. Although these models were robust to the noise in the image, they were also inaccurate in areas of abnormal motion such as where disease had adversely affected the ventilation to sections of the lung. Previous work using finite element modelling (FEM) had demonstrated that such models require accurate segmentation and classification of disease and tissue properties. Methods were sought which could be used to accurately track clearly visible objects and only resort to modelling for regions of low contrast.

The first contribution in this thesis was to develop two novel motion estimation methods which addressed issues specific to respiratory motion in PET/CT.

The first model is presented in chapter 4 and used machine learning to estimate respiratory motion. A machine learning population model was used to find the relationship between motion in a collection of 4DCT scans. The population model was then used to find the most statistically likely motion in a PET/CT even though no 4DCT had been acquired.
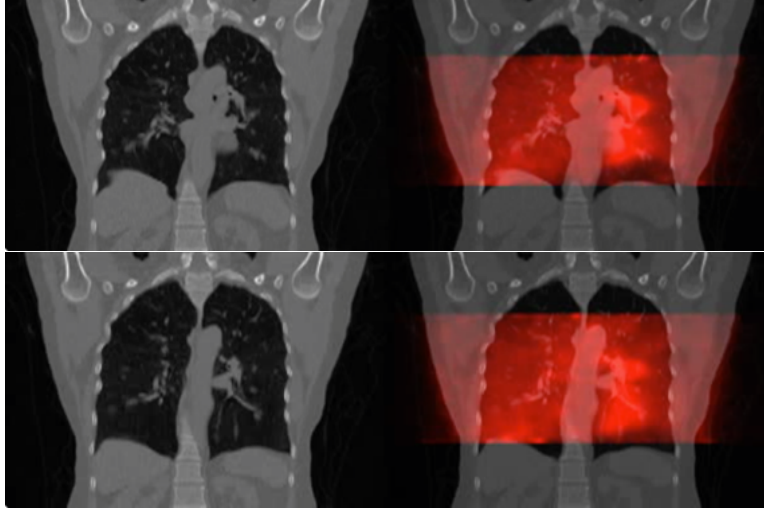
Figure 1.3: Cross population model used to predict motion at End Inspiration (Top) and End Expiration (Bottom).

The second model is presented in chapter 5 and used a 1 dimensional respiratory motion model to constrain the amount that lung tissue expands and contracts. An initial estimate of motion was used to build the 1 dimensional model of air flow and tissue deformation, but it did not require segmentation of lung boundaries. This model was then built into the image registration algorithm to provide the tissue expansion and contraction constrains. This improved motion tracking in areas of low contrast where motion could not be otherwise detected.
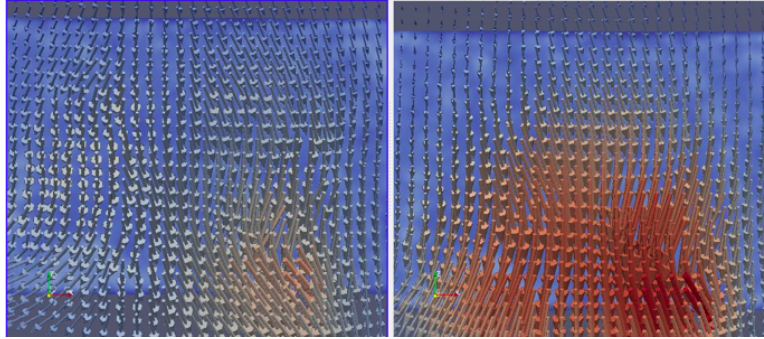
Figure 1.4: A coronal slice of the deformation field from the image registration overlaid onto the target PET frame. The deformation field is represented by arrows of varying colour. Red represents large deformations. Blue and grey represent smaller deformations. The field produced using free-form deformation (left). The field produced after applying wave equation constraints (right). In both cases the bending energy penalty term was 1%.

The second contribution in this thesis was to develop a method for validating motion estimation and is presented in chapter 3. A set of consistency metrics were devised which enabled the optimisation of image registration without the need for human observers. Validation through real scans was seen as necessary because abnormal anatomy was difficult to

model. Models would often fail for highly abnormal patients, highlighting the need to use real patient data to test the robustness of the models.

These contributions demonstrated several approaches to improving attenuation correction of respiratory gated PET without increasing patient radiation dose using a combined PET/CT scanner. Further work is outlined in chapter 6 and it is suspected that these approaches could be enhanced by using the additional information available on modern Time Of Flight (TOF) PET systems.

References

Bai, W., & Brady, M. (2011). Motion correction and attenuation correction for respiratory gated PET images. *Medical Imaging, IEEE Transactions on*, *30*(2), 351–365. doi: 10.1109/TMI.2010.2078514

Barnett, R., Meikle, S., & Fulton, R. (2011). Deformable image registration by regarding respiratory motion as 1d wave propagation in an elastic medium. In *2011 IEEE nuclear science symposium and medical imaging conference (NSS/MIC)* (pp. 2956–2963). doi: 10.1109/NSSMIC.2011.6152528

Corrigan, A. J. G., Schleyer, P. J., & Cook, G. J. (2015). Pitfalls and artifacts in the use of PET/CT in oncology imaging. *Pitfalls in Nuclear Medicine (Part II)*, *45*(6), 481–499. doi: 10.1053/j.semnuclmed.2015.02.006

Greco, C., Rosenzweig, K., Cascini, G. L., & Tamburrini, O. (2007). Current status of PET/CT for tumour volume definition in radiotherapy treatment planning for non-small cell lung cancer (NSCLC). *Lung Cancer*, *57*(2), 125–134.

Liu, C., Alessio, A., Pierce, L., Thielemans, K., Wollenweber, S., Ganin, A., & Kinahan, P. (2010). Quiescent period respiratory gating for PET/CT. *Medical physics*, *37*(9), 5037–5043.

Liu, C., Alessio, A. M., & Kinahan, P. E. (2011). Respiratory motion correction for quantitative PET/CT using all detected events with internal-external motion correlation. *Medical Physics*, *38*(5), 2715–2723. doi: 10.1118/1.3582692

McClelland, J. R., Hawkes, D. J., Schaeffter, T., & King, A. P. (2013). Respiratory motion models: A review. *Medical Image Analysis*, *17*(1), 19–42. doi: 10.1016/j.media.2012.09.005

Osman, M. M., Cohade, C., Nakamoto, Y., & Wahl, R. L. (2003). Respiratory motion artifacts on PET emission images obtained using CT attenuation correction on PET-CT. *European journal of nuclear medicine and molecular imaging*, *30*(4), 603–606.

Tsoumpas, C., Polycarpou, I., Thielemans, K., Buerger, C., King, A. P., Schaeffter, T., & Marsden, P. K. (2013). The effect of regularization in motion compensated PET image reconstruction: a realistic numerical 4d simulation study. *Physics in Medicine and Biology*, *58*(6), 1759. doi: 10.1088/0031-9155/58/6/1759

Wu, J., Aguilera, T., Shultz, D., Gudur, M., Rubin, D. L., Loo Jr, B. W., ... Li, R. (2016). Early-stage non-small cell lung cancer: Quantitative imaging characteristics of 18f fluorodeoxyglucose PET/CT allow prediction of distant metastasis. *Radiology*, 151829.

Chapter 2

Background

PET/CT is a widely used imaging technique used in the diagnosis and staging of many diseases including cancer. This thesis evaluates several attenuation correction methods for PET/CT . This background chapter reviews developments which can offer improvements to the accuracy of attenuation correction. In section 2.1, the design and purpose of standalone PET and standalone CT scanner are introduced. The section discusses the benefit of a combined PET/CT scanner, particularly the availability of CT as a means to produce an attenuation map. It will be argued that despite these benefits, many studies have shown that CT derived attenuation correction is inaccurate due to respiratory motion. Section 2.2 presents respiratory gating as a technique for mitigating respiratory motion in both PET and CT scans.

Several approaches to attenuation correction are reviewed in section 2.3 and it will be shown that the accuracy of attenuation correction could be significantly improved by addressing the misalignment between PET and CT images. Section 2.4 overviews a variety of approaches to image registration. The section will also outline the theoretical framework needed to address PET/CT misalignment using deformable image registration whereby the attenuation map can be deformed such that it is aligned to the PET image.

Although the existing methods for correcting for alignment are promising, they all have shortcomings in the way that they do not explicitly model respiratory motion. Section 2.5 will describe a variety of ways these shortcomings can be addressed by parameterising respiratory motion through modelling. Given the wide variety of options for addressing misalignment, it is conceivable that any of the methods introduced in this review could be the approach which yields the most accurate attenuation correction. Section 2.6 will review two existing approaches to attenuation mismatch and identifies the need for improvement.

2.1   An Introduction to PET/CT Imaging

PET and CT are scanner technologies developed in the 1970s for the purposes of producing 3D images to aid the diagnosis of disease. A standalone CT scanner consists of a rotating X-ray tube and a detector on the opposing side. The patient is fed through the CT scanner on a moving bed and the entire assembly rotates around the patient to take multiple x-ray images of the patient at various angles. A standalone PET scanner consists of a ring of detectors for the purposes of detecting the position of radioactive decay of tracer radio-pharmaceuticals in the body. In a similar way to CT, the patient is fed through the scanner and imaged at various positions. A combined PET/CT scanner is a CT scanner with an adjoining PET scanner such that the patient bed can be fed through both scanners in one episode. In 2001 Townsend produced patient images from a combined PET/CT scanner. The following provides an overview of the benefits of a combined PET/CT scanner.

2.1.1   Helical Computed Tomography

Early designs of CT imaging systems acquired each axial slice separately in sequence with the detector performing a full rotation around the patient before moving to the next bed position. The development of continuously rotating CT systems, known as helical CT, provided much faster acquisition times allowing for rotation times of less than 1 second. With this technical

capability it was possible to operate these systems with the bed moving and sample projection space using a helical geometry. Modern Helical CT systems have multiple rows of detectors, allowing for the simultaneous acquisition of multiple slices. Helical CT is now the most common mode of acquisition, and there has been much investigation into motion artefacts specific to this method (Pan, Lee, Rietzel, & Chen, 2004).

Computed tomography (CT) is the process by which each axial 2D slice is reconstructed from a set of projections. The physics of CT can be described by a basic model of transmission of X-rays through the voxels (ray of length l) representing an object being detected as sinogram values Y_i on the ith detector (equation 2.1). This is known as the Radon transform. The sinogram is also comprised of both noise and scatter which need to be accounted for.

$$Y_i = b_i e^{-\sum_j l_{ij} \mu_j} + s_i + n_i \tag{2.1}$$

Reconstruction of the 2D slice is an inverse problem where the average attenuation in each voxel, μ_j, can be estimated by filtered back projection (FBP). FBP involves back projecting each ray between the X-Ray tube source and the detector and determining the length, l_{ij}, which intersects with the jth voxel. The sensitivity for the detector is b_i, the scatter, s_i, and the noise is n_i. The geometry of helical CT is modified and the volume data must be rebinned to produce axial slices.

A significant problem with CT reconstruction is that sometimes a voxel is not sampled from every sinogram angle. A simple example is when a metal object prevents any transmitted signal being measured at certain angles, in which case 'fan' artefacts will appear. The effect of these artefacts will also increase in the presence of noise, thus increasing the effect for low dose scans. Trade-offs affecting dose and sampling can be measured and tested using phantom experiments, making it possible to predict when artefacts will appear. Typically, every slice is reconstructed with all 360 degrees of sinogram data. Alternatively, it is possible to under sample projection angles using only 180 degrees of sinogram data (half scan reconstruction) however, this also leads to artefacts. The scan duration is an important factor in reducing the chance of motion occurring while the scan is underway. Scan speed can be increased by increasing the pitch. A pitch greater than one stretches out the helical path of the tube such that it skips slices. Increasing scan speed can reduce the chance of the image being corrupted in some way due to motion, however, fast scan speed can often be at the price of poor axial sampling (using high pitch) or high noise (using low mAs).

Image sampling is often assumed to be isotropic; that is, the resolution is the same in every direction. Unfortunately, the sampling scheme employed in CT is usually non-isotropic because the pitch is usually such that axial resolution is reduced in comparison to other directions. Many automated image algorithms must pre-filter the image such that the data has an isotropic sampling and resolution. These effects are also difficult to account for in motion tracking algorithms because motion tracking algorithms rely on the ability of the algorithm to track image features in all directions.

Clinical CT systems allow the operator to adjust the energy of the X-rays by varying the tube voltage (kV) and the amount of X-ray exposure by varying the cumulative tube current (mAs). Whilst these adjustments allow the operator to optimise image quality, careful attention must be paid to the dose deposited by the X-rays. It is important to minimise dose which is the energy deposited per unit volume of tissue (mGy). Table 1 shows typical doses

Table 2.1: CT Dose Reference Levels (DRLs) for various sites. Scanner parameters were interpreted using CT Expo (Stamm & Nagel, 2002).

Dose Reference Level	Dose: $CTDI_{vol}$ (mGy)	Typical scan parameters for 70kg patient
ARPANSA DRL for diagnostic chest CT in adult	15	140 mAs, 140kV 210mAs, 120kV
SNM Clinical Trials. Low dose CT for PET/CT Jallow et al. (2016)	4.9-6.2	100 mAs, 110kV
Example site (Westmead)	1.6	80 mAs, 80kV

Table 2.2: Typical scan parameters for helical CT

Scanner	Dose: $CTDI_{vol}$ (mGy)	Typical scan parameters
GE LightSpeed MSCT Pan et al. (2004)	40mGy	100mAs, 120kV, up to 10 rotations, 5s acquisition per slab.
GE Discovery ST PET/CT Murphy et al. (2011)	160mGy	100mAs, 120kV, up to 30 rotations, 20s acquisition per slab

and scan parameters for a chest CT.

2.1.2   Cine Computed Tomography

Modern multi-slice CT systems have been introduced to increase scan speed. Mutli-slice CTs have large detector coverage and can acquire between 1cm and 16cm of axial data in a single rotation. Cine mode involves scanning at the same slice location for multiple gantry rotations (Low et al., 2003). Cine mode has become possible through the reduction of rotation times to less than 0.5 seconds, enabling the production of a 4D image with a temporal resolution as low as 100ms.

A CT scanner capable of acquiring in cine mode can be combined with a respiratory motion management system to generate a 4DCT. This method of acquisition is covered later in the section on respiratory gating.

2.1.3   Anatomical Imaging of the Thorax using CT

Anatomical imaging of the thorax involves creating an image of internal organs with the purpose of visualising disease. Thoracic computed tomography creates a 3 dimensional image which allows for improved visualisation of disease over a conventional chest X-ray. A chest X-ray of the thorax typically contains many overlapping organs, however, a thoracic CT (Figure 2.1) provides a visual separation of the mediastinum, lungs, pleura, skeleton and liver.

There are many indications requiring a thoracic CT including evaluation of abnormal masses, staging of tumour and metastatic disease, and the evaluation of pulmonary embolus.
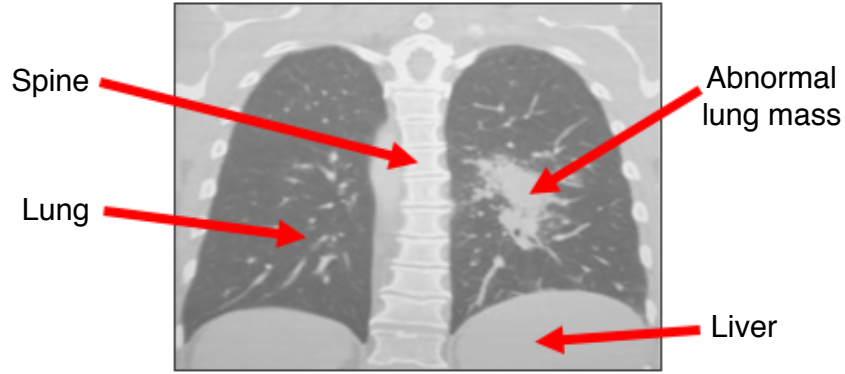
Figure 2.1: A coronal cross-sectional CT image of the thorax

CT in the thoracic area can also be used to evaluate various types of cardiac and aortic disease. Respiratory motion involves the continuous movement of nearly all of these organs in the thorax and is known to cause image artefacts which can lead to misdiagnosis. Both thoracic and cardiac CT imaging are affected by respiratory motion. Numerous scanning protocols have been developed to reduce the appearance of dome artefacts in the liver, however, the outcome of a scan depends on patient compliance and scan speed. A discussion on the mechanics of respiratory motion, as well as methods to estimate motion and assist in the correction of motion for the improvement of diagnosis, is addressed later in this review.

Over the history of CT imaging, the quality of CT images has improved considerably to ensure the reliability of diagnostic CT procedures at correctly detecting and staging disease. The reliable detection and/or staging of disease by an expert radiologist requires a minimum level of image quality. The human observer model is a mathematical construct which describes the image quality in an effort to estimate the reliability of disease detection (Barrett, Abbey, & Clarkson, 1998). In this model anatomical imaging can be thought of as an attempt to identify the existence of a true object which may or may not be visible on the CT image. The impact of CT image quality on disease detection imposes strict limitations on CT acquisition parameters. Improved image quality may come at the expense of a higher required radiation dose. This trade off can often be mitigated using advanced CT dose optimisation methods, however, an investigation into this is beyond the scope of this review. The same trade off also affects methods of respiratory motion detection and correction. This is because detecting respiratory motion is analogous to the human observer model in which we require a minimum image quality to observe an object.

The radio-density of organs in CT images is measured using the Hounsfield Scale, a transformation of linear attenuation coefficient where air is -1000HU and water as 0HU at standard temperature and pressure (STP). Some organs in the thorax are easily identifiable as their HU is significantly different from background. The process of finding anatomical boundaries in CT is known as segmentation. The boundaries between organs of similar HU are difficult to detect and often require expert interpretation, for example, the boundary between the myocardium and liver is difficult to detect because both tissues have a density of approximately 100HU. The visibility of apparent organ boundaries is affected by respiratory gating as described in section 2.2 and this is relevant to image registration methods reviewed in section 2.4.
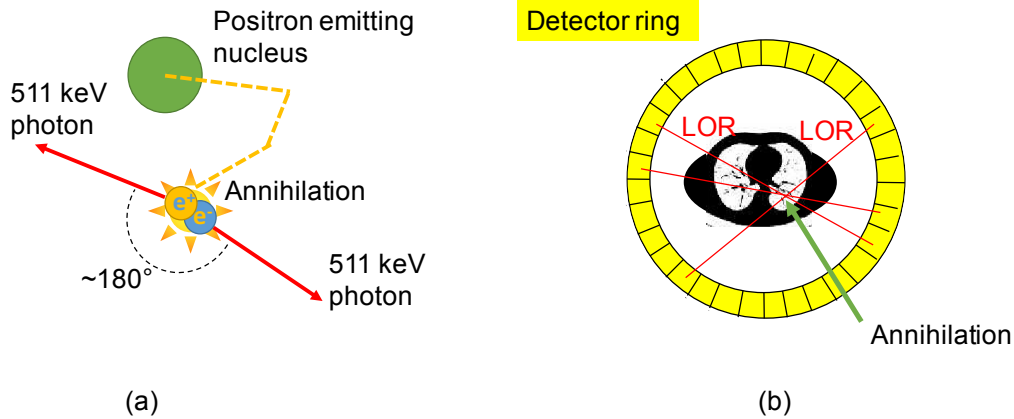
Figure 2.2: The annihilation of a positron leads (left) to a coincidence detection along a line of response (right)

2.1.4 Positron Emission Tomography

Positron Emission Tomography (PET) imaging relies on the radioactive decay of a positron emitting isotope. Emitted positrons travel a short distance through tissue before they annihilate (figure 2.2). Flourine 18 is a common PET isotope and the positrons have a maximum range of 2.6mm in tissue (Levin & Hoffman, 1999).

Each annihilation event creates two 511 keV gamma rays which are emitted in opposite directions. The PET scanner consists of a ring of highly sensitive detectors which are designed to detect both of these 511 keV gamma rays (see figure 2.2). The high sensitivity is achieved by using photo-scintillating crystals which have sufficient density and size to absorb energy from a significant number of 511 keV gamma rays without them just passing through undetected. The crystal is optically coupled to a photosensitive detector such as a photomultiplier tube (PMT) or a solid state photo-detector. If the detection of each 511keV event takes place within a narrow time frame then this is registered as a coincident detection. The position of the detected events can then be used to identify the line-of-response (LOR) upon which the annihilation must have taken place. A PET image is produced through the tomographic reconstruction of the detection events recorded along each LOR.

A significant proportion of detected events are due to the chance of two independent detections happening within the same coincidence time window. These are referred to as random events and need to be estimated and accounted for during tomographic reconstruction. Additionally, approximately 20-40% of events are scattered as a result of interacting with tissue. Scatter alters the path the photon travels to the detector meaning that the LOR no longer represents the positions where the original annihilation took place. Given the numerous potential interactions that lead to detected events, an estimate of the amount of useful detected signal to noise ratio (SNR) is useful in establishing the quality of image which can be produced from the PET scanner. The noise equivalent count rate (NECR) is a benchmark of SNR which is dependent on true detected events, random detected events and the scattered events.

The amount of noise in the image depends upon the number of measured true coincident events. The number of events can be predicted by knowing the sensitivity of the PET

Table 2.3: Typical scan parameters for PET

Scanner	Sensitivity	Dose (MBq)	Duration	Coincidence Events
Biograph 16	4kcps/MBq	200	20mins	10^7
Biograph mCT	8kcps/MBq	250	15min	2×10^7

| 100 msec | 1 sec | 10 sec |
| 4 000 cts | 40 000 cts | 400 000 cts |
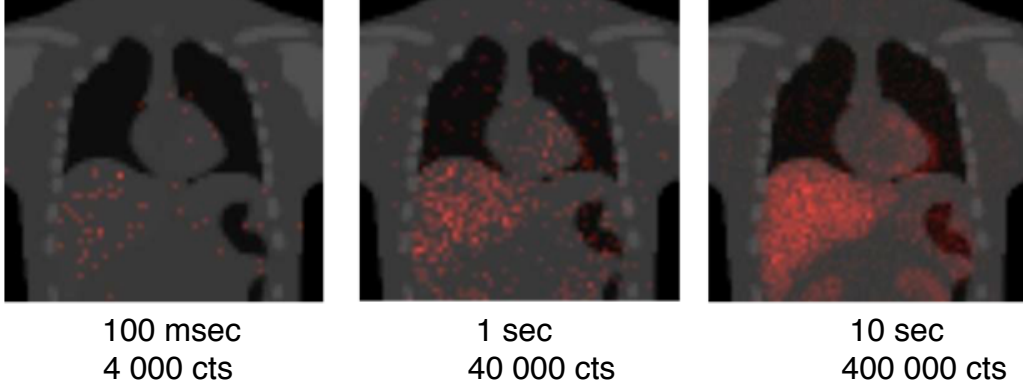
Figure 2.3: Low count PET/CT coronal slices simulated using the X CAT phantom. Shades of red indicate the counts detected using PET. The signal to noise ratio (SNR) is so low that even a large organ such as the liver is not visible until there are at least 4×10^5 counts recorded in the PET field of view (FOV) for low count images and this is why 10^6 counts is really needed to see structures.

scanner, the dose administered to the patient, and the duration of the scan (figure 2.3). The signal to noise (SNR) of the PET system is proportional to the inverse square root of the true coincident rate. The SNR limits the reliability of detecting a tumour in comparison to the surrounding tissue (Caucci & Barrett, 2012). Later in this review it will be shown that the SNR of the PET scanner places significant limitations on estimating respiratory motion. Table 2.3 shows administered dose and acquisition duration and scanner sensitivity suitable for producing a whole body image for clinical use for two commonly used PET/CT scanners.

PET is an intrinsically 3-dimensional (3D) imaging methodology and most modern reconstruction algorithms model coincidence detection in 3D. A 3D model involves ray tracing from the detector element along the line of response to find which image voxels are a potential source of the annihilation event. In order to make improved estimates of the voxel values, $\lambda(x, y, z)$ it is necessary to use the Radon transform. The Radon transform is the forward projection from every voxel to every detector element, $p(s, \phi, z_1, z_2)$, in the PET emission sinogram where s is the detector position, ϕ is the angle and z_1 and z_2 are the axial positions of the detector rings (figure 2.4).

$$p(s, \phi, z_1, z_2) = \mathcal{R}(\lambda(x, y, z)) \tag{2.2}$$

Given that it is not possible to know the exact source of the annihilation event a statistical algorithm can be used to find an estimate of voxel values, $\lambda(x, y, z)$. This estimate is the combination of voxel values which leads to the maximum likelihood of the forward projection matching the emission data (Shepp & Vardi, 1982).

Maximum Likelihood Expectation Maximisation (MLEM) is an iterative algorithm for
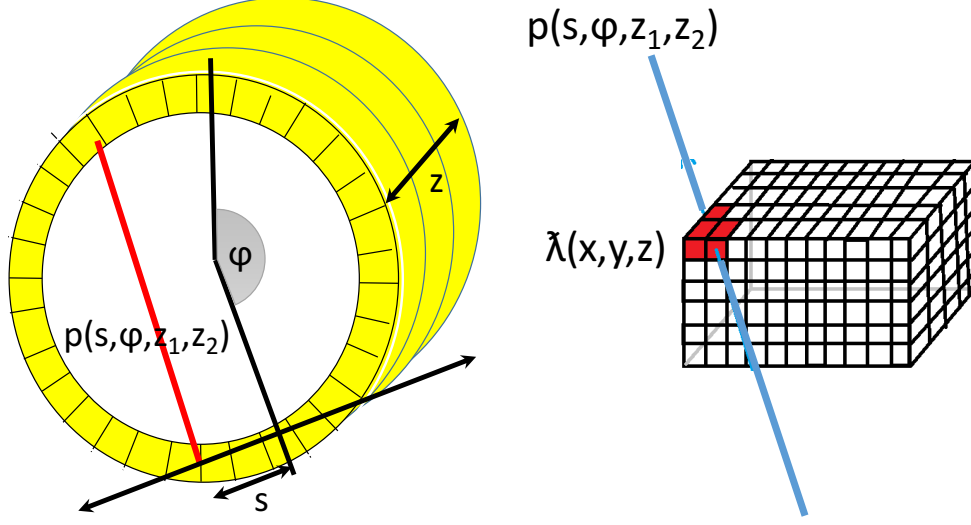
Figure 2.4: Ray tracing in 3D showing detector element $p(s, \phi, z_1, z_2)$ and voxel λ at position x, y, z.

obtaining the jth voxel, λ_j, from the ith detector measurement E_i. Approximately 100 iterations are required for a typical PET image before MLEM converges to an estimate, however, there exist many practical methods to accelerate convergence.

Maximum likelihood reconstruction of low count acquisitions does not necessarily lead to a visually interpretable image (figure 2.3), thus, it is often necessary to perform post reconstruction smoothing for the purposes of visual interpretation. Post reconstruction smoothing may also improve the stability of image registration algorithms discussed in section 2.4. .

PET relies on the accuracy of the system model which incorporates corrections into the MLEM reconstruction. An accurate system model corrects for events which are mispositioned or missing due to photon scatter, photon attenuation, detector resolution and dead time. Accurate corrections for photon attenuation is arguably one of the most critical factors for accurate PET in the thorax.

Equations 2.3 and 2.4 shows the update step including incorporating scatter, S_i and randoms, R_i effects (Iatrou, Ross, Manjeshwar, & Stearns, 2004).

$$\lambda_j^{k+1} = \frac{\lambda_j^k}{\sum_i A_{ii} P_{ij}} \sum_i \frac{A_{ii} P_{ij} E_i}{p_i} \tag{2.3}$$

where

$$p_i = \sum_j A_{ii} P_{ij} \lambda_j^k + R_i + S_i \tag{2.4}$$

2.1.5 Functional Imaging using PET

Positron emission tomography (PET) is an imaging technique aimed at visualising and quantifying the physiological processes. The most common use of PET is to image the glucose metabolism in an effort to either diagnose or stage cancer. In this technique, trace amounts

of the radio-labelled glucose analog, flurodeoxyglucose (FDG), are injected into the blood-stream. A fluorine radioisotope (F18) is used to label flurodeoxyglucose such that it becomes radioactive and undergoes positron emitting decay. The PET system is then used to obtain an image which is representative of FDG metabolism. For thoracic tumour imaging, FDG is distributed to the tumour and metabolised at a much higher rate in comparison to surrounding tissue. The concentration of FDG in the circulatory system decreases over a number of hours as it is absorbed into tissue. FDG remains in the tissue once it has been metabolised, thus the total amount metabolised can be approximated by performing delayed imaging. It can then be assumed that the metabolic rate of FDG does not vary significantly during the scan. The assumption of constant metabolism is useful when attempting to perform motion correction in PET.

It is not practical to wait until all the FDG is metabolised because the F18 also undergoes radioactive decay with a half life of 109.8 minutes. Delayed imaging is performed at approximately 60 minutes post injection because this provides the optimal image when taking into consideration both metabolism and radioactive decay of the tracer. An ideal delayed image provides a high signal at the tumour site in comparison to the surrounding background organs. The system resolution of modern clinical PET scanners is approximately 4mm and acquired images are affected by the partial volume effect (Erlandsson, Buvat, Pretorius, Thomas, & Hutton, 2012). The detection of small tumours (1cm) is limited by the PET system resolution and may only be positively identified if there is a high contrast of FDG-PET uptake in comparison to surrounding organs.

The in vivo concentration of FDG in diseased tissue provides clinical information about the pathology of the disease. Voxel values are directly proportional to in vivo tracer activity when the PET image is normalised for the sensitivity of each detector element pair, random coincidences are subtracted for each line of response and both attenuation and scatter correction are included in MLEM reconstruction. Since the voxel values are proportional, they can be calibrated and converted into units of Becquerels per millilitre (Bq/mL). This technique of estimating the absolute tracer activity is known as quantitative PET. For the FDG-PET tracer it has become popular to scale images according to a standardised uptake value (SUV) corrected for body mass and injected activity. The SUV is frequently used in thoracic imaging to stage the progression of disease, including pulmonary tumour nodules.

Motion affects the quantitation of PET scans and this can be compensated by incorporating motion correction. Motion correction requires knowledge of the motion upfront and this can be achieved by transforming LORs to compensate for the motion of a rigid body. Livieratos et al. (2005) demonstrated the correction of the inferior-superior motion of the heart in respiratory gated PET scans Unfortunately, although it is possible to correct for motion by transforming LORs, this thesis and review will show that this is non-trivial for motion of other organs in the thorax.

2.1.6 Benefits of a Combined PET and CT Scanner

The interpretation of a FDG-PET image involves distinguishing abnormal foci from the metabolism of non-involved organs such as brain, heart, liver, intestines, and the lung. Glucose metabolism is significantly increased by a variety of physiological processes, however, the anatomical location of a foci is needed to interpret whether it is abnormal physiology.

Combined PET/CT imaging allows for better distinction between normal physiology and disease. FDG foci can be localised to anatomical structures identified on CT providing additional interpretation of whether the structure is likely to have disease (Townsend, Carney, Yap, & Hall, 2004). Conversely, interpretation of suspicious structures observed on the CT can potentially be modified or refined by reference to the PET scan. In both cases, the alignment of the CT image and PET images are important.

Attenuation correction in PET imaging requires an accurate estimate of the photon attenuation at every voxel location. This makes it possible to calculate the proportion of photons which were not detected along each LOR due to the linear attenuation of 511keV photons. This proportion is equivalent to the attenuation correction factor (ACF) and can be applied to the counts recorded on a line of response. Attenuation is due to a combination of the photoelectric absorption and Compton scattering within tissue. Although the CT can be used to estimate the radio-density in terms of Hounsfield units, the low energy X-rays with a mean of approximately 70keV have a linear attenuation coefficient which is quite different to that of 511keV annihilation photons. The approach taken by P. E. Kinahan, Townsend, Beyer, and Sashin (1998) was to separate the problem into an approximation of the relationship for two types of tissue leading to a bilinear scaling function. Figure 2.5 shows a proposed relationship between HU and the linear attenuation for 511keV annihilation photons (P. E. Kinahan et al., 1998). Soft tissue such as lung, liver, fat and muscle are assumed to only require a single scaling factor which applies to HU less than 0. Spongiosa and cortical bone contain significant calcium and phosphorous and require an alternate scaling factor. The radio density is dependent on the x-ray energy leading to modified scaling factors for specific CT tube energies (P. E. Kinahan et al., 1998). Soft tissue such as lung, liver, fat and muscle are assumed to only require a single scaling factor which applies to HU less than 0. Spongiosa and cortical bone contain significant calcium and phosphorous and require an alternate scaling factor. The radio density is dependent on the x-ray energy leading to modified scaling factors for specific CT tube energies.

Accurate attenuation correction is essential for quantitative PET/CT, and although the bilinear scaling method can be used to derive the attenuation map from CT, there are challenges and potential improvements to attenuation correction which are detailed in section 2.3.1.

2.1.7   Improving PET Reconstruction with Time of Flight

The principle of time-of-flight PET is to measure the difference in arrival time of a pair of photons along a line of response. Since photons travel at the speed of light, the difference in arrival time can be used to estimate the position of the event along the line of response. Although the TOF resolution is poor, this additional information can be incorporated into the system model for iterative reconstruction to localise events along each LOR. The main benefit of TOF reconstruction is a significant increase in SNR of large body patients (Karp, Surti, Daube-Witherspoon, & Muehllehner, 2008). There were few TOF systems commercially available during initial investigations into the effect of respiratory motion on PET/CT imaging undertaken for this thesis. More recent investigations have shown that TOF can improve reconstruction making it more robust to inconsistencies in attenuation correction (Conti, 2010). The specific benefits of TOF for attenuation correction are detailed in section
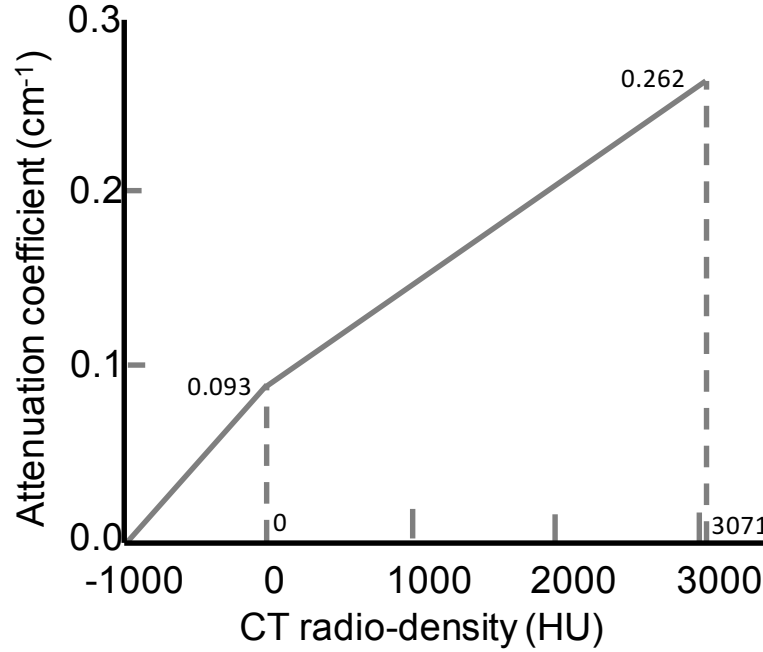
Figure 2.5: Bilinear scaling from HU to attenuation coefficient for 511 keV photons.

2.3.5.

2.2 Respiratory Gating in PET/CT

Respiratory gating in both PET and CT can be used to generate 4D images providing a sequence of images at different phases of the respiratory cycle. The method of generating 4D image is different for each modality. This section describes some of the challenges in acquiring gated images.

Respiratory motion during a PET scan degrades the image, particularly at organ boundaries which may continuously move by 20mm for the duration of the acquisition. Respiratory gating in PET can be used to address motion by partitioning data into bins which correspond to particular phase intervals in the respiratory cycle (figure 2.6). Each gate thus contains only a part of total motion and will have reduced motion artefacts (Dawood, Bther, Lang, Schober, & Schfers, 2007). Gating is often performed over many respiratory cycles whereby the total counts in each bin is added up over many cycles providing sufficient counts to reconstruct an image. Section 2.2.2 outlines key challenges in gated PET including methods for identifying phase intervals in the respiratory cycle, addressing image noise, addressing motion blur and tracking the motion of objects.

Respiratory gating in CT is implemented by taking multiple CT images of the same slice many times. Cine CT mode is the preferred method for acquiring images and is often referred to as 4DCT. The key challenges in generating a 4DCT image are that frames need to be sorted according to their phase interval in the respiratory cycle.
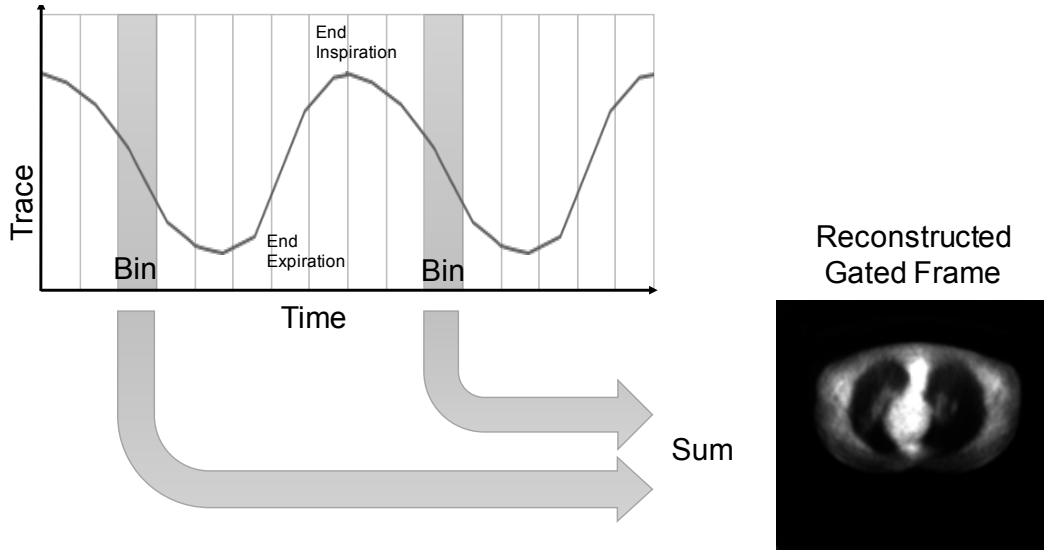
Figure 2.6: A respiratory trace (left) where gated events are used to select events from the list mode stream. These events are summed and reconstructed to create a gated frame (right).

2.2.1   Acquiring a Trace Suitable for Respiratory Gating

The respiratory trace is a time varying function which is an indicator of the depth of respiration. The respiratory trace is analysed to determine the phase interval for each respiratory gated bin. The ideal respiratory trace is the lung air volume and should clearly mark significant time points during the scan such as end inspiration (EI) and end expiration (EE). The only means of directly estimating lung air volume is to use spirometry during imaging, however this is impractical.

Several methods of acquiring a trace by monitoring the external movement of the chest have been proposed such as the use of a strain gauge or reflective markers (Nehmeh et al., 2004), or marker-less motion tracking using video (Noonan et al., 2012) (figure 2.7). The primary advantage of surface data is that it produces a more reliable result and is less dependent on patient positioning. In some cases chest surface tracking can be used to distinguish whole body motion from respiratory motion (Pretorius et al., 2011). Markerless tracking of the chest surface can also be used to construct a patient specific model of respiratory motion and this technique is described in more detail in the section on parametric models for respiratory motion. A disadvantage of external movement tracking is that it is not necessarily representative of internal motion, in particular, because the changes in lung volume lag the external movement of the chest (Liu, Alessio, & Kinahan, 2011). Another method that can be used to monitor internal movement is radio-wave frequency (Gu et al., 2012). In the case of PET it is also possible to interpret a respiratory trace from the listmode data itself. The movement of FDG avid organs across the most sensitive region of the PET detector can be interpreted from variations in the count rate in the list mode data (He et al., 2008). Alternative approaches such as dimensionality reduction or spectral analysis (Thielemans, Schleyer, et al., 2011) are more robust to noise. For example, Thielemans, Rathore, Engbrant, and Razifar (2011) used principal components analysis to interpret the respiratory trace from relative changes between elements in a down-sampled sinogram .
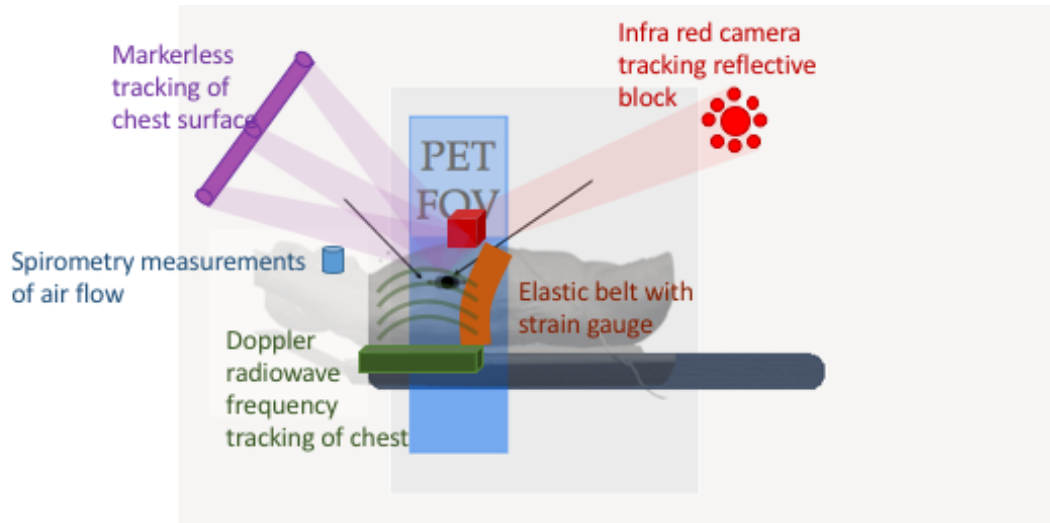
Figure 2.7: A schematic view of all the possible different tracking systems which have been used to create a respiratory trace, including markerless (purple), reflective block (red), spirometery (blue), doppler radar (green) and strain gauge (orange).

2.2.2 Challenges with Respiratory Gating a PET Scan

Respiratory motion during a PET scan tends to corrupt the image causing organ and disease boundaries to blur. Motion blurring dilutes the signal from small areas of increased uptake, thus lung tumour nodules may not be detectable on a PET image. Respiratory gating has been shown to restore the appearance of areas of focal uptake such that they may be more reliably detectable (Liu, Pierce, Larry, Alessio, & Kinahan, 2009). There are various methods for partitioning PET data to produce a respiratory gated scan. Partitioning PET list mode data by phase interval (phase gating) is the most widely used method because it is straightforward to encode timing intervals into the PET list mode data stream (Bruyant, 2007). The list mode stream can be binned into an arbitrary number of sinograms and reconstructed into a sequence of images which each contain less motion than an ungated image.

Although phase gating can be used to address respiratory motion a significant challenge arises due to the Poisson nature of PET imaging. Noise has a significant effect on the identification of anatomical landmarks and the distinction of organ and disease boundaries compromising the interpretation of gated PET images. The noise limitations of gating can be overcome by increasing the acquisition duration however, long duration scans may not be desirable in circumstances where a high patient throughput is required in the PET clinic.

Depending on patient breathing patterns, the quality of images within different gates may not be comparable due to different degrees of motion blur. For patients with irregular breathing patterns alternative methods of gating are necessary. Partitioning PET data into bins is nearly always implemented retrospectively, thus, alternative gating schemes can be applied in PET without the need for an additional scan.

One method of minimising the effect of irregular breathing is to partition events in a PET scan according to the amplitude of the respiratory trace. There are a variety of ways to
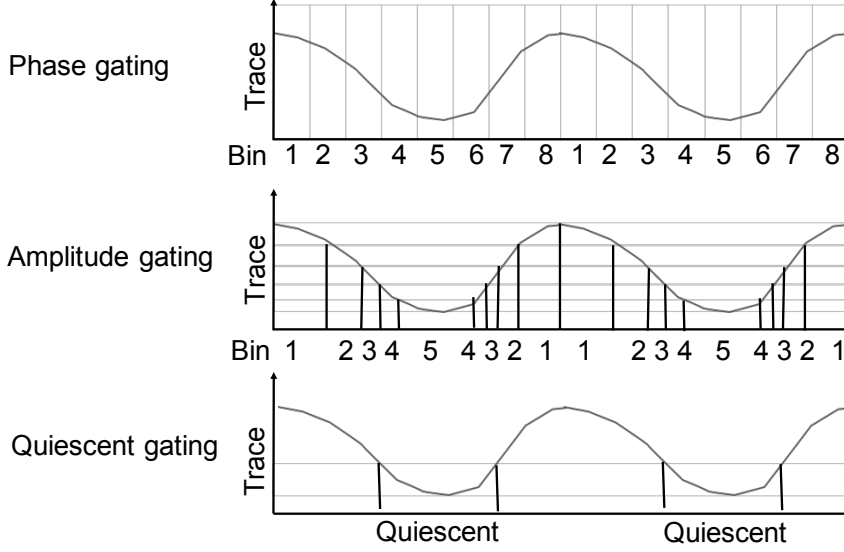
Figure 2.8: Three methods for partitioning the respiratory cycle using gating. The time intervals used for partitioning the list mode data are indicated by vertical lines. The bin numbers are shown on the time axis (horizontal)

gate depending on the trace amplitude cutoffs used for each bin (figure 2.8). Some methods of amplitude gating ensure that every frame has the same number of counts whereas other methods ensure that each frame has a similar degree of motion blur (Dawood et al., 2007). In respiratory gated PET, the optimal gating method is largely influenced by the trade off between the image noise due to limited number of events per frame versus the image blur due to motion within each frame. The effect of noise is so significant that if the number of counts drops below 106 counts then it is questionable whether motion can be interpreted regardless of the method used (Cloquet, Goldman, & Defrise, 2010).

The expiration phase of the respiratory cycle is typically a few seconds long and there is very little motion leading up to the point of end expiration. This phase is referred to as the quiescent part of the respiratory cycle. Quiescent gating is a variation of amplitude gating and is less sensitive to noise and irregular breathing (Liu et al., 2010). In this scheme only events during the quiescent part of the respiratory cycle are binned to obtain a single motion compensated frame (figure 2.8). Despite the benefits, it is not possible to characterise motion from a single frame.

2.2.3 Challenges with Acquiring a Respiratory Cine CT

There are many strategies for producing a 4DCT using a CT scanner in cine mode. The GE LightSpeed multi-slice CT scanner was an early commercial implementation of cine CT. In cine mode, the GE light speed steps through couch positions scanning the body repeatedly with multiple rotations for each table position. Cine mode captures enough image slices at each location to sample the entire respiratory cycle. Post-processing involves sorting CT slices so that the phase interval for each CT slice is similar for each gated frame.

4DCT scans are very susceptible to motion artefacts because they are typically of long

duration being acquired at multiple couch positions. Advances in the methods of sorting slices in cine CT images can reduce the susceptibility to motion. These sorting methods are dependent on the implementation of respiratory gating and are discussed in a later section. Early methods of avoiding motion artefacts in 4DCT were to acquire additional rotations and then discard slices which couldn't be accurately sorted into phase intervals. This involved acquiring 20 to 30 CT rotations per couch position (McClelland et al., 2006) discarding unmatched slices without compromising data sufficiency. State of the art CT scanners with larger axial field of view (FOV) also have fewer artefacts arising from incorrectly sorted slices. For example, a 128 slice CT has the capability to acquire a complete lung cine using only 3 or 4 couch positions. Although 4DCT can be used to capture a motion free image, the purpose of 4DCT is often to increase treatment margins to account for motion during external beam radiotherapy. 4DCT has also been used to track the motion of tumours due to respiration. The tracking can be used to assist in the delivery of highly conformal radiotherapy whereby the treatment field is dynamically adapted to account for tumour motion.

