Surrogate-driven motion models from cone-beam CT for motion management in radiotherapy treatments

James Martin

UCL

A thesis submitted for the degree of

Engineering Doctorate (EngD)

2014

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

For my parents, wife and two beautiful children.

Acknowledgements

The journey taken to developing the content within this thesis has been a long, but exciting one. First and foremost, I should thank my supervisor, David Hawkes, for providing direct and good high-level insight into my research. His opinions and thoughts have ensured that I have kept on track and focussed on areas of higher impact with the time I have had available. Thanks also to Jamie McClelland, who has played a key role in all of my work. His very sharp eye and attention has resulted in my work being considerably more robust than it would otherwise have been. The staff at Guy's and St. Thomas' Hospital have been critical, lending their time and efforts to obtain the various datasets needed. Special thanks to David Landau, Christopher Thomas, Connie Yip, Clare Hartill, Shahreen Ahmad, and Asad Ahmad.

I must extend thanks to Ivan Meir and Norman Smith, from Vision RT, who have been my sponsors throughout my EngD. They have added an important industrial perspective onto my work, and have provided me with invaluable exposure to the inner workings of a highly successful, UK-based, technology company. Thanks to the Royal Commission of the Exhibition of 1851, who awarded me an Industrial Fellowship in the second year of my EngD. Their funding have lended weight behind my research, and their contact network have put me in touch with some of the most promising future scientists and technologists in the UK today.

A variety of events have happened over the course of my work, most impactful of which were the births of my two children, Olivia and Henry. Watching them grow each day has been a privilege, and they have certainly added perspective on my life and encouraged me to focus my time effectively. Finally, an honourable mention goes to my parents and soon-to-be wife, who have been extremely supportive to me over this time.

Abstract

This thesis details a variety of methods to build a surrogate-driven motion model from a cone-beam CT (CBCT) scan. The methods are intended to form a key constituent of a tracked RT treatment system, by providing a markerless means of tracking tumour and organs at risk (OAR) positions in real-time. The beam can then be adjusted to account for the respiratory motion of the tumour, whilst ensuring no adverse effects on the OAR from the adjustment in the beam.

An approach to describe an iterative method to markerlessly track the lung tumour region is presented. A motion model is built of the tumour region using the CBCT projections, which then gives tumour position information during treatment. For simulated data, the motion model was able to reduce the mean L2-norm error from 4.1 to 1.0 mm, relative to the mean position. The model was used to account for the motion of an object placed within a respiratory phantom. When used to perform a motion compensated reconstruction (MCR), measured dimensions of this object agreed to within the voxel size (1 mm cube) used for the reconstruction. The method was applied to 6 clinical datasets. Improvements in edge contrast of the tumour were seen, and compared to clinically-derived positions for the tumour centres, the mean absolute errors in superior-inferior directions was reduced to under 2.5 mm.

The model is then subsequently extended to monitor both tumour and OAR regions during treatment. This extended approach uses both the planning 4DCT and CBCT scans, focusing on the strengths of each respective dataset. Results are presented on three simulated and three clinical datasets. For the simulated data, maximal L2-norm errors were reduced from 14.8 to 4.86 mm. Improvements in edge contrast in the diaphragm and lung regions were seen in the MCR for the clinical data.

A final approach to building a model of the entire patient is then presented, utilising only the CBCT data. An optical-flow-based approach is taken, which is adapted to the unique nature of the CBCT

data via some interesting conceptualisations. Results on a simulated case are presented, showing increased edge contrast in the MCR using the fitted motion model. Mean L2-norm errors in the tumour region were reduced from 4.2 to 2.6 mm.

Future work is discussed, with a variety of extensions to the methods proposed. With further development, it is hoped that some of the ideas detailed could be translated into the clinic and have a direct impact on patient treatment.