2.2.4 Challenges with Tracking Motion from Respiratory Gated Images

In 4DCT tumours are tracked by segmenting every frame in the gated sequence and can then be used to compensate for tumour motion. The same approach can be expanded to track any organ provided their features can be identified through image segmentation. Ideally, the segmentation should be fully automated so that it does not require a human operator to manually locate every image feature. However, automated segmentation is adversely affected due to confounding factors such as inadequate image sampling, excessive image noise and the presence of discontinuities caused by gating inaccuracies. There have been numerous developments in algorithms for automated segmentation of CT images (Heimann et al., 2009; Zheng, Barbu, Georgescu, Scheuering, & Comaniciu, 2008). The automated segmentation of the lungs is straightforward because of the stark difference in density between pulmonary vessels and the bronchiole tree. These soft tissue organs are generally more difficult to segment, even by human experts, due to the similar densities of adjacent tissues. Given the difficulties encountered by human experts, automated algorithms are unlikely to succeed. Despite developments in segmentation based tracking (Keall et al., 2005), fully automated approaches do not perform well in comparison to other methods of respiratory motion estimation (Murphy et al., 2010).

Although a segmentation approach may be suited to cine CT, confounding factors make it unlikely for the technique to be broadly applicable to respiratory gated PET. Simulations demonstrate (Segars, Sturgeon, Mendonca, Grimes, & Tsui, 2010) that the movement of surrounding objects such as the liver, heart and diaphragm are significant, however, the image quality of PET is such that the boundaries of some of these organs are not clearly visible (Mattes, Haynor, Vesselle, Lewellen, & Eubank, 2003). Given the difficulty for a human observer to identify anatomical objects from a PET image it is unlikely that a segmentation approach will be accurate. Alternative methods may be more robust to noise, sampling and image artefacts in gated PET.

2.3 Correction Methods for Attenuation in PET/CT

CT derived attenuation correction (CTAC) is now the most widely used correction method in PET imaging. Despite the wide use of CTAC, recent investigations into the effect of respiratory motion on attenuation correction have led to the development of several alternative methods. Several of these alternative methods to attenuation correction are incorporated into PET image reconstruction. Given the wide use of PET/CT, we have concentrated on methods which can utilise the combination of both CT and PET modalities. A complete study of all methods is beyond the scope of this review and can be found elsewhere (Berker & Li, 2016).

2.3.1 Challenges In CT Based Attenuation Correction

CT derived attenuation correction (CTAC) relies upon the accuracy of the bilinear scaling method for deriving attenuation coefficients from CT. The accuracy of CTAC can be adversely affected because low energy X-rays from CT (less than 80kV) are preferentially absorbed by objects in the body. This effect is referred to as beam hardening and can lead to inaccuracies in voxel wise tissue density from CT because log attenuation is no longer a linear function of object thickness. CT Reconstruction algorithms have been adapted to reduce the effect of beam hardening due to large or dense objects including bone and other highly dense materials such as metal implants. The accuracy of CTAC is also affected because the scan field of view (typically 50cm) does not extend to the full bore size of the PET/CT. The truncated parts of the patient can be accounted for by estimation of the tissue density of voxels which lie outside the scan field of view. Importantly the CT derived attenuation map will not be spatially aligned to the PET image. Spatial mismatch due to respiratory motion is arguably the largest source of inaccuracy in the thorax. Objects visible in the thorax are affected differently by motion in PET and CT. CT is a snapshot of the respiratory cycle whereas non-gated PET is an average image of the respiratory cycle. Respiratory motion can be mitigated in PET using respiratory gating, however, this does not address attenuation mismatch. Unfortunately, respiratory motion is complex and attenuation mismatch is difficult to mitigate because motion is continually occurring and altering during the PET/CT acquisition:

1. *During the CT scan.* Respiratory motion during the CT can lead to artefacts in the attenuation map.

2. *During the PET scan.* Respiratory motion during the gated PET scan leads to misalignment between the stationary attenuation map and moving objects in the gated PET frames. Some gated frames will suffer from increased attenuation mismatch whilst other frames will be better matched.

The propagation of CT artefacts into the attenuation map was a common problem with early model combined PET/CT scanners (Osman, Cohade, Nakamoto, & Wahl, 2003). Improvements in CT scanner speed have reduced the incidence of respiratory motion induced artefacts affecting attenuation correction. In addition to this, it has been shown that it is more accurate to acquire the spiral CT at end expiration to further reduce the chance

of motion artefacts appearing in the image (Killoran et al., 2011). Such respiratory motion artefacts can also be minimised using a deep inspiration breath hold (DIBH) protocol (Nehmeh et al., 2007). When artefacts do appear they tend to affect the PET quantification for the entire set of axial slices which contain the CT artefact, thus leading to the appearance of boundaries in the PET image where there are none (Zaidi, Montandon, & Alavi, 2007).

Although the DIBH technique addresses motion during the CT, it does not address motion between the PET and CT, nor during the PET. If the duration of the gated PET scan is extended, then it is possible to address attenuation mismatch by only sampling list mode events for which the CT and PET match. For example, if the CT was acquired at end expiration then it is feasible to create a single gated PET frame which corresponds to the same phase in the respiratory cycle (Liu et al., 2010). Unfortunately, this method is sensitive to changes in breathing pattern and it is difficult to predict in advance the timing and duration of gate intervals during the PET scan required to match the CT.

Respiratory gated PET typically bins events into 8-12 frames, meaning that each frame represents a 500ms - 1s segment of the respiratory cycle. Intra-gate motion during this time has many implications for attenuation as tumour and organ boundaries may shift and this is not reflected in CTAC. Although this effect can be mitigated by increasing the number of gated frames, intra-gate motion can be addressed using additional information such as a model of respiratory motion (Alessio, Kinahan, Champley, & Caldwell, 2010).

Given the multiple ways in which respiratory motion can affect PET/CT thoracic imaging, it is clear that there is no simple solution to CTAC misalignment. The remainder of this review is dedicated to methods to address this.

2.3.2 Improving the Accuracy of Attenuation Correction using Cine CT

The mismatch between PET and CT can be partially addressed by averaging all frames in a cine CT. By averaging over the complete respiratory cycle the cine CT derived attenuation map is a closer spatial match than CTAC applied to non-gated PET (Alessio et al., 2007). This has an advantage over acquiring CTAC as breath hold or free breathing as the result is highly dependent on the timing and speed of the CT acquisition. The radiation dose of a respiratory cine CT can be reduced by reducing the tube current to as low as 10mAs without compromising its usefulness for attenuation correction (Xia, Alessio, & Kinahan, 2009).

Cine CT can also be used to create a gated CT with the same number of frames as the gated PET. It has been shown that when using this acquisition protocol the breathing pattern was comparable enough between the gated PET and gated CT to enable it to be used for quantitative attenuation correction (Nehmeh et al., 2004). However, subsequent investigations have shown that low dose phase gated CT is even more prone to respiratory motion artefacts due to deficiencies in the sorting of CT slices (Lu, Parikh, Hubenschmidt, Bradley, & Low, 2006). It is difficult to avoid the problem that CT and PET are affected by motion in different ways.

Some studies have shown that there is little difference between phase matched 4DCT and respiratory averaged 4DCT (Killoran et al., 2011; Sun & Mok, 2012). Figure 2.9 shows phantom studies which have been used to quantify of attenuation correction. In the study published by Killoran et al. (2011), spiral CTAC acquired at end-expiration (EE) was com-
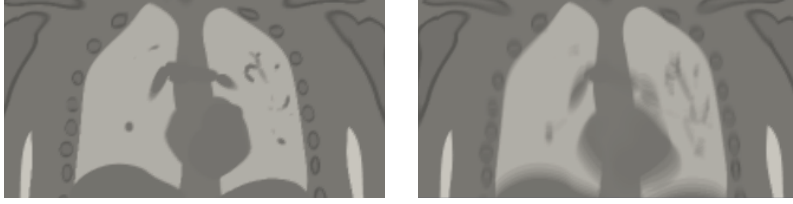
Figure 2.9: Coronal slices of attenuation maps in a simulation study using spiral CTAC (Left) acquired at end-expiration (EE) and respiratory averaged 4DCT (Right) similar to those conducted by Killoran et al. (2011)

parable to respiratory averaged 4DCT. This implies that methods such as quiescent gating may be as accurate as using cine CT (Liu et al., 2010). Partial volume effects may arise due to motion during the spiral CT acquisition, however, this is likely to be a secondary concern in comparison to attenuation mismatch provided the motion during the spiral CT is minimised. The partial volume effect in CT is directly related to scanner speed and would likely be present in very slow spiral CT scans such as those used in early cardiac PET/CT (Di Carli et al., 2007).

2.3.3 Natterer's Formulation of the Consistency Condition Equation

Without attenuation correction PET images reconstructed with iterative MLEM appear noisy due to an inconsistent system model (McQuaid, Lambrou, & Hutton, 2008). Although attenuation correction can be applied as a post reconstruction correction (Bai & Brady, 2011), this does not entirely eliminate the inconsistencies caused by attenuation mismatch. Numerous attempts have been made to quantify this inconsistency as well as to address attenuation correction and attenuation mismatch (Berker & Li, 2016).

PET image reconstruction involves the inverse of the Radon transform, where lines of response (LOR) are iteratively back projected into voxels. The Radon transform has symmetries in the absence of any attenuating objects. The same symmetries hold in the presence of attenuating objects provided accurate attenuation correction is applied. This dependency was exploited to develop a consistency condition for attenuation correction (Natterer, 1999).

It was subsequently proposed that Natterer's consistency condition can be used to address attenuation mismatch (Alessio et al., 2010). In this scheme, (equation 2.5) the attenuation correction factor, $e^{T(s,\phi)}$, for each line of response was used to correct measured sinogram data, $E(s,\phi)$, at each detector position, s, and view angle, ϕ. The consistency condition held for the integral of any combination moment, m, and frequency, k.

$$\int_0^{2\pi} \int_{-\infty}^{\infty} s^m e^{ik\phi} e^{T(s,\phi)} E(s,\phi) ds d\phi = 0, k > n \qquad (2.5)$$

At the beginning of this thesis work it was clear that the potential to exploit the consistency condition to inform the alignment of PET/CT had not been investigated thoroughly.

2.3.4 Reconstruction of Maximum Likelihood Activity and Attenuation

Estimating attenuation from PET sinogram data has been considered as a way to avoid the need for a separate acquisition for obtaining an attenuation map image (Nuyts et al., 1999). Maximum likelihood reconstruction of attenuation and activity (MLAA) was developed as a logical extension of consistency conditions incorporating activity and attenuation estimation into an iterative algorithm (Clinthorne, Fessler, Hutchins, & Rogers, 1991). MLAA is a constrained maximum likelihood approach and takes into account the Poisson nature of detected emissions, where the likelihood of emission, L, is dependent on emission activity, $\vec{\lambda}$, and attenuation, $\vec{\mu}$ (equation 2.6). Nuyts et al. (1999) shows that the maximum likelihood can be solved by an iterative algorithm in which each voxel in the estimate is forward projected along a line of response, i. At each iteration, the forward projection, p_i is compared to the measured events, E_i.

$$L(\vec{\lambda}, \vec{\mu}) = \sum_i (-p_i + E_i \ln(p_i)) \tag{2.6}$$

MLAA is implemented using a nested iterative solver, where the likelihood is maximised by alternating updates of attenuation and emission estimates. Although MLAA can be used to estimate the attenuation map, it relies on a non-zero $\frac{\partial L}{\partial \mu}$ in order to maximise L. This means that MLAA is underdetermined as can be demonstrated by the fact that many different activity distributions can lead to the same Radon transform (sinogram). Prior to the availability of TOF data, simulations showed that MLAA was not accurate for some cases of non-uniform attenuation (Nuyts et al., 1999). After the advent of combined PET/CT, MLAA was perceived as redundant because CTAC was perceived as superior in accuracy. However, it was soon found that CTAC mismatch artefacts had a serious effect and that perhaps MLAA could augment CTAC in some way. This led to a resurgence of interest in alternative methods for improving the estimation of attenuation from CTAC. For example, it was shown that MLAA could be used to address truncation of CTAC derived attenuation maps by estimating the extended field of view (eFOV) of the attenuation map using an adaptation of MLAA (Nuyts et al., 2013).

The shortcomings of MLAA reveal that without additional information assumptions are required for estimating attenuation.

2.3.5 Attenuation Correction using Consistency Conditions For TOF Imaging

Time of Flight (TOF) information localises the position of the annihilation event along the line of response (LOR), however, uncertainties of approximately 15cm (500ps) exist depending on the type and size of photo-scintillating crystals in the scanner (Daube-Witherspoon et al., 2006). The uncertainty in position along the LOR is parameterised using 1D Gaussian. A new consistency equation has been proposed which is an extension of Natterer's consistency equations (Panin, Defrise, & Casey, 2010)

$$t\frac{\partial p}{\partial s} + \frac{\partial p}{\partial \phi} - s\frac{\partial p}{\partial t} + \sigma^2 \frac{\partial^2 p}{\partial s \partial t} = 0 \qquad (2.7)$$

where p are the sinogram values in the absence of attenuation, s is the detector position, ϕ is the angle, t is the photon TOF and σ is the TOF time resolution. The last term in this equation is dependent on overall average attenuation and must be conserved. This consistency equation (2.7) is used to find the derivatives of the Radon transform of the attenuation map, $\frac{\partial \mathcal{R}(\vec{\mu})}{\partial s}$ and $\frac{\partial \mathcal{R}(\vec{\mu})}{\partial \phi}$. These derivatives can be integrated to estimate the ACF sinogram up to a constant (Defrise, Rezaei, & Nuyts, 2012). Although this formulation allowed for the estimation of the ACF factors from sinogram data, it was extremely sensitive to noise (Defrise et al., 2012) and it was suspected that iterative algorithms will lead to significantly superior results.

The availability of time-of-flight (TOF) information enabled an estimation of attenuation correction factors (ACF) up to a constant and the same approach could be extended to address the shortcomings of MLAA (Defrise et al., 2012). MLAA was underdetermined leading to a cross-talk problem where it was not possible to maximise the more likely activity and attenuation (Rezaei et al., 2012). The underdetermined nature of MLAA can also be observed through symmetries in the associated Fisher information matrix (Rezaei et al., 2012).

Maximum Likelihood Reconstruction for TOF-PET with simultaneous estimation of the attenuation factors (MLACF) was developed as an iterative algorithm which uses TOF information to find both the attenuation and emission (Rezaei, Defrise, & Nuyts, 2014). MLACF uses the TOF consistency equations to obtain a ML solution thereby stabilising the joint estimation problem (Rezaei et al., 2014).

Although MLACF is promising, two significant problems stem from the inverse relationship between the emission and attenuation. The first problem is that the attenuation factor is currently only estimated up to a constant, and the overall attenuation factor has to be obtained using an alternative method. The second problem is that LORs with no uptake have no emission signal and thus attenuation cannot be estimated along that LOR. This can be mitigated by excluding low count regions of the sinogram prior to reconstruction.

An advantage of MLACF over CTAC is that even if the PET image is compromised by motion, the MLACF is affected in the same way and there is no need to correct for misalignment. Although MLACF is promising in that it may enable quantitative PET imaging without relying on transmission scans or other anatomical imaging procedures, it is suboptimal to ignore the information obtained from CTAC in a combined PET/CT scanner.

Whilst it would be interesting to further investigate the benefits of TOF, this review indicates that methods of misalignment require a robust method of parameterising respiratory motion. It is also worthwhile considering what information can be obtained about misalignment of the CTAC due to respiratory motion. These issues are covered in the following sections.

2.4 Image Registration Applied to Respiratory Gated PET/CT

There are numerous methods for estimating respiratory motion in both PET and CT. For example, motion has been estimated from the cine CT image data itself by tracking the movement of features during the CT scan (Han, 2010). Data driven methods such as these typically use image registration as a means of finding the correspondence between frames in the cine CT (Ehrhardt, Lorenz, & others, 2013). Although image registration has traditionally been used to align pairs of images it has been extended to sequences of images for the construction of respiratory motion models in both PET and CT. In this case the goal of registration is to estimate the transformations which correspond to the respiratory motion being observed. Several kinds of transforms have been proposed as suitable for modelling respiratory motion. There are also two kinds of image registration techniques, feature-based and intensity-based. In feature-based registration, the similarity between images is determined through the alignment of segmented features in both the fixed and moving image. In intensity-based registration, the similarity between image is determined through voxelwise correspondences between the fixed and moving images. This investigation focused on intensity-based registration because segmenting features in PET images is challenging. Typically only solid large tumours such as single pulmonary nodules (SPN) are the only identifiable features that can be reliably identified and tracked in PET. Thus feature-based registration offers very little for determining the motion of organs other than the tumour.

2.4.1 Objective Function

In intensity-based registration the objective function is a voxelwise measure of similarity between a source and target image. The transformation, $T(\theta)$, is often found by maximising the similarity, S, between the fixed image, I_{Fixed}, and the transformed moving image, I_{Moving} (figure 2.10). Registration is often formulated as the minimisation of an objective function as shown in equation 2.8.

$$\arg\min -S(I_{Fixed}, T(\theta, I_{Moving})) \qquad (2.8)$$

Objective functions based on the direct correlation between voxels are suitable when registering PET to PET and CT to CT. In such cases the transform, $T(\theta)$, is an estimate of the respiratory movement from one frame to the next and is found by minimising the sum of squared differences or correlation coefficient between sequential frames. Direct correlation metrics such as SSD require images with enough counts to distinguish between the variance caused by respiratory movement and background Poisson noise (Huang et al., 2011). Noise corruption from individual frames can be addressed by extending the SSD objective to simultaneously calculate the SSD for all frames using group-wise registration. Cross correlation and correlation coefficient metrics are useful when the appearance of objects changes during the transform. This may be the case in long duration studies where tracer uptake and distribution is changing, however this is not significant for short duration PET studies whereby any effect would be averaged out over many respiratory cycles. Aligning CT to PET images requires information theoretic approaches, where the relationship between grey scale
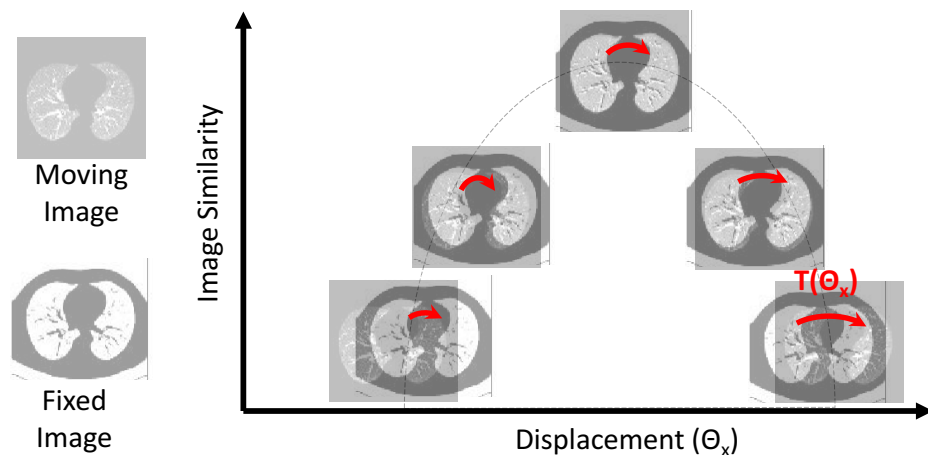
Figure 2.10: An ideal objective function is a smoothly varying function for which a minimum can be reliably estimated using image registration. Many different objective functions have been proposed which are suitable for measuring similarity given the characteristics of PET images and CT images. A summary of approaches is shown in the table 2.4.

Table 2.4: A list of objective functions used for measuring image similarity in image registration.

Direct correlation	
PET to PET	Correlation Coefficient: Bai and Brady (2009a)
	Normalised Correlation: Thorndyke and Xing (2006)
CT to CT	SSD: Vandemeulebroucke, Rit, Kybic, Clarysse, and Sarrut (2011)
PET to CT	*n/a*
Histogram based	
PET to PET	*n/a*
CT to CT	Mutual Information: McClelland et al. (2006)
PET to CT	Mutual Information: Mattes et al. (2003)

values in the PET and grey scale values in the CT are partial (in some cases independent) and non-linear. The relationship between grayscale values in target and source image can be expressed using a joint image histogram between target and source images. This joint histogram is then used to calculate the normalised mutual information (NMI). The greater the NMI, the more significant the relationship between target and source images. This approach has been used as a metric in the past to identify attenuation mismatch (Dawood, Buther, Lang, Jiang, & Schafers, 2006). Image-based registration using NMI is more robust than feature-based image registration because there is no need to track features in the PET image and find corresponding features in the CT. A drawback of image-based registration is that NMI and other histogram based objective functions can be corrupted by noise. The presence of noise leads to local maxima of the objective function confounding the estimate of the transform.

2.4.2   Affine Transforms

Affine transforms are often referred to as global transforms because all voxels are transformed uniformly at all points in the image. A 3D dimensional affine transform allows for rotation, translation and zooming and is described by 9 parameters. Even if the misalignment is non-rigid, an affine transform is often performed as a first step in image registration. In this case affine is a first order approximation of the transform which can be further refined by introducing a deformation model with a large number of parameters.

For intra-modality registration, affine registration has been demonstrated as robust for brains and has been used for motion compensation in PET by realigning the reconstructed PET frames and robust to noise allowing for registration of frames with a duration as short as 30 seconds (Costes et al., 2009). The motion of semi-rigid objects can be tracked using affine transforms and this is useful for localised corrections of motion of the heart (Livieratos et al., 2005) and tumours (Liu et al., 2011). Affine transforms can be applied to reconstructed gated images provided that there is approximately 30 seconds worth of data per sinogram. Alternatively, the transforms can be applied to detected events prior to reconstruction using event by event motion correction without the need for gating (Liu et al., 2011).

For inter-modality registration affine transforms have been used to address attenuation mismatch in the lung. Although, affine transforms are inadequate to capture multiple spatially non-uniform movements, the simple rigid-body approximation was found to account for many types of PET-CT mismatch errors (Alessio et al., 2010). Affine registration can also be incorporated into motion correction of PET scans to align dynamic frames during PET image reconstruction (Verhaeghe et al., 2010). Despite the suitability of affine registration to PET rigid motion correction, it is not able to capture the complex nature of respiratory motion (Lamare et al., 2004). More recent work modelling the lungs expanding and contracting alongside the heart and diaphragm has shown that non-rigid deformations are necessary to address PET-CT mismatch errors (McQuaid, Lambrou, & Hutton, 2011).

2.4.3   Non-Rigid Transforms

Non-rigid image registration requires an objective function, a deformation model and an optimiser to match the target and source images. The deformation model is a transform of

every point in the source image to a location in the target image. The most suitable choice of deformation model is dependent on image quality, tissue properties and physiological variations of organs being imaged. Rather than reviewing all possible representations of the deformation model, this thesis focuses on deformation models which have been widely used for respiratory imaging. A more comprehensive summary of geometric transformations for non-rigid body registration can be found in recent literature such as an excellent review by Holden (2008). Expressing deformation as an independent transform for every voxel in the image requires a large number of parameters, in fact it requires at least 3 transform parameters per voxel in 3 dimensions. Such a scheme is referred to as a local representation because it allows for arbitrary local transforms. Alternatively, the deformation can be expressed as compact function with less degrees of freedom by assuming that the transform is smooth and piecewise within the spatial neighborhood (Buerger, Schaeffter, & King, 2011). The transform of each voxel is represented using a function that is confined to the semi-local area. Transforms need to be resampled to the voxel resolution in order to deform every voxel in the image.

Each deformation model can only be used with a limited set of optimisers and objective functions. For example, inter-modality registration for PET/CT often uses a mutual information cost function, however, there is no straightforward way to implement this cost function for some deformation models (Holden, 2008). In order to succinctly compare deformation models this review only considers several registration algorithms identified in literature (McClelland, Hawkes, Schaeffter, & King, 2013) as applicable to respiratory motion.

2.4.4 Regularisation

Even the most suitable choice of deformation model and objective function in a registration algorithm may not produce the correct deformation. The reason for this is that image registration is an ill-posed inverse problem and it is impossible to distinguish the correct deformation from the other deformations which also minimise the objective function. Sensible limits to the solution of the inverse problem can be imposed through a regularisation penalty term.

Such regularisation can be introduced into image registration by adding a penalty term, R, to equation 2.9.

$$\arg\min -S(I_{Fixed}, T(\theta, I_{Moving})) + R(\theta) \qquad (2.9)$$

A straightforward approach to regularisation is to set R as a least squares penalty term, $\beta||T(\theta)||^2$, known as Tikhonov or L2 regularisation. L2 regularisation allows the image registration to find the smallest deformation that also satisfies minimising the objective function. The amount of regularisation can be tuned with the β which forces the transformation model, T, to stay within sensible limits. Regularisation is also a way to incorporate prior assumptions into the registration framework such as favouring spatially smooth transforms. For example, schemes which model deformations using a local representation often require the use of a Gaussian filter to enforce spatially smooth transforms. In fact, any mathematical or physical property can be penalised, provided a penalty can be devised. The spatially smooth assumption has implications in the kind of deformation field produced and this issue is revisited at the end of this chapter in light of attempts to improve regularisation.

An alternative form of regularisation is to simply avoid capturing the local variation of the deformation field. This can be achieved by limiting the number of parameters used in the deformation model by reducing the resolution of the field. Although this is a suitable method for obtaining an initial estimate of the deformation, there are numerous possible constraints presented later in this chapter that provide an improved model of respiratory motion.

2.4.5 Implementations of Deformable Registration

There are many different methods of implementing deformable registration. This section reviews four deformation models which have been used to perform deformable registration of lung images. These have been chosen based on their usefulness for constructing more elaborate respiratory motion models.

Hierarchical Affine

A straightforward extension of affine registration is to model deformation as different affine transforms at different points in the image space. Hierarchical affine registration decomposes the transform into multiple locally affine transforms. Affine registration is used to find these local transforms which are subsequently combined to form an overall non-rigid deformation. Hierarchical affine registration has been shown to be suitable for estimating respiratory motion from MRI sequences (Buerger et al., 2011). In this scheme a coarse representation of respiratory motion is initially modelled by registering large sections of the image assuming the motion of the section can be described by an affine transform. The image being registered is successively split into smaller rectangular blocks, each associated with an affine registration. Although, this method is highly computationally efficient, the alignment of fine image structures may not be accurate (Buerger et al., 2011)

Demons Method

Another approach of image registration is to use the differences in voxel grayscale values to estimate an imaginary force which would align the source image, s, with the moving image, m (equation 2.10). This method is reasonably straightforward to implement using the gradient of the moving image, ∇m, and the voxelwise difference between images, $s-m$. Each iteration of the algorithm allows each voxel to move one pixel in the direction indicated by $\vec{f}_m$ (Wang et al., 2005).

$$\vec{f}_m = \frac{(s-m)\vec{\nabla}m}{\left|\vec{\nabla}m\right|^2 + (s-m)^2} \tag{2.10}$$

One of the problems with this method is that it only works within a single modality. This has not been typically applied to PET because of the sensitivity to noise. The algorithm is quite sensitive to spatial variations in image counts due to Poisson noise.

The demons method of image registration is adaptable by modifications to the 'active' force equation 2.10. An alternative formulation of equation 2.10 is to define a 'passive' force for which the direction of force is determined from ∇s. A recent review found that using a combination of both forces is the more accurate method (Wang et al., 2005). Demons

registration has remained popular for respiratory imaging in 4DCT because only a small amount of deformation smoothing is required to overcome image noise.

Optical Flow

The optical flow registration method is similar to the demons method in the sense that it models highly local voxelwise transforms. Optical flow is based on a diffusion equation that governs the transport of grayscale values, I, of an image undergoing registration (Dawood, Buther, Jiang, & Schafers, 2008). Images are registered by finding the optical flow velocity, $\vec{v}$, that satisfies equation 2.11.

$$\nabla I \cdot \vec{v} + \frac{\partial I}{\partial t} = 0 \tag{2.11}$$

Minimising the LHS of this equation is an ill posed problem, thus additional regularisation is applied through the use of smoothness constraints (Dawood et al., 2010). Conventional optical flow does not preserve counts upon transformation. In order to preserve counts an additional constraint is added (Dawood et al., 2010). Constraining grayscale values to satisfy equation 2.12 preserves counts even in the presence of the partial volume effect.

$$\nabla \cdot (I\vec{v}) + \frac{\partial I}{\partial t} = 0 \tag{2.12}$$

Optical flow registration has been used to estimate respiratory motion from both simulated and clinical PET scans (Dawood et al., 2008). A significant limitation of the optical flow method is that it cannot be applied to multimodality registration because it relies upon local differences in grayscale values. It is also not straightforward to incorporate prior information about respiration unless the constraint can be expressed in terms of the variables in Equation 2.11 and 2.12.

Free Form Deformation

Free form deformation (FFD) image registration is a versatile approach where basis functions restrict the localised properties of deformations (Holden, 2008). Semi-localised models support basis functions which assume that deformation is smooth and continuous in between control points. For example, sparsely distributed control points, P_{ijk}, are interpolated using cubic B-Splines to estimate and the transform, T, at intervening voxel positions, x, y, z where ψ is the B-Spline basis function at each control point (equation 2.13).

$$T(x, y, z) = \sum_{i,j,k} \psi_i(x)\psi_j(y)\psi_k(z)P_{i,j,k} \tag{2.13}$$

It has been shown that semi-localised models of deformation encourage free form registration to converge to the correct solution (Rueckert, Aljabar, Heckemann, Hajnal, & Hammers, 2006). FFD allows for arbitrarily large deformations, making it suitable in a broad number of circumstances and imaging modalities.

2.4.6 Improvements to Image Registration

Mass Preserving Image Registration

Non-rigid registration in CT lung images can be improved by accounting for local tissue volume changes. Hounsfield Units (HU) are used to measure radiographic-density, therefore, if tissue changes its volume as a result of deformation, then the HU must change, provided that tissue mass is conserved. Changes in density are contributed by two factors. The largest is the change in the regional volume due to deformation and a smaller change is the change in the fractional tissue content within a region due to respiration (Yin, Hoffman, & Lin, 2009). Yin et al. (2009) assumed that the HU consists of two components: tissue, which includes parenchyma and blood, with approximately 55HU, and air with -1000 HU. Any unit of lung volume may be considered a linear combination of these two compartments.

The accuracy of registered PET images is also affected by local tissue volume changes in the lung. These changes can be accounted for by using a method to preserve the total activity in the image. Changes in the volume of a voxel must be accompanied by a change in activity density. The changes in activity due to voxel size can then be corrected for after image registration by finding the amount of stretching and shrinking due to deformation. This change can be estimated from the determinant of the Jacobean of the deformation transform (Thielemans, Asma, & Manjeshwar, 2009). Activity can be preserved by multiplying by the Jacobean at each voxel in the image. Some studies have shown that incorporating activity preservation into the image registration algorithm has the additional benefit of avoiding mis-registration due to the partial volume effect (Dawood et al., 2010). Both free form deformation (Thielemans et al., 2009) and optical flow image registration (Dawood et al., 2010) have been adapted to incorporate activity preservation.

Groupwise Registration

Although it is possible to separately register each frame in a gated acquisition, it is useful to draw upon the similarities of deformations across sequential frames. A groupwise approach to registration simultaneously aligns all frames using a combined registration algorithm which finds the set of all transforms. A groupwise approach allows the use of more robust deformation models and objective functions. The voxelwise variance objective function has been used to visualise the effect of motion for short duration PET scans with a low signal to noise (Huang et al., 2011). In equation 2.14, the voxelwise variance, sigma, is the squared of difference between the voxel in each frame, I_f and the average over all frames, $\bar{I}$.

$$\sigma_i^2 = \sum_f (I_{i,f} - \bar{I}_i)^2 \tag{2.14}$$

In practice, groupwise registration can be accomplished by iteratively aligning all frames to the running average of all registered frames (Costes et al., 2009), thereby minimising the voxelwise variance. Although this form of registration may require significantly more computing resources, it is more robust to noise.

Spatiotemporal Registration

Spatiotemporal registration is a method of modelling both the spatial and temporal properties of a transform for a sequence of dynamic or gated frames. This can be accomplished by parameterising the transform as a function of both space and time (Metz, Klein, Schaap, van Walsum, & Niessen, 2011). The spatiotemporal approach has been reported to be more robust to error caused by CT artefacts in individual frames as well as robust to the influence of noisy or corrupted frames on registration accuracy. Metz et al extended the spatiotemporal approach into a groupwise registration algorithm to jointly parameterise both spatial and time variations of the transform. The trajectory of a point in the spatiotemporal transform can be expressed as either a velocity or a displacement. It is more common to express the trajectory as a displacement from a reference point and this is known as the Lagrangian approach. For example, Castillo, Castillo, Martinez, Shenoy, and Guerrero (2009) used a polynomial to describe the displacement trajectory over time. Vandemeulebroucke et al. (2011) demonstrated a trajectory model based on cubic splines imposing piecewise smooth constraints over the respiratory period. Alternatively, the trajectory of a point can be expressed as the instantaneous velocity between time points, however, this is typically more difficult to parameterise in a way suitable for group-wise registration. Spatiotemporal registration requires fewer temporal parameters than separate pairwise registration for every frame. For example, a study of spatiotemporal registration of frames from thoracic 4DCT scans required only two time points to adequately estimate the trajectory of each point (Metz et al., 2011). The reduction in the number of parameters in spatiotemporal registration increases the robustness to noise. However, fewer parameters can pose challenges with characterising sliding motion, especially at the pleural wall (Ruan, 2009). Inaccurate modelling of sliding motion at the lung boundaries can be partially addressed using a lung mask (Vandemeulebroucke et al., 2011).

Rigidity Based Constraints

An alternative approach to increasing the accuracy of registration is to introduce a soft constraint through regularisation. Regularisation is accomplished such that increased weighting of the penalty term favours local affine transforms in regions where the organ may be considered rigid (e.g bone) (Staring, Klein, & Pluim, 2007). The same constraint also ensures that registration satisfies orthonormality conditions (robust to noise) and suppresses folding in the deformation transform (properness). The usefulness of this constraint is demonstrated through the way it prevents non-physical transforms such as the ribs being disconnected or twisted during the registration process (Ruan, Fessler, Roberson, Balter, & Kessler, 2006).

There are many ways to introduce a rigidity based constraint into image registration. For example, in CT image registration a local rigidity penalty term has been incorporated into registration in order to penalise the deformation of segmented rigid objects (Staring et al., 2007). Another method of enforcing local rigidity is to calculate a rigidity index from the Jacobean determinant. This voxelwise rigidity index is then multiplied by a stiffness weighting factor derived from the Hounsfield units providing a substantial stiffness weighting for bone ($HU > 250$). Equation 2.15 describes the penalty term, R, to be minimised as a function of voxelwise intensity, I, and the Jacobean determinant of the transform, $|J(T)|$.

$$R = I \times (|J(T)|^2 - 1) \qquad (2.15)$$

These approaches are yet to be fully validated, however, preliminary results show that the constraint prevents non-physiological (or impossible) deformations (Ruan et al., 2006). Rigidity based constraints may be especially important for registration of solid lesions (Yip, Chen, Aerts, & Berbeco, 2014).

2.5 Parametric Models for Respiratory Motion

2.5.1 Sources of Variation in Respiratory Motion

Respiration is the action of breathing in (inspiration) and out (expiration) whereby outside air travels in and out of the lungs due to a change in pressure in the thorax. During inspiration, the intercostal muscles (between the ribs) and diaphragm contract to expand the chest cavity (Ehrhardt et al., 2013). The action of breathing is irregular whereby the volume of air which fills the lungs may vary from cycle to cycle. The volume of air during inspiration is the most prominent characteristic of respiration (Ehrhardt, Werner, Schmidt-Richberg, & Handels, 2011). Accurate motion correction requires parameterisation of the deformation of organs such as the rise and fall of the diaphragm. Although breathing is irregular, parametric models can be extended to encompass the variables associated with irregular breathing (e.g. breathing with a deeper amplitude). This requires additional surrogate measurements whose amplitude is related to the volume of air in the lungs, such as the position of the chest surface. Although useful, surrogate measurements are not necessarily accurate because the action of breathing does not instantaneously fill the lungs with air as there is a time lag between diaphragm contraction and when alveoli expand and fill (Kyriakou & McKenzie, 2011). This leads to a phase difference between the external surrogate measurements and the internal lung volume.

Irregularities in breathing can confound attempts to register PET and CT if there are long periods of shallow breathing in the PET but deep inspiration breath hold for the CT. A promising way to ensure that the breathing pattern is the same during the PET and CT scan is to train the patient to breathe consistently using biofeedback (George et al., 2006).

Despite there being multiple sources of variation in respiratory phase and amplitude, several studies have shown that motion can be adequately approximated with remarkably simple parametric models of respiratory motion (McClelland et al., 2013).

2.5.2 A Surrogate Model for Respiratory Motion

A direct correspondence model establishes internal deformation as a function of an external surrogate measurement. For example, a direct correspondence model can relate the observed internal motion in a patient's 4DCT scan to the external movement of a respiratory gauge (Q. Zhang et al., 2007). A direct correspondence approach establishes a relation between coefficient values and an external surrogate measurement. Once a correspondence has been established, it is then possible to extrapolate the relation to predict motion states based on the external surrogate measurement alone. For example, it has been demonstrated that it is

possible to establish a direct correspondence model between the motion observed in a 4DCT scan and the respiratory trace (McClelland et al., 2006). Direct correspondence models may be appropriate for external beam radiotherapy due to the increased use of 4DCT data and surrogate measurements to estimate variations in tumour position during respiration (Keall et al., 2005).

2.5.3 Statistical Models for Respiratory Motion

A further alternative approach to estimating deformations due to respiratory motion is to consider a statistical model of respiration. A statistical motion model is more sophisticated than a surrogate model because it can adapt to individual breathing patterns observed in respiratory imaging such as in 4DCT (R. Li et al., 2011). A statistical model is also able to draw upon both temporal and spatial correlations in the deformation in addition to using smooth basis functions such as B-Splines (Q. Zhang et al., 2007) and polynomials (King, Buerger, Tsoumpas, Marsden, & Schaeffter, 2012).

Patient Specific Modeling

Although there are several techniques for parameterising the time evolution of respiratory motion, an increasingly popular alternative is to use a motion model based on principal components analysis (PCA) (R. Li et al., 2011). The PCA approach has been favoured because it has been shown to adapt to patient specific variations better than a direct correspondence model. PCA can also be implemented as a post-processing step without the need to modify the image registration algorithm. Several authors have demonstrated the use of PCA to find a correspondence between external surrogates and internal deformation (King et al., 2012; R. Li et al., 2011; Q. Zhang et al., 2010). This has wide uses for respiratory motion correction in medical imaging and tumour tracking in external beam radiation therapy.

The application of PCA to make predictions about respiratory motion is an example of machine learning. In this case, the PCA model is supplied with a set of patient specific training samples to isolate commonalities in lung displacement within a sequence of gated frames. The machine learning algorithm can subsequently be used to make a prediction about respiratory motion which is yet to be observed. PCA motion modeling has been successfully used to address attenuation mismatch in gated-PET/CT. In this case, the training samples are derived from a 4DCT scan of the patient and are used to make a prediction about respiratory motion during the gated PET scan. The advantage of machine learning algorithms is that unlike current methods for motion correction (Bai & Brady, 2011) they can predict motion which cannot be otherwise estimated from acquired data (Ehrhardt et al., 2011).

Equations 2.16, 2.17 and 2.18 describe how a patient specific PCA model is computed via the Eigen-decomposition of displacement vectors corresponding to lung motion. The lung motion states are expressed as a matrix, X, of which each row represents the displacement vectors for a particular time sample.

$$X^T = [d_1, d_2, d_3, \ldots d_r] \tag{2.16}$$

The displacement vectors. d, for every voxel in the lung at all time points $(1 \ldots r)$ are obtained from three dimensional deformable image registration. The number of rows in

matrix X is the number of voxels in the lung times three, and the number of columns is the number of sampling points in time. PCA decomposition is performed on matrix X giving rise to a new approximation.

$$X^T \approx [\hat{d}_1, \hat{d}_2, \hat{d}_3, \ldots \hat{d}_r] \tag{2.17}$$

where each motion state, $\hat{d}_t$, is a linear combination of eigenvectors, u_k, with a coefficients vector, v_t, (Q. Zhang et al., 2007).

$$\hat{d}_t = \sum_{k=1}^{r} u_k \cdot v_t \tag{2.18}$$

It is common to approximate lung motion using the first two eigenvectors (R. Li et al., 2011) and this has been found to be adequate to accurately describe respiratory motion in 4DCT scans.

Rather than considering correspondence at each time point individually, a PCA motion model can be improved by taking into account the history of previous surrogate measurements. A Bayesian filter has been incorporated (Smith, Rahni, Jones, & Wells, 2012) into a PCA motion model by estimating the probability density of the relation between PCA coefficients and external surrogates. The Bayesian filter approach may be useful for estimating motion from surface measurements during a PET scan (Abd. Rahni, Lewis, Wells, Guy, & Goswami, 2010) due to its ability to adapt to irregular breathing (Smith et al., 2012).

Indirect Correspondence Models

An indirect correspondence model involves adapting the model in some way to optimise the relationship between either image data or external surrogates and internal deformation. For example a PCA motion model can be optimised by adjusting the PCA weights to improve correspondence. One way to validate such an optimisation is to check that the predicted motion is consistent with the images acquired (figure 2.11). Optimisation of a correspondence model was demonstrated by King et al. (2012). A sequence of 3D-MRI images from a PET-MRI scanner was used to construct a PCA respiratory motion model. A PET-MRI scanner can also simultaneously acquire a time sequence of navigator images, of which each image has a small number of spatial samples in 2 dimensions. King et al utilised the navigator image as a surrogate to find the PCA weights in the motion model and establish a correspondence model between internal deformation and the 2-D navigator. The same correspondence model was then tested with the original 3D-MRI image sequences and the model was continually adapted in an iterative fashion. The advantage of this approach was that the correspondence model was checked for consistency in making motion predictions from the observed surrogate data.

The indirect correspondence approach can be used to eliminate the need for surrogate signals altogether. For example it is possible to construct a PCA motion model in the absence of surrogates and then adapt the motion model to be consistent with another set of acquired images. An example of this is modifying a PCA motion model to be consistent with a sequence of respiratory gated images acquired from a PET scanner. Whilst the PET images may be unsuitably noisy for constructing a motion model, they can be used to adapt
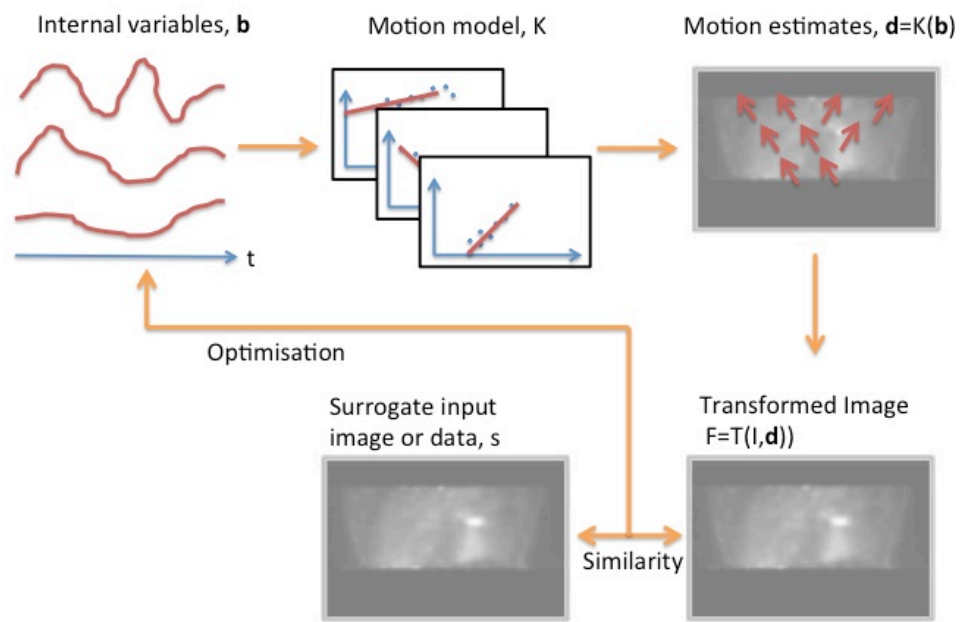
Figure 2.11: A flowchart representation of an indirect correspondence model (Adapted from McClelland et al. (2013))
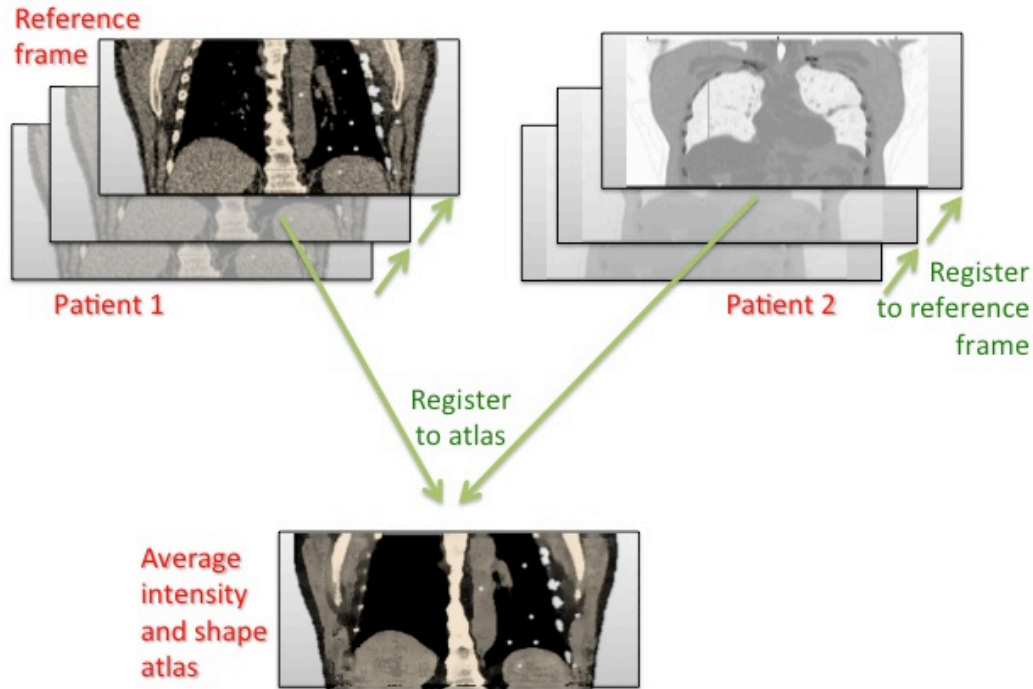
Figure 2.12: Calculation of a mean image which can serve as an atlas in a cross population model (Adapted from Ehrhardt et al. (2011)).

an existing model. A recently published approach constructed a PCA motion model from MR derived motion fields and then adapted the model to be consistent with acquired PET images (Balfour, Polycarpou, Marsden, & King, 2013). Balfour et al. (2013) adjusted the PCA weights to maximise the expectation of the correspondence of the PCA motion model with the PET images.

Cross-Population Models

Cross-population motion models are formed from motion data acquired from many different subjects, and attempt to capture the nature of breathing motion across the population (McClelland et al., 2013). These cross-population models are also referred to as global models (H. Fayad, Clement, et al., 2009).

Unlike patient specific models, cross population models need to also account for the anatomical variation between subjects. This is usually done by registering each subject's image to a population mean image (figure 2.12). The mean image acts as a global atlas and the cross population model parameterises motion at each spatial position in the atlas.

A simple example of the use of a cross-population model is to establish a mean velocity field and assume that the amplitude of motion can be found by scaling the velocity field by an exponential function (Ehrhardt et al., 2011). This model is only approximate because it assumes constant velocity of motion whereas the velocity field changes magnitude as the lungs fill and empty. This limitation partially addressed by Ehrhardt et al. (2011) by calculating

separate velocity fields for end expiration, end inspiration and mid inspiration.

Another type of cross-population model is the statistical shape model. Statistical shape models have a demonstrated use in image segmentation. Such a shape model can be used to refine the boundaries of a segmented region. Shape models have also been used to analyse organ motion. For example, statistical motion model of the heart was created and then used to characterise the most probable deformation of heart tissue during the cardiac cycle. It was then possible to establish a time-dependent deformation field which characterised heart motion of a set of subjects in a population.

Statistical shape models utilise eigenvectors which are derived from statistical analysis of the displacement fields of multiple subjects drawn from a population. A straightforward application of the shape model is to express motion as a linear combination of these component eigenvectors. In the statistical motion model proposed by Chandrashekara, Rao, Sanchez-Ortiz, Mohiaddin, and Rueckert (2003), the direction of the eigenvectors describes the direction of heart motion, whilst the scaling coefficients parameterise the amplitude of motion in each direction. Chandrashekara et al. (2003) use the following equation (2.19) to create a time dependent deformation field, $d^S_{(0,t)}$

$$d^S_{(0,t)}(\dot{x}, \dot{y}, \dot{z}) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u)B_m(v)B_n(w)(\hat{\mathbf{C}} + \mathbf{\Phi b})_{i+l,j+m,k+n} \tag{2.19}$$

where the B-Spline coefficients, $B_l(u), B_m(v) and B_n(w)$ are used to interpolate between positions where the displacement model has been calculated. The initial position of the control points in the displacement model is $\hat{\mathbf{C}}$. The variation in motion within the population is arranged into a matrix of component eigenvectors $\mathbf{\Phi}$. A new set of scaling coefficients, $\mathbf{b}$ is then approximated using a gradient descent solver. The practical advantage of this model is that when a new patient scan is encountered it is only necessary to calculate a new set of scaling coefficients, b^{new}. i.e. this can then be used to build a new correspondence model, $\mathcal{K}$

$$d^{new}_{(0,t)}(\dot{x}, \dot{y}, \dot{z}) = \mathcal{K}(\mathbf{b^{new}}) \tag{2.20}$$

where d^{new} is the new deformation.

McQuaid et al. (2008) proposed that a population model could be used to approximate the shape of the diaphragm in the lung. The population model was created by the analysis of diaphragm shape observed in high resolution images of several patients. The diaphragm vertical position, $s_z(x,y)$, was expressed as a linear combination of component eigenvectors as shown in 2.21.

$$s_z(x,y) = \sum_{k=1}^{r} b_k \cdot u_k(x,y) \tag{2.21}$$

The model was then used to estimate the shape of the diaphragm for patients who were not used to create the population model. This was accomplished by estimating the scaling coefficients, b_k, for each new patient at each time point in the respiratory cycle. Their work presented a method of estimating the coefficients, and analysed the accuracy of the method for tracking the rise and fall of the diaphragm with respiration. The rise and fall of the

diaphragm is the main contributor to attenuation correction inaccuracies when radionuclide-imaging the myocardium. Using the model to implement motion compensation had a clinical significance, even though the expansion and contraction of the chest wall was not modeled.

The PCA approach to a correspondence model exploits the relation between externally measured signals and internal organ deformation. Imaging the deformation of internal organs is challenging, thus the idea of exploiting this correspondence in a cross population model is attractive. H. Fayad, Pan, et al. (2009), proposed a relation between the surface deformation as measured by an optical time of flight (TOF) camera and the internal motion described by a deformation field. A matrix X was constructed in which each row was the concatenation of internal deformation, d, and external surface measurements, s.

$$X^T = [d_1|s_1, d_2|s_2, d_3|s_3, \ldots d_r|s_r] \tag{2.22}$$

X was then decomposed into a linear combination of component eigenvectors, u_k, such that the displacement field, d_p, and surface map, s_p for every patient could be expressed as

$$\hat{d}_p|\hat{s}_p = \sum_{k=1}^{r} u_k \cdot c_{k,p} \tag{2.23}$$

H. Fayad, Pan, et al. (2009) considered cases when the external surface measurements, s, for a patient, p, were known, but the internal deformation, d and scaling coefficients, c_p, were unknown. They manipulated the eigen-decomposition method to infer the internal deformation, d, from the observed surface motion. A limitation of this model is that it relies on very accurate surface tracking measurements and assumes that the internal-external correlation holds true for all patients in the population model, which may not be a valid assumption.

A commonality between cross population models is that they parameterise deformation such that it can be expressed as a function of an external or internal measurement. It even appears possible to construct a cross-population model using a single parameter (Ehrhardt et al., 2011), correlated with lung volume, and achieve surprising accuracy.

There is no consensus about the number of patients required to construct a population model, nor the number of surrogate measurements required. Some models account for individual variability within a population by incorporating the tracking of hundreds of surface location measurements (H. Fayad, Pan, et al., 2009) and relate them to the internal deformation field. A similar technique has been successfully demonstrated with as few as five patient datasets (H. Fayad, Buerger, Tsoumpas, Cheze-Le-Rest, & Visvikis, 2012).

The clinical application of cross-population models remains to be thoroughly validated (McClelland et al., 2013). A cross-population motion model is essentially a way of bypassing the problem of obtaining patient specific motion data. This may be the only option when it is impracticable to construct a patient specific motion model with the imaging modality.

2.5.4   Physiological Models of Respiratory Motion

Explicitly modelling the physics of respiratory motion is a promising approach to improving the accuracy of parametric models. A distinct example is the use of finite element models

Table 2.5: Adaptations of FEM to image registration

Modality	Method
MR	Application of FEM to image registration. Sundaram and Gee (2005)
CT	Application of FEM to 4D image registration. Werner, Ehrhardt, Schmidt, and Handels (2008)
CT	Improvement of deformable image registration using FEM. Zhong et al. (2012)
CT	Spatially varying FEM properties. M. Li et al. (2013)
CT	Novel adaptive mesh refinement. J. Zhang, Wang, Wang, and Feng (2013)

(Sundaram & Gee, 2005) to analytically model lung mechanics. A drawback of such approaches is that it is approximate and some aspects of mechanics or image acquisition are not correctly modeled. This is evident because physiological models of respiratory motion often perform poorly in comparison to voxelwise image registration algorithms (Murphy et al., 2010).

These drawbacks have since been addressed by restricting the modelling of lung mechanics to areas where image based registration may be inaccurate. For example finite element modelling (FEM) has been used to avoid inaccuracies in areas of low image contrast (Zhong et al., 2012). Further adaptive approaches that have been proposed to improve the accuracy of modelling are summarised in table 2.5.

Finite element modelling requires an approximation of the elasticity parameters of lung tissue. A shortcoming of FEM is that it is difficult to predict the model properties needed for accurate registration. FEM registration for respiratory motion generally requires the estimation of lung boundaries and other organ boundaries in the thorax. FEM registration is possible in CT because segmentation of pulmonary cine CT images is fairly straightforward. However, it is difficult to determine lung boundaries accurately from a gated PET image. Although gated PET images can be used to track tumours, there appear to be no investigation tests that accurately capture the deformation of the diaphragm in the absence of a model of respiration. It is possible that FEM may need to be coupled with another form of registration before it could be applied to gated PET images.

2.6 Application of Image Registration to Attenuation Correction in PET/CT

Although many different methods of non rigid registration can be applied to PET/CT, the assumptions and information used within the algorithm determine the accuracy of registration. The next section of this chapter reviews recent attempts at addressing attenuation misalignment including some methods which have focused on combining consistency conditions with image registration.

2.6.1 Using Deformable Registration to Address Attenuation Misalignment

An initial attempt to use image registration to address attenuation misalignment highlighted issues which affect the robustness of the solution (Mattes et al., 2003). An iterative non-linear optimiser was used to maximise the mutual information between the PET and CT using a multi-resolution strategy. Deformations have been modelled using B-splines (Dawood, Lang, Jiang, & Schafers, 2006), however, the resolution of the deformation was limited to spacing the control points by approximately 61mm. Despite the poor resolution of deformation due to coarse spacing of control points, it was shown that attenuation misalignment could be reduced based on a human observer validation scheme. The coarse spacing of control points limiting the degrees of freedom was seen as a necessary compromise for working with very low count data with a large number of degrees of freedom.

Given the wealth of options for parameterising respiratory motion many authors have considered approaches that are more suited to PET imaging. For example, a higher resolution solution has been proposed using optical flow to parameterise respiratory motion (Dawood et al., 2008). Here, the optical flow was regularised by smoothing the image and also smoothing the derivatives used in the optical flow equations. This approach was applied to gated studies by correcting for motion in both emission and attenuation sinograms. Further attempts at addressing attenuation misalignment have focused on parameterising respiratory motion in PET with spatiotemporal splines (Bai & Brady, 2011). A spatiotemporal approach was shown to be more robust because it included temporal smoothing which was a plausible assumption for respiratory motion (Bai & Brady, 2009b).

Both spatiotemporal approaches used the same mechanism, initially reconstructing the image, registering each frame to a reference frame and then adding all frames together to form a motion free image (P. Kinahan et al., 2006). In these methods a gated PET frame was identified as a reference frame in which the timing interval of the frame closely matched the respiratory phase of the spiral CT. An advantage of this approach was the ability to use a wider range of objective functions as it was not necessary to choose an image registration package capable of multimodality PET/CT alignment. The disadvantage of using a reference frame is that mismatch errors could still arise because a CT scanner takes time to traverse the thorax, and although newer scanners are much faster, there is not always a single PET frame which accurately matches the helical CT image. Even after correction for misalignment, any spiral CT motion artefacts encountered (as discussed in section 2.3.1) would still propagate into the attenuation map for each PET frame.

Deformable image registration has been combined with surrogate data to develop a correspondence model (Liu et al., 2011). Although this method is appropriate for tumour tracking, there is difficulty interpreting the motion of other structures in the lung from the gated PET images. It has been suggested in a few studies that regularisation of image registration may improve the tracking of non-involved structures in gated PET (H. J. Fayad, Lamare, Rest, Bettinardi, & Visvikis, 2013) (Dawood, Lang, et al., 2006), however, the level of accuracy in such approaches for the attenuation correction of such structures in clinical scans remains to be fully validated. In place of surrogate data, non-involved structures such as the diaphragm can be tracked using a population model (McQuaid et al., 2011). This has an advantage over other methods as the objects do not need to be clearly visible on PET to obtain an estimation

of motion. The motion of many attenuating organs, such as the dome of the liver, can be interpreted from the population model and the resulting deformation estimates can be used to correct for attenuation misalignment.

Deformable Image registration has been combined with PET imaging in several different ways in an attempt to obtain a motion free image (Blume, Martinez-Moller, Keil, Navab, & Rafecas, 2010). Despite these attempts, there is little certainty about whether PET data contain enough information to perform motion correction (Cloquet et al., 2010). Deformable image registration is typically difficult to validate because the only benchmark is a human observer. In some cases an alternative objective function can be used as an independent metric to validate the image registration. For example, it was shown that the optical flow approach was validated by showing that it reduced the NMI between frames (Dawood et al., 2008), however, the question then arises as to whether NMI would be a better choice of objective function. A lack of validation may be one of the reasons why using image registration for motion correction has not been widely accepted in clinical imaging. What is needed is an independent measure of the accuracy of motion correction other than the objective function used in image registration.

2.6.2 Using Consistency Conditions to Address Attenuation Mismatch

Although consistency conditions can be used to jointly estimate attenuation and emission (MLAA), they can also be used to address attenuation mismatch (Alessio et al., 2010). By combining Natterer's consistency conditions with image registration it is possible to improve the alignment of PET/CT.

Initial attempts to implement this approach used rigid or affine registration in the local area of a tumour (Alessio et al., 2010). An alignment metric was devised as the sum of Fourier transforms of the sinogram. Alessio et al. (2010) used a 9 degree of freedom (DOF) affine model to transform the CT to the PET image. Unlike conventional implementations of image registration, this method uses a simplex algorithm and may not converge since there is no way to incorporate a gradient descent solver. This limitation also means that it is not straightforward to extend the same methodology to deformable registration. The consistency metric is a sum of Fourier transforms of the sinogram making it very susceptible to noise. Despite these issues the method was found to work reliably across a wide range of patients (Alessio et al., 2010). A reason why it was robust may be that it did not rely solely on image registration to match image features.

An alternative to using Natterer's consistency equations as a metric is to consider the maximum likelihood as a function of attenuation alignment. Maximum likelihood has been previously used in MLAA to estimate the attenuation image, however, there some promising studies that indicate it can be used to estimate attenuation alignment. A recent example of this approach is to use the consistency equations for TOF-PET to address attenuation alignment during reconstruction (Rezaei, Michel, Casey, & Nuyts, 2016). This method used a demons registration algorithm, however, the quality of registration was affected by the low counts in PET image. An alternative parameterisation of patient and/or breathing motion may be more appropriate for PET scanners which do not have TOF capability.

Whilst the use of consistency conditions to correct for attenuation misalignment is promising, the implementations are sensitive to noise and are not necessarily practical to incorporate into existing clinical reconstruction algorithms. An alternative possibility is that consistency conditions may be a useful tool for tuning and evaluating the parameterisation of respiratory motion. This thesis proposes that consistency conditions can be used to quantitatively evaluate attenuation mismatch. Thus, different methods of deformable image registration can be compared for their accuracy when used to align PET/CT for attenuation correction.

References

Abd. Rahni, A. A., Lewis, E., Wells, K., Guy, M., & Goswami, B. (2010). Development of a particle filter framework for respiratory motion correction in nuclear medicine imaging. In *Society of photo-optical instrumentation engineers (SPIE) conference series* (Vol. 7623, p. 80).

Alessio, A. M., Kinahan, P. E., Champley, K. M., & Caldwell, J. H. (2010). Attenuation-emission alignment in cardiac PET/CT based on consistency conditions. *Medical physics*, *37*(3), 1191–1200.

Alessio, A. M., Kohlmyer, S., Branch, K., Chen, G., Caldwell, J., & Kinahan, P. (2007). Cine CT for attenuation correction in cardiac PET/CT. *Journal of Nuclear Medicine*, *48*(5), 794–801.

Bai, W., & Brady, M. (2009a). Regularized b-spline deformable registration for respiratory motion correction in PET images. *Physics in Medicine and Biology*, *54*(9), 2719. doi: 10.1088/0031-9155/54/9/008

Bai, W., & Brady, M. (2011). Motion correction and attenuation correction for respiratory gated PET images. *IEEE Transactions on Medical Imaging*, *30*(2), 351–365. doi: 10.1109/TMI.2010.2078514

Bai, W., & Brady, S. M. (2009b). Spatio-temporal image registration for respiratory motion correction in PET imaging. In *2009 IEEE international symposium on biomedical imaging: From nano to macro* (pp. 426–429). doi: 10.1109/ISBI.2009.5193075

Balfour, D., Polycarpou, I., Marsden, P. K., & King, A. P. (2013). Respiratory motion correction of PET imaging using a combination of PET and MR data. In *Nuclear science symposium conference record (NSS/MIC), 2013 IEEE*.

Barrett, H. H., Abbey, C. K., & Clarkson, E. (1998). Objective assessment of image quality. III. ROC metrics, ideal observers, and likelihood-generating functions. *JOSA A*, *15*(6), 1520–1535.

Berker, Y., & Li, Y. (2016). Attenuation correction in emission tomography using the emission dataa review. *Medical physics*, *43*(2), 807–832.

Blume, M., Martinez-Moller, A., Keil, A., Navab, N., & Rafecas, M. (2010). Joint reconstruction of image and motion in gated positron emission tomography. *IEEE transactions on medical imaging*, *29*(11), 1892–1906.

Bruyant, P. P. (2007). A method for synchronizing an external respiratory signal with a list-mode PET acquisition. *Med Phys*, *34*(11).

Buerger, C., Schaeffter, T., & King, A. P. (2011). Hierarchical adaptive local affine registration for fast and robust respiratory motion estimation. *Medical image analysis*, *15*(4),

551–564.

Castillo, E., Castillo, R., Martinez, J., Shenoy, M., & Guerrero, T. (2009). Four-dimensional deformable image registration using trajectory modeling. *Physics in medicine and biology*, *55*(1), 305.

Caucci, L., & Barrett, H. H. (2012). Objective assessment of image quality. v. photon-counting detectors and list-mode data. *JOSA A*, *29*(6), 1003–1016.

Chandrashekara, R., Rao, A., Sanchez-Ortiz, G. I., Mohiaddin, R. H., & Rueckert, D. (2003). Construction of a statistical model for cardiac motion analysis using nonrigid image registration. *Information processing in medical imaging: proceedings of the … conference*, *18*, 599–610.

Clinthorne, N. H., Fessler, J. A., Hutchins, G. D., & Rogers, W. L. (1991). Joint maximum likelihood estimation of emission and attenuation densities in PET. In *Nuclear science symposium and medical imaging conference, 1991., conference record of the 1991 IEEE* (pp. 1927–1932). IEEE.

Cloquet, C., Goldman, S., & Defrise, M. (2010). Cramer-rao bound for gated PET. In *IEEE nuclear science symposuim & medical imaging conference* (pp. 2267–2272). IEEE.

Conti, M. (2010). Why is TOF PET reconstruction a more robust method in the presence of inconsistent data? *Physics in medicine and biology*, *56*(1), 155.

Costes, N., Dagher, A., Larcher, K., Evans, A. C., Collins, D. L., & Reilhac, A. (2009). Motion correction of multi-frame PET data in neuroreceptor mapping: Simulation based validation. *NeuroImage*, *47*(4), 1496–1505. doi: 10.1016/j.neuroimage.2009.05.052

Daube-Witherspoon, M. E., Surti, S., Matej, S., Werner, M., Jayanthi, S., & Karp, J. S. (2006). Influence of time-of-flight kernel accuracy in TOF-PET reconstruction. In *Nuclear science symposium conference record, 2006. IEEE* (Vol. 3, pp. 1723–1727). IEEE.

Dawood, M., Brune, C., Jiang, X., Bther, F., Burger, M., Schober, O., . . . Schfers, K. P. (2010). A continuity equation based optical flow method for cardiac motion correction in 3d PET data. In *International workshop on medical imaging and virtual reality* (pp. 88–97). Springer.

Dawood, M., Buther, F., Jiang, X., & Schafers, K. (2008). Respiratory motion correction in 3-d PET data with advanced optical flow algorithms. *Medical Imaging, IEEE Transactions on*, *27*(8), 1164–1175. doi: 10.1109/TMI.2008.918321

Dawood, M., Buther, F., Lang, N., Jiang, X., & Schafers, K. P. (2006). Transforming static CT in gated PET/CT studies to multiple respiratory phases. *Proceedings of the ICPR 2006*, 1026–1029.

Dawood, M., Bther, F., Lang, N., Schober, O., & Schfers, K. P. (2007). Respiratory gating in positron emission tomography: A quantitative comparison of different gating schemes. *Medical Physics*, *34*(7), 3067–3076. doi: 10.1118/1.2748104

Dawood, M., Lang, N., Jiang, X., & Schafers, K. (2006). Lung motion correction on respiratory gated 3-d PET/CT images. *Medical Imaging, IEEE Transactions on*, *25*(4), 476–485. doi: 10.1109/TMI.2006.870892

Defrise, M., Rezaei, A., & Nuyts, J. (2012). Time-of-flight PET data determine the attenuation sinogram up to a constant. *Physics in medicine and biology*, *57*(4), 885.

Di Carli, M. F., Dorbala, S., Meserve, J., El Fakhri, G., Sitek, A., & Moore, S. C. (2007). Clinical myocardial perfusion PET/CT. *Journal of Nuclear Medicine*, *48*(5), 783–793.

Ehrhardt, J., Lorenz, C., & others. (2013). *4d modeling and estimation of respiratory motion for radiation therapy.* Springer.

Ehrhardt, J., Werner, R., Schmidt-Richberg, A., & Handels, H. (2011). Statistical modeling of 4d respiratory lung motion using diffeomorphic image registration. *IEEE Transactions on Medical Imaging*, *30*(2), 251–265. doi: 10.1109/TMI.2010.2076299

Erlandsson, K., Buvat, I., Pretorius, P. H., Thomas, B. A., & Hutton, B. F. (2012). A review of partial volume correction techniques for emission tomography and their applications in neurology, cardiology and oncology. *Physics in medicine and biology*, *57*(21), R119.

Fayad, H., Buerger, C., Tsoumpas, C., Cheze-Le-Rest, C., & Visvikis, D. (2012). A generic respiratory motion model based on 4d MRI imaging and 2d image navigators. *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2012 IEEE*, 4058–4061. doi: 10.1109/NSSMIC.2012.6551927

Fayad, H., Clement, J., Pan, T., Roux, C., Le Rest, C., Pradier, O., & Visvikis, D. (2009). Towards a generic respiratory motion model for 4d CT imaging of the thorax. In *Nuclear science symposium conference record (NSS/MIC), 2009 IEEE* (pp. 3975–3979).

Fayad, H., Pan, T., Roux, C., Le Rest, C., Pradier, O., Clement, J., & Visvikis, D. (2009). A patient specific respiratory model based on 4d CT data and a time of flight camera (TOF). In *Nuclear science symposium conference record (NSS/MIC), 2009 IEEE* (pp. 2594–2598).

Fayad, H. J., Lamare, F., Rest, C. C. L., Bettinardi, V., & Visvikis, D. (2013). Generation of 4-dimensional CT images based on 4-dimensional PETderived motion fields. *Journal of Nuclear Medicine*, *54*(4), 631–638. doi: 10.2967/jnumed.112.110809

George, R., Chung, T. D., Vedam, S. S., Ramakrishnan, V., Mohan, R., Weiss, E., & Keall, P. J. (2006). Audio-visual biofeedback for respiratory-gated radiotherapy: Impact of audio instruction and audio-visual biofeedback on respiratory-gated radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, *65*(3), 924–933. doi: 10.1016/j.ijrobp.2006.02.035

Gu, C., Li, R., Zhang, H., Fung, A. Y., Torres, C., Jiang, S. B., & Li, C. (2012). Accurate respiration measurement using DC-coupled continuous-wave radar sensor for motion-adaptive cancer radiotherapy. *IEEE Transactions on Biomedical Engineering*, *59*(11), 3117–3123.

Han, X. (2010). Feature-constrained nonlinear registration of lung CT images. *Medical image analysis for the clinic: a grand challenge*, 63–72.

He, J., O'Keefe, G. J., Gong, S. J., Jones, G., Saunder, T., Scott, A. M., & Geso, M. (2008). A novel method for respiratory motion gated with geometric sensitivity of the scanner in 3d PET. *IEEE Transactions on Nuclear Science*, *55*(5), 2557–2565.

Heimann, T., Van Ginneken, B., Styner, M. A., Arzhaeva, Y., Aurich, V., Bauer, C., ... Bekes, G. (2009). Comparison and evaluation of methods for liver segmentation from CT datasets. *IEEE transactions on medical imaging*, *28*(8), 1251–1265.

Holden, M. (2008). A review of geometric transformations for nonrigid body registration. *IEEE transactions on medical imaging*, *27*(1), 111–128.

Huang, S.-C., Ye, H., Wardak, M., Wong, K.-P., Dahlbom, M., Shao, W., ... Barrio, J. (2011). A bootstrap method for identifying image regions affected by intra-scan body movement during a PET/CT scan. In *2011 ieee nuclear science symposium and medical imaging conference (nss/mic)* (pp. 2905 –2908). doi: 10.1109/NSSMIC.2011.6152516

Iatrou, M., Ross, S. G., Manjeshwar, R. M., & Stearns, C. W. (2004). A fully 3d iterative image reconstruction algorithm incorporating data corrections. In *Nuclear science symposium conference record, 2004 IEEE* (Vol. 4, pp. 2493–2497). IEEE.

Jallow, N., Christian, P., Sunderland, J., Graham, M., Hoffman, J. M., & Nye, J. A. (2016). Diagnostic reference levels of CT radiation dose in whole-body PET/CT. *Journal of Nuclear Medicine*, *57*(2), 238–241.

Karp, J. S., Surti, S., Daube-Witherspoon, M. E., & Muehllehner, G. (2008). Benefit of time-of-flight in PET: experimental and clinical results. *Journal of Nuclear Medicine*, *49*(3), 462–470.

Keall, P. J., Joshi, S., Vedam, S. S., Siebers, J. V., Kini, V. R., & Mohan, R. (2005). Four-dimensional radiotherapy planning for DMLC-based respiratory motion tracking. *Medical Physics*, *32*(4), 942–951.

Killoran, H., J, Gerbaudo, V. H., Mamede, M., Ionascu, D., Park, S., & Berbeco, R. (2011). Motion artifacts occurring at the lung/diaphragm interface using 4d CT attenuation correction of 4d PET scans. *Journal of applied clinical medical physics*, *12*(4), 261–274.

Kinahan, P., MacDonald, L., Ng, L., Alessio, A., Segars, P., Tsui, B., & Pathak, S. (2006). Compensating for patient respiration in PET/CT imaging with the registered and summed phases (RASP) procedure. In *Biomedical imaging: Nano to macro, 2006. 3rd ieee international symposium on* (pp. 1104–1107).

Kinahan, P. E., Townsend, D. W., Beyer, T., & Sashin, D. (1998). Attenuation correction for a combined 3d PET/CT scanner. *Medical physics*, *25*(10), 2046–2053.

King, A., Buerger, C., Tsoumpas, C., Marsden, P., & Schaeffter, T. (2012). Thoracic respiratory motion estimation from MRI using a statistical model and a 2-d image navigator. *Medical Image Analysis*, *16*(1), 252–264. doi: 10.1016/j.media.2011.08.003

Kyriakou, E., & McKenzie, D. R. (2011). Dynamic modeling of lung tumor motion during respiration. *Physics in Medicine and Biology*, *56*(10), 2999–3013. doi: 10.1088/0031-9155/56/10/007

Lamare, Cresson, Savean, Cheze-Le Rest, Turzo, Bizais, ... Visvikis (2004). Affine transformation of list mode data for respiratory motion correction in PET. In *IEEE symposium conference record nuclear science 2004.* (Vol. 5, pp. 3151–3155 Vol. 5). doi: 10.1109/NSSMIC.2004.1466349

Levin, C. S., & Hoffman, E. J. (1999). Calculation of positron range and its effect on the fundamental limit of positron emission tomography system spatial resolution. *Physics in medicine and biology*, *44*(3), 781.

Li, M., Castillo, E., Zheng, X.-L., Luo, H.-Y., Castillo, R., Wu, Y., & Guerrero, T. (2013). Modeling lung deformation: A combined deformable image registration method with spatially varying young's modulus estimates. *Medical physics*, *40*(8).

Li, R., Lewis, J. H., Jia, X., Zhao, T., Liu, W., Wuenschel, S., ... Jiang, S. B. (2011). On a PCA-based lung motion model. *Physics in Medicine and Biology*, *56*(18), 6009–6030. doi: 10.1088/0031-9155/56/18/015

Liu, C., Alessio, A., Pierce, L., Thielemans, K., Wollenweber, S., Ganin, A., & Kinahan, P. (2010). Quiescent period respiratory gating for PET/CT. *Medical physics*, *37*(9), 5037–5043.

Liu, C., Alessio, A. M., & Kinahan, P. E. (2011). Respiratory motion correction for quantitative PET/CT using all detected events with internalexternal motion correlation.

Medical Physics, *38*, 2715. doi: 10.1118/1.3582692

Liu, C., Pierce, I. I., Larry, A., Alessio, A. M., & Kinahan, P. E. (2009). Impact of respiratory motion on tumor quantification and delineation in static PET/CT imaging. *Physics in Medicine and Biology*, *54*, 7345–7362.

Livieratos, L., Stegger, L., Bloomfield, P. M., Schafers, K., Bailey, D. L., & Camici, P. G. (2005). Rigid-body transformation of list-mode projection data for respiratory motion correction in cardiac PET. *Phys Med Biol*, *50*(14), 3313–22. (Journal Article England)

Low, D. A., Nystrom, M., Kalinin, E., Parikh, P., Dempsey, J. F., Bradley, J. D., . . . Christensen, G. (2003). A method for the reconstruction of four-dimensional synchronized CT scans acquired during free breathing. *Medical physics*, *30*(6), 1254–1263.

Lu, W., Parikh, P. J., Hubenschmidt, J. P., Bradley, J. D., & Low, D. A. (2006). A comparison between amplitude sorting and phase-angle sorting using external respiratory measurement for 4d CT. *Medical physics*, *33*(8), 2964–2974.

Mattes, D., Haynor, D. R., Vesselle, H., Lewellen, T. K., & Eubank, W. (2003). PET-CT image registration in the chest using free-form deformations. *IEEE Transactions on Medical Imaging*, *22*(1), 120–128. doi: 10.1109/TMI.2003.809072

McClelland, J. R., Blackall, J. M., Tarte, S., Chandler, A. C., Hughes, S., Ahmad, S., . . . Hawkes, D. J. (2006). A continuous 4d motion model from multiple respiratory cycles for use in lung radiotherapy. *Medical Physics*, *33*(9), 3348. doi: 10.1118/1.2222079

McClelland, J. R., Hawkes, D. J., Schaeffter, T., & King, A. P. (2013). Respiratory motion models: A review. *Medical Image Analysis*, *17*(1), 19–42. doi: 10.1016/j.media.2012.09.005

McQuaid, S. J., Lambrou, T., & Hutton, B. F. (2008). Statistical shape modeling of the diaphragm for application to rb-82 cardiac PET-CT studies. In *Nuclear science symposium conference record, 2008. NSS '08. IEEE* (pp. 3651–3655).

McQuaid, S. J., Lambrou, T., & Hutton, B. F. (2011). A novel method for incorporating respiratory-matched attenuation correction in the motion correction of cardiac PETCT studies. *Physics in medicine and biology*, *56*(10), 2903.

Metz, C. T., Klein, S., Schaap, M., van Walsum, T., & Niessen, W. J. (2011). Non-rigid registration of dynamic medical imaging data using nD + t b-splines and a groupwise optimization approach. *Medical Image Analysis*, *15*(2), 238–249. doi: 10.1016/j.media.2010.10.003

Murphy, K., Ginneken, B. v., Reinhardt, J. M., Kabus, S., Ding, K., Deng, X., . . . Pluim, J. P. W. (2011). Evaluation of registration methods on thoracic CT: The EMPIRE10 challenge. *IEEE Transactions on Medical Imaging*, *30*(11), 1901–1920. doi: 10.1109/TMI.2011.2158349

Murphy, K., van Ginneken, B., Reinhardt, J., Kabus, S., Ding, K., Deng, X., & Pluim, J. (2010). Evaluation of methods for pulmonary image registration: The EMPIRE10 study. *Grand Challenges in Medical Image Analysis*, *2010*.

Natterer, F. (1999). Numerical methods in tomography. *Acta Numerica*, *8*, 107–141.

Nehmeh, S. A., Erdi, Y. E., Meirelles, G. S., Squire, O., Larson, S. M., Humm, J. L., & Schoder, H. (2007). Deep-inspiration breath-hold PET/CT of the thorax. *J Nucl Med*, *48*(1), 22–26.

Nehmeh, S. A., Erdi, Y. E., Pan, T., Yorke, E., Mageras, G. S., Rosenzweig, K. E., . . . Pevsner, A. (2004). Quantitation of respiratory motion during 4d-PET/CT acquisition.

Medical physics, *31*(6), 1333–1338.

Noonan, P. J., Howard, J., Tout, D., Armstrong, I., Williams, H. A., Cootes, T. F., ... Hinz, R. (2012). Accurate markerless respiratory tracking for gated whole body PET using the microsoft kinect. In *Nuclear science symposium and medical imaging conference (NSS/MIC), 2012 IEEE* (pp. 3973–3974). IEEE.

Nuyts, J., Bal, G., Kehren, F., Fenchel, M., Michel, C., & Watson, C. (2013). Completion of a truncated attenuation image from the attenuated PET emission data. *IEEE transactions on medical imaging*, *32*(2), 237–246.

Nuyts, J., Dupont, P., Stroobants, S., Benninck, R., Mortelmans, L., & Suetens, P. (1999). Simultaneous maximum a posteriori reconstruction of attenuation and activity distributions from emission sinograms. *Medical Imaging, IEEE Transactions on*, *18*(5), 393–403.

Osman, M. M., Cohade, C., Nakamoto, Y., & Wahl, R. L. (2003). Respiratory motion artifacts on PET emission images obtained using CT attenuation correction on PET-CT. *European journal of nuclear medicine and molecular imaging*, *30*(4), 603–606.

Pan, T., Lee, T.-Y., Rietzel, E., & Chen, G. T. (2004). 4d-CT imaging of a volume influenced by respiratory motion on multi-slice CT. *Medical physics*, *31*(2), 333–340.

Panin, V. Y., Defrise, M., & Casey, M. E. (2010). Restoration of fine azimuthal sampling of measured TOF projection data. In *Nuclear science symposium conference record (NSS/MIC), 2010 IEEE* (pp. 3079–3084). IEEE.

Pretorius, P. H., King, M. A., Johnson, K. L., Mukherjee, J. M., Dey, J., & Konik, A. (2011). Combined respiratory and rigid body motion compensation in cardiac perfusion SPECT using a visual tracking system. In *Nuclear science symposium and medical imaging conference (NSS/MIC), 2011 IEEE* (pp. 2768–2773). IEEE.

Rezaei, A., Defrise, M., Bal, G., Michel, C., Conti, M., Watson, C., & Nuyts, J. (2012). Simultaneous reconstruction of activity and attenuation in time-of-flight PET. *IEEE transactions on medical imaging*, *31*(12), 2224–2233.

Rezaei, A., Defrise, M., & Nuyts, J. (2014). ML-reconstruction for TOF-PET with simultaneous estimation of the attenuation factors. *Medical Imaging, IEEE Transactions on*, *33*(7), 1563–1572.

Rezaei, A., Michel, C., Casey, M. E., & Nuyts, J. (2016). Simultaneous reconstruction of the activity image and registration of the CT image in TOF-PET. *Physics in medicine and biology*, *61*(4), 1852.

Ruan, D. (2009). Directionally selective regularization for sliding preserving medical image registration. In *2009 IEEE nuclear science symposium conference record (NSS/MIC)* (pp. 2936–2939). doi: 10.1109/NSSMIC.2009.5401607

Ruan, D., Fessler, J. A., Roberson, M., Balter, J., & Kessler, M. (2006). Nonrigid registration using regularization that accomodates local tissue rigidity. In *Medical imaging* (pp. 614412–614412–9). International Society for Optics and Photonics.

Rueckert, D., Aljabar, P., Heckemann, R. A., Hajnal, J. V., & Hammers, A. (2006). Diffeomorphic registration using b-splines. In R. Larsen, M. Nielsen, & J. Sporring (Eds.), *Medical image computing and computer-assisted intervention MICCAI 2006* (Vol. 4191, pp. 702–709). Springer Berlin Heidelberg.

Segars, W. P., Sturgeon, G., Mendonca, S., Grimes, J., & Tsui, B. M. W. (2010). 4d XCAT phantom for multimodality imaging research. *Medical Physics*, *37*(9), 4902.

doi: 10.1118/1.3480985

Shepp, L. A., & Vardi, Y. (1982). Maximum likelihood reconstruction for emission tomography. *IEEE transactions on medical imaging*, *1*(2), 113–122.

Smith, R., Rahni, A., Jones, J., & Wells, K. (2012). Recursive bayesian estimation for respiratory motion correction in nuclear medicine imaging. *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2012 IEEE*, 2942–2945. doi: 10.1109/NSS-MIC.2012.6551672

Stamm, G., & Nagel, H. D. (2002). CT-expo–a novel program for dose evaluation in CT. *RoFo: Fortschritte auf dem Gebiete der Rontgenstrahlen und der Nuklearmedizin*, *174*(12), 1570–1576.

Staring, M., Klein, S., & Pluim, J. P. W. (2007). A rigidity penalty term for nonrigid registration. *Medical Physics*, *34*(11), 4098–4108. doi: 10.1118/1.2776236

Sun, T., & Mok, G. S. (2012). Techniques for respiration-induced artifacts reductions in thoracic PET/CT. *Quantitative imaging in medicine and surgery*, *2*(1), 46.

Sundaram, T. A., & Gee, J. C. (2005). Towards a model of lung biomechanics: pulmonary kinematics via registration of serial lung images. *ITKOpen science - combining open data and open source software: Medical image analysis with the Insight Toolkit*, *9*(6), 524–537. doi: 10.1016/j.media.2005.04.002

Thielemans, Rathore, Engbrant, & Razifar. (2011). Device-less gating for PET/CT using PCA. In *2011 IEEE nuclear science symposium conference record* (pp. 3904–3910). doi: 10.1109/NSSMIC.2011.6153742

Thielemans, Schleyer, Marsden, Manjeshwar, Wollenweber, & Ganin. (2011). Comparison of different methods for data-driven respiratory gating of PET data. In *2013 IEEE nuclear science symposium and medical imaging conference (2013 NSS/MIC)* (pp. 1–4). doi: 10.1109/NSSMIC.2013.6829055

Thielemans, K., Asma, E., & Manjeshwar, R. M. (2009). Mass-preserving image registration using free-form deformation fields. In *2009 IEEE nuclear science symposium conference record (NSS/MIC)* (pp. 2490–2495). doi: 10.1109/NSSMIC.2009.5402070

Thorndyke, K., Schreibmann, & Xing. (2006). Reducing respiratory motion artifacts in positron emission tomography through retrospective stacking. *Med Phys*, *33*(7), 2632–41. (Journal Article United States)

Townsend, D. W., Carney, J. P., Yap, J. T., & Hall, N. C. (2004). PET/CT today and tomorrow. *Journal of Nuclear Medicine*, *45*(1), 4S–14S.

Vandemeulebroucke, J., Rit, S., Kybic, J., Clarysse, P., & Sarrut, D. (2011). Spatiotemporal motion estimation for respiratory-correlated imaging of the lungs. *Medical Physics*, *38*(1), 166–178. doi: 10.1118/1.3523619

Verhaeghe, J., Gravel, P., Mio, R., Fukasawa, R., Rosa-Neto, P., Soucy, J. P., ... Reader, A. J. (2010). Motion compensation for fully 4d PET reconstruction using PET superset data. *Physics in medicine and biology*, *55*(14), 4063.

Wang, H., Dong, L., O'Daniel, J., Mohan, R., Garden, A. S., Ang, K. K., ... Cheung, R. (2005). Validation of an accelerated demons algorithm for deformable image registration in radiation therapy. *Physics in medicine and biology*, *50*(12), 2887.

Werner, R., Ehrhardt, J., Schmidt, R., & Handels, H. (2008). Modeling respiratory lung motion: a biophysical approach using finite element methods. In *Proc. spie 6916, medical imaging 2008* (Vol. 6916, pp. 69160N–69160N–11). doi: 10.1117/12.769155

Xia, T., Alessio, A. M., & Kinahan, P. E. (2009). Limits of ultra-low dose CT attenuation correction for PET/CT. In *Nuclear science symposium conference record (NSS/MIC), 2009 IEEE* (pp. 3074–3079). IEEE.

Yin, Y., Hoffman, E. A., & Lin, C.-L. (2009). Mass preserving nonrigid registration of CT lung images using cubic b-spline. *Medical Physics*, *36*(9), 4213–4222. doi: 10.1118/1.3193526

Yip, S., Chen, A. B., Aerts, H. J., & Berbeco, R. (2014). Sensitivity study of voxel-based PET image comparison to image registration algorithms. *Medical physics*, *41*(11), 111714.

Zaidi, H., Montandon, M., & Alavi, A. (2007). Advances in attenuation correction techniques in PET. *PET clinics*, *2*(2), 191–217.

Zhang, J., Wang, J., Wang, X., & Feng, D. (2013). The adaptive FEM elastic model for medical image registration. *Physics in medicine and biology*, *59*(1), 97.

Zhang, Q., Hu, Y.-C., Liu, F., Goodman, K., Rosenzweig, K. E., & Mageras, G. S. (2010). Correction of motion artifacts in cone-beam CT using a patient-specific respiratory motion model. *Medical Physics*, *37*(6), 2901–2909. doi: 10.1118/1.3397460

Zhang, Q., Pevsner, A., Hertanto, A., Hu, Y. C., Rosenzweig, K. E., Ling, C. C., & Mageras, G. S. (2007). A patient-specific respiratory model of anatomical motion for radiation treatment planning. *Medical Physics*, *34*, 4772.

Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., & Comaniciu, D. (2008). Four-chamber heart modeling and automatic segmentation for 3-d cardiac CT volumes using marginal space learning and steerable features. *IEEE transactions on medical imaging*, *27*(11), 1668–1681.

Zhong, H., Kim, J., Li, H., Nurushev, T., Movsas, B., & Chetty, I. J. (2012). A finite element method to correct deformable image registration errors in low-contrast regions. *Physics in medicine and biology*, *57*(11), 3499.

Chapter 3

Improving The Accuracy of Attenuation Correction in Gated PET Images Using Image Registration

3.1 Introduction

3.1.1 Compensating for Attenuation Misalignment

The misalignment of the attenuation correction map often degrades the quantitative accuracy of PET images in the thorax. Respiratory motion is the most significant cause of misalignment and there are several approaches to compensating for this (Berker & Li, 2016). Although misalignment can be addressed in many ways, parameterising respiratory motion is a useful approach to compensation since it can be used to model the deformation which causes the PET and CT misalignment. In addition to that, the models can be used to estimate deformation in all gated PET frames and subsequently used to generate a motion free image whilst also compensating for attenuation mismatch (Bai & Brady, 2011).

This investigation builds on previous attempts at using deformable image registration to address attenuation misalignment by systematically comparing each method and devising an optimal approach. The comparison is restricted to methods which use deformable image registration of non-attenuated corrected (NAC) gated PET images to estimate deformations to compensate for PET/CT misalignment. In this case it is necessary to perform an additional image reconstruction for each frame using the aligned CT for attenuation correction (Dawood, Lang, Jiang, & Schafers, 2006). Methods for estimating deformations generally draw from one of two distinct approaches. One is to use multimodality registration to align the PET image to the CT image prior to AC correction (Dawood et al., 2006). The second approach is to find the gated PET reference frame which is best aligned to the CT and then use mono modality registration to align the remaining gated PET frames to the reference frame. There are many methods of motion compensation in respiratory gated PET which use mono-modality registration (Kinahan et al., 2006) and these methods can also be extended to gated PET/CT for compensating for attenuation misalignment. Provided that breathing style doesn't change significantly during the PET scan or CT scan this approach to motion compensation can be extended to attenuation correction. This is achieved by deforming the CT image or attenuation map with the same deformation used for motion compensation (Dawood et al., 2006) and then reconstructing each frame with the aligned attenuation map. In both mono-modality and multimodality approaches the NAC gated PET images have more noise than attenuation corrected images (McQuaid, Lambrou, & Hutton, 2011), although the noise in the NAC images can be addressed through improved forms of registration (Bai & Brady, 2009). Simulated NAC PET scans have been used to investigate the impact of image noise on aligning attenuation maps using image registration (Fayad, Lamare, Rest, Bettinardi, & Visvikis, 2013).

The accuracy of image registration is largely dependant on the quality of the images which are to be aligned. Although registration accuracy is difficult to determine, the robustness of image registration can however be determined through the simulation of potential deformations (Kybic, 2010; Tsui et al., 2001). Despite this, it is difficult to realistically simulate the effect of respiratory motion on the acquisition of both PET and CT images. It is also challenging to simulate abnormal anatomy which may be observed in diseased patients. Rather than simulations, real patient data has been used to compare image registration algorithms where the ground truth is specified through the localisation of landmarks by an expert operator (Murphy et al., 2011). Such comparisons evaluate the robustness of image registration

to real abnormalities and artefacts which are found in routine clinical images. Even the more realistic simulations may not include image corruption due to changes in breathing style. Using such gated PET images may lead to perturbations in the apparent displacement of organs. For example, the ratio of abdominal to thoracic displacement may change and this cannot be compensated for with respiratory phase gating.

The aim of this thesis is to improve the accuracy of attenuation correction, yet most methods of compensating for misalignment make use of deformable image registration methods for which the accuracy cannot be determined. Given the significant challenges in validating deformable image registration, it is suspected that the accuracy of attenuation correction is largely affected by the accuracy of the registration method used.

Although there is evidence that image quality can be improved using amplitude based gating, phase gating remains the predominant approach for clinical gated PET/CT scans . Finding the timing interval of the gated PET frame which matches the respiratory phase of the spiral CT is technically challenging. The surrogate respiratory signals from both PET and CT modalities must be initially synchronised (Bruyant et al., 2007) and then amplitude or quiescent gating of the PET data must be adapted to produce a PET image which is aligned with the CT image (Dawood et al., 2006; Liu et al., 2010). An alternative to this method is a data driven approach whereby the phase interval of the CT acquisition in the respiratory cycle is found by selecting the gated PET frame which has the maximum mutual information (MI) with the CT image (Nam, Ahn, Kim, Kim, & Ra, 2013). A shortcoming of using the mutual information metric for gated interval selection is that it may be affected by noise from the gated PET image. Without investigations into the accuracy of this method it is not possible to rule out a negative impact of interval selection on methods for addressing attenuation mismatch.

3.1.2 Tumour Quantification

Accurate tumour quantification requires the alignment of the attenuation map and emission image. Although phase gating partially compensates for motion in the PET image, it does not guarantee accurate alignment with the attenuation map. In respiratory gated PET there are two main expected contributions to such inaccuracies:

1. *Emission:* Distortion of tumour shape due to motion blur of the emission image. In phase gating, the amount of tumour blur is proportional to the range of respiratory amplitude in each frame.

2. *Attenuation:* Application of incorrect attenuation correction factors due to mismatch of the tumour mass in the attenuation map with the apparent tracer uptake in the emission image.

Both effects can be mitigated using either amplitude gating (Dawood, Buther, Lang, Schober, & Schafers, 2007) or *Registered and Summed Phases* (Kinahan et al., 2006). In addition to this as the lung expands density decreases and objects may appear brighter at end inspiration and darker at end expiration. Correcting for density of the attenuation map should reduce this effect (Thielemans, Asma, & Manjeshwar, 2009) leading to a reduced variance in the tumour uptake from the average uptake over the respiratory cycle.

3.1.3 Validating Attenuation Correction Accuracy

Validation of a PET attenuation map can be achieved prior to PET image reconstruction by exploiting symmetries in Natterer's formulation of the consistency equation that was presented in chapter 2. The attenuation correction factors (ACF) and the emission sinogram are said to be consistent when Natterer's equation holds true (Alessio, Kinahan, Champley, & Caldwell, 2010). Alessio et al used this equation to develop a metric which measures the alignment of both gated PET image and static PET image with a single helical CT image. Simulation studies have demonstrated that there is a global minimum in the metric that corresponds to the affine translation which aligns the two images (Alessio et al., 2010), however, this approach has not been extended to deformable image registration.

Maximum likelihood of activity and attenuation (MLAA) was reviewed in chapter 2 as an alternative to CT derived attenuation correction through the use of consistency conditions. MLAA has also been extended to compensate for attenuation misalignment in the myocardium (Presotto et al., 2015) without the need for image registration, however, there was no quantitative analysis of image accuracy. Although MLAA is promising as a method to compensate for attenuation misalignment, there are as yet no investigations which use consistency to validate existing methods of compensating for misalignment.