Contents

List of Figures

List of Tables

Chapter 1

Background

1.1 Lung cancer and strategies for its treatment

Lung cancer has been estimated to be the most common cancer in the world, with the highest number of cases and deaths [21]. In 2008, 1.61 million new cases were recorded, accounting for 12.7% of all new cancers. 1.38 million deaths were also reported in the same year, accounting for 18.2% of total cancer deaths. Since 1980, there has been an increase in the proportion of lung cancer patients from the developing world, from 30% to 55% in 2008. Lung cancer is associated with a high fatality rate, with over 86% of those diagnosed dying of the disease.

The most common forms of cancer treatment are surgery, chemotherapy and radiotherapy (RT). This thesis is concerned with gaining and using knowledge of respiratory motion to improve treatment, and since RT treatments are the biggest benefactor of this knowledge, RT-based treatment will be discussed further.

RT is the medical use of ionising radiation to treat cancer. The intention is to build up a lethal exposure to X-rays on the cancerous region, causing irrepairable damage to the cancer cells' DNA. At the same time, it is essential to spare as

much of the surrounding healthy tissue as possible. An X-ray beam delivers dose throughout its path through the tissue, hence X-ray beams are delivered from a variety of directions, overlapping where the greatest dose is desired. The X-ray beam can also be collimated across its width to further improve the dose profile.

The delivery strategy is determined during planning. A series of CT scans are taken, from which key regions are identified. They include the tumour region (TR) and organs at risk (OAR). Given the clinically-specified radiation tolerances of each organ of the OAR, the delivery strategy is optimised such that the dose on the tumour region is maximised whilst that delivered to the OAR is minimised, or at least kept below a clinically acceptable level. One issue where knowledge of respiratory motion can be of use is with the definition of the TR. Ideally [71], the TR would be formed from the clinical target volume (CTV), which includes the gross tumour volume (GTV) delineated from the planning CT scans plus a margin to account for microscopic spread. However, in practice, further margins need to be added to statistically compensate for daily positioning errors and organ motion, resulting in the planning target volume (PTV). There is growing evidence emphasising the importance of these additional margins, with a variety of studies showing that changes in patient anatomy and breathing pattern occur between planning and treatment fractions [60, 54, 64, 65].

The introduction of Image-guided radiotherapy (IGRT) has allowed daily position errors to be reduced. An important technology facilitating IGRT has been the incorporation of cone-beam CT (CBCT) onto the treatment gantry, allowing scans immediately prior, and during, a fraction of treatment. See Figure 1.1 for a picture of a typical RT treatment beam with orthogonally mounted CBCT scanner. As a result of IGRT, PTV margins can be substantially de-

Omitted for copyright reasons.

Figure 1.1: Image from Medical Physics website (www.medicalphysicsweb.org). RT treatment beam (MV) with on-board CBCT scanner (kV). The CBCT scanner is orthogonally mounted onto the gantry.

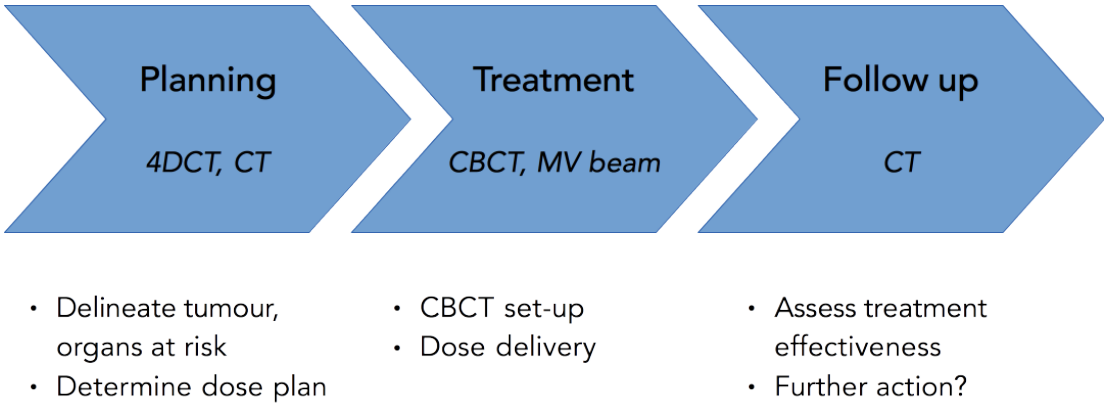

Figure 1.2: A typical treatment RT workflow.