This chapter examines whether validation of attenuation correction accuracy can be achieved through an extension of image reconstruction. Given the multitude of possible approaches to image reconstruction, both maximum likelihood approaches and least squared approaches have been considered. The proposed validation method is loosely based on findings that inconsistent system models in iterative reconstruction often do not converge (Alessio et al., 2010). Attenuation correction mismatch will lead to inaccurate attenuation correction factors (ACF) and it is suspected that this will result in a significant inconsistency in the system model. The remainder of this chapter outlines the evaluation of some novel metrics which are able to detect this inconsistency and thus provide a means of validating and comparing methods of compensating for misalignment.

3.2 Methods

3.2.1 Consistency Conditions and Comparison Metrics

We propose a novel method which builds upon conventional approaches to assessing attenuation alignment including Natterer's consistency equation (Alessio et al., 2010) and the image similarity of PET and CT images (Nam et al., 2013). The proposed method involves initially reconstructing attenuation corrected gated PET images. The PET images are then forward projected with all corrections applied and it is hypothesised that the sinogram should best match the detected emissions when the attenuation map is correctly aligned. In iterative reconstruction the difference between projected and the detected emissions is often used to calculate an update term, however, in this method the difference was summed as a metric for the assessment of attenuation alignment.

Each gated PET frame was reconstructed with attenuation correction for 3 iterations and 24 subsets of Expectation Maximisation (OSEM), as further iterations did not improve image quality. All reconstruction settings were kept consistent across all studies as any variations
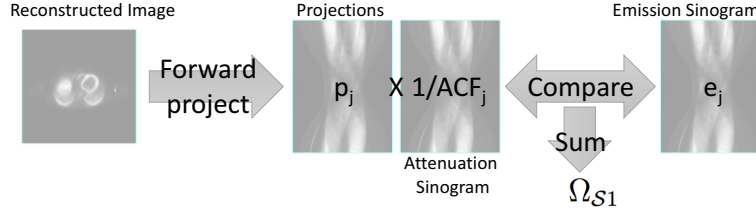
Figure 3.1: Comparison of forward projections with emission data applying attenuation factor following the projection of the reconstructed image

would have a significant impact on the derived metrics. Lines of response which did not pass through the body were assumed not to contain any information about attenuation alignment. In a similar approach to Nuyts et al, these LORs were identified by regions on the sinogram with low counts (approximate background level from LSO crystal) and were thresholded to zero. The Biograph mCT has 400 detector elements, 168 views and 621 sinogram planes which were rebinned into a smaller 2D sinogram to ensure that the signal due to attenuation mismatch outweighed noise. The down-sampled sinogram had 100 elements, 42 views and 109 axial planes and each element in the emission sinogram, e_j, was compared to the forward projected data, p_j using equation 3.1. Our devised metrics were very sensitive to inaccuracies in the system model, thus we produced each projected sinogram element, p_j, using the same ray tracing algorithm used for OSEM image reconstruction and we applied scatter correction, randoms subtraction and detector normalisation. Incorrect attenuation factors, ACF_j caused by attenuation mismatch were anticipated to contribute to an increased value of the summed metric, Ω_{S1} provided that S increased due to the difference between emission sinograms and projected data (figure 5.8).

$$\Omega_{S1} = \sum_j S(e_j, \frac{p_j}{ACF_j}) \tag{3.1}$$

Some implementations of OSEM pre-correct the sinogram with the attenuation correction factor, ACF_j, although this should make minimal difference in an iterative reconstruction it does change the statistical distribution of the estimated image values. We investigated pre-correction because it weights values in the sinogram which contribute to the metric and could be favourable for detecting misalignment (equation 3.2). The precorrected metrics identified using a subscript for the summed metric, Ω_{S2}.

$$\Omega_{S2} = \sum_j S(e_j * ACF_j, p_j) \tag{3.2}$$

The log likelihood metric was calculated by summing the log likelihood comparator for all LORs which had a positive number of net trues in every gated frame. Both the unweighted, Ω_{LL1}, and weighted, Ω_{LL2}, log likelihood metrics were calculated using the equations 3.3 and 3.4.

$$\Omega_{LL1} = \sum_j (e_j \log(\frac{p_j}{ACF_j}) - \frac{p_j}{ACF_j}) \tag{3.3}$$

$$\Omega_{LL2} = \sum_j (e_j ACF_j \log(p_j) - p_j) \tag{3.4}$$

Although maximum likelihood is suited to image reconstruction of PET images, alternative metrics which might be more sensitive to detecting attenuation mismatch were investigated. For example, the weighted least squares (WLS) iterative reconstruction is suited to iterative estimation of images affected by Gaussian noise (Jing Wang, Tianfang Li, Hongbing Lu, & Zhengrong Liang, 2006). It was suspected that a least squares approach may be a sensitive metric as it was calculated from the sum of squared differences (SSD) between the emission and projection data using equations 3.5 and 3.6.

$$\Omega_{SSD1} = \sum_j (\frac{p_j}{ACF_j} - e_j)^2 \tag{3.5}$$

$$\Omega_{SSD2} = \sum_j (p_j - e_j ACF_j)^2 \tag{3.6}$$

We also tested the Chi square statistic as a metric using equations 3.7 and 3.8 because it is an unbiased estimator when testing the goodness of fit of a Poisson distributed variable.

$$\Omega_{ChiSq1} = \sum_j \frac{(\frac{p_j}{ACF_j} - e_j)^2}{\frac{p_j}{ACF_j}} \tag{3.7}$$

$$\Omega_{ChiSq2} = \sum_j \frac{(p_j - e_j ACF_j)^2}{p_j} \tag{3.8}$$

We calculated these metrics in addition to calculating Natterer's consistency conditions which have already been shown to be sensitive to attenuation misalignment (Alessio et al., 2010).

3.2.2 Acquisition of Respiratory Gated PET/CT

Scans were acquired and reconstructed using a Biograph mCT PET/CT. Phase based respiratory gating was performed following a conventional whole body static PET and spiral CT. Each patient had been injected with approximately 350 MBq of F18-FDG. The count statistics achieved with the Biograph mCT PET/CT permitted short duration gated acquisitions with a total duration of 3 minutes (sensitivity was 8 cps/kBq). The phase based respiratory gating system AZ733 V (Anzai Medical, Tokyo, Japan), was employed as previously described by Pönisch (Pönisch, Richter, Just, & Enghardt, 2008). The system utilises a pressure sensor inserted into the pocket of an elastic belt to detect external respiratory motion. Previously acquired clinical gated PET/CT scans were selected for which list mode PET data was available from patients who had undergone PET/CT investigations for the staging of single pulmonary nodules. The set of 11 scans were used to compare methods of reducing attenuation mismatch. Patients were instructed to breathe out during the CT, increasing the likelihood that the CT would image the quiescent phase of the respiratory cycle.
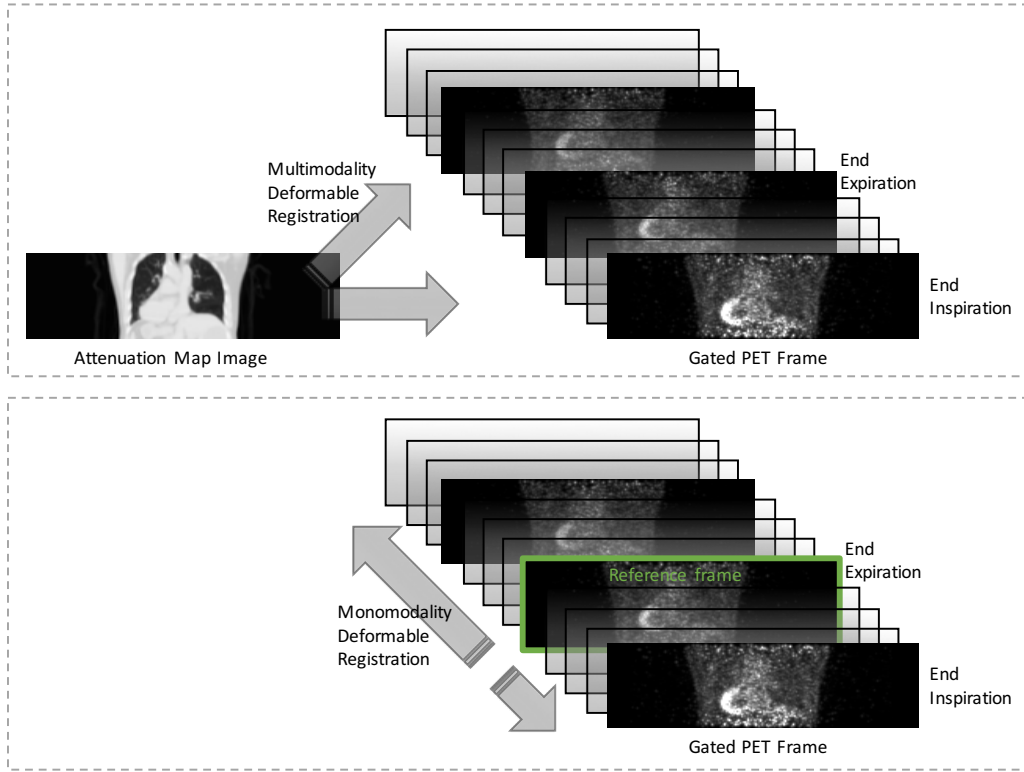
Figure 3.2: Multi-modality deformable registration (top) independently aligns the attenuation map image to each gated PET frame. Mono-modality deformable registration (bottom) aligns a reference frame (green) to each gated PET frame. In both cases the deformable transform is subsequently used to deform the attenuation map.

3.2.3 Compensation for Attenuation Misalignment

Each respiratory gated PET scan was initially reconstructed without attenuation correction (NAC). Some organ boundaries are visible on the NAC gated PET images indicating that it may be feasible to find the deformation between a gated PET frame and the attenuation map. Each approach to compensating for attenuation misalignment used either mono-modality deformable image registration (PET only) or a multimodality registration (PET/CT) or a combination of both (figure 3.2). A variety of different options for image registration were used for both mono-modality and multimodality approaches including mass preserving registration, similarity metrics, masks and regularisation tuning parameters. Mass-preservation was implemented through a weighting which was applied to attenuation maps following registration (P15 in table 3.2, P16 in table 3.1, P18 and P20 in table 3.3). The effect of each combination of options for image registration were initially investigated by using the deformed attenuation maps in a final OSEM reconstruction and subsequently assessing the final images for accuracy of attenuation correction.

A Multimodality Approach to Compensation

The first approach taken was to align the attenuation map image to the emission image using multimodality image registration with various options summarised in table 3.1. In this study, the respiratory gated PET scan contained multiple frames and thus the attenuation map image was independently registered to each frame. The deformable image registration algorithm involved aligning the attenuation map with each gated PET frame by maximising the NMI between them (Modat et al., 2010). Two registration packages were investigated including two sets of regularisation tuning parameters, P5 and P6, for niftyreg image registration (Modat et al., 2010) and two sets of regularisation tuning parameters, P51 and P52, for elastix image registration (Metz, Klein, Schaap, Walsum, & Niessen, 2011). The niftyreg registration package was extended to correct for the redistribution of attenuating tissue during deformation. Mass preservation required the use of a different set of tuning parameters, P16, whereby the attenuation map was corrected by weighting the map with the Jacobean of the transform.

Table 3.1: Registration options and parameters used in independent multi-modality alignment of the attenuation map to each frame.

Method	P5	P6	P16	P51 & P52
Package	Niftyreg	Niftyreg	Niftyreg	Elastix
Spatiotemporal	No	No	No	No
Mass preserving	No	No	Yes	No
Similarity metric	NMI	NMI	NMI	NMI
Mask	None	None	None	None

Identifying a Reference Frame

The reference frame was the frame in the gated PET scan with the most accurate attenuation map. Several data driven methods were devised for finding a suitable reference frame in the gated PET scan (labelled green in figure 3.2). Initially a similar approach to Nam et al (Nam et al., 2013) was used to choose the reference frame by finding the PET frame that maximised the NMI with respect to the CT-derived attenuation map.

Several approaches were developed based on metrics for sinogram comparison including Natterers consistency, SSD, Chi Squared as well as the negative Log Likelihood. It was suspected that choosing a reference frame by finding the PET frame where these metrics were minimised may be a more accurate approach. The movement of activity in and out of the field of view may lead to a fluctuation in the total number of counts in each sinogram (He et al., 2008). We analysed each metric for its independence from total count level and its suitability to select the best visually matching frame. The number of patients, N_Q, for which the reference frame fell within the quiescent phase was recorded. The quiescent phase of the respiratory cycle was determined from the respiratory trace using a 33% window centred at end expiration.Each metric was compared to the amplitude of the respiratory trace, although the evaluation was limited because the respiratory trace amplitude during the CT was not recorded.

A Mono-Modality Approach to Compensation

A mono modality approach to attenuation alignment was taken by estimating deformation due to respiratory motion during the gated PET scan. In a similar way to Bai and Brady (2011), deformation was estimated through the registration of PET frames reconstructed without attenuation correction (NAC). The reference frame was assumed to have have the most accurate attenuation map. The transformation of each remaining PET frame was found with respect to the reference frame (Figure 3.2). Attenuation mismatch was then compensated for by deforming the attenuation map according to the transform to each NAC PET frame relative to the reference frame.

The accuracy of image registration was compared with two measurements of similarity between NAC PET frames. Image similarity was maximised in the mono-modality approach using the NMI objective for methods P1, P2 and P15 and sum of squared differences (SSD) objective for method P3 (table 3.2). These methods were also compared to spatiotemporal group-wise registration method, P3, which finds a set of transforms which simultaneously minimises the voxelwise variance across all frames (Metz et al., 2011). Due to the computational complexity of spatiotemporal registration, a larger pixel size (4mm) was used and a coarse control point spacing of 26mm x 26mm x 26mm with a temporal control point spacing of 3 frames. Mass preservation was used for method P15, to correct for changes in density due to the redistribution of attenuating tissue.

Table 3.2: Registration options and parameters used in mono-modality registration of each gated PET frame to the reference frame and subsequent deformation of the attenuation map.

Method	P1 & P2	P3	P15	P53
Package	Niftyreg	Niftyreg	Niftyreg	Elastix
Spatiotemporal	No	No	No	Yes
Mass preserving	No	No	Yes	No
Similarity metric	NMI	SSD	NMI	Variance
Mask	None	None	None	None

Combining Mono-Modality and Multimodality Approaches

An approach was devised which comprised of both a multimodality and a mono-modality step (figure 3.3). The first step was to align the attenuation map image to the reference NAC PET frame using multi-modality image registration with a NMI image similarity metric. The reference frame was anticipated to be a close initial match to the attenuation map, thus, the regularisation tuning parameters in multi-modality registration were chosen to allow only for minimal deformation. The second step was to align the reference NAC PET frame to all other PET frames using mono-modality registration. Image similarity was maximised in the second step using the NMI objective for methods P7, P8, P9, P10, P12, P18 and P20 and SSD objective for method P11 (table 3.3). Mass preservation was used for two methods, P18 and P20, to correct for changes in density due to the redistribution of attenuating tissue. Deformation due to respiratory motion was assumed not significantly affect patient posture or be related to the movement of surrounding objects such as the patient bed palate. To avoid spurious estimates of deformation in areas with little or no motion, a mask was applied
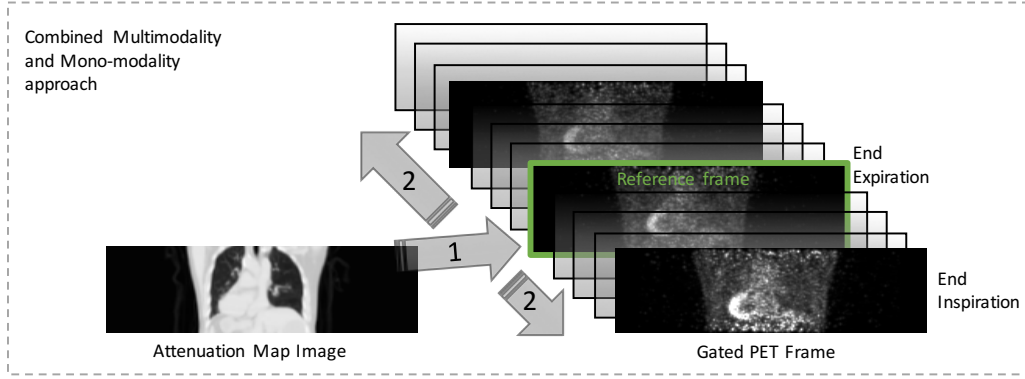
Figure 3.3: Multimodality registration (step 1) is performed in combination with mono modality registration (step 2). The multi modality registration finds the transform between the attenuation map and the PET reference frame. The mono-modality registration finds the transform from the reference frame to each other frames in the gated PET.

in methods P9, P10, P12 and P20 to ensure that any objects outside of the mask did not contribute to the deformation estimate (in table 3.3)

Table 3.3: Registration options and parameters used a combination of multi-modality and mono-modality registration subsequent deformation of the attenuation map.

Method	P7 & P8	P9 & P10	P11
Package	Niftyreg	Niftyreg	Niftyreg
Spatiotemporal	No	No	No
Mass preserving	No	No	No
Similarity metric	NMI	NMI	NMI & SSD
Mask	None	Lung	None

Method	P12	P18	P20
Package	Niftyreg	Niftyreg	Niftyreg
Spatiotemporal	No	No	No
Mass preserving	No	Yes	Yes
Similarity metric	NMI	NMI	NMI
Mask	Whole body	None	Lung

'

3.2.4 Segmentation and Tracking of Lesions

Each PET/CT study had visible lung lesions with increased FDG uptake on the respiratory gated PET image of the thorax. The lesions were segmented from the spiral CT and the volume of the lesion was recorded in millilitres. The position of the lesion was tracked

by transforming the region with the displacements estimated from image registration of sequential gated PET frames. This moving region, referred to as the tumour tracking region, was analysed by recording both the mean (Bq/mL) and total (Bq) activity segmented from the sequence of gated PET frames. Tumour uptake was quantified using a simplified approach of aligning all frames and recording the spatial mean of the activity (Bq/mL) in the tracking region. This approach differed from the registered and summed phases approach (RASP) as the texture and appearance of the lesions were not considered. The effect of respiratory motion on quantification was assessed by measuring the variance of the activity in comparison to the mean of all frames. In addition to this, the tumour motion extent was determined from combining the sequence of tumour tracking regions to form a single tumour volume, TV, which encompasses tumour in every frame. The variation in quantitative parameters of the TV were assessed including the mean (Bq/mL) and total (Bq) TV activity (a_t) through the gated PET sequence. A high standard deviation of the TV activity, σ_{TV}, was also indicative of variations which could be caused by respiratory motion.

3.3 Results

3.3.1 Metrics for Identifying a Reference Frame

Initially the criterion for the reference frame was identified by the maximum NMI. For patients with the reference frame at end expiration we expected correlation of the NMI with the respiratory trace amplitude, however this was not consistent across all patients. Figure 3.4 shows that both SSD2 and NMI were correlated with the amplitude of the respiratory trace for study 16 with Pearson's correlation being $R = 0.85$ for NMI and $R = 0.65$ for SSD2. However, when study 14 was analysed negative correlation between NMI and respiratory trace amplitude was observed as shown in figure 3.5, with Pearson's correlation being $R = -0.51$ for NMI and $R = 0.83$ for SSD2. The reason for the opposite correlation in study 14 was not known, however, it is known that using NMI to match image features can be misleading because organ boundaries on the attenuation map do not always appear on the PET image (Alessio et al., 2010). Despite the shortcomings of NMI as a metric, alternative methods for identifying the reference frame were suspected to produce similar results to using NMI in a majority of studies. Figure 3.6 shows a systematic analysis of the NMI between the emission image and attenuation map for alternative methods of finding the reference frame. Finding the reference frame using the SSD2, ChiSq2 and Natterers metrics resulted in the greatest average NMI. Reference frames identified with the SSD2 metric had close to the maximum NMI with an average NMI of 0.227. It was expected that methods would frequently select a reference frame from the quiescent phase of the respiratory cycle, yet this was only observed in 50% of the studies for the NMI metric (see table 3.4). It was suspected that the NMI metric may be inaccurate due to the lack of similarity between emission images and the attenuation map. The SSD2 metric frequently selected the most quiescent frame making it a suitable alternative for selecting a reference frame in cases where there was limited mutual information.

The number of events per sinogram in a gated study varies from frame to frame and this was found to impact metrics used to identify a reference frame since a high count rate
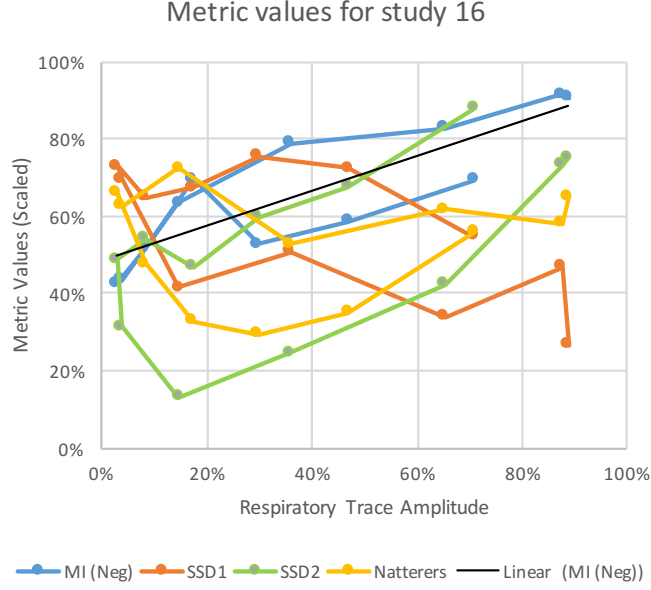
Figure 3.4: The NMI (Negative), Natterers, SSD1 and SSD2 metrics scaled between 0% and 100% for study 16. The metric values were temporally smoothed with a 3 frame boxcar kernel. The attenuation map similar to the emission image when negative NMI is minimum. The NMI (Negative) metric was correlated with respiratory amplitude $R = 0.85$ and a line of best fit (black) shows the direction of the correlation.

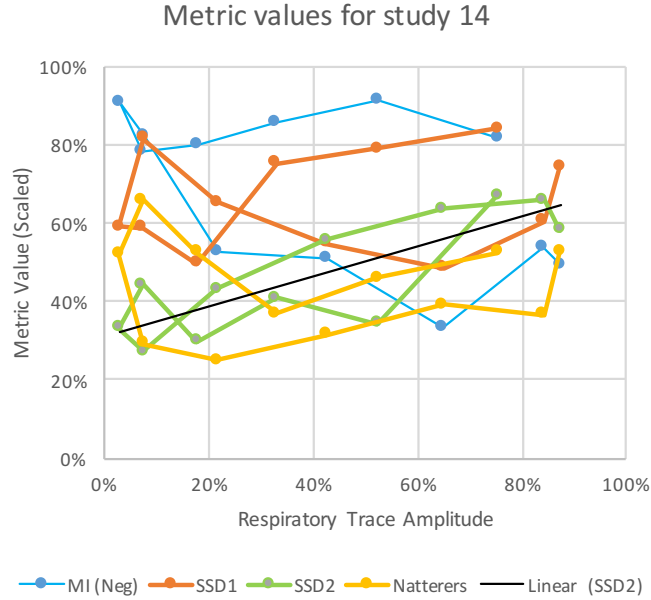Figure 3.5: The NMI (Negative), Natterers, SSD1 and SSD2 metrics scaled between 0% and 100% for study 14. The metric values were temporally smoothed with a 3 frame boxcar kernel. The attenuation map similar to the emission image when negative NMI is minimum. The SSD2 metric was correlated with respiratory amplitude $R = 0.82$ and a line of best fit (black) shows the direction of the correlation.
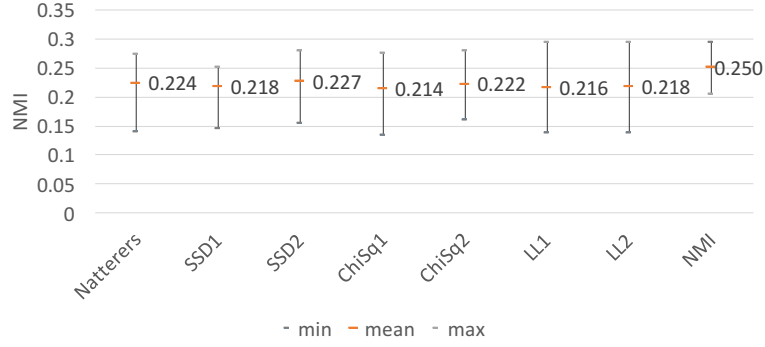
Figure 3.6: The NMI of the designated reference frame with the attenuation map when using various metrics to find the reference frame. The NMI was compared for all 11 studies and the minimum (least similar) and maximum (most similar) results are shown. The SSD2 metric had the closest overall performance to the NMI metric.
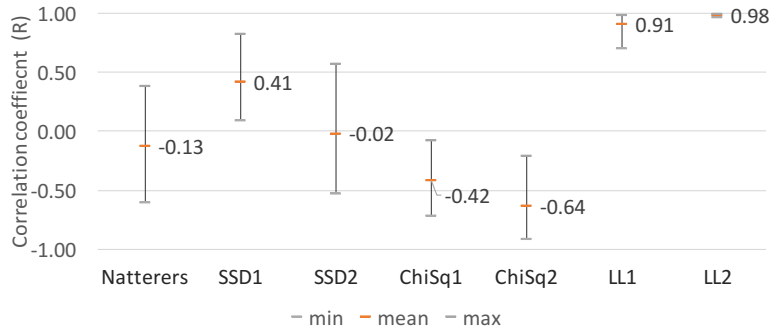

Figure 3.7: The correlation of metrics with the total number of emissions in each respiratory gated bin. The Pearson's correlation (R) was calculated separately for each study (n=11) and the mean, minimum and maximum correlation coefficient was plotted. The Natterers and SSD2 similarity metrics had the lowest correlation with total number of emissions.

frame might result in a higher than expected metric value. It was assumed that there should be minimal correlation between the number of events and the metric values, however, it was found that some metrics such as the log likelihood metrics were highly correlated with the total number of events in each frame. Pearson's correlation between the count level for each frame versus the metric value was calculated for each frame for all 11 studies (Figure 3.7). A significant correlation ($R > 0.9$) ruled out log likelihood metrics because they were disproportionately likely to identify frames with low counts as a reference frame.

The quiescent phase of the respiratory cycle was defined as the portion of the cycle when the respiratory trace amplitude was less than 33% of the end inspiration value. It was expected that the attenuation map would be a reasonably accurate representation of the anatomy during the quiescent phase because the CT acquisition protocol involved the operator directing the patient to breathe out during the scan. Table 3.4 shows that the SSD2 metric designated the reference frame as during the quiescent phase of the respiratory cycle in 88% of studies. 3 studies were excluded from this analysis due to difficulties in synchronising the respiratory trace with the gated PET scan. The study for which the reference frame was

65

not in the quiescent phase had a respiratory motion artefact on the whole body PET/CT image (Figure 1.2). The NMI, Natterers, and ChiSq2 metrics were ruled out for identifying the reference frame because they designated the reference frame as during quiescent phase of the respiratory cycle in only 50% of studies.

Table 3.4: A comparison of 8 studies analysed with various metrics. N_Q is the number of studies where the metric designated a frame in the quiescent phase of the respiratory cycle. The SSD2 metric was most likely to identify a frame in the quiescent phase.

Metric	N_Q
NMI	4
SSD1	5
SSD2	7
ChiSq1	6
ChiSq2	3
LL1	6
LL2	5
Natterers	4

3.3.2 Metrics for Evaluating the Accuracy of Attenuation Alignment

The same metrics used for identifying a reference frame were used to compare registration methods, however, some metrics needed to be excluded. It was shown in section 3.2.1 that the metrics defined by equations 3.1 - 3.8 increase due to attenuation mismatch. Upon further analysis, we excluded the use of log likelihood metrics based on the previously identified correlation with the total number of counts per frame because the correlation may confound attempts to compare studies with significantly different count levels.

When an attenuation map is deformed, the average ACF value per frame is likely to change. If the metric is sensitive to changes in the magnitude of the average ACF, then an overall change in the ACF as a result of deformation would potentially be mistaken for a misalignment. Analysis of the equation 3.6 for SSD2 and equation 3.8 for ChiSq2 revealed a linear dependance of the metric on the magnitude of the average ACF. The correlation of the metrics with the average ACF value was confirmed for all studies and a scatter plot for the SSD metric for study 14 is shown in figure 3.8. Although SSD2 was the optimal metric for finding the reference frame it was necessary to rule out both SSD2 and ChiSq2 metrics for measuring the accuracy of the alignment of emission images and deformed attenuation maps.

Although it is possible to use NMI as a measure of the alignment of attenuation maps with emission images, it is not possible use NMI as an independent measurement of accuracy if NMI has already been used for estimating the deformation. We used NMI extensively in registration methods presented in tables 3.1, 3.2 and 3.3, thus NMI was excluded as an evaluation metric.

This analysis of metrics ruled out the use of log likelihood, SSD2, ChiSq2 and NMI metrics. The results of SSD1, ChiSq1 and Natterer's as a metric for evaluating alignment
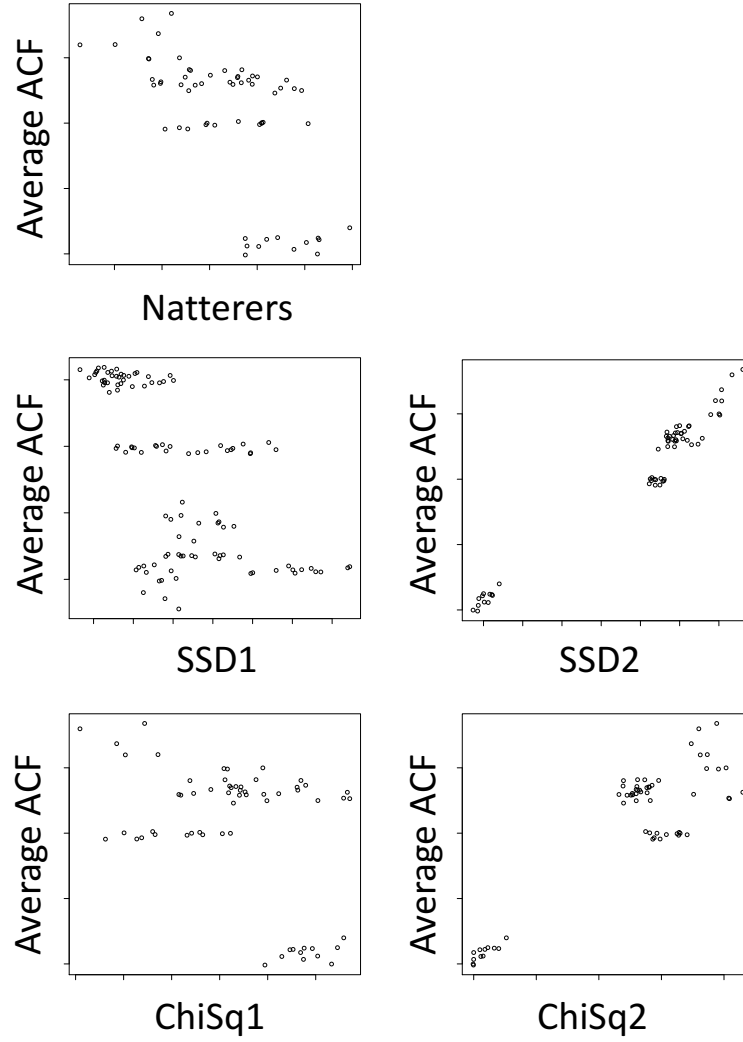
Figure 3.8: A scatter plot showing the relationship between SSD, Chisq and Natterer's metrics and the average attenuation correction factor for study 14. A correlation with the ACF was observed for both the SSD2 and ChiSq2 metrics.
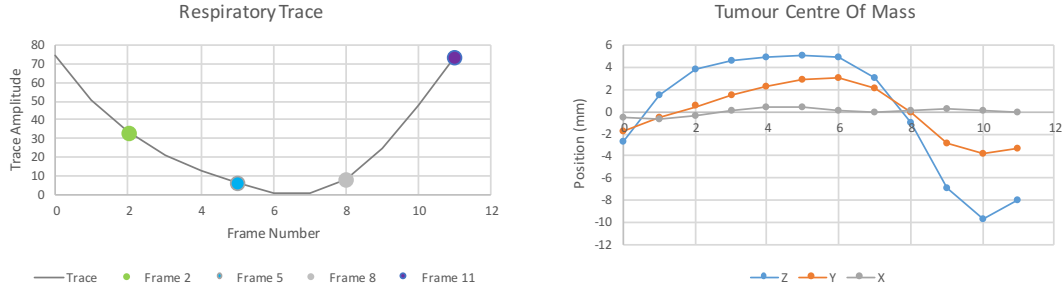
Figure 3.9: The respiratory amplitude where end inspiration is at frame 11 and end expiration is between frame 6 and 7. Frames 2, 5, 8 and 11 are highlighted with the same colours as the segmented regions in figure 3.10. The plot on the right shows centre of mass (COM) of the segmented tumour region for each frame. The tumour moves predominantly in the superior and inferior direction (z direction) over a 15mm range. There is also anterior and posterior motion (y direction) over a range of 7mm.

methods are presented in the final section.

3.3.3 Analysis of Lesion Deformation

An analysis of lesion deformation was completed following the compensation for attenuation misalignment. Figure 3.9 shows that the centre of mass of segmented lesions moved in the superior direction at end expiration (EE) and to the inferior extent at end inspiration (EI). Although it was possible to trace the tumour volumes separately for each frame, we found it reliable to use a larger region (TV), shown in figure 3.10, to measure total activity (in kBq) using the same segmented volume in every frame.

An analysis of the apparent activity concentration (Bq/mL) in the tracking region was performed for multimodality, mono-modality and combined approaches (figure 3.11). Mass preserving image registration was applied by correcting for density changes following the deformation of the attenuation map and the effect on apparent activity concentration was recored. In each approach to registration it was found that mass preservation decreased the variance and decreased the mean activity in the emission image. Mass preserved attenuation maps at end inspiration (EI) appeared to have a lower density in comparison to end expiration (EE). The decrease in density caused the average amount of attenuation to decrease and was likely responsible for the corresponding decrease in average tumour activity.

3.3.4 Optimal Methods for Compensating for Attenuation Misalignment

Table 3.12 summarises the results of compensating for attenuation misalignment using several methods of image registration for study 14. The change in Natterers, SSD1 and ChiSq1 metrics using compensated attenuation maps varied significantly in comparison to CTAC. In cases where image registration accurately compensated for alignment it was expected that the metrics would decrease relative to the CTAC value. For example, multimodality
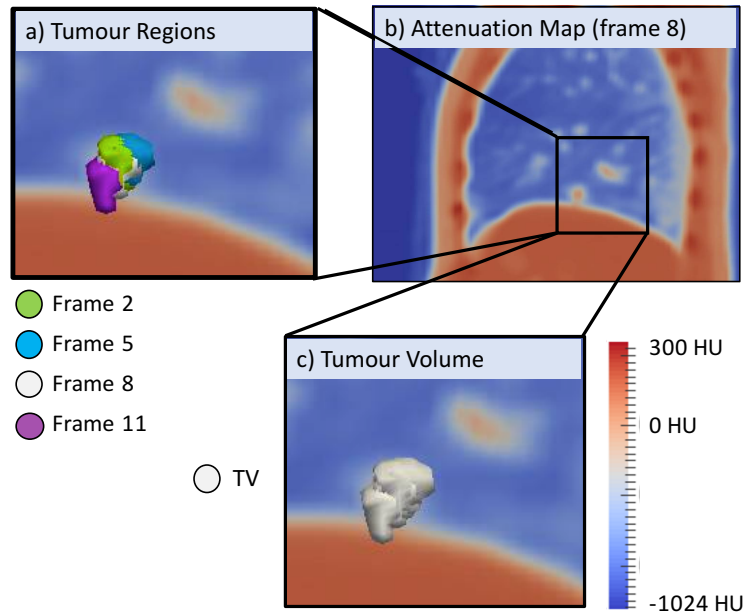
Figure 3.10: The segmented tumour in study 16 is superimposed on a coronal cross section of the attenuation map for PET/CT. Subfigure (b) shows the reference frame (frame 8) with no transform applied to the tumour or the attenuation map. Subfigure (a) shows the deformed tracking region for frame 2 (Green), frame 5 (Light blue), frame 8 (White) and frame 11 (Purple). Subfigure (c) shows the tumour volume (TV) extent which is a larger region that encapsulates the movement of the tumour over the entire respiratory cycle.

Figure 3.11: Tumour lesion activity in the tracking region for selected registration parameter sets for study 14. The activity is a function of respiratory amplitude with results for no attenuation alignment shown in gray (CTAC). The blue plots show the lesion activity without mass preservation using multimodality (P6), mono-modality (P2) and combined multi-modality and mono-modality (P10) parameters. The orange plots show the lesion activity with mass preservation using multimodality (P16), mono-modality (P15) and combined multi-modality and mono-modality (P20) parameters leading to a decrease in average tumour activity.

Table 3.5: Attenuation alignment metrics calculated from the median metric difference for each frame study 14. Differences are expressed as a percentage of the metric when CTAC was used without attenuation alignment. The last two columns compare tumour lesions using the standard deviation of the TV activity and the volume of the TV. Each row has metric values for a different registration parameter set described in tables 3.1, 3.2 and 3.3.

Multimodality	Natterers	SSD1	ChiSq1	σ_{TV} (kBq)	TV vol (mL)
P5	-7.1%	-1.6%	2.2%	2.92	9.4
P6	-11.9%	-2.1%	2.4%	3.60	9.3
P16	45.8%	-0.7%	3.9%	3.30	9.3
P51	-31.4%	0.3%	1.5%	6.04	12.8
P52	-48.3%	-8.6%	-3%	5.79	16.6
mono-modality	Natterers	SSD1	ChiSq1	σ_{TV} (kBq)	TV vol (mL)
P1	1.9%	-0.3%	1%	5.47	11.3
P2	7.1%	0.2%	1%	6.16	10.9
P3	-1.1%	1.8%	1%	6.54	11.8
Combined	Natterers	SSD1	ChiSq1	σ_{TV} (kBq)	TV vol (mL)
P7	71.0%	0.2%	3.5%	5.82	16.1
P8	27.2%	-2.0%	1.4%	4.83	15.4
P9	-94.7%	-2.1%	-0.1%	4.03	21.4
P11	-23.9%	5.1%	2.2%	7.03	14.5
CTAC				3.53	9.1

registration method, P52, reduced both Natterers and SSD1 metrics by 48.3% and 8.6% respectively, yet alternative methods such as P7 increased all metrics. Compensating for misalignment did not significantly affect the variation of radioactivity concentrations across frames as seen in methods P5, P6 and P16 in table 3.12. However, for two methods, P51 and P52, the registration algorithm estimated large deformations leading to variations in TV activity (6.04kBq and 5.79kBq) which were large in comparison to methods P5, P6 and P16. Multi-modality method P52 led to estimates of substantial deformation, hence motion, of the tumour. This was evident by the size of the TV which had increased from 9.1mL (for CTAC) to 16.6mL (for P52). Our analysis of table 3.5 indicated that method P52 was the best performing multi-modality method reducing Natters and SSD1 metrics by 48.3% and 8.6% respectively. The mean activity was noticeably higher for method P52 (figure 3.12), however, this was thought to be related to the size of the TV and not related to the uptake of the tumour.

Compensating for attenuation misalignment using a combination of both mono-modality and multi-modality registration led to the most significant improvements. The best performing method reduced Natterer's and SSD1 metrics by 94.7% and 2.1% respectively. Method P9 in table 3.5 shows that size of the TV was larger than any other method (21.4mL). However, despite the large estimation of tumour motion, the variance in TV activity was lowest ($\sigma_{TV} = 4.03kBq$).

The median metric differences are shown in table 3.6 and expressed as a percentage of the metric for CTAC. We found that three multimodality registration methods P5, P6,
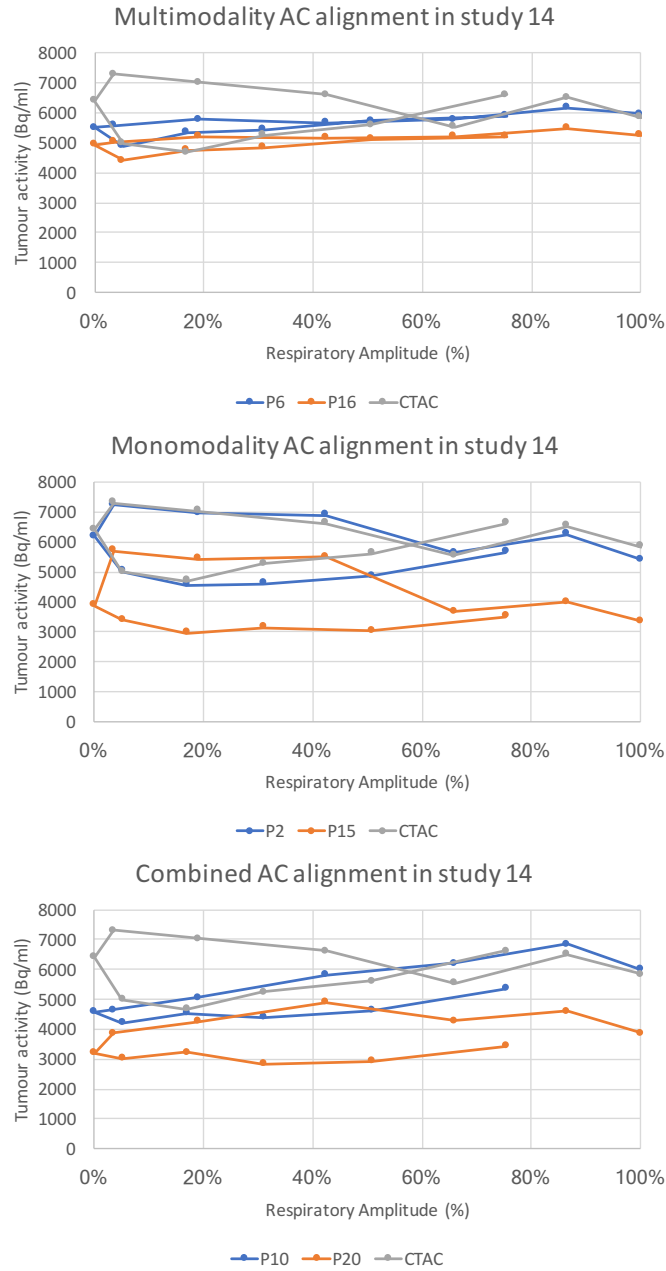
Figure 3.12: Tumour lesion activity in the TV for selected registration parameter sets for study 14. The activity is a function of respiratory amplitude with results from no attenuation alignment shown in black (CTAC). A comparison of lesion activity using multimodality parameters P5 (blue), P6 (orange), P16 (yellow) and P51 (green) is top left. A comparison of lesion activity using mono-modality parameters P1 (blue), P2 (orange), P3 (yellow) and P52 (green) is top right. A comparison of lesion activity using a combined multi-modality and mono-modality P7 (blue), P8 (orange), P9 (yellow) and P11 (green) is bottom.
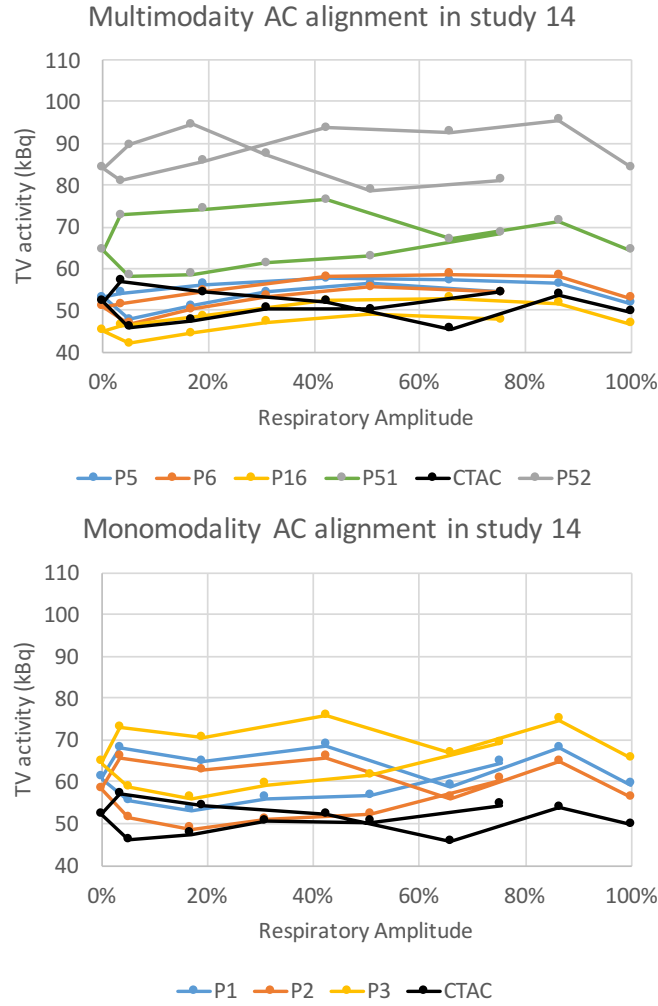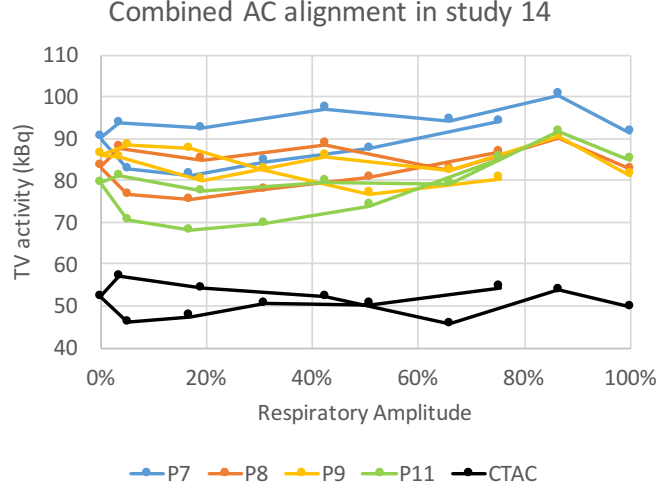
Figure 3.13: Tumour lesion activity in the TV for selected registration parameter sets for study 14. The activity is a function of respiratory amplitude with results from no attenuation alignment shown in black (CTAC). A comparison of lesion activity using multimodality parameters P5 (blue), P6 (orange), P16 (yellow) and P51 (green) is top left. A comparison of lesion activity using mono-modality parameters P1 (blue), P2 (orange), P3 (yellow) and P52 (green) is top right. A comparison of lesion activity using a combined multi-modality and mono-modality P7 (blue), P8 (orange), P9 (yellow) and P11 (green) is bottom.

P16 did not reduce the median Natterers metric after attenuation alignment. The metric differences between attenuation alignment versus CTAC were analysed for all 11 studies. The optimal multimodality method was P52 in table 3.6 which reduced Natters metric by 10.1% and the SSD1 metric by 8.4%. Using mono-modality registration to align frames resulted in a reduction in Natterers metric by up to 9.5%. Combining both mono-modality and multi modality registration led to almost no reduction in Natterers consistency and SSD1 metric. P9 in table 3.8 was the best performing combined approach leading to a reduction in the Natterers consistency and SSD1 metric by 1.8% and 3.0% respectively. The optimal multimodality (P52), mono-modality (P1) and combined (P9) approaches are compared in figure 3.14).

Attenuation image alignment using multi-modality registration method P52 reduced Natterer's consistency, SSD1 metric and ChiSq1 metric for mono-modality by the largest magnitude. A boxplot analysis (figure 3.14) showed that the SSD1 metric and ChiSq1 metrics had the most significant statistical difference when comparing methods. Although the magnitude of the changes in Natterer's consistency were larger, the results appeared equivocal when comparing different methods for multiple studies.