creased (cm to mm), substantially reducing the volume of radiation prescribed [5, 6, 70, 74, 73, 2, 13, 24]. This large decrease is due to margins being expressed in one dimension, which has a cubic relationship to the volume, on which the dose volume is related. For example [71], reducing a 2 to a 0.5 cm margin around a spherical tumour of diameter 5 cm results in a decrease in the irradiated volume of the surrounding organs from 316 to 48 cm^3. Figure 1.2 shows a typical RT workflow utilising an on-board CBCT scanner.

Methods to better shape the dose are essential to realise the benefits of a better defined target region. Methods such as conformal radiotherapy (CRT) allow the dose to be shaped to that defined at planning. Other approaches, such as Intensity Modulated RT (IMRT) go beyond CRT by varying the intensity of irradiation within a radiation field. IMRT is achieved in practice via multi-leaf collimators (MLCs); a series of beam shaping blocks arranged in opposed pairs.

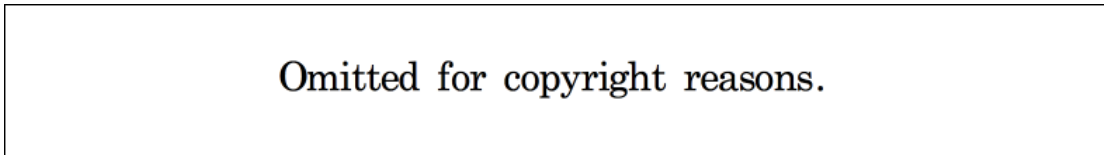

Figure 1.3: Image from Varian website (www.varian.com). Visualisation of how a multi-leaf collimator (MLC) shapes the treatment beam (highlighted yellow).

Two MLC-strategies include delivering IMRT statically (beam off whilst MLCs move to new shape) or dynamically (beam on throughout radiation field delivery, with MLCs exposing areas requiring higher dose for a longer period, within the field). Figure 1.3 shows a visualisation of how the MLCs shape the treatment beam.

Other methods of delivering IMRT include helical tomotherapy [39, 40, 20], Cyberknife [1] and volumetric modulated arc therapy (VMAT) [72, 53]. With helical tomotherapy, the dose is delivered in rotational fashion on a ring-based gantry, with the patient is slowly translated through the bore of the machine. Cyberknife achieves accurate delivery of dose by mounting a LINAC on a car-production robot, which allows a large number of narrow beams to accurately build a dose distribution within the target volume. The robot allows for dextrous beam positioning, delivering beams from a large number of angles and locations around the patient. VMAT uses a similar set-up to SABR except the dose is delivered in one continuous arc, with the MLCs dynamically shifting to deliver the required dose distribution over the target region. Delivering the dose in this manner, as opposed to a series of static fields, decreases the time taken to deliver the dose from 8-12 minutes (SABR) to around 2 minutes.

1.2 Accounting for respiratory motion

Faster treatment times and smaller margins are important drivers behind further developments of next generation RT treatments. Accounting for respiratory motion can offer a promising way to provide further improvements in these areas. Gated [34] and tracked [62, 59] treatments utilise knowledge of tumour position throughout the fraction to treat either at specific parts of the respiratory cycle, or throughout it, respectively. Margins added around the tumour due to respiratory motion can potentially be decreased from around 1 cm to a few mm [59]. Due to a reduced treatment duty cycle, gated treatments will most likely increase treatment times. Breath-hold techniques or equipment [41, 31, 55] may reduce this impact, by allowing the patient to spend more time in the required part of the breathing cycle. Given treatment throughout the duty cycle, tracked treatments may offer the optimal combination of reduced margins and shorter treatment time. For gated treatments, it is possible to choose a treatment window such that the tumour lies close to its average position. This makes planning the treatment simpler. For tracked treatments, the patient may have a variety of respiratory states, and simply moving the beam based on tumour position may give rise to complications. For example, if the tumour is at an extremal state, the organs at risk (OAR) may also be in one, hence simply moving the beam may also expose the OAR to additional dose. Strict [7] limits are detailed for the maximum dosages the OAR can be exposed to, and this unplanned exposure may result in these limits being breached. During arc-therapy treatments, margins are added to the tumour to account for the full seen respiratory motion during planning. Additional margins are added to account for the expected variation of this mea-