For both multi-modality and combined approaches, mass preservation decreases the variance and decreases the mean activity. This makes sense as most CTs are taken at EE and then stretched to EI. In this case the density of the AC map should decrease overall.

Table 3.6: Attenuation alignment metrics for multi-modality alignment calculated from the median metric difference over 11 studies. Differences are expressed as a percentage of the metric when CTAC was used without attenuation alignment. Each row has metric values for a different registration parameter set described in table 3.1.

Multi-modality	Natterers	SSD1	ChiSq1
P5	17.3%	-3.2%	-1.9%
P6	18.2%	-2.4%	-1.3%
P16	22.1%	-1.4%	-0.7%
P51	-5.0%	0.6%	0.5%
P52	-10.1%	-8.4%	-4.9%

Table 3.7: Attenuation alignment metrics for mono-modality alignment calculated from the median metric difference over 11 studies. Differences are expressed as a percentage of the metric when CTAC was used without attenuation alignment. Each row has metric values for a different registration parameter set described in table 3.2

Mono-modality	Natterers	SSD1	ChiSq1
P1	-9.3%	-1.9%	-1.4%
P2	-7.1%	-2.1%	-1.6%
P3	-9.5%	-1.8%	-1.4%
P15	-5.3%	-1.4%	-0.8%
P53	-4.2%	0.0%	-0.2%

Table 3.8: Attenuation alignment metrics for combined alignment calculated from the median metric difference over 11 studies. Differences are expressed as a percentage of the metric when CTAC was used without attenuation alignment. Each row has metric values for a different registration parameter set described in table 3.3.

Combined	Natterers	SSD1	ChiSq1
P7	24.5%	-1.9%	-1.4%
P8	17.1%	-3.5%	-2.6%
P9	-1.8%	-3.0%	-1.6%
P10	2.3%	2.5%	0.9%
P11	7.1%	-1.3%	-1.0%
P12	42.4%	5.9%	2.3%
P18	20.5%	-2.0%	-1.1%
P20	1.3%	1.5%	1.5%

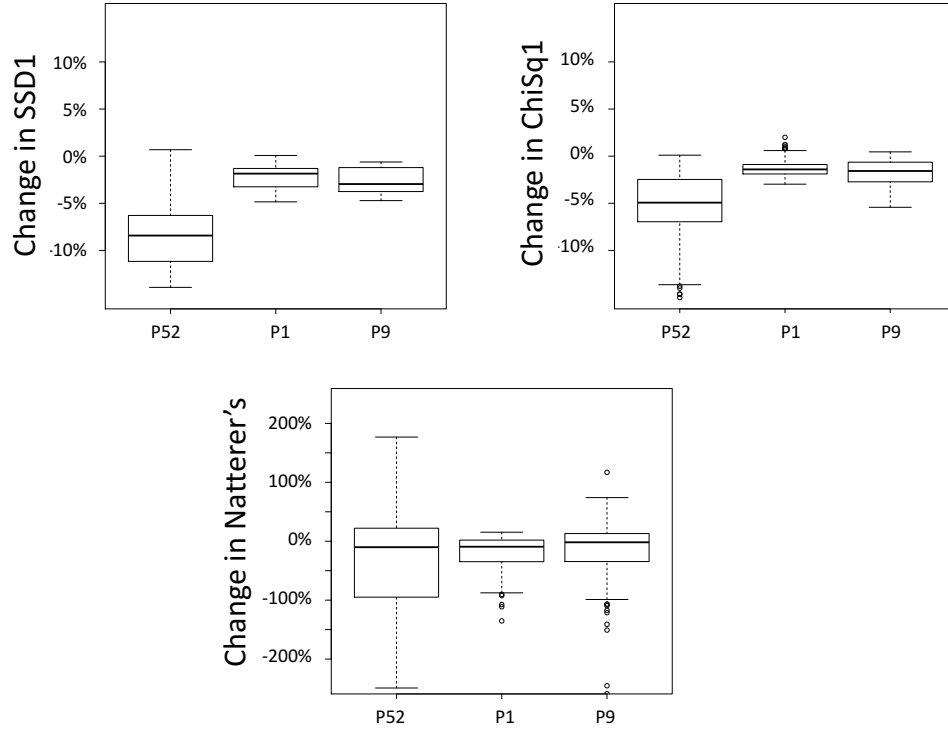Figure 3.14: A comparison of multimodality (P51), mono-modality (P52) and combined (P9) approaches with attenuation alignment metric differences. Comparisons are made using the Natterer's consistency (bottom), SSD1 (top left) metric and ChiSq1 (top right) metrics. Each box and whisker shows the metric difference in 11 studies expressed as a percentage of the metric when CTAC was used without attenuation alignment.

3.4 Discussion

Seven different metrics were devised for the purposes of validating methods of image registration to address attenuation mismatch. The SSD2 metric was found to perform best at finding the reference frame in a gated study when CTAC is applied. The SSD1 or ChiSq1 metrics were suited to validation of the alignment of attenuation maps to each frame in a gated study. We used these metrics to investigate attenuation alignment in an example gated study and found that attenuation alignment using multi-modality registration was the optimal method. We analysed 11 studies to find that multi-modality elastix registration of the attenuation map to the emission image reduced the SSD1 and ChiSq1 metrics (p <0.05). Mono-modality and combination approaches were also promising, but did not result in significant metric improvements. It is suspected that mono-modality registration performed poorly because the respiratory motion signal in the PET data was confounded by Poisson noise. Further work is presented in chapter 5 in improving the estimation of motion from the PET data by regularising image registration so that it is robust to the poor SNR in gated PET images. The analysis of tumour lesions revealed that some alignment methods led to an underestimation of the tumour volume extent. A smaller than expected TV volume was likely due to image registration inadequately modelling the deformation due to respiratory motion. Further work on modelling respiratory motion is presented in chapters 4 and 5. A complete investigation of mass preservation could not be completed since mass preservation had not been fully implemented for all cost functions in the Elastix and Niftyreg registration packages. Further work may reveal more promising results with this approach, however, it would require a significant amount of software development. Although a limited investigation into mass-preserving registration was performed (Barnett, Fulton, & Meikle, 2013), those results were not conclusive enough to rule out that the encountered benefits didn't arise due to unrelated differences in the implementation of the registration algorithm. In some cases the Natterer's consistency condition result conflicted with the SSD1 and ChiSq1 metrics, however, further analysis showed that Natterer's consistency had a wide distribution of values when comparing multiple studies in figure 3.14. This could be due to Natterers consistency condition being oversensitive to certain symmetries in the sinogram (Alessio et al., 2010) meaning that the metric may not be adequate as a method of validation of attenuation correction. The work in this chapter forms a basis for using the SSD1 metric in chapter 5 to pursue further improvements.

References

Alessio, A. M., Kinahan, P. E., Champley, K. M., & Caldwell, J. H. (2010). Attenuation-emission alignment in cardiac PET/CT based on consistency conditions. *Medical physics*, *37*(3), 1191–1200.

Bai, W., & Brady, M. (2009). Regularized b-spline deformable registration for respiratory motion correction in PET images. *Physics in Medicine and Biology*, *54*(9), 2719–2736. doi: 10.1088/0031-9155/54/9/008

Bai, W., & Brady, M. (2011). Motion correction and attenuation correction for respiratory gated PET images. *Medical Imaging, IEEE Transactions on*, *30*(2), 351–365. doi:

10.1109/TMI.2010.2078514

Barnett, R., Fulton, R., & Meikle, S. (2013). Robust estimation of respiratory motion from PET image sequences through minimisation of voxel-wise variance. In *Nuclear science symposium and medical imaging conference (NSS/MIC), 2013 IEEE.*

Berker, Y., & Li, Y. (2016). Attenuation correction in emission tomography using the emission dataa review. *Medical Physics*, *43*(2), 807–832.

Bruyant, P. P., Le Rest, C. C., Turzo, A., Jarritt, P., Carson, K., & Visvikis, D. (2007). A method for synchronizing an external respiratory signal with a list-mode PET acquisition. *Medical physics*, *34*(11), 4472–4475.

Dawood, M., Buther, F., Lang, N., Schober, O., & Schafers, K. P. (2007). Respiratory gating in positron emission tomography: A quantitative comparison of different gating schemes. *Medical Physics*, *34*(7), 3067. doi: 10.1118/1.2748104

Dawood, M., Lang, N., Jiang, X., & Schafers, K. (2006). Lung motion correction on respiratory gated 3-d PET/CT images. *Medical Imaging, IEEE Transactions on*, *25*(4), 476–485. doi: 10.1109/TMI.2006.870892

Fayad, H. J., Lamare, F., Rest, C. C. L., Bettinardi, V., & Visvikis, D. (2013). Generation of 4-dimensional CT images based on 4-dimensional PETderived motion fields. *Journal of Nuclear Medicine*, *54*(4), 631–638. doi: 10.2967/jnumed.112.110809

He, J., O'Keefe, G. J., Gong, S. J., Jones, G., Saunder, T., Scott, A. M., & Geso, M. (2008). A novel method for respiratory motion gated with geometric sensitivity of the scanner in 3d PET. *IEEE Transactions on Nuclear Science*, *55*(5), 2557–2565.

Jing Wang, Tianfang Li, Hongbing Lu, & Zhengrong Liang. (2006). Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose x-ray computed tomography. *IEEE Transactions on Medical Imaging*, *25*(10), 1272–1283. doi: 10.1109/TMI.2006.882141

Kinahan, P., MacDonald, L., Ng, L., Alessio, A., Segars, P., Tsui, B., & Pathak, S. (2006). Compensating for patient respiration in PET/CT imaging with the registered and summed phases (RASP) procedure. In *Biomedical imaging: Nano to macro, 2006. 3rd ieee international symposium on* (pp. 1104–1107).

Kybic, J. (2010). Bootstrap resampling for image registration uncertainty estimation without ground truth. *IEEE Transactions on Image Processing*, *19*(1), 64–73.

Liu, C., Alessio, A., Pierce, L., Thielemans, K., Wollenweber, S., Ganin, A., & Kinahan, P. (2010). Quiescent period respiratory gating for PET/CT. *Medical physics*, *37*(9), 5037–5043.

McQuaid, S. J., Lambrou, T., & Hutton, B. F. (2011). A novel method for incorporating respiratory-matched attenuation correction in the motion correction of cardiac PETCT studies. *Physics in medicine and biology*, *56*(10), 2903.

Metz, C. T., Klein, S., Schaap, M., Walsum, T. v., & Niessen, W. J. (2011). Non-rigid registration of dynamic medical imaging data using nD + t b-splines and a groupwise optimization approach. *Medical Image Analysis*, *15*(2), 238–249. doi: 10.1016/j.media.2010.10.003

Modat, M., Ridgway, G. R., Taylor, Z. A., Lehmann, M., Barnes, J., Hawkes, D. J., ... Ourselin, S. (2010). Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine*, *98*(3), 278–284. doi: 10.1016/j.cmpb.2009.09.002

Murphy, K., Van Ginneken, B., Reinhardt, J. M., Kabus, S., Ding, K., Deng, X., ... others (2011). Evaluation of registration methods on thoracic CT: the EMPIRE10 challenge. *IEEE transactions on medical imaging*, *30*(11), 1901–1920.

Nam, W. H., Ahn, I. J., Kim, K. M., Kim, B. I., & Ra, J. B. (2013). Motion-compensated PET image reconstruction with respiratory-matched attenuation correction using two low-dose inhale and exhale CT images. *Physics in medicine and biology*, *58*(20), 7355.

Pönisch, F., Richter, C., Just, U., & Enghardt, W. (2008). Attenuation correction of four dimensional (4d) PET using phase-correlated 4d-computed tomography. *Physics in Medicine and Biology*, *53*(13), N259. doi: 10.1088/0031-9155/53/13/N03

Presotto, L., Busnardo, E., Perani, D., Gianolli, L., Gilardi, M. C., & Bettinardi, V. (2015). Simultaneous reconstruction of attenuation and activity in cardiac PET can remove CT misalignment artifacts. *Journal of Nuclear Cardiology*, 1–12.

Thielemans, K., Asma, E., & Manjeshwar, R. M. (2009). Mass-preserving image registration using free-form deformation fields. In *2009 IEEE nuclear science symposium conference record (NSS/MIC)* (pp. 2490–2495). doi: 10.1109/NSSMIC.2009.5402070

Tsui, B. M., Lalush, D. S., Frey, E. C., King, M. A., Manocha, D., & Segars, W. (2001). Development and application of the new dynamic nurbs-based cardiac-torso (NCAT) phantom. *School of Nursing Faculty Publication Series*, *42*(5).

Chapter 4

Cross Population Respiratory Motion Modelling for Improving Attenuation Correction

The following chapter is published as, Barnett R, Meikle S, Fulton R . *Cross population motion modeling applied to attenuation correction of respiratory gated F18-FDG PET*. IEEE Transactions on Nuclear Science, Volume: 63, Issue: 1, 170-179, Feb. 2016 with minor editions to reduce length.

4.1 Introduction

Attenuation correction relies on the PET and attenuation maps being spatially aligned and thus any mismatch between the PET and CT affects the diagnostic accuracy of the PET/CT scan. Positron emission tomography and computed tomography are typically acquired using a single gantry (PET/CT). Although mismatch is significantly reduced because the patient is not required to move from the bed between each scan, compensation for mismatch due to involuntary respiratory motion was demonstrated to be beneficial in chapter 3. A typical clinical implementation of PET respiratory gating in PET/CT is to use the same spiral CT image to generate an identical attenuation map for all respiratory gated PET frames. CT imaging takes only a few seconds to traverse the thoracic region so it is affected by respiration to a lesser degree. In this case respiratory motion leads to the appearance of attenuation mismatch artifacts in the reconstructed PET image (Hamill, Bosmans, & Dekker, 2008; McQuaid, Lambrou, & Hutton, 2008).

There are many promising methods for motion compensation during a PET scan (Chan et al., 2013; Lamare et al., 2007), however, few address the effect of respiratory motion on CT-based attenuation correction (Bai & Brady, 2011). One approach to mitigating CT motion artifacts in the thoracic region is to perform a breath hold CT image (Nehmeh et al., 2007). Alternatively, it is possible to use 4DCT to capture the motion of organs by acquiring a cine of the respiratory cycle. Organ motion can then be corrected through analysis of the 4DCT frames (Zhang et al., 2010). A common challenge to all of these approaches is that they require identifying a reference frame when the gated PET image matches the CT derived attenuation map.

A more accurate attenuation map can be obtained by acquiring a 4DCT which is phase correlated with the gated PET such that each frame in the 4DCT matches the corresponding frame in the gated PET scan (Pönisch, Richter, Just, & Enghardt, 2008). Unfortunately, the 4DCT frames may not match the gated PET frames because the 4DCT is only acquired over 4-5 respiratory cycles and the breathing motion during these cycles may be significantly different to the breathing motion during the PET scan which is averaged over many respiratory cycles. The 4DCT is also affected by motion causing the appearance of image artifacts such as distorted organ boundaries. Such artifacts are typically mitigated through additional time sampling (Pan, Lee, Rietzel, & Chen, 2004), alternatively it is possible to address the same issue using motion compensated reconstruction (Ehrhardt et al., 2007).

4DCT derived attenuation correction may be promising if the radiation dose from the scan can be reduced significantly such that it is preferable to a conventional low dose spiral CT scan. Even if the dose can be reduced, 4DCT acquisitions that have been localized to the thoracic region usually require an additional spiral CT acquisition to apply attenuation correction if it is to be combined with a whole body PET scan.

An alternative approach to 4DCT is to merely apply accurate attenuation correction

for a single PET frame (Daouk, Fin, Bailly, & Meyer, 2008). In this approach attenuation mismatch is avoided by finding a single frame in the gated PET which best matches the spiral CT and discarding the remainder of the PET data. However, it would be preferable to be able to apply attenuation correction for all PET frames. Chapter 2 proposed and compared methods of deformable image registration as a promising solution for this problem whereby the attenuation mismatch is compensated by deforming the spiral CT image in such a way that it matches each frame of a respiratory gated PET.

An alternative approach to correcting for respiratory motion is to consider a patient specific model of respiration in which the motion model is tailored to the patient. This approach has been proposed as a suitable model for identifying individual breathing patterns (Li et al., 2011; Zhang et al., 2007). Patient specific models find underlying patterns in respiratory motion and can be used for the prediction of arbitrary phase and amplitude of respiration (Thielemans et al., 2013). Without a model, motion estimation would normally require accurate time sequence images for many respiratory cycles.

A direct correspondence model can relate the observed internal motion in a patients 4DCT to the external movement of a respiratory gauge. Although correspondence has been demonstrated in PET imaging for single tumors (Chan et al., 2013), it is yet to be reliably demonstrated for attenuation correction for which motion needs to be modeled in the entire thoracic region.

Cross population motion models are formed from motion data acquired from many different subjects, and attempt to capture the nature of breathing motion across the population (McClelland, Hawkes, Schaeffter, & King, 2013). A cross population model may be considered in cases where it is impracticable to construct a patient specific correspondence model. Although this method was shown to allow comparison of breathing motion between subjects (Ehrhardt, Werner, Schmidt-Richberg, & Handels, 2011), it was not applied to predict the motion of the entire thorax.

Cross population models have a demonstrated use in image segmentation. Once such type of cross population model is the statistical shape model used for refining the boundaries of a segmented region (Heimann & Meinzer, 2009). Cross population models can also be used to analyze the motion of a segmented region. For example, a statistical shape model was used to approximate the shape of the diaphragm in imaging the radionuclide uptake of the myocardium (McQuaid et al., 2008). The rise and fall of the diaphragm is the main contributor to attenuation correction inaccuracies. Using the model to implement motion compensation improved diagnostic accuracy, even though the model could have been improved by considering expansion and contraction of the chest wall.

The cross population model this chapter builds on a model proposed by Chandrashekara, Rao, Sanchez-Ortiz, Mohiaddin, and Rueckert (2003) which was presented in section 2.5.3. Rather than parameterize at every voxel location, B-Spline coefficients were used to interpolate between parameterized positions (equation 2.19). The direction of heart motion in the population was captured through a set of eigenvectors which were then scaled with time dependent coefficients (**b**). The practical advantage of this model was that the nature of

heart motion was captured in a single function $\mathcal{K}$ (equation 4.1) and when a new patient scan was encountered it was only necessary to calculate a new set of scaling coefficients ($\mathbf{b}$).

$$d_{(0,t)}^{S}(\dot{x}, \dot{y}, \dot{z}) = \mathcal{K}(\mathbf{b}) \tag{4.1}$$

The parameterization overcame issues caused by the poor SNR in images acquired using fast sequences in dynamic MR. Our implementation of a cross population model for lung motion is an adaptation of this parameterization with the purpose of overcoming the poor signal to noise ratio in PET imaging.

Our aim is to address attenuation mismatch due to respiratory motion in order to improve the accuracy of attenuation correction. In this paper we build on existing methods for deformable image registration to estimate motion from respiratory gated images. Rather than using deformable image registration alone, we explore the possible use of a cross population motion model to calculate a set of patient specific deformation fields due to respiratory motion during a gated PET scan. We build on previous work which utilizes a cross population model for the heart (Chandrashekara et al., 2003) and the diaphragm (McQuaid et al., 2008). We combine our cross population model with a correspondence model to estimate motion of the whole thorax from a sequence of gated images in the PET scan.

4.2 Methods

Cross population modeling involves analyzing a collection of input scans to find similarities of breathing motion across the population. Our method of utilizing a cross population model is summarized in Figure 4.1. We pre-processed a collection of 4DCT images from which frame by frame deformations were estimated to establish a cross population model (denoted by $\mathcal{K}$). The cross population model, $\mathcal{K}$, was used to address the issue that estimating deformation from gated PET using frame-by-frame image registration (denoted by Q) was adversely affected by poor signal to noise (SNR). We extended the cross population model to ensure correspondence with the initial deformation estimate, Q, and produced an improved deformation estimate, V. The following sections describe our proposed joint correspondence and cross population model and how we used the deformation estimate, V, to address mismatch between each gated image and the spiral CT.

4.2.1 Data Acquisition and Image Reconstruction

A set of 10 individual 4DCT thorax scans were acquired in cine mode using GE LightSpeed RT 16 scanner in a similar way to the method adopted by Nemeh (Nehmeh et al., 2004). Cine mode was implemented by acquiring repeated axial CT images for a specified period of time, at each table position. The 10 4DCT thorax scans were from a cohort of oncology patients undergoing a respiratory gated CT for the purposes of radiation treatment planning. Breathing motion was tracked by monitoring the vertical displacement of two infrared reflective markers
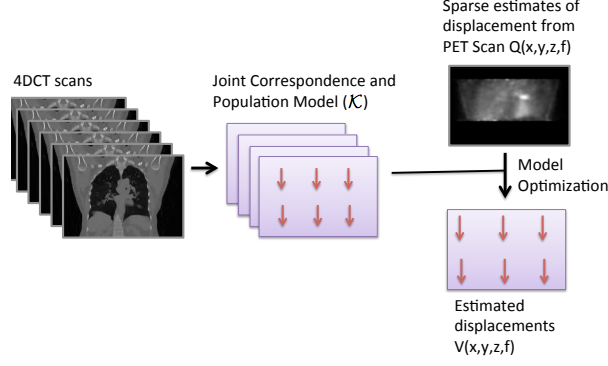
Figure 4.1: A collection of 4DCT scans are used to construct a motion model. We use a combination of the motion model and an sparse estimation from the gated PET scan to construct a refined deformation field.

on the patient's chest using the Varian real-time position management (RPM). The phase in the respiratory cycle was measured relative to the time when the RPM marker had maximum displacement (End inspiration). Each axial CT slice was retrospectively binned into 10 frames based on the phase in the respiratory cycle.

Eleven PET scans were acquired and analyzed with respiratory gating on a Biograph 16 PET/CT using a phase based gating system. All of these gated PET scans were independent from the 4DCT scans used to train the population model. The phase based respiratory gating system AZ733V (Anzai Medical, Tokyo, Japan) was employed as previously described by Pönisch et al. (2008). The system utilizes a pressure sensor inserted into the pocket of an elastic belt to detect external respiratory motion.

Each patient was injected with approximately 350MBq of F18-FDG. A single bed list mode scan with total duration of 5 minutes and an approximate sensitivity of 4cps/kBq was performed after a conventional whole body static PET and spiral CT. The list mode data were binned into 6 frames to produce a set of respiratory gated images. Conventional attenuation correction of each respiratory gated image was implemented from the spiral CT for comparison with our proposed method.

4.2.2 Image Pre-processing

Segmentation

Before being used as training data in the cross population model the 4DCT images were segmented to mask out internal features within organs. Due to the significant variation of organ appearance across the population, ignoring such features was expected to minimize the number of required training scans. A simple classification scheme based on a threshold of Hounsfield Units (HU) was used to select lung (-1024HU to -800HU), abdominal tissues (-800HU to -100HU) and bone (above 200HU).
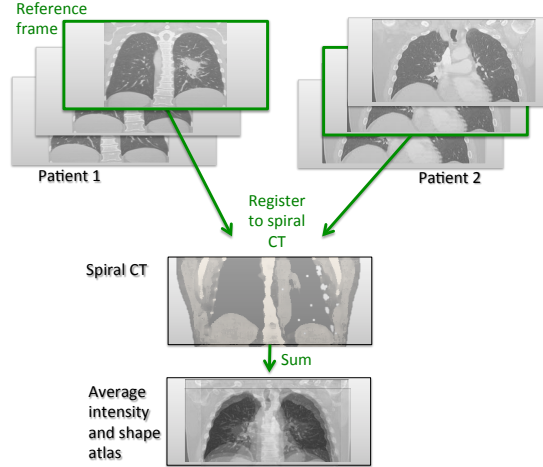
Figure 4.2: A reference frame was selected from each 4DCT scan. The reference frame image from each subject was registered to the spiral CT (common target) and summed to create an average intensity and shape atlas.

Global Atlas

To build a cross population model of respiratory motion, anatomical correspondence between different scans had to be established. In a similar way to Ehrhardt (Ehrhardt et al., 2011), an atlas was used to provide an anatomical reference for the comparison between different scans. This provided an average snapshot of the lung anatomy for a particular time range during the respiratory cycle.

The atlas image was constructed by selecting a single reference frame from each 4DCT scan and registering that frame to the common target. The common target image was chosen as the spiral CT from the PET/CT scan which is to be corrected for attenuation. The reference frame from each 4DCT was selected such that it reflected the same phase offset of the common target in the respiratory cycle. In order to capture non-rigid respiratory motion we chose to construct the atlas image by using free-form deformation, cubic B-Splines interpolation and a mutual information cost function. Performing an affine registration as an initial step accelerated the registration to the common target. All registered frames were added together to form an average image which was then referred to as the atlas image (Figure 4.2). Although it is usually desirable to be able to re-use the same atlas for many applications, we found it more practical for the common target to be the Spiral CT from the newly encountered PET/CT scan and because it was not severely affected by respiratory motion since it passes over the thoracic region in a few seconds.

In a similar way to Ehrhardt et al. (2011), the reference frame from each 4DCT scan was registered for a second time to the atlas image and the deformation field was stored. For each 4DCT scan the stored deformation was then applied to all other frames in each 4DCT. The deformed frames of each 4DCT then represented one of the many possible ways the lung could expand or contract relative to the atlas image. The deformed 4DCT scans could then be used in a cross population model to make an inference about how the spiral CT could deform during the newly encountered PET/CT scan.
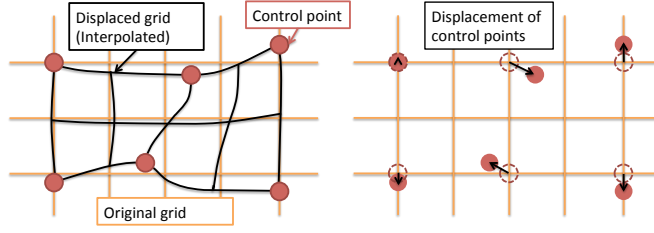
84

Figure 4.3: Parameterisation of displacement using free form image registration. A grid represents the original position of all voxels before deformation (left). Control points (solid red circles) were displaced and the new position of all voxels (black grid) was calculated using cubic B-Splines interpolation. This displacement of control points (right) was parameterized as displacement vectors (black arrows).

Estimating Frame by Frame Deformations

The deformation between consecutive frames in each 4DCT training scan was used as an estimate of respiratory motion. Deformation was represented as displacement vectors to describe the vector mapping of a voxel in a frame at end inspiration to a voxel in a frame at a later time. Estimation of every displacement vector at every voxel was found to be computationally expensive, thus we used a free form image registration algorithm which estimated the deformation on a coarse grid with a grid spacing of 5mm. The free form image registration was initialized by placing the freely adjustable control points at the grid vertices (Figure 4.3). The displacement at intervening voxels was approximated using cubic B-Splines interpolation. For the purposes of training the cross population model frame-by-frame respiratory motion was parameterized as $U(x, y, z, f)_p$, where U was the displacement of the control points from their initial position at the end inspiration frame, x, y, z was the coarse sampling position, f was the frame and p was the training scan number.

We also used this method of frame-by-frame respiratory motion parameterization to establish an estimate of the motion in the gated PET scan. We used free form image registration to estimate the displacement of every control point in each frame in the gated PET with respect to the end inspiration frame. The estimation was expected to only be reliable for a sparse set of positions with sufficient radio-tracer uptake to identify features in each gated PET frame (Mattes, Haynor, Vesselle, Lewellen, & Eubank, 2003), thus limiting the accuracy of the registration. We identified the sparse positions using a mask based on the signal to noise ratio (SNR), described in a later section, to retain only those estimates which were reliable. Respiratory motion in the gated PET scan was then parameterized as the displacement of control points, $Q(\grave{x}, \grave{y}, \grave{z}, f)$, where the sparse positions were $\grave{x}, \grave{y}, \grave{z}$ in each frame f.

4.2.3 Correspondence and Cross Population Model

A joint correspondence and cross population model was devised to improve the sparse estimates Q, to estimate a complete set of control point displacements $V(x, y, z, f)$. As described in equation 4.2, we expressed our model as a function $\mathcal{K}$, which produced the estimate, $V(x, y, z, f)$, when supplied with a set of cross population model coefficients, λ, and corre-
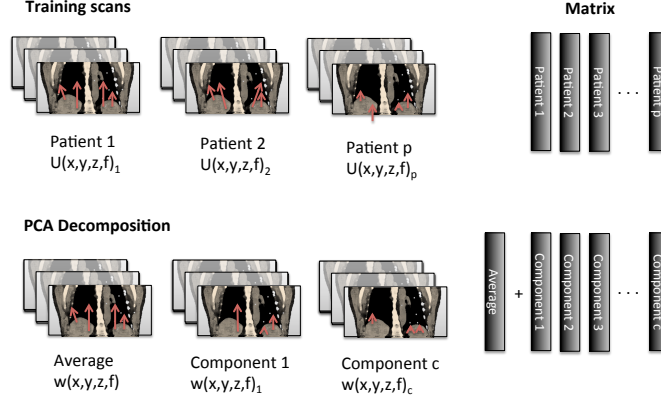
Figure 4.4: The parameterized motion fields for each training scan (Top Left) were arrange into a matrix (Top Right). Principal Components Analysis (PCA) produced a decomposed matrix whose eigenvectors were motion components (Bottom right). Motion was then parameterized as an average vector plus a linear combination of the motion components (Bottom left).

spondence model coefficients, μ.

$$V(x, y, z, f) = \mathcal{K}(\mu_0, \mu_1, \mu_2, ...\lambda_0, \lambda_1, \lambda_2...) \tag{4.2}$$

We found the coefficients for $\mathcal{K}$ such that $V(x, y, z, f)$ best matched the sparse and noisy estimates, $Q(\hat{x}, \hat{y}, \hat{z}, f)$, derived from the acquired PET data from a patient. The following sections describe how $\mathcal{K}$ was constructed using a linear combination of eigen-vectors.

Cross Population Model

The cross population model $\mathcal{K}$, was devised from N_p individual 4DCT scans from different patients to provide prior knowledge of respiration necessary for motion prediction. The displacement of control points $U(x, y, z, f)$ were used to parameterize the motion for each patient (p) in the cross population model. The parameterized motion field $U(x, y, z, f)_p$ was arranged into a 2 dimensional matrix where each column represented the displacement of all control points for a given patient at all positions x, y, z, f (Figure 4.4). A set of PCA component vectors was found by performing the eigen-decomposition of the covariance matrix $(U^T U)$ which represents the variance of control point displacements over each training scan (p). We have chosen to represent these vectors as continuous on x, y, z, f as the control point displacements can be discretized arbitrarily depending on the required resolution.

We decomposed the control point displacement field, U (Equation 4.3), for every training scan (p) into an average displacement field $(\bar{w})$ plus a linear combination (κ_c) of the components of motion $(w(x, y, z, f)_c)$.

$$\bar{w}(x, y, z, f) + \sum_c^{N_c} \kappa_c w(x, y, z, f)_c \approx U(x, y, z, f)_p \tag{4.3}$$

where

$$\bar{w}(x, y, z, f) = \frac{1}{N_p} \sum_{p}^{N_p} U(x, y, z, f)_p \tag{4.4}$$

The first N_c components of motion $(w(x, y, z, f)_c)$ and their coefficients (κ_c) were calculated from the Principal Components Analysis (PCA) algorithm. In practice it was found that $(w(x, y, z, f)_c)$ was unchanged regardless of whether motion was parameterized as the control point displacement (U) optimized by the registration software or the displacement vector mapping used for image deformation.

We created an individualized displacement field, W (Equation 4.5), which was adapted to the estimation of deformation between end-inspiration and end-expiration in the PET/CT scan.

$$W(x, y, z, f) = \bar{w}(x, y, z, f) + \sum_{c}^{N_c} \lambda_c w(x, y, z, f)_c \tag{4.5}$$

The coefficients (λ_c) for each patients individualized displacement field were derived by a function minimization procedure. A downhill simplex algorithm was used to vary λ_c until the individualized displacement field (W) was a close fit to the sparse control point displacements (Q) estimated from the gated PET frames (Equation 4.6).

$$\arg\min_{\lambda_c} \left\{ \sum_{\grave{x}, \grave{y}, \grave{z}} |W(\grave{x}, \grave{y}, \grave{z}, f_{max}) - Q(\grave{x}, \grave{y}, \grave{z}, f_{max})| \right\} \tag{4.6}$$

We computed the coefficients (λ_c) only at the extreme of motion (f_{max}) meaning that we only estimated the displacement (W) between end-inspiration and end-expiration during the PET scan. This constraint was necessary to ensure that the amplitude of displacement was larger than errors induced by image noise, but had the disadvantage of not optimizing control point displacement for intervening frames.

Correspondence Model

A correspondence model was used in combination with the estimated displacement from the cross population model to estimate control point displacement for all frames (including intervening frames) in the gated PET scan. W was expressed as a matrix with N_f columns where each column represented the control point displacement for a particular frame. In a similar way to the cross population model, principal components analysis (PCA) was performed on the matrix to isolate commonalities of control point displacements across different frames within W. The PCA algorithm produced eigenvectors (v_m) in order of their prevalence. Thus, the first few eigenvectors represented the most prevalent modes of motion that were common to all frames. We selected the N_m most prominent eigenvectors from the PCA algorithm to be the modes of motion (v_m). The outcome of using PCA to decompose W is described in equation 4.7 where W was approximated as the average displacement $(\bar{v})$ plus a linear combination $(\eta(f)_m)$ of the modes of motion (v_m).

$$\bar{v}(x, y, z) + \sum_{m}^{N_m} \eta(f)_m v(x, y, z)_m \approx W(x, y, z, f) \tag{4.7}$$

where

$$\bar{v}(x,y,z) = \frac{1}{N_f} \sum_f^{N_f} W(x,y,z,f) \tag{4.8}$$

The correspondence model described by Equation 4.7 differs from the cross population model described in Equation 4.3 because it is used to analyze control point displacement of all frames (W) in an individual patient. The cross population model was not used to predict displacement of all frames because of differences in respiratory phase and amplitude between the 4DCT training scans used to construct the cross population model and the gated PET scan. The correspondence model was found to be more suited to account for the variation in respiratory phase and amplitude between the 4DCT training scans and the gated PET scan. Equation 4.9 describes how we used the correspondence model to create a refined displacement field (V) which could have an arbitrary respiratory phase and amplitude.

$$V(x,y,z,f) = \bar{v}(x,y,z) + \sum_m^{N_m} \mu_m(f) v(x,y,z)_m \tag{4.9}$$

The coefficients $\mu(f)_m$ represented the contribution of displacement from each frame in the selected modes of motion, m, to each frame, f, in the refined motion field V. We estimated the control point displacement for all frames in the gated PET scan by varying $\mu(f)_m$ until V fitted our sparse estimate Q using a minimization procedure shown in Equation 4.10. This method was similar to finding optimal coefficients for the cross population model, however, there were fewer variables, thus we expected the function minimization procedure to be less affected by noise.

$$\arg\min_{\mu_m(f)} \left\{ \sum_{\dot{x},\dot{y},\dot{z},f} |V(\dot{x},\dot{y},\dot{z},f) - Q(\dot{x},\dot{y},\dot{z},f)| \right\} \tag{4.10}$$

All displacements had been calculated on a coarse grid to reduce computational demands. The refined displacement field (V) was interpolated to produce a voxel wise deformation which was at the resolution of the CT image.

Signal to Noise Ratio Mask

Estimation of motion at the axial edges of the PET image appeared inaccurate due to image noise. We avoided these areas of low camera sensitivity by only considering areas where the axial sensitivity of the PET scanner had adequate signal to noise (SNR) to estimate motion. Estimation of small displacements were also inaccurate due to PET image noise, leading to a situation where the model was fitted to image noise rather than useful signal (over-fitting). We overcame both effects by establishing a voxel-wise metric of the SNR as a product of magnitude of displacement and camera sensitivity. We used this metric by applying a threshold to determine the masked region where the SNR was adequate to use to the PET to measure organ deformation. The mask was then applied to the PET displacement field (Q) so that it was only defined for a restricted set of positions $\dot{x}, \dot{y}, \dot{z}$ in each frame f. We calculated the mean amplitude of displacement in the SNR masked region for the PET displacement field, $\sum Q(\dot{x},\dot{y},\dot{z})$, to the fitted displacement field, $\sum V(\dot{x},\dot{y},\dot{z})$,

and then adjusted the SNR mask cut-off until the ratio of $\sum Q(\dot{x}, \dot{y}, \dot{z})$ to $\sum V(\dot{x}, \dot{y}, \dot{z})$ was 95%.

4.2.4   Model Evaluation

Model Cross-Validation

In order to quantitatively assess the robustness of the joint correspondence and cross population model ($\mathcal{K}$), we determined how well the model could predict changes in lung volume. The lung volume was estimated directly from a 4DCT via the segmentation algorithm used to construct the global atlas. We measured the change in the segmented lung volume in milliliters and expressed it as a percentage of the volume at end-inspiration. We used a leave-one-out cross-validation method (McQuaid et al., 2008) for which one 4DCT data set was removed from the cross population model ($\mathcal{K}$). A single static frame was chosen from the removed 4DCT data set. We then used the model ($\mathcal{K}$) to estimate the deformation and applied it to the single static frame to predict the deformation of the remaining frames. We determined the change in lung volume due to the deformation by segmenting the remaining frames and comparing it to lung volume at end-inspiration. We compared the changes in lung volume predicted by the cross population model to the changes observed from the segmented 4DCT.

Attenuation Correction

The CT image was deformed to account for mismatch between the spiral CT and each PET frame. We used the cross population model to find the deformation of each PET frame and then applied the deformation to the CT image to produce phase correlated PET and CT frames. Each correlated CT frame was scaled from Hounsfield Units (HU) to a voxel-wise attenuation map. Each phase correlated attenuation map was then used to correct for photon attenuation during the reconstruction of each gated PET frame.

We evaluated PET image contrast as a measure of whether attenuation mismatch had been addressed. Inadequate correction for counts lost due to attenuation can be responsible for a reduction in image contrast. We specifically investigated attenuation mismatch at the diaphragm boundary because it is known to lead to a reduction in image contrast of the myocardial wall (McQuaid et al., 2008).

In each of the 11 gated PET studies we either segmented the heart (n=4) or a pulmonary nodule (n=7) and measured the counts per voxel in the segmented regions for each gated PET frame. We calculated the maximum range of count levels over all gated frames in the segmented regions of each gated PET. We compared the maximum range after applying no correction (NAC), applying correlated attenuation correction (Corr. AC) and using a static spiral CT for attenuation correction (CT AC). A large range has been shown to be associated with attenuation mismatch (Hamill et al., 2008).
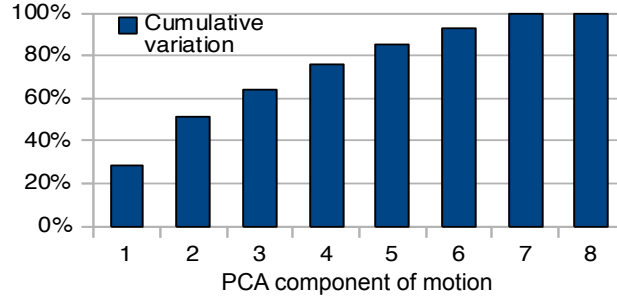
Figure 4.5: The bars show the cumulative variance for each component of motion expressed as a percentage of the total variance in the population.
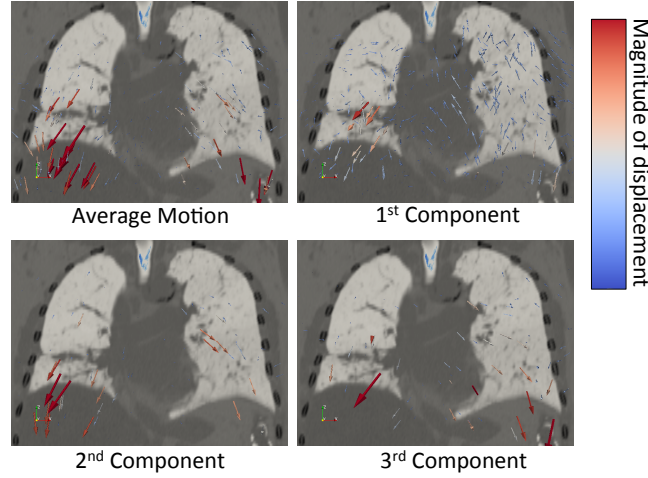

Figure 4.6: The components of motion $(w(x, y, z, f)_c)$ for frame 3. The displacement field is represented using a sparse vector field with arrow color indicating the amplitude of motion (red - large, blue small).

4.3 Results

4.3.1 Training the Cross Population Model

The results of performing principal components analysis on 10 patients who underwent a 4DCT scan (Fig. 4.5) showed that the first eigenvector accounted for 29% of the total variation of motion and that we must use at least the first 5 eigenvectors to model 85% of the variation in patient scans.

The average motion, $\bar{w}$, for all 10 patients was a continuous rise from the base of the diaphragm in the superior direction. Figure 4.6 shows the magnitude and direction for the first few principal components, $w(x, y, z, f)_c$, which can be interpreted as a departure from the population average motion, $\bar{w}$.
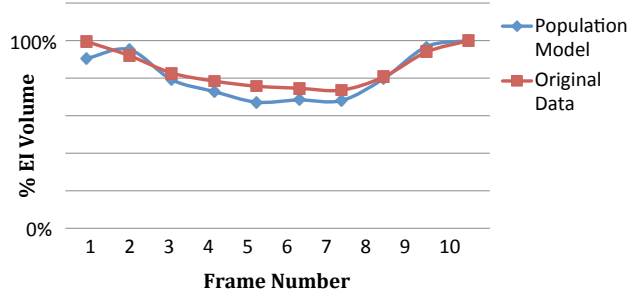
Figure 4.7: The segmented lung volume for a selected data set compared to the lung volume estimated using the cross population model, expressed as a percentage of the lung volume at end inspiration using the original CT data. The model volume curve was reasonably correlated with the original removed data set (R=0.95).

4.3.2   Population Model Cross-Validation

Our results of cross-validation using the leave-one-out method demonstrated consistency between the cross population model and new motion data not previously encountered. This consistency can be seen in Figure 4.7 where there was good correspondence of lung volume between each frame in the model and each frame in the data. We quantified the correspondence in lung volume using the Pearson correlation coefficient and identified an R value for each patient data set ranging from R=0.68 to R=0.99 with a median of R=0.97. 9 of 10 datasets had acceptable consistency (R > 0.85) implying that the cross population model may account for variations in patient anatomy and breathing patterns for approximately 90% of patients.

4.3.3   Estimates of Control Point Displacement

Without cross population modeling the PET displacement field (Q) was found to be exceptionally noisy. The noise produced a dampening effect which reduced the mean displacement of control points and thus an underestimate of motion. This was in contrast to the estimate produced using the joint correspondence and cross population model (V) which had a much larger amplitude of displacement. Prior to motion modeling the PET deformation field (Q) typically consisted of large displacements in the center of the field of view as opposed to small and seemingly random displacements at the edge of the field of view (Figure 4.8 top). This was addressed using the SNR masked region and ensured that the cross population model was much less affected by noise (Figure 4.8 bottom).

4.3.4   Attenuation Correction

The 11 PET/CT scans for which we were comparing methods of attenuation correction were assigned the labels A-K.

We initially investigated the use of a cross-population model to address attenuation mismatch by measuring myocardial uptake from the images in Figure 4.9. An axial profile was drawn on the image in Figure 4.9 to determine the ratio of counts at the myocardial wall

PET displacement field

Fitted displacement field

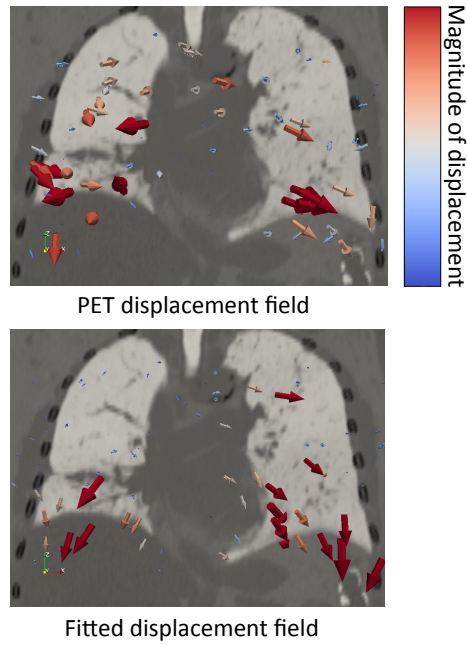Figure 4.8: The displacement field estimated from a single gated PET (Study D) and the result after fitting it to a cross population model. The displacement field is represented using a sparse vector field with arrow color indicating the amplitude of motion (red - large, blue - small). The fitted displacement appears to be less affected by noise and the directions of the vectors make more physiological sense.

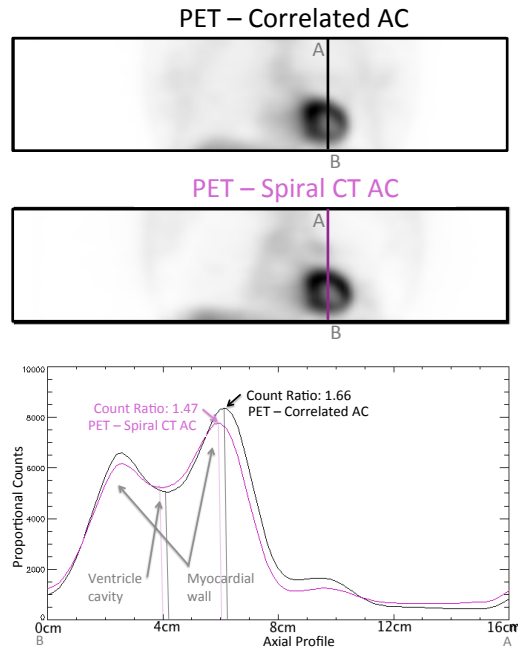Figure 4.9: A coronal slice of a single respiratory PET scan (Study B) with myocardial uptake to show a comparison of correlated AC (top) with using spiral CT AC (bottom). All frames in the respiratory gated FDG PET were summed after Attenuation Correction (AC). An axial profile was taken through the myocardium (point B to point A). The profile was used to determine the ratio of counts at myocardial wall to the ventricle cavity.
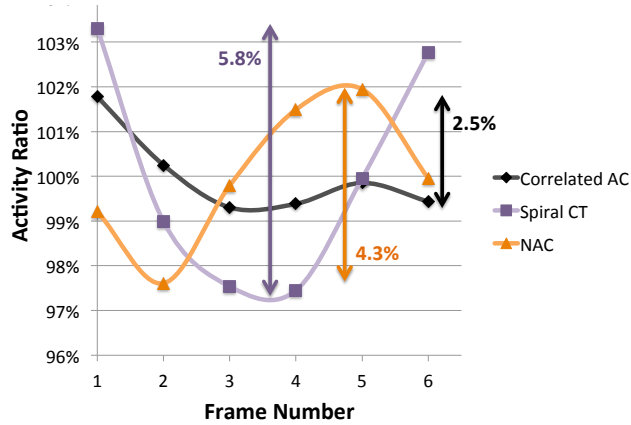
Figure 4.10: The ratio of counts per voxel in the segmented heart region for each frame of Study D. The counts are scaled as a ratio of the mean counts per voxel over the entire respiratory cycle (100%). The vertical lines show the maximum range of each ratio when using Correlated AC (2.5%), Spiral CT based AC (5.8%) and No AC (4.3%).

Table 4.1: Heart: Maximum range (%) of variation over respiratory cycle.