surement. If the actual respiratory motion were known, this additional margin could be removed, instead delivering the planned treatment at specific parts of the duty cycle to ensure the tumour positions are within the limits of that seen during planning.

A variety of techniques and modalities have been used to assess tumour motion and shape. Weekly 4DCT scans have been used to assess changes in tumour motion and shape during treatment [10]. Tumour motion was found to vary widely between patients, but could be up to centimetres in extent. The majority of this motion occurs in the superior-inferior direction. These scans can be used to modify treatment before each fraction of treatment and take inter-fraction variation into account. However, giving additional CT scans has significant hospital resource requirements, as a 4DCT scanner must be made available at each treatment, followed by rapid processing of the scans by medical physicists and clinicians immediately prior to the fraction. Given the extent of the additional requirements, as well as the patient dose implications from the extra scans, it is unlikely that prior 4DCT-based modifications to the plan will be a practical approach. CBCT is available immediately prior to each fraction of treatment, making this modality a much more suitable candidate for monitoring inter-fraction variations and adjusting treatment accordingly. 4D-CBCT reconstructions can be made by binning projections according to their phase [66], which is calculated using a surrogate signal extracted directly from the projections. A 4D-CBCT scan can replace a standard CBCT scan, and hence can be integrated into the clinical workflow without the need for additional scans. 4D-CBCT is able to provide an indication of tumour motion, but because each reconstruction is made from a reduced number of projections, the quality of each reconstruction is limited.

This could be improved with longer acquisitions, but would impact on treatment times and delivered dose. Due to binning, it is also possible for 4D-CBCT to underestimate [9] the true extent of tumour motion.

Given the current popularity of integrated CBCT scanners on RT treatment equipment, approaches to measuring respiratory motion which utilise this modality have immediate access to actual patient datasets. Given the large number of problems with the current clinical state-of-the-art, 4D-CBCT, there is clinical interest in better approaches to determine respiratory motion. For CBCT-based approaches which can offer improvements upon 4D-CBCT, there is the possibility to make a meaningful impact in this area within a relatively short time-scale.

1.3 Overview of thesis

This thesis details a variety of methods to build a surrogate-driven motion model from a CBCT scan. The methods are intended to form a key constituent of a tracked RT treatment system, by providing a markerless means of tracking tumour and OAR position in real-time. The beam can then be adjusted to account for the respiratory motion of the tumour, whilst ensuring no adverse effects on the OARs from the adjustment in the beam.

Chapter 1 provides some background information on lung cancer and how it is treated with RT-based treatments. Chapter 2 looks in more detail at how accounting for respiratory motion can be advantagous in RT treatments. Relevant literature, including some state-of-the-art methods to determine respiratory motion during a CBCT scan are discussed.

Chapter 3 describes an iterative method to markerlessly track the lung tu-

mour region. Parameters of a surrogate-driven motion model are fitted during the CBCT scan. Simply by monitoring the surrogate signal during treatment, the model can then be used to determine the tumour position during treatment. The model is robust against many of the issues affecting CBCT, which is of a relatively poor quality, such as limited field-of-view and scatter. The model can also account for complex variations in breathing, such as hysteresis and varying depth and length of breathing cycle. With each iteration, the cost function which is optimised has been structured in a manner whereby it can be optimised quickly, with a unique, global minimum. With suitable improvements, this also means that the approach should fit the model within a clinically useful timeframe. Results are given on a phantom case, simulated patient case and six clinical datasets.

Chapter 4 builds upon this and extends the model to a fully deformable one. With this approach, it is now possible to monitor both tumour and OAR regions during treatment. This extended approach uses both the planning 4DCT and CBCT scans, focusing on the strengths of each respective dataset. The high quality planning 4DCT is used to describe how the various parts of the patient anatomy deform relative to each other. The CBCT, taken immediately prior to treatment, is used to understand the inter-fraction variations that have occurred between planning and the day of treatment. Many of the benefits with the tumour-only approach are retained, such as a cost function which can be optimised quickly, with the potential to fit the model fast enough to be of clinical use. Results are presented on three simulated and three clinical datasets.

A different approach is taken in Chapter 5, which presents an approach to building a fully deformable motion model, but exclusively from the CBCT. An optical-flow-based approach is taken, which is adapted to the unique nature of

the CBCT data via some novel conceptualisations. For this latest work, results on a simulated case are presented.

Chapter 6 explores a variety of extensions to the approaches discussed in earlier Chapters. Briefly, some of these include dividing the number of regions into regions (e.g. lung and non-lung), each of which can have a different relationship to the surrogate signal. Another interesting area to explore includes using the fully deformable model to augment the constrained non-rigid motion model, especially in areas of large inter-fraction variation. Connections between the motion models used in each of these methods are discussed with this in mind.

It is hoped that this thesis can provide a contribution to the literature in this field. With further development, the authors also hope that some of the ideas detailed could be made clinically feasible and hence be used to improve patient outcome.

Chapter 2

Literature review

2.1 FDK-based reconstruction

The Feldkamp-Davis-Kress (FDK) reconstruction [19] algorithm is currently the most popular method used in clinical CBCT scanners. Assuming that enough projections have been taken, this reconstruction approach can provide reasonable quality reconstructions faster than iterative approaches. With a low number of projections, streaking and other artefacts are present in the reconstructed volume. Iterative approaches, although taking longer to produce a reconstruction, may offer a better alternative in this situation. However, lack of projections will not be an issue with actual patient data, as the CBCT aquisitions are designed for use with an FDK reconstruction approach, hence will sample the patient volume enough to achieve a reasonable reconstruction quality with this algorithm. Therefore, although this is also compatible with iterative methods, given the speed of the reconstruction and popularity with clinical scanners, FDK-based reconstruction will be used in the first instance.

The FDK reconstruction approach to CBCT can be visualised by considering

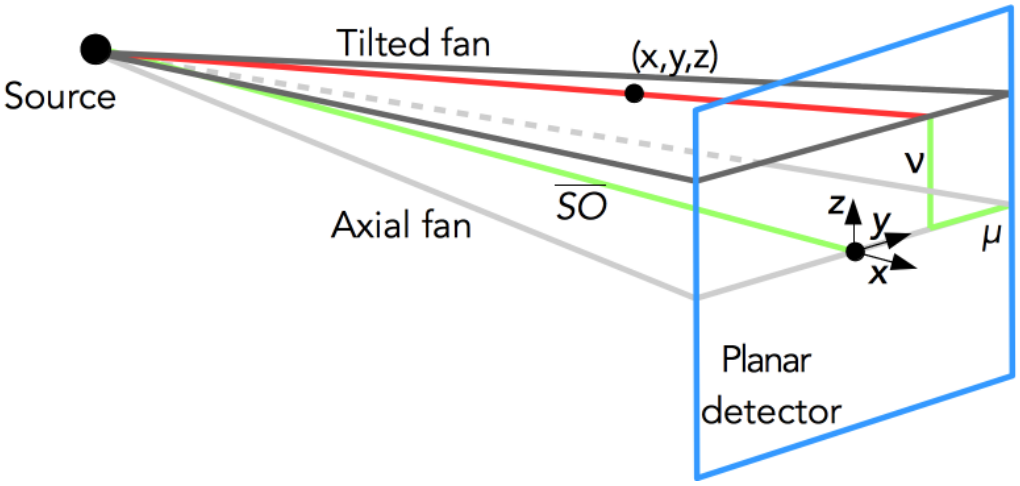

Figure 2.1: Diagram of CBCT geometry, where the source is assumed to have a circular orbit, around a fixed point on line $\overline{SO}$ in the plane of the axial fan. Each projection has an associated projection angle β. Useful for the FDK-based reconstruction approach, the geometry has been visualised in terms of tilted (off-axis) fan-beams.

each imaging cone as a series of tilted fans. Each fan spreads from the source to the edges of the planar detector, passing through the volume. Figure 2.1 shows an illustration of the arrangement. Apart from the axial fan beam, all the others are angled out of plane. The FDK approach approximates each angled fan by pre-weighting the fan-beam formula.