Study	B	D	F	J
NAC	3.3%	4.3%	4.2%	3.4%
CT AC	7.4%	5.8%	6.7%	5.5%
Corr. AC	3.9%	2.5%	2.8%	3.0%

to the ventricle cavity when using both phase correlated attenuation correction and spiral CT alone for attenuation correction. The results demonstrated an increase in contrast (1.47 → 1.66) through the myocardium when using the phase correlated attenuation map. These images are a summation of all gated frames because individual gated PET frames are highly variable and very noisy.

We further assessed the accuracy of attenuation correction near the heart in a subset of the PET/CT scans with significant apparent myocardial uptake (studies B, D, F and J). We calculated the counts per voxel in the segmented heart region and determined the ratio of counts per voxel in each frame versus the mean counts per voxel for the entire respiratory cycle. Figure 4.10 compares count ratios for Study D when using Correlated AC, Spiral CT AC and No Attenuation Correction (NAC). The maximum range of the count ratio over the entire respiratory cycle for Study B, D, F and J were between 2.5% and 7.4%. Table 4.1 shows that, relative to applying no correction (NAC), the maximum range decreased when applying correlated attenuation correction (Corr. AC) but increased when using a static spiral CT for attenuation correction (CT AC).

The images in Figure 4.11 demonstrate an increase in contrast near the lung-diaphragm boundary in Study A. A line profile has been drawn crossing that boundary. Attenuation correction utilising a phase correlated attenuation map resulted in higher contrast between the lung and diaphragm.

We observed significant uptake of single pulmonary nodules in 7 PET studies included in
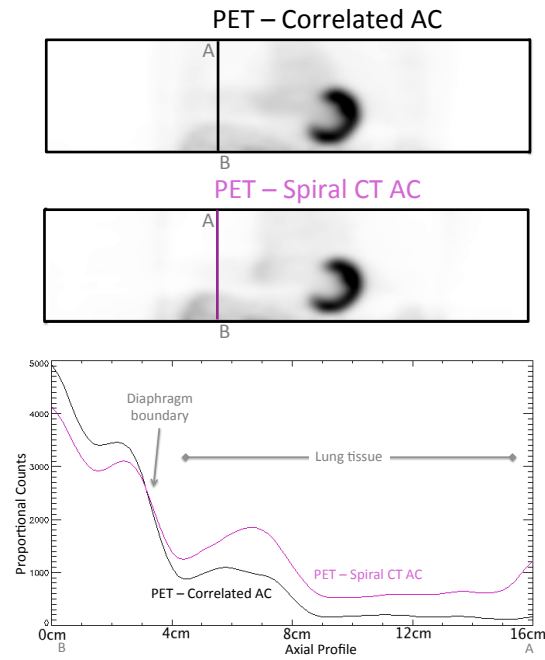
Figure 4.11: A coronal slice of the respiratory PET scan (Study A) with diaphragm and lung tissue uptake to show a comparison of correlated AC (top) with using spiral CT AC (bottom). All frames in the respiratory gated FDG PET were summed after Attenuation Correction (AC). An axial profile was taken through the lung-diaphragm boundary to present the changes in contrast (point B to point A).

Table 4.2: Single Pulmonary Nodule: Maximum range (%) of variation over respiratory cycle.

Study	A	C	E	G	H	I	K
NAC	8.6%	2.6%	7.5%	13.8%	24.3%	6.5%	3.2%
CT AC	13.2%	2.2%	7.3%	17.5%	15.5%	7.1%	2.9%
Corr. AC	4.8%	4.1%	4.7%	6.6%	11.3%	6.4%	2.3%

our investigation. The segmented counts per voxel in the nodules varied over the respiratory cycle partially because they altered position relative to surrounding tissue and thus were subject to variations in photon attenuation. We calculated the counts per voxel in the segmented nodule region as a ratio of the mean counts per voxel over the entire cycle and reported the maximum range. The maximum range over the respiratory cycle was between 2.2% and 24.3%. Table 4.2 lists 6 studies (A,E,G,H,I,K) for which the maximum range was lower with correlated attenuation correctionthan for other methods of correction. In Study C, the diaphragm was outside the PET field of view and there was an increase in the maximum range when using correlated AC in comparison to simply using the spiral CT for AC.

4.4   Discussion

These results indicate that attenuation mismatch during a gated PET scan may be addressed using a cross population respiratory motion model trained with 4DCT. In gated PET, static spiral CT AC fails to account for variations in photon attenuation over the respiratory cycle. The resulting variations of count levels between different gated frames were mitigated using our method of correlated attenuation correction. McQuaid et al. (2008) suggest that the accuracy of a population model is affected by the uptake of moving organs, such as the diaphragm, being visible in both the end-inspiration and end-expiration PET frames. Further investigation is also required to determine if our method is still accurate in less than optimal imaging conditions such as when there is inaccurate respiratory gating.

Some authors (Bai & Brady, 2011) argue that incorporating regularization into the non-rigid image registration algorithm is the best way of addressing image noise, however, some questions have been raised as to whether it is realistic to penalize sudden changes in displacement (Ruan, 2009). Our results using the proposed respiratory cross population model demonstrate that the noise may be significantly reduced such that the model can be used to deform the spiral CT for the purposes of attenuation correction. We are yet to evaluate promising registration methods which establish a more physically realistic model of motion such as regularized mass-preserving image registration (Gigengack et al., 2010).

Although our leave-one-out analysis results are promising, further work is required to determine the optimal number of training data sets. There is a large variation in respiratory motion across patients, especially those with lung pathology (Liu, Pierce II, Alessio, & Kinahan, 2009). Although our training data included oncology patients with lung pathology, our leave-one-out cross validation suggested that the population model may successfully predict motion in about 90% of cases. We are optimistic since population models have been used for model-based prediction of tumour motion with a very limited number of patients (Ehrhardt et al., 2011).

The proposed method of constructing a cross population model of respiratory motion is a way of estimating respiratory movement from gated PET without the need for a 4DCT from the same patient. The model was successfully adapted to produce a phase correlated attenuation map. The cross population approach to attenuation alignment demonstrated less variation in count density due to respiratory motion. The correlated attenuation correction also produced an improvement in image contrast in some circumstances.

References

Bai, W., & Brady, M. (2011, February). Motion correction and attenuation correction for respiratory gated PET images. *Medical Imaging, IEEE Transactions on*, *30*(2), 351–365. doi: 10.1109/TMI.2010.2078514

Chan, C., Jin, X., Fung, E. K., Naganawa, M., Mulnix, T., Carson, R. E., & Liu, C. (2013, November). Event-by-event respiratory motion correction for PET with 3d internal-1d external motion correlation. *Medical physics*, *40*(11), 112507. doi: 10.1118/1.4826165

Chandrashekara, R., Rao, A., Sanchez-Ortiz, G. I., Mohiaddin, R. H., & Rueckert, D. (2003, July). Construction of a statistical model for cardiac motion analysis using nonrigid image registration. *Information processing in medical imaging: proceedings of the ... conference*, *18*, 599–610.

Daouk, J., Fin, L., Bailly, P., & Meyer, M.-E. (2008, October). Improved attenuation correction via appropriate selection of respiratory-correlated PET data. *Computer Methods and Programs in Biomedicine*, *92*(1), 90–98. doi: 10.1016/j.cmpb.2008.06.014

Ehrhardt, J., Werner, R., Schmidt-Richberg, A., & Handels, H. (2011, February). Statistical modeling of 4d respiratory lung motion using diffeomorphic image registration. *IEEE Transactions on Medical Imaging*, *30*(2), 251–265. doi: 10.1109/TMI.2010.2076299

Ehrhardt, J., Werner, R., Sring, D., Frenzel, T., Lu, W., Low, D., & Handels, H. (2007, January). An optical flow based method for improved reconstruction of 4d CT data sets acquired during free breathing. *Medical Physics*, *34*(2), 711–721. doi: 10.1118/1.2431245

Gigengack, F., Ruthotto, L., Burger, M., Wolters, C., Jiang, X., & Schafers, K. (2010, November). Motion correction of cardiac PET using mass-preserving registration. In *Ieee nuclear science symposium conference record (nss/mic)* (pp. 3317 –3319). doi: 10.1109/NSSMIC.2010.5874418

Hamill, J. J., Bosmans, G., & Dekker, A. (2008, February). Respiratory-gated CT as a tool for the simulation of breathing artifacts in PET and PET/CT. *Medical Physics*, *35*(2), 576–585.

Heimann, T., & Meinzer, H.-P. (2009, August). Statistical shape models for 3d medical image segmentation: A review. *Medical Image Analysis*, *13*(4), 543–563. doi: 10.1016/j.media.2009.05.004

Lamare, F., Carbayo, M. J. L., Cresson, T., Kontaxakis, G., Santos, A., Rest, C. C. L., ... Visvikis, D. (2007, September). List-mode-based reconstruction for respiratory motion correction in PET using non-rigid body transformations. *Physics in Medicine and Biology*, *52*(17), 5187. doi: 10.1088/0031-9155/52/17/006

Li, R., Lewis, J. H., Jia, X., Zhao, T., Liu, W., Wuenschel, S., ... Jiang, S. B. (2011, September). On a PCA-based lung motion model. *Physics in Medicine and Biology*,

56(18), 6009–6030. doi: 10.1088/0031-9155/56/18/015

Liu, C., Pierce II, L. A., Alessio, A. M., & Kinahan, P. E. (2009). The impact of respiratory motion on tumor quantification and delineation in static PET/CT imaging. *Physics in medicine and biology*, *54*(24), 7345.

Mattes, D., Haynor, D. R., Vesselle, H., Lewellen, T. K., & Eubank, W. (2003, January). PET-CT image registration in the chest using free-form deformations. *IEEE Transactions on Medical Imaging*, *22*(1), 120–128. doi: 10.1109/TMI.2003.809072

McClelland, J., Hawkes, D., Schaeffter, T., & King, A. (2013, January). Respiratory motion models: A review. *Medical Image Analysis*, *17*(1), 19–42. doi: 10.1016/j.media.2012.09.005

McQuaid, S. J., Lambrou, T., & Hutton, B. F. (2008). Statistical shape modeling of the diaphragm for application to rb-82 cardiac PET-CT studies. In *Nuclear science symposium conference record* (pp. 3651–3655).

Nehmeh, S. A., Erdi, Y. E., Meirelles, G. S., Squire, O., Larson, S. M., Humm, J. L., & Schoder, H. (2007, January). Deep-inspiration breath-hold PET/CT of the thorax. *J Nucl Med*, *48*(1), 22–26.

Nehmeh, S. A., Erdi, Y. E., Pan, T., Yorke, E., Mageras, G. S., Rosenzweig, K. E., . . . Humm, J. L. (2004, June). Quantitation of respiratory motion during 4d-PET/CT acquisition. *Medical Physics*, *31*(6), 1333–1338. doi: 10.1118/1.1739671

Pan, T., Lee, T.-Y., Rietzel, E., & Chen, G. T. Y. (2004, January). 4d-CT imaging of a volume influenced by respiratory motion on multi-slice CT. *Medical Physics*, *31*(2), 333–340. doi: 10.1118/1.1639993

Pönisch, F., Richter, C., Just, U., & Enghardt, W. (2008, July). Attenuation correction of four dimensional (4d) PET using phase-correlated 4d-computed tomography. *Physics in Medicine and Biology*, *53*(13), N259. doi: 10.1088/0031-9155/53/13/N03

Ruan, D. (2009, October). Directionally selective regularization for sliding preserving medical image registration. *Nuclear Science Symposium Conference Record (NSS/MIC), 2009 IEEE*, 2936–2939. doi: 10.1109/NSSMIC.2009.5401607

Thielemans, K., Schleyer, P., Marsden, P. K., Manjeshwar, R. M., Wollenweber, S. D., & Ganin, A. (2013). Comparison of different methods for data-driven respiratory gating of PET data. In *Nuclear science symposium and medical imaging conference (NSS/MIC), 2013 IEEE* (pp. 1–4). IEEE.

Zhang, Q., Hu, Y.-C., Liu, F., Goodman, K., Rosenzweig, K. E., & Mageras, G. S. (2010, May). Correction of motion artifacts in cone-beam CT using a patient-specific respiratory motion model. *Medical Physics*, *37*(6), 2901–2909. doi: 10.1118/1.3397460

Zhang, Q., Pevsner, A., Hertanto, A., Hu, Y. C., Rosenzweig, K. E., Ling, C. C., & Mageras, G. S. (2007). A patient-specific respiratory model of anatomical motion for radiation treatment planning. *Medical Physics*, *34*, 4772.

Chapter 5

1D Respiratory Motion Models for Improving Attenuation Correction

5.1 Introduction

Alignment of the attenuation correction map in PET/CT could be addressed if there was a way to ascertain the internal organ deformation during both the PET and CT scans. For this reason it is desirable to estimate motion during a PET scan in an attempt to compensate for the contribution of respiratory motion to attenuation misalignment. Although a single cycle of respiratory motion can be captured using CT, this is not practical in PET imaging. In PET each frame may not contain enough counts for accurate estimation of respiratory motion thus, assumptions about respiratory motion are required to construct a model of respiration. Respiratory motion during a medical imaging procedure is complex and leads to challenges in developing a model of respiratory motion (McClelland, Hawkes, Schaeffter, & King, 2013). These challenges were discussed in chapter 4 and were mainly due to the variations between different patients.

Respiratory motion can be estimated from respiratory gated images whereby temporal samples are partitioned into distinct phases of the respiratory cycle. Respiratory gating is common in CT, MR and PET and can be used to produce a sequence of images, each of which are a snapshot of the phase position at various parts of the respiratory cycle. It is a useful technique which means that the PET image can be acquired over multiple respiratory cycles thereby avoiding the trade off between image quality and temporal resolution. Unfortunately gating is significantly affected by breathing irregularities (George et al., 2006) and this can further complicate attempts to estimate internal organ motion from gated PET data.

The technical limitations of capturing respiratory-gated images contribute significantly to the problem of modeling the motion of lung tissue. There has been significant investigation into respiratory models which can be created from medical images and the practicalities of using such models to accomodate for unexpected variation in the breathing (Liu, Alessio, & Kinahan, 2011).

Although the motion of lung tissue is difficult to model, the physics of respiration can be considered in a very simple and rudimentary way. In a model proposed by Kyriakou, McKenzie, Suchowerska, and Fulton (2007) the lung was represented by a simple mechanical analogue in which a piston drives the lung tissue, expelling air, considered as incompressible, from the top. Given this assumption it was found that respiratory motion approximately satisfied a one-dimensional equation in the superior-inferior direction.

$$\mu \Delta V \frac{\partial^2 u}{\partial t^2} = \Delta V \left(E \frac{\partial^2 u}{\partial z^2} - \frac{\partial^2 P}{\partial^2 z} \right) \tag{5.1}$$

u is the superior displacement due to respiratory motion, z is the height above the diaphragm, t is the time point of the frame in the respiratory phase, c is a constant determined by the speed of propagation of a pressure wave in air, μ is the lung mass density, ΔV is the air volume element size, and P is the pressure at position z. The meaning of these terms are further explained in the methods section in figure 5.4. This simple model has an advantage over the population model presented in chapter 4 in that each patient can be independently modelled by solving for the coefficients in equation 5.1. However, this equation doesn't account for the time lag between the displacement, u, and the external surrogate measurements used for phase or amplitude gating.

The motion of lung tissue can be analysed by finding the deformation of the image from

one time point to the next. By convention, the deformation map is often implemented as a spatial mapping of voxel locations in a reference frame to the voxel locations in a subsequent frame. The mapping from one frame to the next may not be a 1:1 map, because the size of tissue features will change as a result of the deformation. The Jacobean of the transform is the amount by which tissue expands or compresses due to deformation and this attribute has been used to consider whether the deformation of tissue is physiologically plausible.

There are a variety of approaches to image registration that have been applied to respiratory motion (Murphy, 2011). The basic principles of image registration have been covered in chapter 2, however, a complete review of all registration methods that have been proposed as applicable to respiratory motion is beyond the scope of this thesis. There are several methods which include potentially useful assumptions which constrain image deformation to ensure the robustness of image registration. In general, all of these methods seek to both reduce the number of parameters being estimated whilst also ensuring that the constraints are physically realistic.

The simplest approach to constraining image registration is to penalise non-smooth deformations. The reason why this may be advantageous is because the registration algorithm is very sensitive to noise. It has been suggested (Mattes, Haynor, Vesselle, Lewellen, & Eubank, 2003) noise significantly corrupts the measurement of the similarity between image features and thus the deformation field becomes corrupted in the process of trying to match features. This effect can be mitigated through a simple penalty term which favours smooth deformation fields. Despite its wide use, penalising non-smooth deformations has been shown to be ineffective (Ruan, 2009). In chapter 3 it was found that regularising the registration of gated PET images provided limited improvements to the accuracy of tracking the motion of attenuating lung tissue.

It appears to be plausible that following the deformation of tissue, there should also exist an inverse deformation which would restore the tissue to its former state. This property of a deformation is known as diffeomorphism, however, this is only one of many possible assumptions that can be made about a deformation. Some other assumptions, also known as priors include spatial smoothness of deformation, continuity of deformation (Holden, 2008), direction dependent smoothness (Ruan, 2009) and mass preservation (Thielemans, Asma, & Manjeshwar, 2009) after deformation. Not only are there many possible priors, there is also variability in the way priors are implemented. For example, spatial continuity of deformation can be modeled using a constant velocity framework (Ashburner, 2007), however, spatial continuity of deformation can also be encouraged by penalising transforms such that the Jacobean is always positive (Holden, 2008).

Temporal constraints are possible when registering multiple frames from a gated or dynamic study. Example constraints include using a time based spline parameterisation (Metz, Klein, Schaap, van Walsum, & Niessen, 2011) and enforcing periodic motion. Time based parameterisation makes the assumption that motion is smooth between frames. The low temporal resolution of scanners means that the sampling rate is not sufficient to fully capture respiratory motion and time based spline parameterisation therefore can only provide an approximate model for continuous respiratory motion. Temporal constraints are especially useful in situations where it is necessary to make approximations about the motion in between frames, such as when applying transforms to list mode events (Liu et al., 2011).

Investigation into the methods of constraining image registration has been predominantly

focused on the accuracy of deformation for high resolution and high contrast images such as those attained from CT (Murphy, 2011). These approaches are general and not specific to the characteristics of lung tissue or respiratory motion. It was suspected these methods could be improved by finding unique characteristics of lung tissue which are not present in other tissues.

Modeling the physics of respiration has been proposed as an adjunct to relying on image features used for image registration (Sundaram & Gee, 2005). Finite Element Models (FEM) have been demonstrated as a way of estimating respiratory motion from CT images (Werner, Ehrhardt, Schmidt, & Handels, 2008). A shortcoming of the FEM method is that the accuracy of the model is largely dependent on the material properties used to construct the model in the first place. Thus, it is difficult to reproduce similar results even if two very similar FEM models are used.

Previous investigations have cited considerable difficulties in attaining accurate image registration in areas of low image contrast (Mattes et al., 2003). A recent study proposed that this problem can be mitigated by using FEM in areas where there is low image contrast. More recently it has been proposed to combine FEM modeling with image registration, so as not to rely wholly on the accuracy of the FEM model (Samavati, Velec, & Brock, 2015). Although FEM models are promising, they require the adjustment of many different tuning parameters. It is difficult to explain why FEM based registration may fail in some circumstances and thus difficult to determine what model properties are necessary to get an accurate registration.

The review of existing state-of-the-art registration packages in chapter 2 demonstrated that there are several methods which are applicable to PET imaging. In chapter 3 image registration was marginally improved by smoothing the deformations and PET images. In chapter 4 a population approach was used to estimate a large number of internal variables which could then be used to model deformation. In contrast to the previous two chapters the work described in this chapter focuses on developing a method of registration which incorporates a simple 1 dimensional physiological model of respiratory motion. The purpose of the model is to improve motion estimation from PET frame data such that it can be used to address attenuation mismatch.

5.2   Method

This section introduces a novel method of image registration which was initially developed by modelling respiratory motion in 4DCT. This new method was then used to estimate motion in respiratory gated PET images and subsequently used to compensate for attenuation correction misalignment.

5.2.1   Data Acquisition and Processing

This section describes the data acquired for the evaluation of models (section 5.2.4) and the optimisation of registration for both CT and PET registration (sections 5.2.6 and 5.2.7)

4DCT scans of the thorax were selected from a cohort of oncology patients undergoing a respiratory gated CT for the purposes of radiation treatment planning. The selection of

8 4DCT scans was based on images for which the lung was acquired from apex to base for a complete respiratory cycle. The 4DCT scans had been acquired in cine mode using a GE LightSpeed RT 16 scanner and were subsequently used in this investigation and optimisation of several candidate respiratory motion models. The scans were originally intended for the purposes of radiation treatment planning and were provided for the purposes of retrospective analysis. The phase in the respiratory cycle was measured relative to the time when the chest surface had maximum displacement (End inspiration). Cine mode was implemented by acquiring repeated axial CT images for a specified period of time, at each table position. Each axial CT slice was retrospectively binned into 10 frames based on the phase in the respiratory cycle. There are several publications comparing the performance of image registration for respiratory motion estimation from 4DCT images. In order to analyse the candidate respiratory motion models the displacements estimated from state of the art registration of 4DCT were considered as a gold standard. The motion estimates were obtained using free form image registration which is described in the next section.

Methods of image registration were compared by retrospectively analysing 27 PET/CT scans. Phase based gating was performed after a conventional whole body static PET and spiral CT. The phase based respiratory gating system AZ733 V (Anzai Medical, Tokyo, Japan) was employed as previously described by Ponisch, Richter, Just, and Enghardt (2008). The system utilises a pressure sensor inserted into the pocket of an elastic belt to detect external respiratory motion. The first 16 PET scans were acquired and reconstructed using respiratory gating on a Biograph 16 PET/CT the remaining 11 PET scans were acquired and reconstructed using respiratory gating on a Biograph mCT PET/CT. Each patient was injected with approximately 350 MBq of F18-FDG. The Biograph 16 PET/CT had an approximate sensitivity of 4 cps/kBq and the duration of the scan was 5 minutes. Improved count statistics were achieved with the Biograph mCT PET/CT which had approximate sensitivity of 8 cps/kBq and the duration of the scan was 3 minutes. CT derived attenuation correction was used as a baseline and compared to proposed methods of compensating for attenuation mismatch.

5.2.2   Estimation of Motion Using Free Form Image Registration

The motion between successive gated frames was approximated as a transform at each point (x, y, z). The approximate deformation transform $\vec{T}$ was found using image registration software (Modat et al., 2010). The deformation transform $\vec{T}$ was the new position, (T_x, T_y, T_z), of a voxel an frame number f after being displaced by (D_x, D_y, D_z), where the original position in the reference gated frame was (x, y, z). The voxel size in each direction was s_x, s_y and s_z .

$$\vec{T}(x, y, z, f) = \{\vec{T_x}, \vec{T_y}, \vec{T_z}\} = \{D_x + s_x x, D_y + s_y y, D_z + s_z z\} \tag{5.2}$$

Free form registration allowed for arbitrary deformation $\vec{T}$, through the adjustment of the position of control points (Modat et al., 2010). The deformation of individual voxels is interpolated using a spline basis approximation (figure 5.1).

The estimation of deformation was significantly influenced by noise, thus a first step was to apply image smoothing before using the frames for registration. Frames were pre-smoothed using a 3D gaussian kernel, however, to avoid missing image features, the kernel FHWM
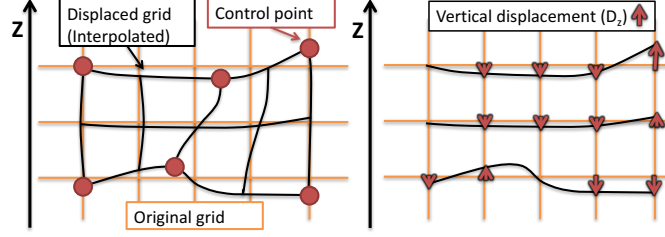
Figure 5.1: Left: A grid of voxels (orange) is displaced using free from registration. The new transformed grid positions (black) are determined through spline interpolation between control points (red). Right: The vertical displacement of each voxel is D_z and corresponds to the displacement of the grid from the original location.

was chosen to be smaller than the approximate system resolution of the PET (4mm) or CT (1mm). Although pre-smoothing could have been avoided using sparsely spaced control points, our experience was that noise at frequencies greater than the system resolution propagated into the objective function depending on how the optimisation had been implemented. Pre-smoothing ensured that the estimation was less dependent on the implementation of registration algorithm.

The predominant motion during respiration was displacement, D_z, in the superior and inferior direction. Image registration was used to make an initial estimate of 3D transforms allowing for motion in all directions, but the investigation of respiratory motion was confined to a single direction. The directional model was used in a fully 3D registration framework which allowed for fully 3D transforms.

A velocity framework was used to address the issue that displacement from end inspiration to end expiration was much larger than the small displacements between adjacent frames (Ashburner, 2007). Large displacements were approximated with a series of smaller displacements between successive frames and then composed them together to estimate motion with respect to the reference frame (Figure 5.2). In the velocity framework, the velocity ($\vec{v_D}$) was defined as the displacement between sequential frames. The velocity was assumed to be constant for the duration of the frame (Δt). The equation for obtaining the velocity from a displacement was

$$\vec{v}_{D,z} \approx \frac{\partial \vec{D}_z}{\partial t} \Delta t \tag{5.3}$$

where Δt was the frame duration and $\vec{v}_{D,z}$ was the vertical displacement between sequential frames. Image registration was used to estimate the velocity at each frame. The displacement for a given frame was then estimated by composing together each transform. When successive transforms were composed together, the spatial location of each velocity vector ($\vec{v}_{D,z}$) was itself transformed (Figure 5.2). For example, the transform from frame number 3 to the reference frame number 0 was approximated as

$$\vec{T}_{0\leftarrow 3} = \vec{T}_{0\leftarrow 1} \circ \vec{T}_{1\leftarrow 2} \circ \vec{T}_{2\leftarrow 3} \tag{5.4}$$

The velocity vector was always expressed as a displacement from the reference frame (Figure 5.3). The current frame was aligned to the reference frame and then the transform
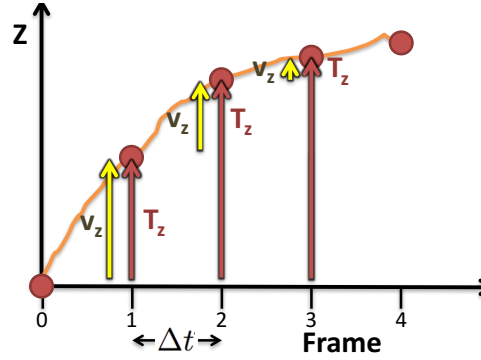
Figure 5.2: The vertical displacement for successive frames where frame 0 is the reference frame. T_z is the vertical displacement from the reference frame. v_z is the vertical velocity
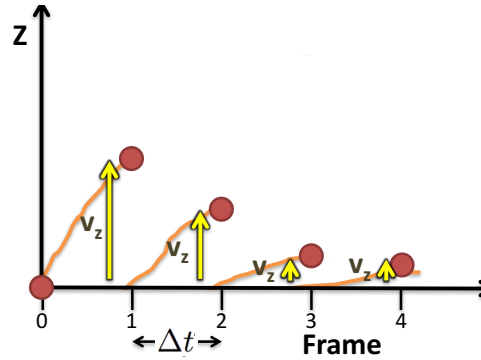

Figure 5.3: The vertical displacement for successive frames where every frame has been deformed to the reference frame before calculating the displacement to the subsequent frame. v_z is the vertical velocity

to the next frame was determined by image registration. The alignment to the reference frame was found by combining successive velocity transforms in reverse order.

$$\vec{T}_{0 \leftarrow 3} = \vec{T}_{2 \leftarrow 3} \circ \vec{T}_{1 \leftarrow 2} \circ \vec{T}_{0 \leftarrow 1} \tag{5.5}$$

5.2.3 Candidate Models

A range of methods were developed to incorporate a simple 1 dimensional physiological model of respiratory motion into image registration. The initial candidate model considered was the piston model of breathing in an elastic permeable medium as proposed by Kyriakou et al. (2007). In this model the vertical position (D_z) of a permeable barrier in the lung was determined by the forces applied to the barrier (Figure 5.4). In this model, only two acting forces were considered:

1. the force of the air flow through the permeable barrier and

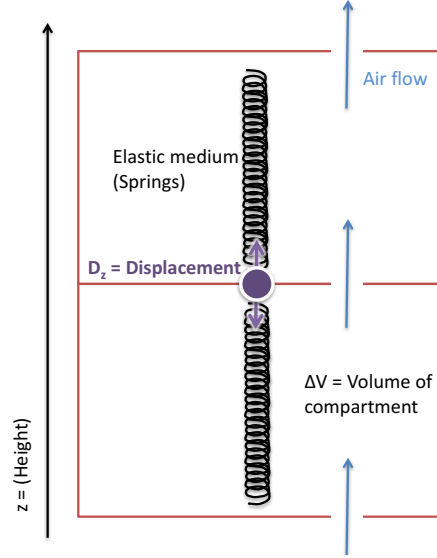2. the elastic force of the medium arising from volume change.

105

Figure 5.4: A representation of respiratory motion using a piston model. A piston is made up of a series of vertically stacked compartments.

The vertical position of the permeable barrier, D_z. was substituted into the derivative of equation 5.1:

$$\mu \Delta V \frac{\partial^3 D_z}{\partial t^2 \partial z} = \Delta V \left(E \frac{\partial^3 D_z}{\partial z^3} - \frac{\partial^2 P}{\partial^2 z} \right) \tag{5.6}$$

where E was the energy (related to tissue elasticity), P was the pressure, μ was the density and ΔV was the volume of the compartment. Equation 5.6 was simplified by assuming that there was no change in the derivative of the pressure and by integrating both sides along the vertical direction (z).

$$\mu \frac{\partial^2 \vec{D}_z}{\partial t^2} = E \frac{\partial^2 \vec{D}_z}{\partial^2 z} \tag{5.7}$$

A first order solution to the wave equation (equation 5.7) was found through the separation of variables into a linear combination of spatial ($A(x, y, z)$) and temporal factors ($B(t)$).

$$\vec{D}_z = \sum_i A_i(x, y, z).B_i(t) \approx A_0(x, y, z).B_0(t) \tag{5.8}$$

A higher order solution may have permitted more complicated modes of motion, however, this was avoided to keep the model simple. One solution to the wave equation while holding $A_0(x, y, z)$ constant was a harmonic oscillator where the first order equation had coefficients ($\frac{\lambda}{\mu}$) which were position dependant.

$$\frac{\partial^2 \vec{D}_z}{\partial t^2} = \left(\frac{\lambda}{\mu} \right) \vec{D}_z \tag{5.9}$$

Another solution to the wave equation while holding $B_0(t)$ constant was to solely model the elasticity of the medium where the first order equation had coefficients ($\frac{\kappa}{E}$) which were

time dependant.

$$\frac{\partial^2 \vec{D}_z}{\partial x^2} = \left(\frac{\kappa}{E}\right) \vec{D}_z \tag{5.10}$$

The harmonic oscillator and elasticity models were investigated in isolation because they were also solutions to the wave equation. Substituting these solutions into the wave equation led to the simplification.

$$(\lambda - \kappa)\vec{D}_z = 0 \tag{5.11}$$

Although the original formulation of the wave equation was in terms of displacement, it was also possible to express each solution as the instantaneous velocity over the frame duration (Δt):

$$\vec{v}_{D,z} \approx A_0(x, y, z)\frac{\partial B_0(t)}{\partial t}\Delta t \tag{5.12}$$

Since $\vec{D}_z$ was separable into $A_0(x, y, z)B_0(t)$, it was straightforward to take the time derivative of both sides of the wave equation and substitute terms A_0 and B_0 giving:

$$\mu\frac{\partial^2 \vec{v}_{D,z}}{\partial t^2} = E\frac{\partial^2 \vec{v}_{D,z}}{\partial^2 z} \tag{5.13}$$

The advantage of expressing the wave equation in terms of $\vec{v}_{D,z}$ was that it was an estimate of the transform between successive frames. Fitting $\vec{D}_z$ required the choice of an arbitrary reference frame for which the magnitude of transforms would change depending on the choice of reference frame.

The piston model of breathing simplified to the wave equation with the assumption that E and μ were homogenous and did not vary from voxel to voxel. If E and μ were permitted to vary in this way then it may have been accurate to not integrate equation 5.6. In this case it was necessary to calculate the Jacobean (J_z) rather than using the displacement (D_z).

$$J_z = \frac{\partial}{\partial z}\left(\vec{D}_z + s_z z\right) \tag{5.14}$$

Where z_0 was the voxel position before being transformed. J_z was a constant for the identity transform. Furthermore, the velocity of the Jacobean was computed as

$$\vec{v}_{J,z} = \frac{\partial}{\partial t}\left(\frac{\partial \vec{T}_z}{\partial z}\right)\Delta t \tag{5.15}$$

Thus, another set of solutions to the piston model of breathing was found by solving the wave equation of the Jacobean:

$$\mu\frac{\partial^2 \vec{v}_{J,z}}{\partial t^2} = E\frac{\partial^2 \vec{v}_{J,z}}{\partial^2 z} \tag{5.16}$$

In order to adequately investigate all possible variations of the piston model of breathing, nine different candidate model equations were considered.

107

Table 5.1: Candidate model equations

	Displacement	Velocity	Velocity of Jacobean
Oscillator	$\frac{\partial^2 \vec{D}_z}{\partial t^2} = \left(\frac{\lambda}{\mu}\right)\vec{D}_z$	$\frac{\partial^2 \vec{v}_{D,z}}{\partial t^2} = \left(\frac{\lambda}{\mu}\right)\vec{v}_{D,z}$	$\frac{\partial^2 \vec{v}_{J,z}}{\partial t^2} = \left(\frac{\lambda}{\mu}\right)\vec{v}_{J,z}$
Elastic	$\frac{\partial^2 \vec{D}_z}{\partial x^2} = \left(\frac{\kappa}{E}\right)\vec{D}_z$	$\frac{\partial^2 \vec{v}_{D,z}}{\partial x^2} = \left(\frac{\kappa}{E}\right)\vec{v}_{D,z}$	$\frac{\partial^2 \vec{v}_{J,z}}{\partial x^2} = \left(\frac{\kappa}{E}\right)\vec{v}_{J,z}$
Wave	$\frac{\partial^2 \vec{D}_z}{\partial t^2} = \left(\frac{\mu}{E}\right)\frac{\partial^2 \vec{D}_z}{\partial^2 z}$	$\frac{\partial^2 \vec{v}_{D,z}}{\partial t^2} = \left(\frac{\mu}{E}\right)\frac{\partial^2 \vec{v}_{D,z}}{\partial^2 z}$	$\frac{\partial^2 \vec{v}_{J,z}}{\partial t^2} = \left(\frac{\mu}{E}\right)\frac{\partial^2 \vec{v}_{J,z}}{\partial^2 z}$

The purpose of these models was to address inaccuracies in the estimated deformation across frames due to image noise which corrupts image registration. A statistical analysis of these candidate models would reveal whether the constraints on estimates of $\vec{D}_z$ or $\vec{v}_D$ or $\vec{v}_J$ are plausible.

5.2.4 Evaluation of Models

Pearson's goodness of fit to gold standard data allowed us to rank the appropriateness of each model. Both Pearson's R and Chi squared were intended to reveal if the candidate model could make a prediction about unknown displacements given a set of existing displacements and thus was suitable for incorporating into image registration.

The fitting process of candidate equations in table 5.1, involved rearranging terms into the form $y = mx$. For example, the wave equation was rearranged into a linear equation $y = mx$ where $y = \frac{\partial^2 \vec{D}_z}{\partial t^2}$, $x = \frac{\partial^2 \vec{D}_z}{\partial^2 z}$ and the model coefficient $m = \frac{\mu}{E}$. Two scenarios were considered for fitting the model:

1. The model coefficient is the same for the **Whole organ**

2. The model coefficients are independent for each voxel, thus allowing for **Voxelwise** variations in the air flow and tissue properties.
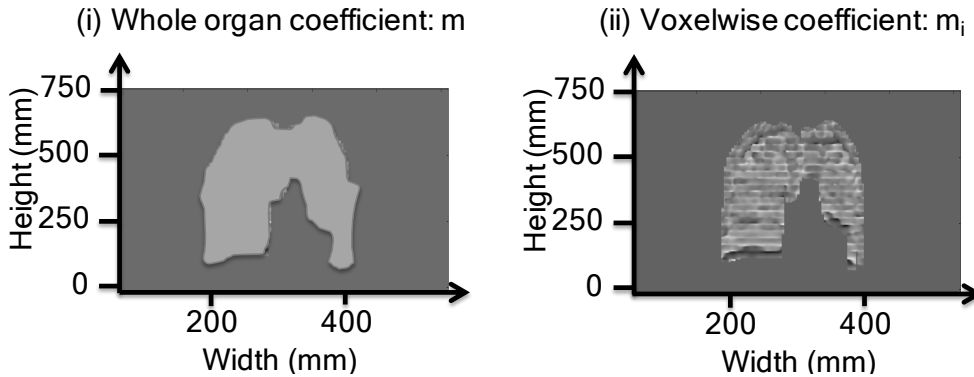

Figure 5.5: A coronal cross-section of the lung showing fitting coefficients, m, for the whole lung (left) and voxelwise (right).

An individual fit for every voxel was achieved by finding the coefficient m_i for every voxel i through least squares minimisation.

$$\arg\min_{m_i} \sum_t (y_{i,t} - m_i x_{i,t})^2 \tag{5.17}$$

where $x_{i,t}$ and $y_{i,t}$ were functions of displacement at voxel i and frame t.

Outlier detection was used to only include voxels which fell within a range, δ^2, of the linear model.

$$\frac{(y_{i,t} - m_i x_{i,t})^2}{(y_{i,t} - \bar{y}_i)^2} < \delta^2 \tag{5.18}$$

This accounted for the possibility that some voxels completely depart from the linear model $y = mx$ even though a majority of voxels demonstrate a pattern. Less than 1% of data was excluded by setting δ^2 to 8.

The chi squared statistic, χ^2, was used to test the assumption that $x_{i,t}$ and $y_{i,t}$ satisfy the model. χ^2 was used to determine the probability that $x_{i,t}$ and $y_{i,t}$ could have been obtained by chance alone (The null hypothesis). When reporting χ^2, it was assumed that the number of degrees of freedom as $N_t - 1$, where N_t is the number of time points being used for the fit. A candidate model was considered as plausible if χ^2 was small:

$$\frac{\chi^2}{N_t - 1} < 1 \tag{5.19}$$

χ^2 relies on an accurate estimate of variance, however, such estimates were difficult within voxelwise models because each voxel was fitted independently over a small number of time points, N_t. The χ^2 for each voxel, i, was:

$$\chi_i^2 = \frac{\sum_t (y_{i,t} - m_i x_{i,t})^2}{\frac{1}{N_t} \sum_t (y_{i,t} - \bar{y}_i)^2} \tag{5.20}$$

To produce a single statistic for the voxelwise model, the average χ_i^2 was computed. In order to prevent inaccuracies caused by a poor estimate of the variance of individual χ_i^2 the average variance over the whole organ was computed rather than individual variances for each χ_i^2.

$$\chi_{voxelwise}^2 = \frac{\frac{1}{N_i} \sum_i \sum_t (y_{i,t} - m_i x_{i,t})^2}{\frac{1}{N_i} \sum_i \frac{1}{N_t} \sum_t (y_{i,t} - \bar{y}_i)^2} \tag{5.21}$$

Pearson's correlation coefficient, R^2, was used to test whether the assumptions satisfied the model and considered a candidate model as plausible if $R^2 > 0.6$. Pearson's correlation for the voxelwise model, r_i, was used to test that the hypothesis that the $x_{i,t}$ and $y_{i,t}$ were independent.

$$r_i = \frac{\sum_t (y_{i,t} - \bar{y}_i)(x_{i,t} - \bar{x}_i)}{\sqrt{\sum_t (y_{i,t} - \bar{y}_i)}\sqrt{\sum_t (x_{i,t} - \bar{x}_i)}} \tag{5.22}$$

The voxelwise pearsons correlation was collated into a single statistic, $r_{voxelwise}$ using a method inspired by (Nikolic, Murean, Feng, & Singer, 2012). In order to prevent inaccuracies caused by a poor estimate of the covariances of individual r_i, the average covariance over the whole organ was calculated.

$$r_{voxelwise} = \frac{\sum_i (\sum_t (y_{i,t} - \bar{y}_i)(x_{i,t} - \bar{x}_i))}{\sqrt{\sum_i \sum_t (y_{i,t} - \bar{y})} \sqrt{\sum_i \sum_t (x_{i,t} - \bar{x})}} \tag{5.23}$$

The $r_{voxelwise}$ and $\chi^2_{voxelwise}$ statistics were computed for all candidate models in table 5.1 and used the results to decide if the models were plausible and thus suitable for implementing as a regulariser.

5.2.5 Regularisation of Image Registration

Each frame in the gated study was registered to the reference frame, for the purposes of estimating the transform $(\vec{T})$ corresponding to respiratory motion. Each implementation of image registration found the transform for which the intensity values of a frame (I_{float}) corresponded to the intensity values of a reference frame I_{ref}. There were many different possible transforms $\vec{T}$ which can account the same apparent motion between I_{float} and I_{ref}. This problem has been previously identified as ill-posed since it is not possible to distinguish which transform $\vec{T}$ is a more accurate estimate of motion (Zitov & Flusser, 2003). Initially, a fairly common method for addressing the ill-posedness was used by adding a regularisation weighting term l_2 which penalises large displacements $(\vec{D})$. The implementation image registration algorithm minimised the following:

$$\arg\min_{\vec{T}} \left\{ \mathcal{C}(\vec{T}; I_{float}, I_{ref}) + l_2 \left\| \vec{D} \right\|^2 \right\} \tag{5.24}$$

where $\mathcal{C}$ was a measure of the voxelwise similarity between I_{ref} and the transformed I_{float}.

Our chosen implementation of image registration used a gradient descent algorithm to minimise $\mathcal{C}$. We suspected that it would be inaccurate to use image registration to resolve displacements which were smaller than the system resolution. Image I_{float} and I_{ref} were smoothed, as well as the gradient of $\mathcal{C}$ using a gaussian filter which was of similar width to the system resolution.

Motion modelling was introduced into image registration by altering equation 5.24 to include $\Omega(\vec{T})$, which was a penalty term due to a departure of $\vec{T}$ from one of the candidate models from table 5.1. The altered implementation of image registration algorithm minimised:

$$\arg\min_{\vec{T}} \left\{ \mathcal{C}(\vec{T}; I_{float}, I_{ref}) + \left| \Omega(\vec{T}) \right|^2 \right\} \tag{5.25}$$

Initially the wave equation (equation 5.7) was used as a penalty term, $\Omega(\vec{T})_{wave}$.

$$\left| \Omega(\vec{T})_{wave} \right|^2 = \left(E \frac{\partial^2 \vec{D}_z}{\partial t^2} - \mu \frac{\partial^2 \vec{D}_z}{\partial^2 z} \right)^2 \tag{5.26}$$

The penalty term was simplified by substituting λ and κ from equations 5.9 and 5.10 after dividing both sides by $\vec{D}$. The penalty term became the sum of two independent variables, one which varied with time (κ) and one which varied with space (λ). The penalty term was weighted by z_2.

110

$$\left|\Omega(\vec{T})_{wave}\right|^2 = z_2(\lambda - \kappa)^2 \vec{D}_z^2 \tag{5.27}$$

It was proposed that if $\Omega(\vec{T})_{wave}$ was small then it did not vary significantly with $\vec{D}$ because it was the sum of the two independent variables. Using this approximation allowed us to avoid recalculating $\Omega(\vec{T})_{wave}$ for every estimate of $\vec{T}$.

Previous work on directionality constraints (Ruan, 2009) led to the suspicion that over-penalising $\vec{D}_z$ could lead to slipping of $\vec{T}$ in the $\vec{D}_x$ and $\vec{D}_y$ directions. In order to avoid this, orthogonal directions were penalised with a uniform weighting factor l_2 and combined into a single voxelwise penalty term.

$$\left|\Omega(\vec{T})_{wave}\right|^2 = l_2(\vec{D}_x^2 + \vec{D}_y^2) + z_2(\lambda - \kappa)^2 \vec{D}_z^2 \tag{5.28}$$

A multi resolution approach was used, whereby image registration was iteratively applied, but the spacing between control points was decreased at each iteration (Table 5.2). The gated frames were grouped into time intervals and summed. The number of intervals (n_i) was increased at each iteration of resolution until the number of intervals equaled the number of frames (Figure 5.6).

Table 5.2: Control point spacing and the number for intervals using 3 multi resolution levels

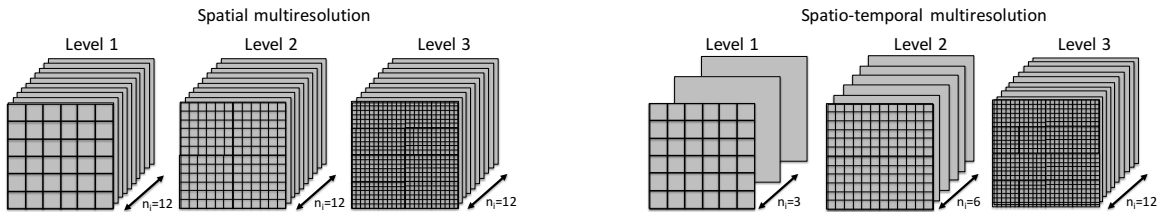| Spacing	X	Y	Z	n_i
Level 1	$80mm$	$80mm$	$40mm$	3 time intervals
Level 2	$40mm$	$40mm$	$20mm$	6 time intervals
Level 3	$20mm$	$20mm$	$10mm$	12 time intervals

Figure 5.6: The control point grid using 3 spatial multi resolution levels (left) and 3 spatiotemporal multi resolution levels (right). The first spatiotemporal level had 3 time intervals with coarse control point spacing, the second level had 6 time intervals and the third level had 12 time intervals with fine control spacing

The algorithm for the multi resolution approach began with the first level of registration which was initialised at $z_2 = 0$. $\lambda - \kappa$ was fitted and then used as a penalty term in the subsequent iteration in the loop. Table 5.3 shows the algorithm used for resampling the control point positions at each resolution level.

Table 5.3: The multi-resolution algorithm

Level 1
Registration with $z_2 = 0$
Loop 3 times:
Use fitting to calculate $\lambda - \kappa$
Level 1 Registration with $z_2 > 0$

$$\Downarrow$$

Level 2
Resample control points to Level 2
Loop 3 times:
Use fitting to calculate $\lambda - \kappa$
Level 2 Registration with $z_2 > 0$

$$\Downarrow$$

Level 3
Resample control points to Level 3
Loop 3 times:
Use fitting to calculate $\lambda - \kappa$
Level 3 Registration with $z_2 > 0$

As the number of time intervals increased, the size of the transform between each individual interval decreased and the noise within each summed frame increased. To minimise the effect of these changes the transform penalty was multiplied by n_i such that a larger penalty was applied to small transforms between closer time intervals.

$$\left|\Omega(\vec{T})\right|^2 = n_i^2 \left(l_2(\vec{D}_x^2 + \vec{D}_y^2) + z_2(\lambda - \kappa)^2 \vec{D}_z^2 \right) \tag{5.29}$$

It was suspected that in some cases a very large $(\lambda - \kappa)^2$ could lead to sideways shearing. This issue was prevented by making the penalisation of $\vec{D}_z$ conditional such that $\vec{D}_z$ was not penalised more than $\vec{D}_x$ or $\vec{D}_y$.