The volume $V(x, y, z)$ can be reconstructed as follows [42]:

1. Preweighting with

$$\omega(\mu, \nu) = \frac{\overline{SO}}{\sqrt{\overline{SO}^2 + \mu^2 + \nu^2}} = \cos\gamma \cdot \cos\kappa, \tag{2.1}$$

where μ and ν are the detector coordinates, γ is the fan angle and κ the cone angle (in the plane of the fan angle specified by γ).

2. Row-wise ramp filtering with 1D ramp filter $h(\mu)$:

$$\tilde{p}(\beta, \mu, \nu) = (\omega(\mu, \nu) \cdot p(\beta, \mu, \nu)) * h(\mu), \tag{2.2}$$

3. 3D back-projection according to:

$$V(x,y,z) = \int_0^{2\pi} U(x,y,\beta)\,\tilde{p}\left(\beta,\mu(x,y,\beta),\nu(x,y,z,\beta)\right)d\beta, \tag{2.3}$$

where

$$U(x,y,\beta) = \frac{\overline{SO}^2}{\left(\overline{SO} + x\cos\beta + y\sin\beta\right)^2} \tag{2.4}$$

Note that this formulation is only exact in the midplane. For tilted fan angles, incomplete Radon data is collected as the source trajectory is circular and does not meet the sufficiency condition for exact reconstruction, given by [63].

Describing equations (2.1,2.2,2.3,2.4) in words, each projection is back-projected after pre-weighting and the application of a ramp filter. The pre-weighting term can be thought of as the geometrical magnification caused by the CBCT geometry (i.e. X-ray point source combined with flat panel detector). The ramp filter is important to avoid severe resampling artefacts when moving from Cartesian sampled projection (radon) space of the detector, to the Cartesian sampled patient volume. This is due to the back-projection formula being descibed essentially in cylindrical coordinates (the directions of x and y depend on the angle β). The back-projection maps the pixel values in the 2D projections onto a host of lines connecting the centre of the pixels to the X-ray source position within the 3D volume. The line is described using the Radon transform, specially modified for CBCT geometry. For further detail, the interested reader is recommended to read the original paper [19] and other good summaries of this approach, for example [69]. The back-projections are then summed together and renormalised (divided by the number of projections), producing the final reconstruction. This final step

is also called accumulation.

2.2 FDK-based motion compensated reconstruction

In this work, we use a modified FDK-based reconstruction algorithm [57] to calculate the motion free image, called the motion compensated reconstruction (MCR). The FDK algorithm can be conveniently modified to perform motion compensation in the reconstruction. For a given motion estimate, each back-projection is deformed with the inverse of the estimated motion, calculated from the motion model. This results in all the back-projections being moved into the space of the reference image ($s_n = \dot{s}_n = 0$). Upon the final summation and renormalisation of the back-projections, the final reconstruction should show a reduction in motion artefacts, assuming the motion estimate during the CBCT is accurate. Figure 2.2 shows a single back-projection before and after motion correction, plus the improvements in motion artefacts and edge contrast attainable with motion correction.

For the work presented, an open-source CBCT reconstruction toolkit [56] is used for FDK-reconstructions. The FDK reconstruction algorithm was modified to perform motion compensation after each back-projection step. To give an example of reconstruction time, for a volume of size 300x300x150 voxels, an MCR can be determined (from around 350 projections) in under 3 minutes. If motion correction were not necessary, a GPU-accelerated version is available, which can produce a reconstruction in under 15 seconds.

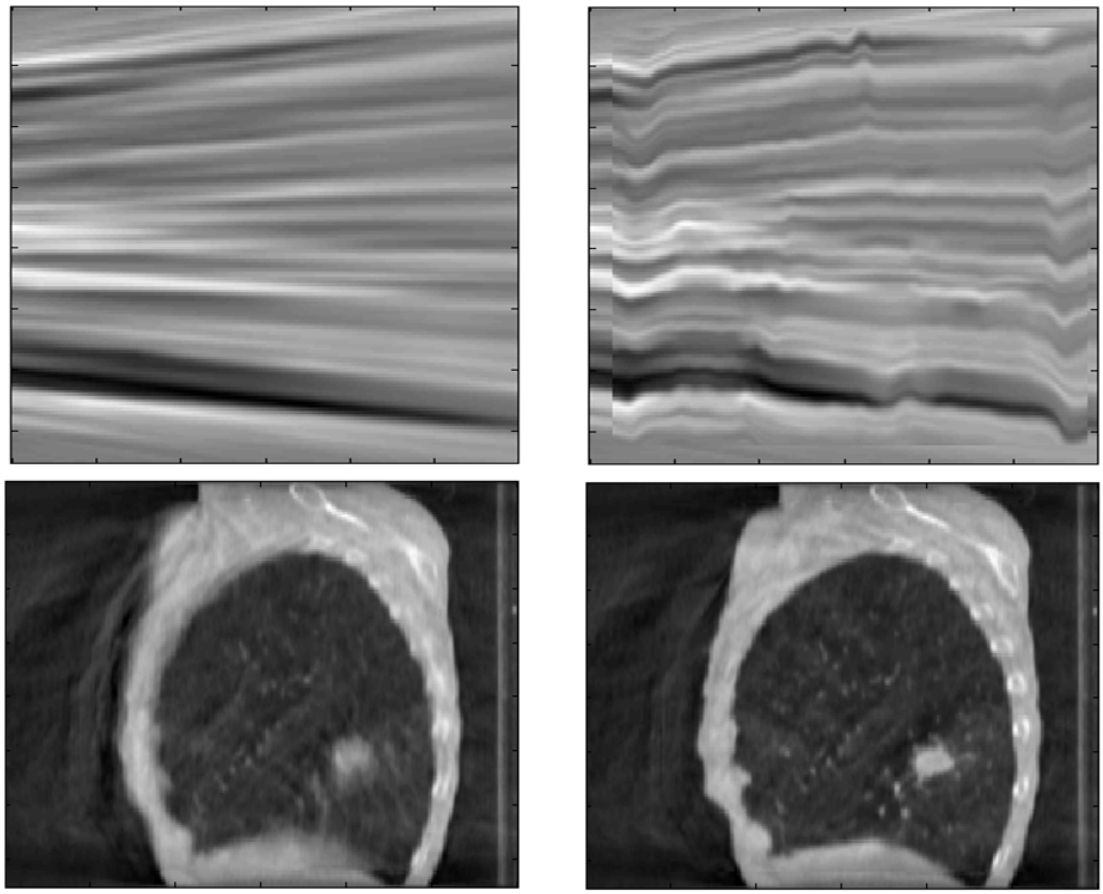

Figure 2.2: An example back-projection before (left) and after (right) application of deformation field from motion model during reconstruction. An example slice of resulting reconstructions using contributions from all back-projections are shown in the bottom row. With suitable motion correction, a visible reduction in motion artefacts can be seen with higher edge contrast around anatomical features.