Where $z_2(\lambda - \kappa)^2 < l_2$:

$$\left|\Omega(\vec{T})\right|^2 = n_i^2 \left(l_2(\vec{D}_x^2 + \vec{D}_y^2) + z_2(\lambda - \kappa)^2 \vec{D}_z^2 \right) \tag{5.30}$$

otherwise:

$$\left|\Omega(\vec{T})\right|^2 = n_i^2 l_2 \left(\vec{D}_x^2 + \vec{D}_y^2 + \vec{D}_z^2 \right) \tag{5.31}$$

The regulariser was altered to include a variation of the piston model of breathing using the wave equation of the Jacobean from table 5.1. The Jacobean was estimated using the derivative of the displacement from equation 5.14. A penalty term, $\Omega(\vec{T})_{wave}$, was devised by integrating equation 5.15 from the base to the height (h) of the voxel undergoing respiratory motion

$$\left|\Omega(\vec{T})\right|^2 = \left| \frac{1}{\vec{D}_z} \int_0^h E\frac{\partial^2 J_z}{\partial t^2} - \mu\frac{\partial^2 J_z}{\partial^2 z} dz \right|^2 \vec{D}_z^2 \tag{5.32}$$

$\Omega(\vec{T})_{wave}$ was expressed as a weighting of the displacement $\vec{D}_z^2$ to permit the use of the same method of regularising estimates of the displacement from equations 5.26 and 5.27.

5.2.6 Optimisation of Image Registration for Cine CT

The novel method of using the wave equation regularisation was incorporated into CT image registration. Although the registration of frames within cine CT images is applicable to techniques for tumour tracking it is necessary for the purposes of attenuation correction. The purpose of analysing motion estimates from cine CT was to exclude poorly performing models and devise a range of potentially optimal parameters.

A trial and error approach was used for finding the optimal image registration parameters. Registration of each cine CT frame to the first CT frame was achieved by maximising the normalised mutual information (NMI) between sequential frames. The negative NMI, $-\mathcal{N}$, was substituted as the cost function $\mathcal{C}$ to be minimised through image registration (Pluim, Maintz, & Viergever, 2003),

$$\mathcal{C} = -\mathcal{N}\left(I_t, \vec{T}_t(I_{t+1})\right) = -\frac{H\left(\vec{T}_t(I_{t+1})\right) + H\left(I_t\right)}{H\left(I_t, \vec{T}_t(I_{t+1})\right)} \quad (5.33)$$

where $H(...)$ was the Shannon entropy for a joint distribution, I_t is an image of the first frame and I_{t+1} the next frame and $\vec{T}_t$ is the transform estimate from image registration.

The inter-frame displacement $\vec{v}_D$ was initially found using frame-by-frame registration. The image registration was regularised by penalising the cost function with the size of the displacement squared $\|\vec{v}_D\|^2$ weighted by l_2 (equation 5.24). The proposed additional penalty was compared to the cost function for departure from the wave equation model weighted by z_2 (equation 5.28). The registration parameters are summarised in table 5.4 and labelled as C1 ... C8.

Table 5.4: CT image registration weighting parameters

	C1	C2	C3	C4	C5	C7	C8
l_2	10^{-9}	10^{-7}	10^{-11}	10^{-9}	10^{-11}	10^{-9}	10^{-11}
z_2	0	0	0	10^{-5}	10^{-5}	10^{-3}	10^{-3}

The goodness of fit to the wave equation model ($\chi^2_{voxelwise}$) and the NMI were used to benchmark the robustness of registration when using weighting parameters from table 5.4. The voxelwise average of the goodness of fit, $\chi^2_{voxelwise}$, was calculated using equation 5.21. The NMI between each sequential frame after registration was found and the mean NMI, $\bar{\mathcal{N}}$ in equation 5.34, was also calculated by averaging over all frames, n_f.

$$\bar{\mathcal{N}} = \frac{1}{n_f} \sum_{n_f} \mathcal{N}\left(I_t, \vec{T}_t(I_{t+1})\right) \quad (5.34)$$

The mean NMI was used as a metric to choose optimal registration weighting parameters from table 5.4. Image registration is an iterative process which minimises a cost function, however, local minima may prevent the global minimum from being reached. Previous experience suggested local minima is often caused by noise and discontinuities in the cost function as well as incorrect or large penalty terms. The mean NMI was used to check that registration was optimal and that the maximum NMI has been reached even after the inclusion of additional constraints.
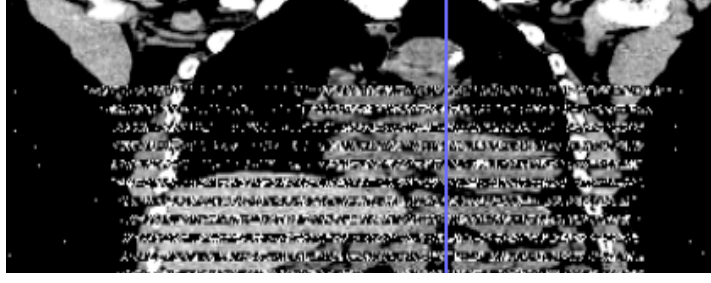
Figure 5.7: A coronal cross-section of the corrupted reference frame used for image registration. Approximately 50% of the slices have been replaced with Poisson noise. The SNR for the corrupted slices is approximately 0.5.

The absence of a ground truth for the accuracy of displacement estimates $\vec{v}_D$ meant that it was not possible to isolate the many causes of potential error. However, it was possible to test for the robustness of registration by intentionally corrupting the CT. Image corruption was simulated for the first frame in the cine CT to observe which combination of weighting parameters were least effected by the corruption. Image corruption was simulated (figure 5.7) by removing features from the image and replacing slices with poisson noise equivalents. The NMI_{mean} was recorded to measure the effect of image corruption on registration when using each of the weighting parameters in table 5.4.

5.2.7 Optimisation of Image Registration for Respiratory Gated PET

Two retrospective studies were used to investigate image registration applied to gated PET acquisitions from different models of PET/CT scanner. The initial study used 6 frames of data from a Biograph 16 scanner, and then it was expanded testing to 12 frames of data from a mCT scanner. Each study in table 5.5 and 5.6 required a different set of image registration parameters because of difference in image quality. The optimised registration parameters have been identified with the labels P1, P2 ... P170 and are presented in the results section.

Table 5.5: Registration options for study 1 using a Biograph 16

Registration Parameters	P1 - P20
Number of frames	6
Number of patients	16
Multiresolution	Spatial
Cyclic	No
$\chi^2_{voxelwise}$ Cutoff	5

Table 5.6: Registration options for study 2 using a Biograph mCT

Registration Parameters	P101 - P111	P151 - P161	P164 - P170
Number of frames	12	12	12
Number of patients	11	11	11
Multiresolution	Spatial	Spatiotemporal	Spatiotemporal
Cyclic	No	No	Yes
$\chi^2_{voxelwise}$ Cutoff	11	11	11

Image registration in gated PET was regularised by penalising the cost function (Ω from equation 5.24) with a weighting of l_2 and z_2. Registration of each floating gated PET frame to the reference gated PET frame was achieved by minimising the sum of square differences between each voxel.

$$\mathcal{C} = \sum_{allvoxels} \left(I_{ref} - \vec{T} \left(I_{float} \right) \right)^2 \tag{5.35}$$

The noise in the PET images were addressed by smoothing both the reference and floating gated PET image with a gaussian kernel of width w_I. The gradient of the cost function was also smoothed with a gaussian kernel of width $w_{\mathcal{C}}$. Given the large number of permutations of registration parameters, optimising every combination of parameters l_2, z_2, w_I and $w_{\mathcal{C}}$ was impractical. The parameter search was stated by optimising l_2 weighting factor in isolation without applying any wave equation regularisation (P1, P2, P3 from table 5.10). The parameter, $w_{\mathcal{C}}$, was then modified to see whether the smoothing complimented (lower $\chi^2_{voxelwise}$) or hindered (higher $\chi^2_{voxelwise}$) fitting the model to the data (P4, P5, P6, P9 from table 5.10).

Unlike previous approaches of spatiotemporal image registration (Metz et al., 2011), wave equation regularisation did not explicitly constrain displacements to follow a smooth cyclic variation. In order to ensure that displacements follow a continuous periodic function the wave equation regularisation was extended to include a cyclic constraint by ensuring that the composite of all sequential frame-to-frame transforms were the identity transform. The rationale that displacements are a periodic function has been justified in previous investigations on respiratory gating (Metz et al., 2011). This approach imposed a cyclic constraint on the 1 dimensional regulariser by assuming that $\vec{v}_{D,z}$ for the last frame must be the same as the inverse of the cumulative displacement to that time point $\vec{D}_z$.

$$(\vec{v}_{D,z})_{n_f - 1} = \int_{t=0}^{t=n_f-1} (\vec{v}_{D,z})_t^{-1} dt = \vec{D}_z^{-1} \tag{5.36}$$

In addition to penalising the cost function, the effect of a multiresolution strategy was investigated. A multi-resolution strategy applied in only the spatial domain was compared to multi-resolution in both a spatial and temporal domain (as illustrated in figure 5.6). A new set of optimal registration parameters were found for each multi-resolution strategy (as shown in table 5.3). The parameters optimised from the retrospective studies were evaluated based on the average of the goodness of fit ($\chi^2_{voxelwise}$) of the voxelwise respiratory motion model to the PET data. The cutoff for average goodness of fit was the number of degrees of freedom for each voxel, which was calculated as the number of frames - 1. Both the median $\chi^2_{voxelwise}$ and the standard deviation of $\chi^2_{voxelwise}$ across all patients were considered in the
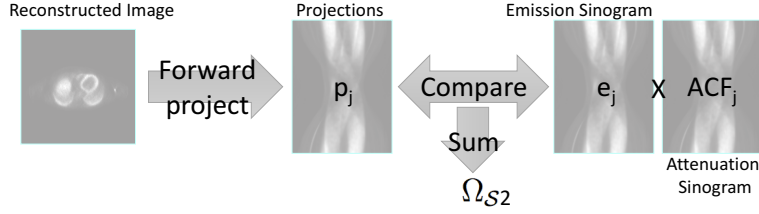
Figure 5.8: The metric for consistency is the sum of squared differences (SSD) between emission and projection sinograms. e_j are the elements in the emission sinogram. ACF_j are the attenuation correction factors. p_j are elements in the reprojected sinogram. Each emission sinogram is corrected for attenuation during reconstruction using 3D OSEM.

investigation. Registration parameters were rejected if the median was above the $\chi^2_{voxelwise}$ cutoff and these models were not considered as robust enough to apply motion estimation and attenuation correction.

5.2.8 Motion Estimation and Attenuation Correction for Respiratory Gated PET

The attenuation correction map was initially approximated using the bilinear scaling of the spiral CT from Hounsfield Units (HU) to voxelwise linear attenuation coefficients (Kinahan, Townsend, Beyer, & Sashin, 1998). This approach was known to be sensitive to attenuation mismatch artefacts, however, at least one frame in the respiratory gated PET was assumed to be aligned to the spiral CT and thus, attenuation mismatch would be minimal for that frame. A metric for the consistency of the attenuation map with the emission data from each frame was devised. The metric described in figure 5.8 was used because of its robustness to noise and reliability to identify the reference frame. The frame with the lowest metric was assumed to have the minimal amount of attenuation mismatch. The frame with minimal metric was identified as the reference frame because it was assumed to be aligned to the CT and thus provide a spatial reference for transforms between the CT and other frames in the PET. The sum of squared differences (SSD) between emission and projection sinograms was investigated in chapter 3 as a metric for attenuation alignment.

Motion between sequential gated frames in respiratory gated PET were estimated using free form image registration. This registration of sequential frames was applied without attenuation correction (NAC) and was used to find the displacements of each PET frame with respect to the reference PET frame. Optimal image registration parameters were selected from the results of motion model analysis and these were used to improve the accuracy of the estimated displacements. The displacements were then used to transform the attenuation map. The attenuation map was transformed using these displacements, so as to produce a sequence of registered attenuation maps which were aligned to their corresponding PET frame. However, the attenuation map for the reference PET frame was not transformed because it was already aligned. The consistency between emission PET data and projection sinograms was dependant on the accuracy of the ACF and therefore an indirect measure of the alignment of the registered attenuation map to each corresponding PET frame. The ratio
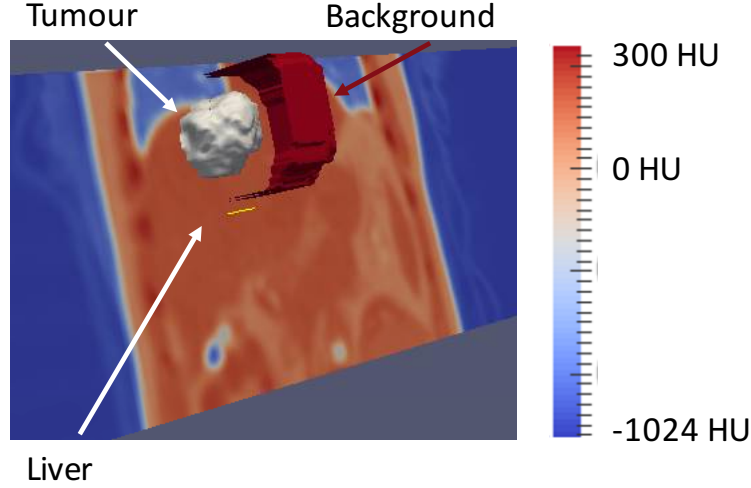
Figure 5.9: Tumour segmentation

of the sum of squared differences (SSD) was used to compare the difference between using the registered attenuation maps to just using a single spiral CT for attenuation correction. The ratio (ΔSSD) was the squared difference between emission sinogram elements ($e_{j,aligned}$) and projection sinogram elements ($p_{j,aligned}$) using an aligned attenuation map divided by the squared difference between sinogram elements using a single spiral CT ($e_{j,ctac}$ and $p_{j,ctac}$).

$$\Delta SSD = \frac{\sum_j (e_{j,aligned} - p_{j,aligned})^2}{\sum_j (e_{j,ctac} - p_{j,ctac})^2} - 1 \tag{5.37}$$

Each PET/CT study acquired on the Biograph mCT had visible lesions of FDG uptake on the respiratory gated PET image of the thorax. The outline of the suspected tumours were identified and segmented from the spiral CT and the volume of the tumour was recorded in millilitres. The outline on CT was likely to also include necrotic regions of tumour, and although it is often desirable to track the active portions of the tumour, segmenting using a threshold for FDG uptake was likely to be affected by the magnitude of respiration, thus introducing an undesired source of variation between patients. The position of the tumour was tracked by transforming the region with the displacements estimated from image registration of sequential gated PET frames. The tumour tracking region was assessed for the reliability of tracking moving FDG uptake by recording both the mean (Bq/mL) and total (Bq) activity segmented from the sequence of gated PET frames. A sequence of centre of mass (COM) measurements for tumour tracking region were recorded. A background region with approximately the same volume as the tumour region was automatically drawn encasing the tumour tracking region (figure 5.9). Tumour activity was calculated relative to both the mean background (Bq/mL) and total background (Bq) activity.

The sequence of tumour tracking regions were combined into a single tumour volume, TV, which encompasses the full motion of the tumour. The variation of the mean (Bq/mL) and total (Bq) TV activity (a_t) were assessed for all frames in the gated PET sequence. In order to compare studies with varying tumour uptake, the standard deviation of the TV activity was then calculated as a fraction of the mean TV activity over the entire cycle.

$$CV_{act} = \frac{\sqrt{\sum_t (a_t - 1/n_f \sum_t a_t)^2}}{1/n_f \sum_t a_t} \qquad (5.38)$$

The variation of the tracking region volume, v_t, was assessed. In order to compare studies with varying tumour size, the standard deviation of v_t was calculated as a fraction of the mean tracking region volume over the entire cycle.

$$CV_{size} = \frac{\sqrt{\sum_t (v_t - 1/n_f \sum_t v_t)^2}}{1/n_f \sum_t v_t} \qquad (5.39)$$

5.3 Results

5.3.1 Evaluation of Candidate Models

An initial evaluation was performed to see if the candidate oscillator, elastic and wave models described in section 5.2.3 arose a real trends observed in patient data Pearson's R^2 correlation coefficient revealed the cases where the models made consistently inaccurate predictions in the CT scans of 8 patients. At least one model was assumed to be a candidate for predicting the Jacobean, Velocity of Jacobean and the Velocity. In order to ensure that at least one model was considered, a cutoff for R^2 was chosen as $R^2 < 0.6$ ($|R| < 0.77$). A value lower than the cutoff of was interpreted as meaning no relationship between the actual values and the values predicted by the model (Cutoff show by red vertical line in Fig. 5.10) .

Pearson's R^2 correlation for the whole lung between the actual measurements of the Jacobean and velocity and the predictions made by the wave model and oscillator model were affected by the spatial correlation of deformation. Fewer control points for the spline parameterisation of the deformation field favoured a whole lung approach because there was less spatial variance between adjacent voxels. Figure 5.11 is a visual representation of measurements from 1000 voxels chosen at random. There was good correlation between adjacent voxels for the oscillator model because the frequency of oscillation should be very similar across all voxels (rate of respiration).

Pearson's correlation coefficient (Fig. 5.10) revealed that the Oscillator Model was a candidate for predicting the Velocity ($R^2 = 0.83$) and that the Wave Equation Model was a candidate for predicting the Velocity of the Jacobean ($R^2 = 0.62$). Outliers in the approximation of the Jacobean led to a value of Pearson's R^2 correlation coefficient which was either disproportionally high or low, particularly for the Elastic Model. The interpretation of outliers was avoided by just considering the median R^2 for all patients (Table 5.7).

Scaling Pearson's correlation coefficient reduced the effect of the outliers, particularly where R^2 was disproportionally high in the elastic model. Scaling Pearson's correlation coefficient excluded the covariance between voxels which increased R^2 on average in the oscillator and wave equation models. Scaling Pearson's correlation coefficient (Table 5.8) revealed that the Wave Equation Model was a candidate for predicting both the Jacobean ($R^2 = 0.61$) and Velocity of Jacobean ($R^2 = 0.66$) whilst the Oscillator Model was a candidate for predicting the Velocity ($R^2 = 0.87$).

Chi squared does not measure covariance, so most models were identified as potential candidates (Table 5.9). A p-value cut off of p=0.01 was used to identify models which
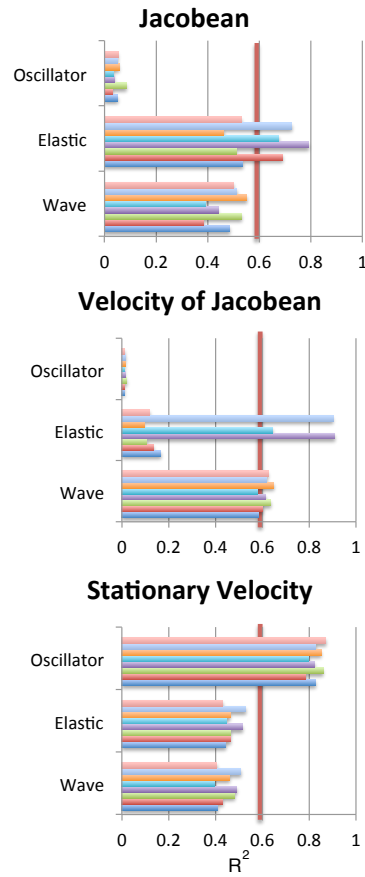
Figure 5.10: R^2 for Oscillator, Elastic and Wave models applied to the CT scans of 8 patients. The R^2 value for each patient is represented by a different colour bar. The vertical red line indicates the cutoff for R^2 chosen to ensure that at least one model was considered
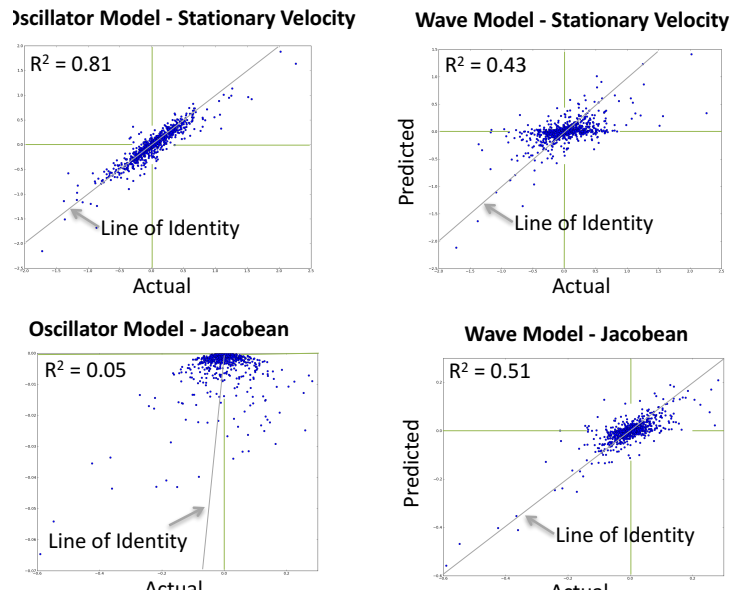

Figure 5.11: Scatter plot of models for patient 9

Table 5.7: Median values of Pearson's R^2. *Candidate models were identified by $R^2 > 0.6$

	Pearsons R^2		
	Wave Model	Elastic Model	Oscillator Model
Jacobean	0.49	**0.60***	0.05
Velocity Of Jacobean	**0.62***	0.15	0.01
Velocity	0.44	0.46	**0.83***

Table 5.8: Median Values of Scaled Pearson's R^2. *Candidate models were identified by $R^2 > 0.6$

	Scaled Pearsons R^2		
	Wave Model	Elastic Model	Oscillator Model
Jacobean	**0.61***	0.32	0.46
Velocity Of Jacobean	**0.66***	0.06	0.48
Velocity	0.51	0.47	**0.87***

could be ruled out as candidates. The number of degrees of freedom (ν) was the number of frames per voxel (10) minus the number of model coefficients per voxel (1). Candidate models were those for which χ^2/ν which fell within the confidence interval $(0.19, 2.62)$. The model with the lowest Chi squared was the Elastic model for predicting both the Jacobean ($\chi^2/\nu = 0.76$), Wave Equation model for predicting the Velocity Of Jacobean ($\chi^2/\nu = 0.69$) and the Oscillator model for predicting the Velocity ($\chi^2/\nu = 0.19$). χ^2 for using the Oscillator model to predict the Velocity was exceptionally low implying that the number of degrees of freedom (ν) may actually be reduced due to some unaccounted symmetry.

Table 5.9: Median Values of Reduced $\chi^2 (\nu = 9)$. *Candidate models were identified by the 99% $CI = (0.19, 2.62)$

	χ^2/ν		
	Wave Model	Elastic Model	Oscillator Model
Jacobean	1.15*	**0.76***	27.79
Velocity Of Jacobean	**0.69***	6.50	134.85
Velocity	1.12*	1.02*	**0.19***

Although a high chi squared was useful for ruling out candidate models, there was a bias when comparing the chi squared between velocity and displacement approximations. $\vec{D}_z$ already contained a correlation in that each successive frame was a product of all the transforms leading up to that frame. Thus, a lower chi squared was expected for the velocity model making direct comparison to the Jacobean model difficult.
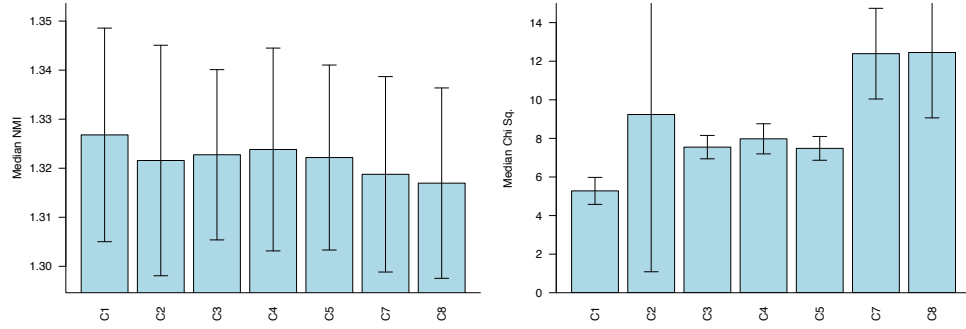
Figure 5.12: Comparison of NMI (left) and χ^2 (right) for various tuning parameters in all 8 patients. The barplot on the left shows median NMI for all patients and all frames and the error bars are the standard deviation from the median value. The barplot on the right shows the median χ^2 for all patients and the error bars are the standard deviation from the median value. The median gives an overall comparison of tuning parameters C1, C2, C4, C5, C7 and C8. C1 and C4 produce very the highest NMI, whereas C2, C7 and C8 produce a slightly reduced NMI and a corresponding increase in χ^2.
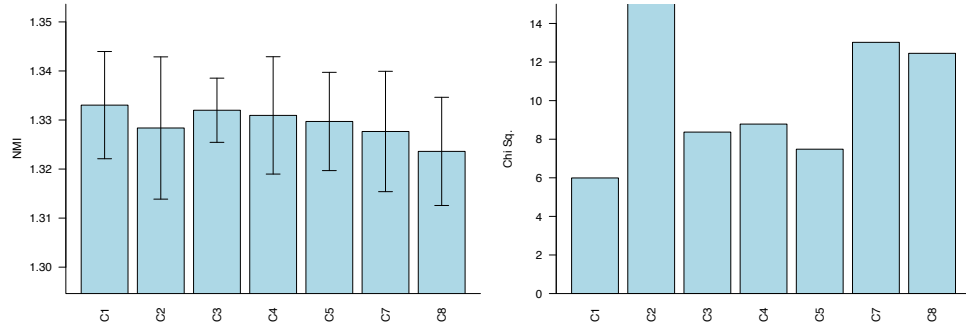

Figure 5.13: Comparison of NMI (left) and χ^2 (right) for various tuning parameters in patient 2. The barplot on the left shows median NMI over all frames of patient 2 and the error bars are the standard deviation from the median value. The barplot on the right shows χ^2 for patient 2. C2, C7 and C8 produce a slightly reduced NMI and a corresponding increase in χ^2.
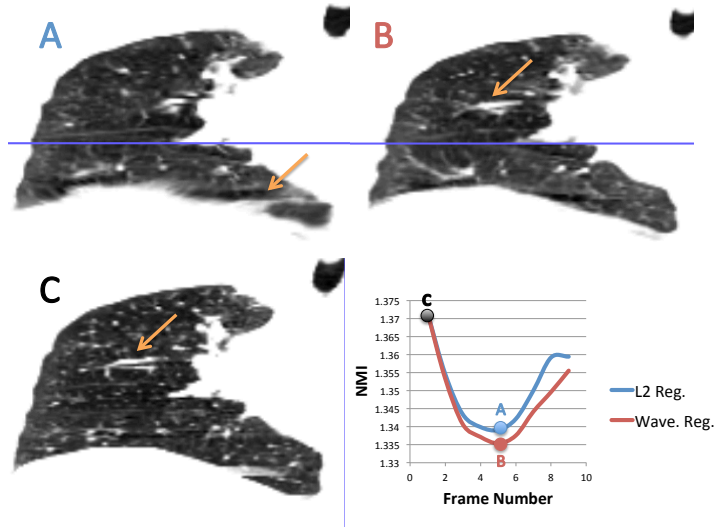
121

Figure 5.14: Coronal slices of frame number 5 (A,B) after being registered to the reference frame (C). The plot shows the Normalised Mutual Information (NMI) for each frame in comparison to the reference frame. Regularisation was applied (A) using a weighting of $l_2 = 10^{-9}$. Wave equation regularisation was applied (B) using a weighting of $z_2 = 10^{-3}$. l_2 weighted regularisation led to smearing in the image (arrow in A). Some anatomical features were only present after w_2 weighted regularisation using the wave equation (orange arrow in B and C).

5.3.2 Modelling Respiratory Motion in Cine CT

The results from modelling respiratory motion in cine CT allowed us to exclude poorly performing models and devise a range of potentially optimal parameters. The results indicated that NMI was not significantly affected provided that the tuning parameter, z_2, was small. It was found that the NMI was highest if the penalty term was set to $z_2 = 10^{-5}$ (C4 in figure 5.12) and that the NMI was decreased marginally when z_2 was increased to $z_2 = 10^{-3}$ (C7 in figure 5.12). It was optimal to set $l_2 = 10^{-9}$ (C1 in figure 5.12) as using a larger or smaller l_2 led to a decrease in NMI (C2 and C3 in figure 5.12). It was suspected that NMI was not sensitive to model fitting because large segments of each image remained stationary and were not affected by respiratory motion. Changes in the NMI were very marginal making the results difficult to interpret, thus the changes in the NMI were compared to the corresponding changes in χ^2.

The χ^2 was observed to increase in circumstances where the NMI decreased for individual patients (figure 5.13) as well as for the median across all patients (figure 5.12). In most cases, χ^2 had a much smaller variation (standard deviation) across patients and thus was a more sensitive metric than NMI. In one case χ^2 varied significantly due to poor fitting for a single patient (C2 in 5.13). It was found that increasing z_2 increased the χ^2, thus, over-weighting the z_2 penalty did not enforce the wave equation constraint.

Prior to corrupting the first frame of the cine CT, a qualitative assessment of robustness of registration was completed. Using l_2 regularisation in isolation led to image smearing in figure 5.14 (A). The smearing of the image near significant structures was potentially due to
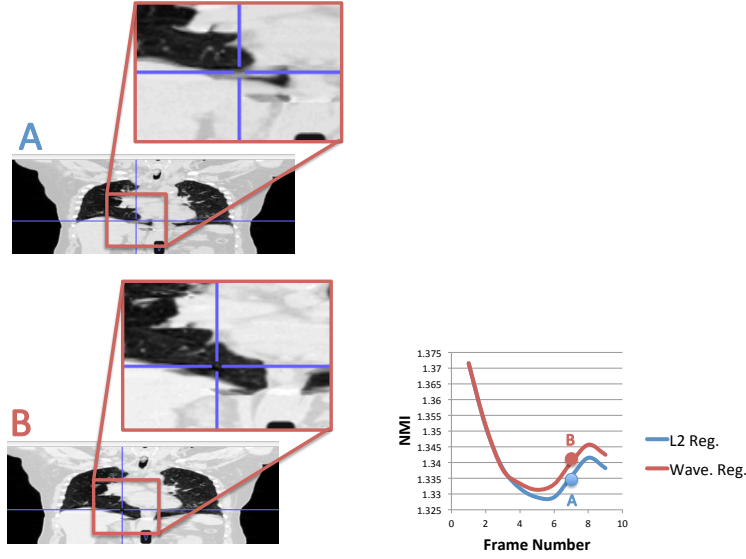
Figure 5.15: Coronal slices of frame number 7 (A,B) after being registered to a corrupted version of the reference frame. The plot shows the Normalised Mutual Information (NMI) for each frame in comparison to the uncorrupted reference frame. Regularisation was applied (A) using a weighting of $l_2 = 10^{-9}$. Wave equation regularisation was applied (B) using a weighting of $z_2 = 10^{-3}$

over penalisation of l_2. Using wave equation regularisation through z_2 weighted penalisation appeared to preserve image sharpness in figure 5.14 (B). Object boundaries were preserved, especially at the diaphragm and cysts within the lung. Interestingly, the NMI when using the z_2 weighted regularisation was less than the NMI for l_2 weighted regularisation alone, however, the former produced a sharper image.

A qualitative assessment showed that after corrupting the first frame of the cine CT z_2 weighted regularisation limited artefacts introduced from a corrupted target image. Figure 5.15 (B) shows that smearing artefacts were removed and boundary edges were preserved. The z_2 weighted regularisation led to an increase in NMI which may have been due to the improved definition of boundary edges. It was suspected that wave equation regularisation guided the registration in the corrupted regions, thus improving the robustness of the registration algorithm.

5.3.3 Modelling Respiratory Motion in Gated PET

Various registration parameter combinations were used in the first retrospective PET/CT study (gated into 6 frames). χ^2 was measured and fell within the median $\chi^2 < 5$ for 50% of the parameter combinations. Table 5.10 shows the median $\tilde{\chi}^2$ over all patients when applying the l_2 weighting factor in isolation. Although many parameter combinations were ruled out as unsuitable ($\chi^2 > 5$), the optimal χ^2 could be achieved with $l_2 = 10^{-7}$ and a smaller l_2 was optimal if further smoothing (w_C) was applied.

123

Table 5.10: Parameter combinations using l_2 weighting factor and spatial smoothing. The bottom row shows the median $\tilde{\chi}^2$ fit of the wave equation model to 6 time point estimates of $\vec{v}_{J,z}$.

	P1	P2	P3	P4	P5	P6	P9
l_2	10^{-9}	10^{-7}	10^{-5}	10^{-9}	10^{-7}	10^{-5}	10^{-9}
$w_{\mathcal{C}}$	0	0	0	2	2	2	4
w_I	6.4	6.4	6.4	6.4	6.4	6.4	6.4
$\tilde{\chi}^2(\nu = 5)$	**4.22**	**3.86**	5.32	**3.44**	5.98	5.31	**3.39**

When wave equation modelling of $\vec{v}_{J,z}$ was used the fitting analysis led to a small increase in the median $\tilde{\chi}^2$ as shown in table 5.11. It was found that a similar χ^2 could be obtained by combining a smaller l_2 weight and $z_2 = 10^{-5}$ weight , however, increasing z_2 above 10^{-3} was ruled out as it led to $\vec{v}_{J,z}$ vanishing to zero. It was not clear whether it was useful to apply further smoothing ($w_{\mathcal{C}}$).

Table 5.11: Parameter combinations using l_2 and z_2 weighting factors in combination with spatial smoothing. The bottom row shows the median χ^2 fit of the wave equation model to 6 time point estimates of $\vec{v}_{J,z}$.

	P12	P13	P14	P15	P16	P17	P18	P19	P20
l_2	10^{-9}	10^{-11}	10^{-9}	10^{-9}	10^{-11}	10^{-9}	10^{-9}	10^{-11}	10^{-9}
z_2	10^{-3}	10^{-5}	10^{-5}	10^{-3}	10^{-5}	10^{-5}	10^{-3}	10^{-5}	10^{-5}
$w_{\mathcal{C}}$	0	0	0	2	2	2	4	4	4
w_I	6.4	6.4	6.4	6.4	6.4	6.4	6.4	6.4	6.4
$\tilde{\chi}^2(\nu = 5)$	6.05	**3.95**	**3.96**	5.82	**4.00**	**4.06**	5.22	**4.22**	**3.86**

Despite measuring the mean frame to frame NMI for each parameter combination, it was found that the NMI was not sensitive enough to distinguish the robustness of the motion model. Figure 5.16 shows a patient where smaller χ^2 was weakly correlated with a larger NMI, but the variation of NMI from frame to frame exceeded this difference.

Various registration parameter combinations were used in the second retrospective PET/CT study (gated into 12 frames). χ^2 was measured and fell within the median $\tilde{\chi}^2 < 11$ for 74% of the parameter combinations. A spatial multi-resolution approach was used to optimising the wave equation motion model of $\vec{v}_{J,z}$, however, difficulties were experienced in finding parameter combinations which consistently had a median of $\tilde{\chi}^2 < 15$. It was necessary to increase z_2 weighting substantially and the lowest χ^2 in table 5.12 was achieved with $z_2 = 10^{-5}$ and smoothing the image with a $6.4mm$ gaussian (w_I) and the cost function with a $4mm$ gaussian ($w_{\mathcal{C}}$).
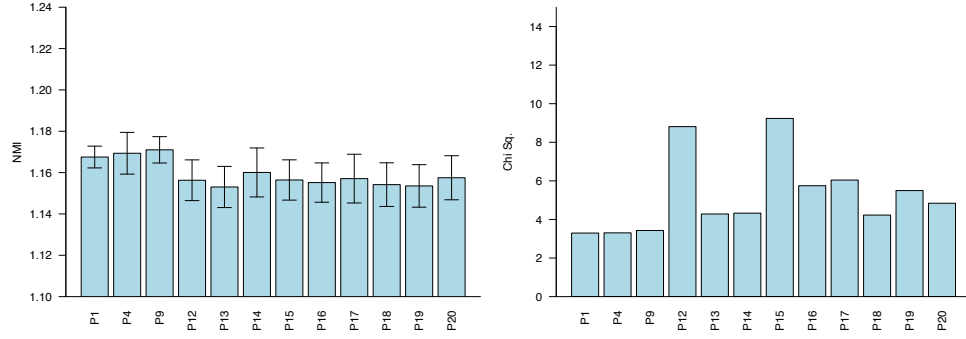
Figure 5.16: A barplot of mean NMI (left) and χ^2 (right) values for a single patient in the retrospective PET/CT study. The error bar is the standard deviation in values obtained from the frame to frame NMI in the gated PET. Analysis of the χ^2 revealed that parameter combinations P12 and P15 should be rejected, however, there was no significant pattern in the NMI values.

Table 5.12: Parameter combinations using l_2 and z_2 weighting factors in combination with spatial smoothing and spatial multiresolution. The bottom row shows the median $\tilde{\chi}^2$ fit of the wave equation model to 12 time point estimates of $\vec{v}_{J,z}$.

	P101	P104	P107	P110
l_2	10^{-9}	10^{-7}	10^{-9}	10^{-5}
z_2	0	10^{-4}	10^{-6}	10^{-2}
$w_{\mathcal{C}}$	2	2	2	2
w_I	5	5	5	5
$\tilde{\chi}^2(\nu = 11)$	16.27	**13.79**	17.71	**12.34**

	P102	P105	P108	P111
l_2	10^{-9}	10^{-7}	10^{-9}	10^{-5}
z_2	0	10^{-4}	10^{-6}	10^{-2}
$w_{\mathcal{C}}$	4	4	4	4
w_I	6.4	6.4	6.4	6.4
$\tilde{\chi}^2(\nu = 11)$	17.23	**13.65**	16.61	**10.00**

Applying a spatio-temporal multi-resolution approach improved fitting of the wave equation motion model of $\vec{v}_{J,z}$ (Table 5.13) and $\vec{D}_z$ (Table 5.14). In both cases it was optimal to apply no weight to the model using the l_2 weighting factor in isolation, however, it also found that χ^2 was very sensitive to the magnitude of l_2. It was possible to use a wider range of weights when l_2 and z_2 were used in combination, thus this form of regularised registration may be less sensitive to the chosen tuning parameters.
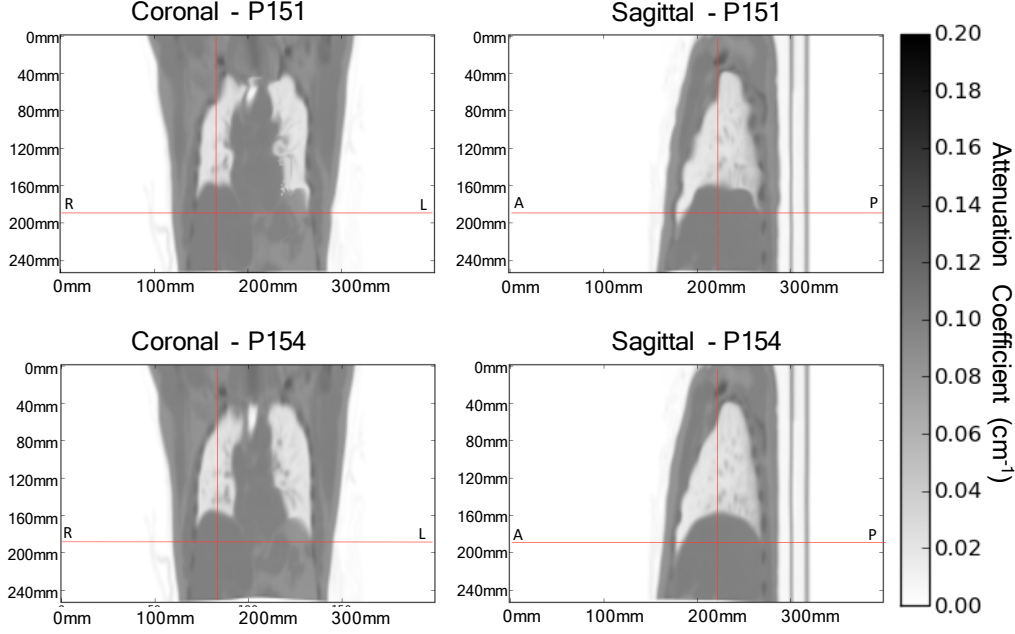
Figure 5.17: Coronal and sagittal slices of the attenuation map after transforming frame 3 of patient 21. The transforms were estimated with spatiotemporal multi-resolution registration. Each column compares fitting the wave equation model using parameter combination P154 (bottom) to no fitting using parameter combination P151 (top). The position of the dome of the liver was tracked and is marked indicated by the red crosshairs.

Table 5.13: Results of fitting the model to the velocity Jacobean, $\vec{v}_{J,z}$. Parameter combinations using l_2 and z_2 weighting factors in combination with spatial smoothing and spatiotemporal multiresolution. The bottom row shows the median $\tilde{\chi}^2$ fit of the wave equation model to 12 time point estimates of the velocity transform. The median is calculated over all 11 patients in the study.

	P151	P154	P157	P160
l_2	10^{-4}	10^{-5}	10^{-5}	10^{-6}
z_2	0	10^{-6}	10^{-7}	10^{-7}
$w_{\mathcal{C}}$	0	0	0	0
w_I	6.4	6.4	6.4	6.4
$\tilde{\chi}^2(\nu = 11)$	**6.55**	**10.21**	**9.06**	**8.04**

Although χ^2 was lowest when setting weighting for the wave equation model to zero (P151 $z_2 = 0$ in table 5.13), a visual inspection of the transformed attenuation map revealed discontinuities near the lung boundaries for this case (Figure 5.17).

Applying a spatiotemporal multi-resolution approach reduced noise when fitting larger transforms, thus it became viable to consider every frame as an independent transform from the reference frame. Fitting the wave equation model to $\vec{D}_z$ (Table 5.14) resulted in a higher median $\tilde{\chi}^2$ than $\vec{v}_{J,z}$ (Table 5.13).
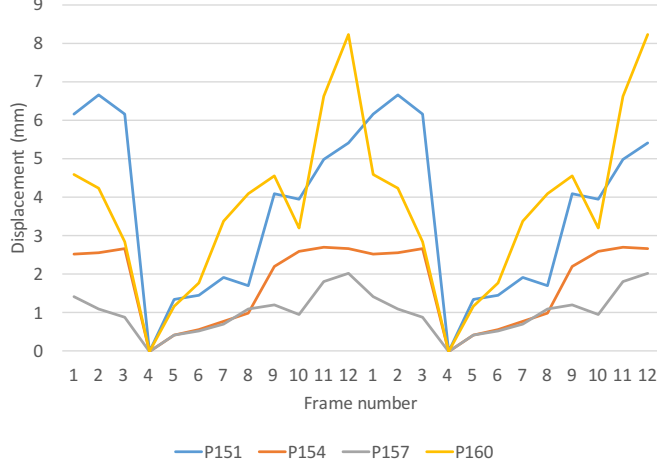
Figure 5.18: Displacement of the right dome of the liver using a velocity approximation for patient 21. Displacements are relative to the position at the reference frame (frame 4). The trace has been repeated to visualise its cyclic features. P151 has no wave equation constraint. Each colour represents a varying level of constraint P154 (Orange), P157 (Gray) and P160 (Yellow)

Table 5.14: Results of fitting the model to the displacement transform, $\vec{D}_z$. Parameter combinations using l_2 and z_2 weighting factors in combination with spatial smoothing and spatiotemporal multi-resolution. The bottom row shows the median $\tilde{\chi}^2$ fit to the wave equation model. The median is calculated over all 11 patients in the study.

	P152	P155	P158	P161
l_2	10^{-5}	10^{-5}	10^{-5}	10^{-6}
z_2	0	10^{-6}	10^{-7}	10^{-7}
$w_{\mathcal{C}}$	0	0	0	0
w_I	6.4	6.4	6.4	6.4
$\tilde{\chi}^2(\nu = 11)$	**9.36**	17.78	**13.50**	**11.75**

A comparison of fitting the model to the velocity displacement (Table 5.13) versus fitting to the displacement (Table 5.14) was achieved by recording the magnitude of estimated displacements at the dome of the liver. It was found that increasing z_2 or l_2 decreased the amplitude of motion, although the amplitude of motion increased with less regularisation (P151 and P160). It was suspected that displacements which were smoothly varying and cyclic over the respiratory period may be preferable (P154 and P157). The large displacements at the end of the cycle (frame 3 and 15 in figures 5.18 and 5.19) were due to the cumulation of errors in successive transforms after the reference frame. Figure 5.18 shows that the effect of cumulative errors was minimised by determining each transform independently as a displacement from the reference frame. It was suspected that the apparent lack of smoothness over time is due to the poor statistics of a 12 frame gated study. Patient 21 was chosen in figures 5.18, 5.19 and 5.20 because χ^2 was close to the median $\tilde{\chi}^2$ for all patients.

Applying a cyclic constraint in combination with fitting the model to the velocity jacobian,
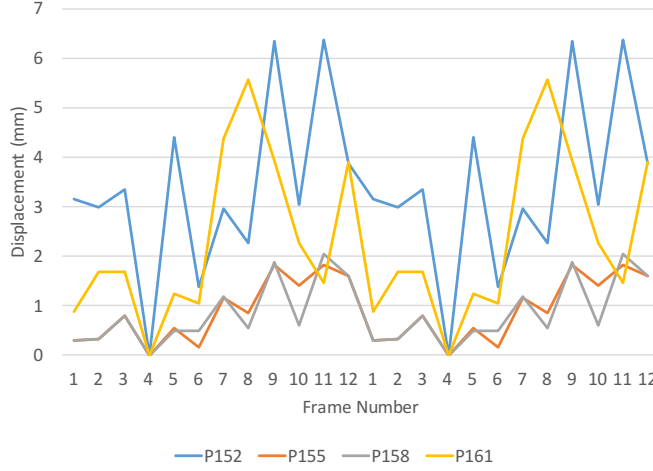
Figure 5.19: Displacement of the right dome of the liver using independent registration for patient 21. Displacements are relative to the position at the reference frame (frame 4). There are 12 unique time points, however, the trace has been repeated over 2 cycles. P152 has no wave equation constraint. Each colour represents a varying level of constraint P155 (Orange), P158 (Gray) and P161 (Yellow)

$\vec{v}_{J,z}$ led to a lower median $\tilde{\chi}^2$ (Table 5.15). The lower median $\tilde{\chi}^2$ was expected because many cyclic functions are also solutions to the wave equation. The magnitude of estimated displacements at the dome of the liver were measured to observe the effect of the cyclic constraint in figure 5.20 in comparison to no cyclic constraint in figure 5.18. Applying a cyclic constraint led to displacements that were temporally smoother and more symmetrical about the reference frame.

Table 5.15: Results of fitting the model to the velocity Jacobian, $\vec{v}_{J,z}$ with a cyclic constraint. Parameter combinations using l_2 and z_2 weighting factors in combination with spatial smoothing and spatio-temporal multiresolution and a cyclic constraint. The bottom row shows the median $\tilde{\chi}^2$ fit of the wave equation model to 12 time point estimates of the velocity transform. The median is calculated over all 11 patients in the study.

	P164	P167	P170
l_2	10^{-5}	10^{-5}	10^{-6}
z_2	10^{-6}	10^{-7}	10^{-7}
w	6.4	6.4	6.4
$\tilde{\chi}^2(\nu = 11)$	**6.94**	**6.85**	**7.41**

The transforms were visually assessed using an image of the magnitude of the displacements shown in figure 5.21. The images showed that less regularisation using parameter combinations P151, P152, P160 and P170 resulted in noticeable discontinuities in the deformation field.
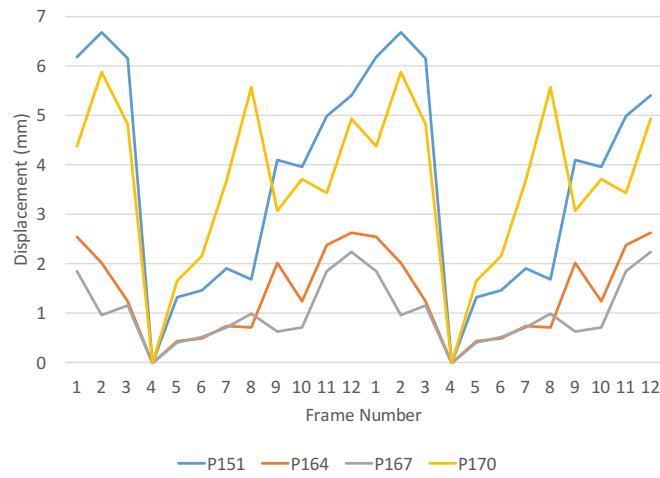
128

Figure 5.20: Displacement of the right dome of the liver using a cyclic constraint for patient 21. Displacements are relative to the position at the reference frame (frame 4). There are 12 unique time points, however, the trace has been repeated over 2 cycles. P151 has no wave equation constraint. Each colour represents a varying level of constraint P164 (Orange), P167 (Grey) and P170 (Yellow).
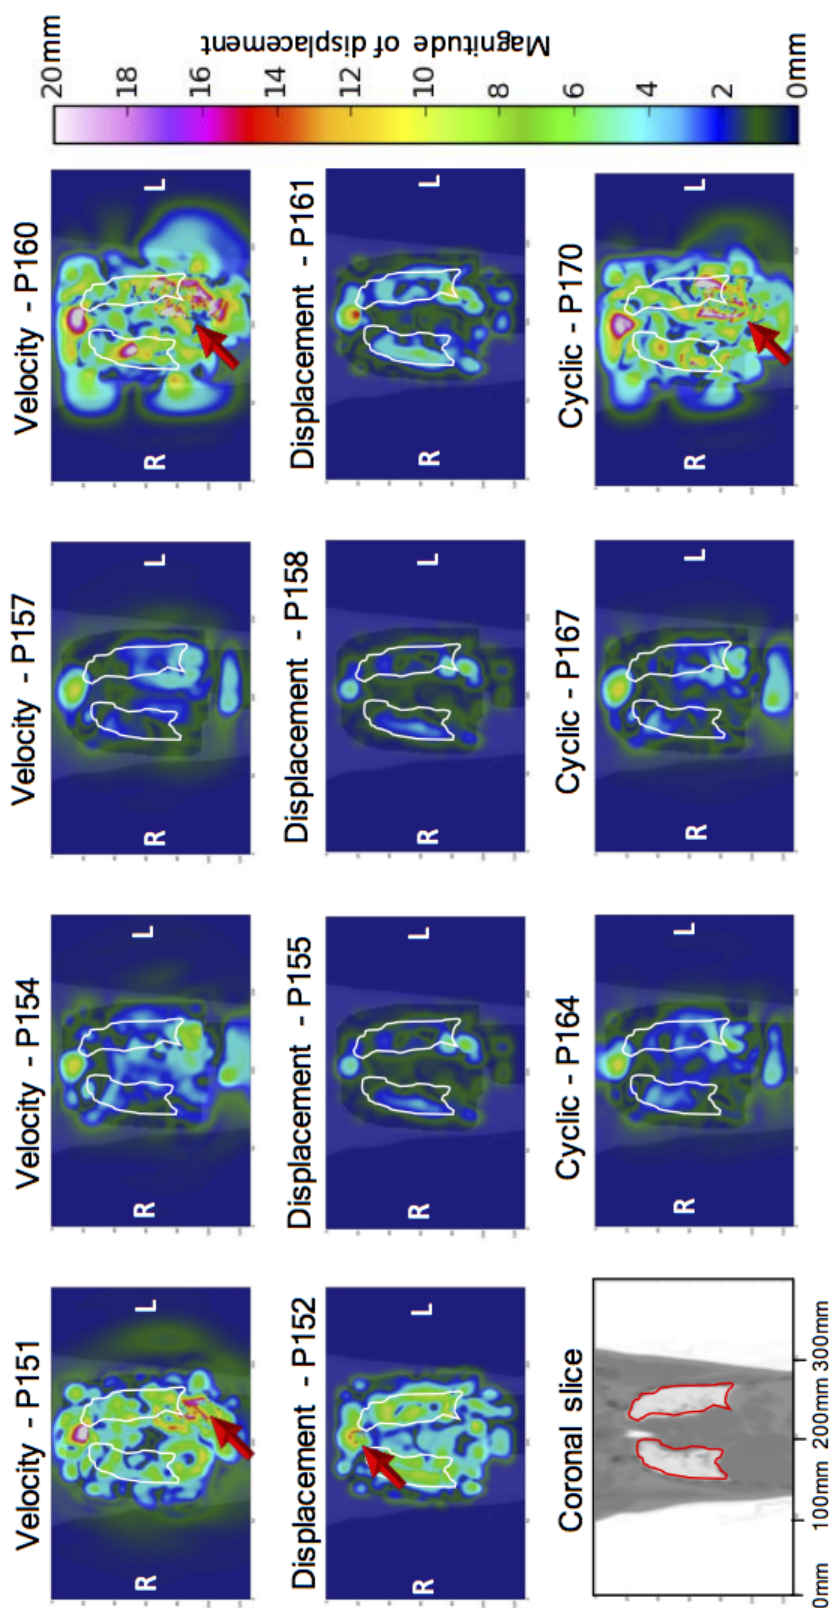
Figure 5.21: Coronal slices of an image where each voxel is coloured with the magnitude of displacement for frame 3 in patient 21. The first row shows coronal slices when using a velocity model (P151, P154, P157, P160). The second row shows coronal slices when using the displacement model (P152, P155, P158, P161). The final row shows coronal slices using a cyclic model (P164, P167, P170). A coronal slice of the attenuation map is shown in the bottom left for comparison. The displacements are relative to the position at the reference frame (frame 4). The outline of the lung in frame 4 has been superimposed on the displacement images. Discontinuities in the deformation field are marked with a red arrow.

5.3.4 Attenuation Correction for Respiratory Gated PET

Compensating for misaligned attenuation correction improved the consistency of the attenuation map with the emission image in comparison to CTAC alone. The consistency was quantified using the SSD metric between emission and reprojected sinograms for each frame in the respiratory gated PET. Following compensation there was no change in SSD for the reference frame, however, the SSD metric of other frames in the respiratory gated PET was decreased. The reason for the decrease was solely attributed to the deformation of the attenuation map in accordance with the respiratory motion model. Respiratory motion models used to deform the attenuation map and were considered based on when the SSD was less than that for CTAC alone. Figure 5.22 shows a reduction in the SSD for models P154, P157, P164 and P167 in patient 21 indicating that these are useful models for the deformation of the attenuation map.

A systematic analysis of the SSD metric for all patients and all frames was used to quantify the effect of 12 different motion models on the SSD. Figure 5.23 shows an overall decrease (ΔSSD) in the SSD when using various motion models versus CTAC alone. The SSD for some models had a significant number of outliers from the median of the population and a visual inspection of the SSD revealed sharp increases corresponding to inconsistencies in the reconstruction. The respiratory motion models which had significant outliers in the SSD were excluded on the basis that these models did not reliably provide an improved method of attenuation correction.

As shown in table 5.16, the median ΔSSD for all patients and all frames was lowest when using both the velocity model and the cyclic model of the wave equation. Using $l_2 = 10^{-5}$ in models P154, P157, P164 and P167 produced the lowest median and the least number of outliers.

The activity concentration (uptake) of suspected tumour regions was highly variable across all 12 frames and all patients. The standard deviation of uptake in the tumour region was between 10% and 20% depending on the respiratory motion model applied. The TV region encompassed the full motion of the tumour and had a smaller standard deviation of between 8.2% and 10.9% depending on the respiratory motion model. The large standard deviation was attributed to the many factors, including attenuation, which influence the measured activity concentration. CV_{act} was expressed as a fraction of mean uptake in the TV region because it was less affected by tumour size. Table 5.16 shows which respiratory motion models produced a smaller variation in uptake (CV_{act}).

Considering that lung tumours are typically less compressible than surrounding tissue (Kyriakou et al., 2007), there were far greater changes in tumour size than expected. The standard deviation in region size, CV_{size}, was greater than 20% for models P151, P152, P160, P161 and P170 and shown in figure 5.24. Increasing the penalty of z_2 to 10^{-5} decreased the standard deviation in region size for models P154, P155 and P160 in table 5.16. Figure 5.25 shows a tumour region which changed size from 13.1mL in frame 10 to 32.0mL in frame 12 when using no wave equation constraint (P151), however, this variation was reduced when using a wave equation constraint (P154). Figure 5.26 compared the effect of including the wave equation constraint in the motion model for all frames in patient 10. The wave equation model is based on a physical assumption that tissue cannot compress without a corresponding flow of air to allow that compression to occur. The wave equation model constraint generally
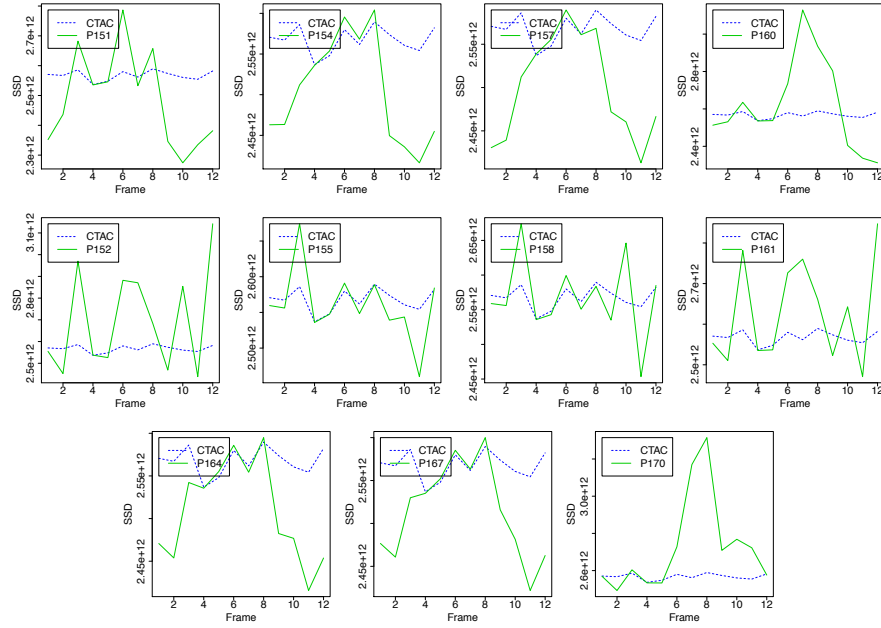
131

Figure 5.22: Sum of squared differences (SSD) metric for consistency in gated PET reconstruction for patient number 21. Each plot shows the total SSD between projection and emission sinogram after attenuation correction has been applied. The SSD has been calculated for all 12 frames in the respiratory cycle. Each model of attenuation deformation (P151 - P170, green) has been compared to spiral CT attenuation correction (CTAC, blue - dotted). The first row shows consistency when using a velocity model (P151, P154, P157, P160). The second row shows consistency when using the displacement model (P152, P155, P158, P161). The final row shows consistency when using a cyclic model (P164, P167, P170).
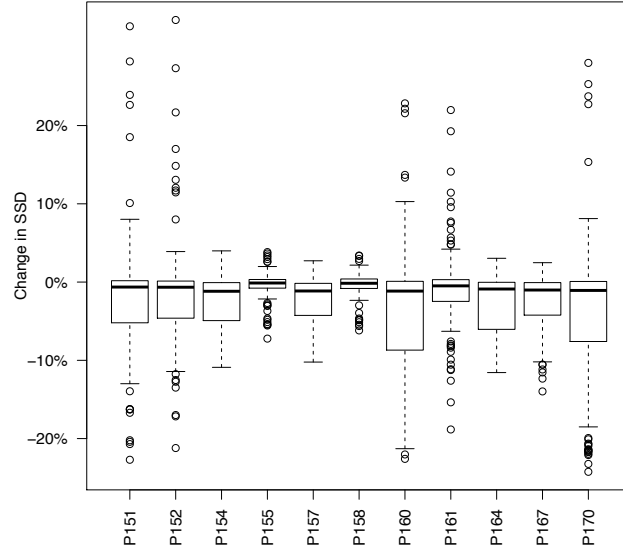
Figure 5.23: A box and whisker plot of changes in sum of squared differences (ΔSSD) metric for consistency in gated PET reconstruction for every patient. Changes have been expressed as a percentage difference in total SSD when using a model (P151 - P170) versus using CTAC alone. A negative change in SSD corresponds to a more consistent reconstruction. The bounding box contains 50% of the values and the outliers have been individually plotted as circles.
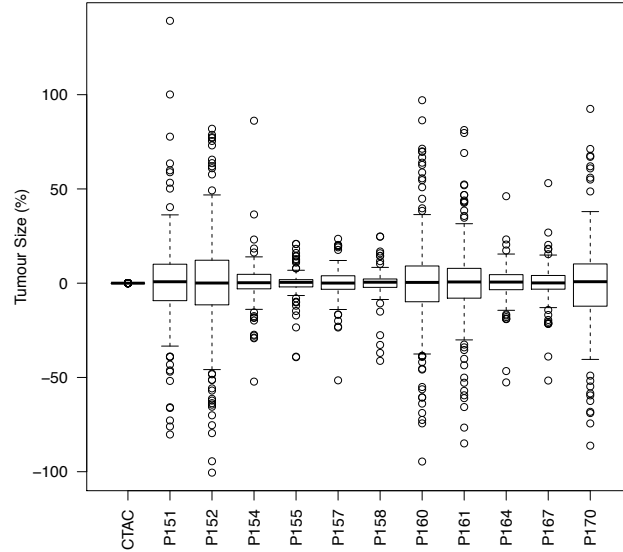

Figure 5.24: A box and whisker plot of the tracking region size v_t. The region size in each gated PET has been plotted for every frame and for every patient. Changes have been expressed as a percentage of the mean (CV_{size}). The bounding box contains 50% of the values and the outliers have been individually plotted as circles.

Table 5.16: A comparison of metrics to evaluate the use of each model (P151 - P170) versus using CTAC alone. The first row is the median ΔSSD and contains the median values of changes in sum of squared differences metric for consistency in gated PET reconstruction from figure 5.23 where changes in SSD have been expressed as a percentage difference in total SSD. The second row is the variance in the TV activity CV_{act} and is expressed as the standard deviation of the activity concentration in the TV region as a percentage of the mean uptake in the TV region. The third row is the variance in the region size CV_{size} and is expressed as the standard deviation of the size of the tracking region as a percentage of the mean size of the tracking region.

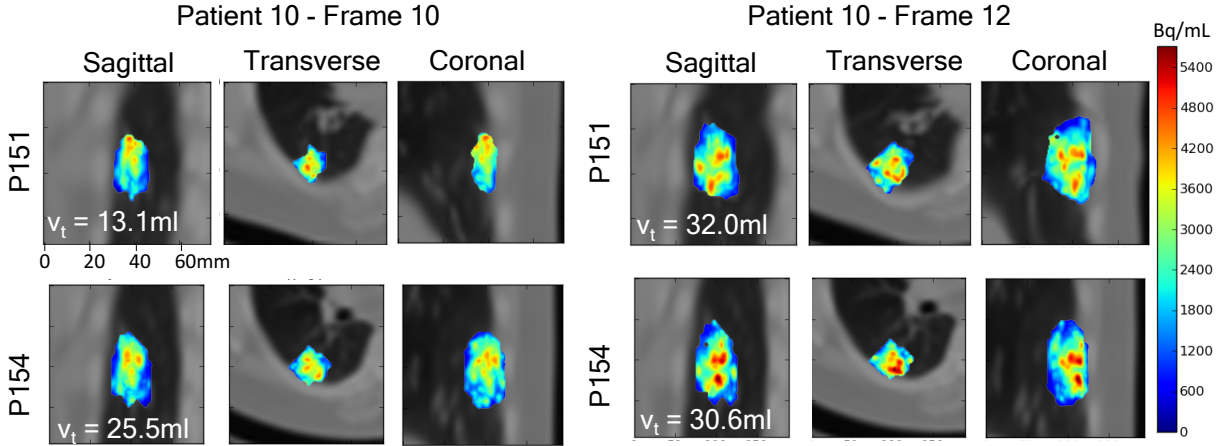| Velocity Model	P151	P154	P157	P160
Median ΔSSD	-0.6%	-1.2%	-1.1%	-1.2%
CV_{act}	8.7%	8.4%	8.5%	9.4%
CV_{size}	29.2%	12.7%	9.2%	31.8%
Displacement Model	P152	P155	P158	P161
Median ΔSSD	-0.7%	-0.1%	-0.2%	-0.5%
CV_{act}	8.4%	9.0 %	9.0%	8.4%
CV_{size}	34.9%	7.8%	8.3%	26.8%
Cyclic Model		P164	P167	P170
Median ΔSSD		-0.9%	-1.0%	-1.1%
CV_{act}		8.2%	8.4%	10.9%
CV_{size}		10.8%	10.9%	28.5%
CTAC only	$CV_{act} = 10.3\%$			

Figure 5.25: A maximum intensity projection of the activity in the tracking region volume, v_t, superimposed on the attenuation correction map. The attenuation correction map has been sliced in the sagittal, transverse and coronal directions at the centre of mass of the tumour. The first two rows are the 10th frame from Patient 10, and the second two rows are the 12th frame from patient 10. The volume of the tumour is approximately 30mls in the 12th frame, but decreases significantly (P151, first row), if not using the wave equation model. The wave equation velocity model preserves the size of the tumour (P154, second row).

constrained the expansion and compression of tumour regions for the cyclic and velocity model.

The standard deviation of the TV region uptake (CV_{act}) was between 8.2% and 10.9% and was slightly decreased to about 8.4% when we applied the velocity model of the wave equation. The decrease in the standard deviation of the TV region uptake was not significant. Models P154, P157, P164 and P167 had the lowest median ΔSSD and also had a CV_{act} between 8.2% and 8.5%. It was suspected that a contributor to both the decreased consistency metric (ΔSSD) and the reduced variation (CV_{act}) was due to improved modelling of motion and hence a more accurate attenuation map.

5.4   Discussion

The initial analysis of candidate models revealed that some models were more suited to estimating spatial and temporal derivatives of the displacement. Many regularisers are expressed as the first or second derivative of the displacement in order to enforce temporal and spatial smoothness of the deformation field (Holden, 2008). The results showed that the model which was most suited to predicting the velocity of the Jacobean was the wave equation model. The results identified tuning parameters of the wave equation model for which regularising the velocity of the Jacobean ensured that the deformation was both spatially and temporally continuous but also consistent with the physiological assumptions of the wave equation model.

The implementation of the regulariser was an adaptive approach, penalising deformations
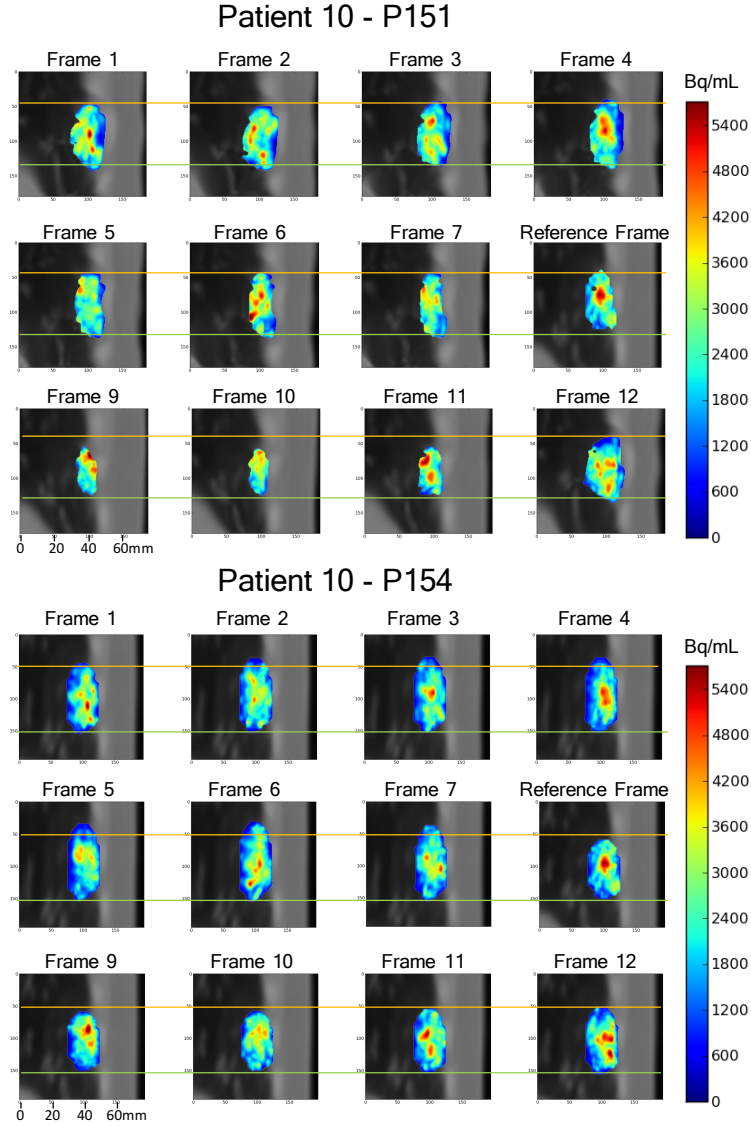
Figure 5.26: A maximum intensity projection of the activity in the tumour region super-imposed on the the attenuation correction map. The attenuation correction map has been sliced in the coronal direction at the centre of mass of the tumour. The 3 rows on the left are 12 frames from patient 10 when using no wave equation constraint (P151). The 3 rows on the right are 12 frames from patient 10 when using the wave equation constraint (P154).

in a voxelwise manner. The results demonstrated an advantage over conventional regularisation which treats all point displacements equally by either smoothing or penalising displacement. The novel regulariser was also tested in combination with smoothing, however, it was not clear whether smoothing displacements with a Gaussian of wider than 6.4mm was beneficial. Alternatives to smoothing have been proposed such as using a Markov model solver (Bai & Brady, 2009) and demons image registration (Vercauteren, Pennec, Perchant, & Ayache, 2009). Future work could involve investigating whether smoothing could be avoided by using a solver for image registration other than the conjugate gradient descent.

This novel approach is mainly beneficial for imaging modalities similar to PET with a low signal to noise. The results from CT imaging showed that if lung boundaries or features were well defined then the benefits of wave equation regularisation were reduced. In such cases it might be worthwhile investigating other forms of image registration as well an assessment of the accuracy of registration of fine structures within the lungs.

The potential benefits of a cyclic constraint in addition to the wave equation model were investigated. Initially, a cyclic constraint was avoided by allowing for non-periodic expansion, although this is not physiologically possible in a gated study such expansion may have been observed due to inaccuracies in respiratory gating or acquisition because the complete respiratory cycle may not have been fully captured. The results indicated that this was not a concern as including a cyclic constraint had little effect.

The investigation of tumour regions revealed that the tuning parameters for the model had a significant effect on the expansion and contraction of the region over the respiratory cycle. It was found that a larger penalty term reduced the variance in region size. Tumours are much less compressible than normal lung tissue (Kyriakou et al., 2007), so it was expected that realistic models would limit the expansion and contraction of tumour regions. Another limitation of this investigation was that it did not reveal local changes, such as changes in activity distribution within the tumour.

The chi squared statistic was used to identify the most suitable model for the physics of respiration. Although the wave model had some similarities to finite element modelling (FEM) a key difference was that this method did not require the detection of lung boundaries or features (Samavati et al., 2015). Further work on developing a FEM model which is less reliant on the accuracy of lung boundaries or features may be of interest.

The wave equation is an analytical approach to respiratory motion modelling. The analysis of candidate models leaves scope for future work for using the same models in contexts other than image registration alone, for example, jointly estimating deformation and emission. Although 1 dimensional modelling is clearly useful, extending this approach to 3 dimensions should lead to improved fitting because of the additional degrees of freedom. A 3D model may also be less prone to discontinuities in the deformation field.

This work has demonstrated that a velocity model of the wave equation was optimal for addressing attenuation mismatch. The consistency was estimated by measuring the SSD between emission data and reprojected sinograms. This estimate of the consistency conditions revealed whether the gated PET frames and attenuation map may be misaligned. It was also found that a velocity model of the wave equation was more robust, with reduced variations in results across patients. The population variance in tumour region size, TV uptake and SSD consistency were all reduced when using a velocity model of the wave equation. It was also found that Chi Squared can be used to provide online feedback on performance of motion

model fitting and could be used to predict whether the model was robust.

References

Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *NeuroImage*, *38*(1), 95–113. doi: 10.1016/j.neuroimage.2007.07.007

Bai, W., & Brady, M. (2009). Regularized b-spline deformable registration for respiratory motion correction in PET images. *Physics in Medicine and Biology*, *54*(9), 2719. doi: 10.1088/0031-9155/54/9/008

George, R., Chung, T. D., Vedam, S. S., Ramakrishnan, V., Mohan, R., Weiss, E., & Keall, P. J. (2006). Audio-visual biofeedback for respiratory-gated radiotherapy: Impact of audio instruction and audio-visual biofeedback on respiratory-gated radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, *65*(3), 924–933. doi: 10.1016/j.ijrobp.2006.02.035

Holden, M. (2008). A review of geometric transformations for nonrigid body registration. *IEEE Transactions on Medical Imaging*, *27*(1), 111–128. doi: 10.1109/TMI.2007.904691

Kinahan, P. E., Townsend, D. W., Beyer, T., & Sashin, D. (1998). Attenuation correction for a combined 3d PET/CT scanner. *Medical physics*, *25*(10), 2046–2053.

Kyriakou, E., McKenzie, D. R., Suchowerska, N., & Fulton, R. R. (2007). Breathing as a low frequency wave propagation in nonlinear elastic permeable medium. *Physica B: Condensed Matter*, *394*(2), 311–314. doi: 10.1016/j.physb.2006.12.056

Liu, C., Alessio, A. M., & Kinahan, P. E. (2011). Respiratory motion correction for quantitative PET/CT using all detected events with internal-external motion correlation. *Medical Physics*, *38*(5), 2715–2723. doi: 10.1118/1.3582692

Mattes, D., Haynor, D. R., Vesselle, H., Lewellen, T. K., & Eubank, W. (2003). PET-CT image registration in the chest using free-form deformations. *IEEE Transactions on Medical Imaging*, *22*(1), 120–128. doi: 10.1109/TMI.2003.809072

McClelland, J. R., Hawkes, D. J., Schaeffter, T., & King, A. P. (2013). Respiratory motion models: A review. *Medical Image Analysis*, *17*(1), 19–42. doi: 10.1016/j.media.2012.09.005

Metz, C. T., Klein, S., Schaap, M., van Walsum, T., & Niessen, W. J. (2011). Non-rigid registration of dynamic medical imaging data using nD + t b-splines and a groupwise optimization approach. *Medical Image Analysis*, *15*(2), 238–249. doi: 10.1016/j.media.2010.10.003

Modat, M., Ridgway, G. R., Taylor, Z. A., Lehmann, M., Barnes, J., Hawkes, D. J., . . . Ourselin, S. (2010). Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine*, *98*(3), 278–284. doi: 10.1016/j.cmpb.2009.09.002

Murphy, K. (2011). Evaluation of registration methods on thoracic CT: The EMPIRE10 challenge. *IEEE Transactions on Medical Imaging*, *30*(11), 1901–1920. doi: 10.1109/TMI.2011.2158349

Nikolic, D., Murean, R. C., Feng, W., & Singer, W. (2012). Scaled correlation analysis: a better way to compute a crosscorrelogram. *European Journal of Neuroscience*, *35*(5),

742–762.

Pluim, J. P., Maintz, J. A., & Viergever, M. A. (2003). Mutual-information-based registration of medical images: a survey. *IEEE transactions on medical imaging*, *22*(8), 986–1004.

Ponisch, F., Richter, C., Just, U., & Enghardt, W. (2008). Attenuation correction of four dimensional (4d) PET using phase-correlated 4d-computed tomography. *Physics in Medicine and Biology*, *53*(13), N259. doi: 10.1088/0031-9155/53/13/N03

Ruan, D. (2009). Directionally selective regularization for sliding preserving medical image registration. In *2009 IEEE nuclear science symposium conference record (NSS/MIC)* (pp. 2936–2939). doi: 10.1109/NSSMIC.2009.5401607

Samavati, N., Velec, M., & Brock, K. (2015). A hybrid biomechanical intensity based deformable image registration of lung 4dct. *Physics in Medicine and Biology*, *60*(8), 3359. doi: 10.1088/0031-9155/60/8/3359

Sundaram, T. A., & Gee, J. C. (2005). Towards a model of lung biomechanics: pulmonary kinematics via registration of serial lung images. *ITKOpen science - combining open data and open source software: Medical image analysis with the Insight Toolkit*, *9*(6), 524–537. doi: 10.1016/j.media.2005.04.002

Thielemans, K., Asma, E., & Manjeshwar, R. M. (2009). Mass-preserving image registration using free-form deformation fields. In *2009 IEEE nuclear science symposium conference record (NSS/MIC)* (pp. 2490–2495). doi: 10.1109/NSSMIC.2009.5402070

Vercauteren, T., Pennec, X., Perchant, A., & Ayache, N. (2009). Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*, *45*(1), S61–S72. doi: 10.1016/j.neuroimage.2008.10.040

Werner, R., Ehrhardt, J., Schmidt, R., & Handels, H. (2008). Modeling respiratory lung motion: a biophysical approach using finite element methods. In *Proc. spie 6916, medical imaging 2008* (Vol. 6916, pp. 69160N–69160N–11). doi: 10.1117/12.769155

Zitov, B., & Flusser, J. (2003). Image registration methods: a survey. *Image and Vision Computing*, *21*(11), 977–1000. doi: 10.1016/S0262-8856(03)00137-9

Chapter 6

Conclusion

The purpose of this thesis was to develop methods to compensate for misalignment of the attenuation map and the emission image in PET/CT. The validation of methods for compensation were benchmarked against CT-derived attenuation. Although there are many ways to improve gated-PET image reconstruction, the literature review in the background chapter of this thesis provides evidence that the most significant image artefacts in respiratory gated PET arise from attenuation correction mismatch.

Despite our efforts to implement previously published approaches as outlined in chapter 3 there was varying accuracy between different methods. Whilst this was concerning, it was not unexpected given the variety of approaches to image registration. This varying accuracy highlighted the need for feedback on the accuracy of image registration and guided decisions on how to best model respiratory motion. For example, if feedback indicates that a respiratory motion model lacks precision then the assumptions made by the model about respiratory motion can be tuned to produce an improved result.

6.1 Compensating for Attenuation Correction Misalignment

Multiple approaches to compensating for attenuation correction misalignment using image registration have been presented in chapters 3, 4 and 5. Since much anatomy is not visible on FDG PET, methods have been developed which are tailored to the accurate registration of images in regions where organs have minimal FDG uptake. This has been accomplished through improvements to image registration using statistical and physiological models as presented in chapters 4 and 5. These models either favour plausible deformations during the process of image registration or infer the most plausible motion in regions where image registration could not derive an accurate deformation. It would be ideal to develop a method which is not reliant on image features or assumptions about respiratory motion, however, chapters 4 and 5 showed this was necessary in order to maintain accuracy. A limitation of this study is that non-attenuated gated images have been used to estimate motion and while recent investigates have shown that this limitation can be addressed using joint reconstruction and registration (Rezaei, Michel, Casey, & Nuyts, 2016), the approach taken in this thesis was to consider practical approaches to compensate for misalignment that do not require significant modifications to the PET image reconstruction.

In chapter 3, multiple different image registration approaches to aligning the emission and attenuation image were taken. The significant variation in results for each approach was taken as an indication that the accuracy of attenuation correction is highly dependent on the accuracy of registration. For example, the results in chapter 3 showed that multimodality registration is more accurate than a mono-modality approach, or even a combination of both of these approaches. This suggested that there could be a limitation with the PET only approach that cannot be overcome. In some cases, poor accuracy may be an indication of a problem with the implementation of registration rather than the approach taken. For example, a spatiotemporal approach to mono-modality registration performed the worst out of all approaches, yet, previous investigations suggested that this method should be robust to noise (Metz, Klein, Schaap, van Walsum, & Niessen, 2011). This may have been an indication that results could be improved if the spatiotemporal approach to registration was

investigated with more rigour.

In chapter 5 the robustness of the mono-modality approach was improved by incorporating a wave equation model into registration. Whilst the results were promising this approach was only implemented by extending the niftyreg registration package. Some issues could be specific to this package, particularly considering niftyreg registration performed poorly as described in chapter 3. The limitation of using a single software package (niftyreg) as described in chapter 5 could be overcome by using other packages such as elastix.

A cross-population approach to modelling respiratory motion was investigated in chapter 4. This method extended mono-modality registration as a means of fitting the population model to the patient image. Cross-population models generally need a large cohort of patients (more than 100) to be accurate, however, our results with 10 patients support findings that a large cohort may not be required to model respiratory motion (Fayad, Buerger, Tsoumpas, Cheze-Le-Rest, & Visvikis, 2012). Our approach to modelling respiratory motion was similar to the statistical motion modelling introduced by Li et al. (2011) which was based on the dimensionality reduction of displacement vector fields. This approach has since been improved upon using manifold alignment strategies (Baumgartner et al., 2014), but the effect on attenuation mismatch has not been investigated. A limitation of our approach is that it does not consider other methods of measuring respiratory motion. Measurements of the chest surface have been shown to be a useful surrogate which can be incorporated into the population model.

6.2 Further Work on Compensating for Misalignment

The aim of this thesis has been to compensate for misalignment, however, the methods presented in chapters 3, 4 and 5 are similar to methods used for the motion compensation of emission images. In addition to compensating for misalignment this method could also be extended the correction of emission images.

Further work on the image registration methods introduced in chapter 3 is required in refining the combined mono-modality and multi-modality approach. The combined approach involves the calculation of the inverse deformation and this was found to be dependent on the implementation provided by the image registration package. Further investigation of alternative reconstruction packages may reveal more positive results

Given that inaccuracies in image registration may be specific to the niftyreg package, further work is required to convincingly rule out that a mono-modality approach cannot be improved upon. Some packages such as elastix provide the ability to combine mono-modality registration in conjunction with other improvements such as groupwise registration

The calculation of SSD metrics were useful in evaluating the improvements to image registration. Although this approach has merit, further work on developing online statistics such as the goodness of fit introduced in chapters 4 and 5 may be more practicable in a clinical situation where computational resources and time are limited.

6.3 Evaluation of Attenuation Correction Methods

Compensating for attenuation correction misalignment was evaluated using both Natterer's and SSD consistency. The results in chapter 3 indicated that the SSD was sensitive to the registration inaccuracies such as the misalignment of lung boundaries. In chapter 3, the SSD decreased in comparison to CTAC following the use of multimodality registration to compensate for the misalignment in 11 patients.

A method for identifying the reference frame was developed in chapter 3 using a weighted SSD measure of consistency. The weighted SSD method was the best approach at identifying the emission frame which was most similar to the CT. Unfortunately, there was no direct way to measure the accuracy of this method without having prior knowledge of the timing and position of the spiral CT relative to time intervals used for gating the PET image.

Further work could be to expand the evaluation of misalignment to a novel method for the evaluation of motion compensation. A limitation of the current method is that it is not sensitive to inaccuracies in the deformation field in areas of constant HU. Whilst this had negligible effect on the accuracy of attenuation correction, it may have a significant effect on the accuracy of motion compensation.

Online statistics on the goodness of fit were used in chapter 5 to investigate registration parameter settings, however, it was challenging to compare alternative methods such as modeling, versus multi-modality, versus mono-modality. There were also difficulties evaluating the population model using leave-one-out analysis as it did not reveal the accuracy of correlated attenuation correction. In chapters 4 and 5 various methods were used to estimate the accuracy of deformation fields including leave-one-out analysis, simulations, and goodness of fit. A limitation of these methods was that they were not particularly sensitive to attenuation mismatch at organ boundaries.

The appearance of tumours and the variance of radio-activity in the tumour region were used to evaluate attenuation correction. In some cases compensating for misalignment increased this variance. Although this was unexpected, it is possible that the variance could be introduced from other sources such as fluctuations in motion blur during phase based respiratory gating. . The wave equation model esnforced the rigid motion of tumours and this reduced variations in the apparent size of the tumour as a result of inaccurate image registration. A limitation to the analysis of tumours was that it was not clear how much variance needed to be mitigated.

6.4 Further Work on Evaluation Methods

The evaluation methods presented in this thesis were useful for detecting misalignment, however, they did not indicate the extent of misalignment. This could be addressed by further work to relate the SSD consistency to the physical extent of misalignment or the quantitative uptake of tracer (SUV). For example, it would be feasible to simulate motion in scans and establish a relation between the simulated motion and the SSD consistency values.

Natterer's consistency had a large variation in values across the 11 patients evaluated in chapter 3. This could be explained by the sensitivity to noise arising from high frequencies in the consistency equation (equation 2.5), however, further work might involve investigating

the source of the variance and why the alternative SSD consistency had a lower variation.

The information available in Time-Of-Flight (TOF) PET has been used in recent methods for addressing attenuation correction (Rezaei, Defrise, & Nuyts, 2014). Our approach to evaluating consistency has essentially summed all TOF bins into a single sinogram. The sensitivity of our proposed consistency metrics could be significantly enhanced by using a consistency which incorporates the TOF signal.

References

Baumgartner, C. F., Kolbitsch, C., Balfour, D. R., Marsden, P. K., McClelland, J. R., Rueckert, D., & King, A. P. (2014). High-resolution dynamic MR imaging of the thorax for respiratory motion correction of PET using groupwise manifold alignment. *Medical Image Analysis*, *18*(7), 939–952. doi: 10.1016/j.media.2014.05.010

Fayad, H., Buerger, C., Tsoumpas, C., Cheze-Le-Rest, C., & Visvikis, D. (2012). A generic respiratory motion model based on 4d MRI imaging and 2d image navigators. *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2012 IEEE*, 4058–4061. doi: 10.1109/NSSMIC.2012.6551927

Li, R., Lewis, J. H., Jia, X., Zhao, T., Liu, W., Wuenschel, S., . . . Jiang, S. B. (2011). On a PCA-based lung motion model. *Physics in Medicine and Biology*, *56*(18), 6009–6030. doi: 10.1088/0031-9155/56/18/015

Metz, C. T., Klein, S., Schaap, M., van Walsum, T., & Niessen, W. J. (2011). Non-rigid registration of dynamic medical imaging data using nD + t b-splines and a groupwise optimization approach. *Medical Image Analysis*, *15*(2), 238–249. doi: 10.1016/j.media.2010.10.003

Rezaei, A., Defrise, M., & Nuyts, J. (2014). ML-reconstruction for TOF-PET with simultaneous estimation of the attenuation factors. *IEEE transactions on medical imaging*, *33*(7), 1563–1572.

Rezaei, A., Michel, C., Casey, M. E., & Nuyts, J. (2016). Simultaneous reconstruction of the activity image and registration of the CT image in TOF-PET. *Physics in medicine and biology*, *61*(4), 1852.

Appendix A

Appendix

A.1 Anatomical Phantom Simulations

The following is an excerpt of the results of using human anatomical phantom simulations to validate the 1D Respiratory Motion Model in Barnett R, Meikle S, Fulton R. *Deformable Image Registration By Regarding Respiratory Motion As 1D Wave Propagation In An Elastic Medium.* Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE

These simulations have not been included in the thesis body because the implementation of the 1D motion model in chapter 5 used a form of regularisation which improved upon this method. Thus, these results have been excluded from the discussion and conclusion.

I. METHOD

A. Simulation of respiratory gated PET

A human anatomical phantom, NCAT [13], was used to generate realistic emission images. The images were forward projected into singorams whilst applying poisson noise and attenuation modeling [14]. Six emission images were produced throughout a normal respiratory cycle of six seconds, thus each frame had a temporal resolution of one second. The sinograms were reconstructed using an implementation of OSEM [15] without attenuation correction (NAC). The performance of various image registration algorithms were tested on the NAC images to establish the accuracy of the registration with respect to the ground truth values used for the simulation. We chose not to simulate a lesion in the lung because we wished to test the performance of the image registration with limited distinguishable features when the image similarity metric is extremely noisy.

II. RESULTS

A. Validation using simulated respiratory gated PET

The NCAT simulation synthesises a nearly constant tracer value in the lungs. Fig 2. demonstrates the absence of visible features in the lung.
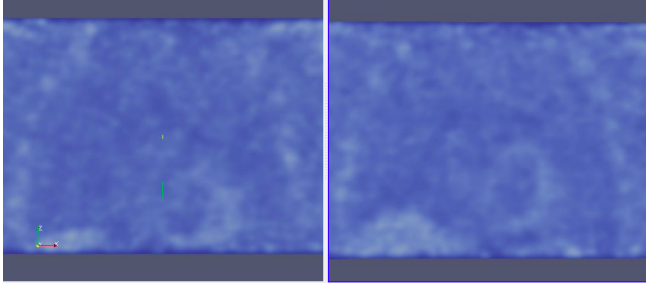

Fig. 2. Coronal cross sections from noisy PET simulated data. The heart and diaphragm are barely visible. The source frame (left) is to be registered to the target frame (right). These are frames 1 and 4 in a 6 frame simulated gated data set.
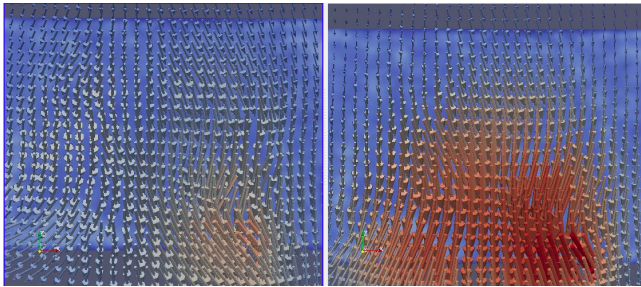

Fig. 3. A coronal slice of the deformation field from the image registration overlaid onto the target PET frame. The deformation field is represented by arrows of varying color. Red represents large deformations. Blue and gray represent smaller deformations. The field produced using free-form deformation (left). The field produced after applying wave equation constraints (right). In both cases the bending energy penalty term was 1% .

We applied conventional free-form registration and found that the presence of noise in the PET image caused image registration to randomly fail. In figure 3 the deformation field for the entire left lung is grossly underestimated due to noise. Applying the wave equation constraint produced increased superior-inferior deformation without requiring a change in the penalty term for bending energy. There is a good visual match between the registered frame and the target frame as a result of the increased deformation.

The amplitude of motion can be recovered by decreasing the bending energy from 1% to 0.1%. The problem with decreasing the bending energy penalty term is that sharp local deformations are favoured.
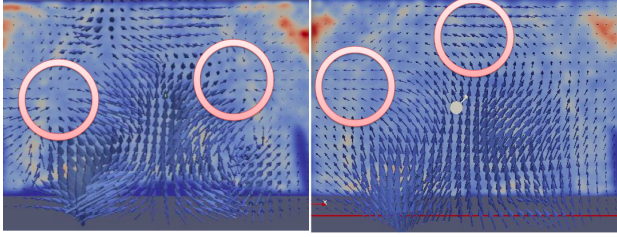
Fig 4. A coronal slice of the deformation field from the image registration using a lower bending energy (be) penalty term. Setting the be=0.1% (left) and be=0.1% with smoothing applied (right). The pink circles indicate sudden changes in the direction of motion to what is expected.

In figure 4 the bending energy has been reduced to 0.1% and we observed sudden changes in the direction of motion in comparison to what was expected. A practical solution was to apply regularisation such as smoothing the gradient of the objective function. Figure 4 (right) demonstrates that smoothing had a dampening effect on the deformation field but it did not completely remove the undesired motion.

The ground truth deformation field was obtained from the motion parameters of the NCAT simulation. We measured the amplitude of motion by summing the absolute value of all superior-inferior motion, thereby creating a curve which is analogous to a respiratory motion trace. We compared the amplitude of the deformation field created from free form image registration with the amplitude in the ground truth results. Figure 5 shows a comparison of the amplitudes.
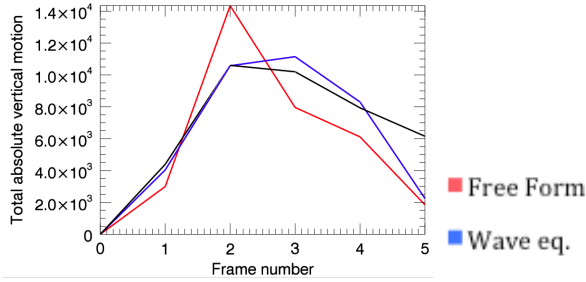

Fig. 5. The sum of the superior-inferior deformation for each frame. Ground truth superior-inferior motion (black). Free form deformation (red) with wave equation constraints (blue)

We measured the per-voxel difference between the deformation field and the ground-truth for each frame. The per-voxel difference followed a predictable Gaussian distribution as shown in Figure 6. We then tabulated the shift of this distribution from zero.
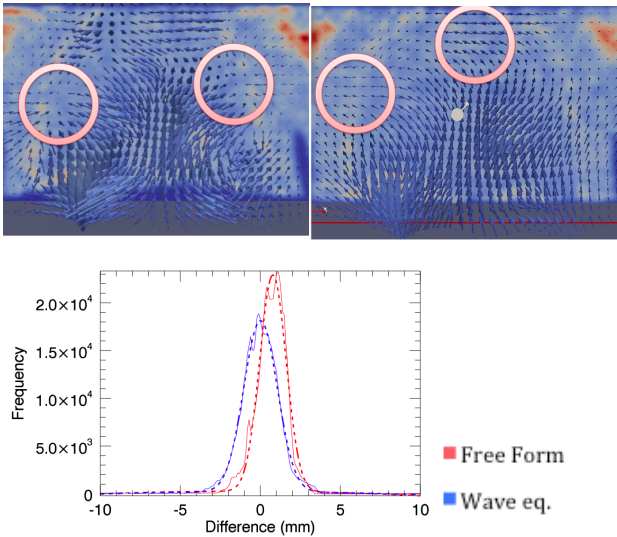

Fig. 6. The difference between free-form deformation (red) and ground truth. The difference between free form deformation using wave equation constraints (blue) and ground truth.

Motion between frames	Free form		Wave equation	
	Difference (mm)	FWHM (mm)	Difference (mm)	FWHM (mm)
1 -> 2	1.6	3.9	0.0	4.9
2 -> 3	-2.2	11.1	3.6	7.7
3 -> 4	**8.5**	11.0	-0.1	4.7
4 -> 5	-3.7	5.2	2.3	6.6
5 -> 6	**-7.1**	4.7	-0.6	5.4

Table I presents a summary of the differences from ground truth. The large differences have been highlighted in bold to demonstrate where the image registration is inaccurate. In frame 4 and frame 6 the deformation field for the entire left lung is grossly underestimated unless wave equation constraints are applied.

[13] B. M. Tsui, D. S. Lalush, E. C. Frey, M. A. King, D. Manocha, and W. Segars, "Development and application of the new dynamic Nurbs-based Cardiac-Torso (NCAT) phantom.," *School of Nursing Faculty Publication Series*, vol. 42, no. 5, Jan. 2001.

[14] C. Tsoumpas et al., "Fast generation of 4D PET-MR data from real dynamic MR acquisitions," *Physics in Medicine and Biology*, vol. 56, pp. 6597-6613, Oct. 2011.

[15] K. Thielemans, S. Mustafovic, and C. Tsoumpas, "STIR: software for tomographic image reconstruction release 2," in *Nuclear Science Symposium Conference Record, 2006. IEEE*, 2006, vol. 4, pp. 2174–2176.