2.3 Determining tumour and other respiratory-correlated motion from a CBCT scan

Critical to delivering tracked RT treatment is knowledge of tumour position throughout the fraction. The time taken for each individual projection (milliseconds) is such that they can be essentially regarded as free of motion blurring. If the tumour can be identified in each projection, it is possible to estimate its likely trajectory. Implanted markers are one way to easily identify tumour position via thresholding, from which methods [52, 4] have been published to determine the

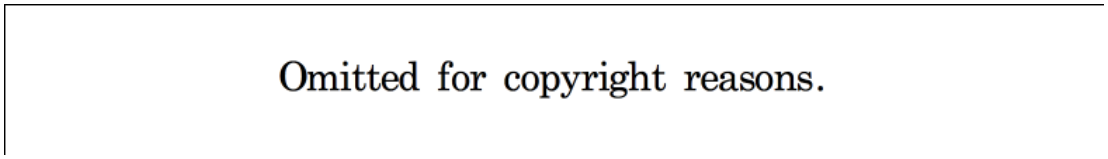

Figure 2.3: From Poulsen et al. [52]. Probability-based estimation of tumor trajectory from cone-beam CT projections. The tumor is projected into imager points (u_1, v_1), (u_2, v_2), (u_3, v_3) for imaging angles a_1, a_2, a_3. (a) First, the Gaussian probability distribution that maximizes the total probability $P(u_1, v_1, a_1) \times P(u_2, v_2, a_2) \times P(u_3, v_3, a_3) \times \cdots$. is determined by maximum likelihood esti- mation (indicated by the gray ellipse). The lines represent the ray lines from the focus point of the imager system to the projected tumor position on the imager. The tumor is known to be located somewhere on these lines. (b) Next, the tumor position along these lines (and thus the three-dimensional [3D] position) is estimated for each projection as the mean position according to the 3D Gaussian, that is, the midpoint of the lines cross section with the gray ellipse (indicated by 1, 2, 3).

likely average and complete trajectory during the scan. Poulsen et al. [52] use a probabilistic appoach to determining the tumour position. A Gaussian distribution is fitted to a series of lines connecting the tumour and source positions during the CBCT scan, such that the total probability is maximised (see Figure 2.3). Given a new tumour position, the estimated 3D position will be the intersection of a line connecting this tumour position in the projection to the source, and the fitted Gaussian distribution. This approach does not take into account hysteresis effects which may lead to errors for patients where the hysteresis variation is large. In order to have real-time tumour position estimation, new projections will be required continuously, increasing dose delivered to the patient during treatment. For practical reasons, this continuous exposure is also likely to be from one direction only (orthogonal to treatment beam) possibly introducing errors in the direction parallel to the CBCT angle.

Becker et al. [4] determine an average position of the marker over the CBCT according to phase. Each projection of the CBCT is binned according to phase.

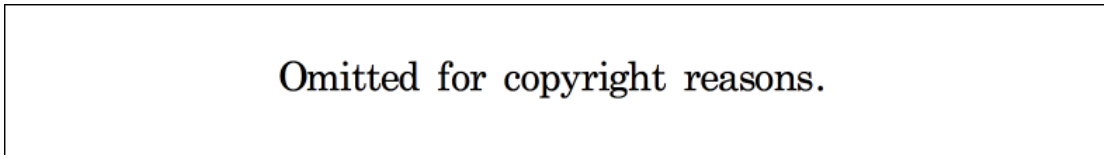

Figure 2.4: From Becker et al. [4]. (a) The average intersection of the back-projected rays from a given phase represents the average 3D coordinate for the marker at that phase. (b) The 3D coordinate for each image (R) is estimated as the point along the backprojection ray closest to the average trajectory ($R_{ave}(\Psi)$).

For each phase, lines connecting the marker in the projections to the source position are then drawn, with the average position chosen to minimise the distances to the lines. Over all phases, an average trajectory therefore results. To determine the specific path for the CBCT (rather than the average path), the point on each line closest to its phase's average position was chosen as the 3D position for that particular projection. Figure 2.4 shows a visualisation of how the average and projection-specific positions are determined. Becker et al.'s approach is retrospective hence can only be used to determine tumour motion throughout the scan. Assumming that the average positions for each phase remain constant throughout treatment, it may be possible to slightly extend the technique to allow real-time monitoring. However, as with Poulsen et al. [52], it would require continuous X-ray projections from the CBCT (increasing the dose delivered to the patient) and to be practical at a direction perpendicular to the treatment beam (introducing errors when estimating where on the line to assume the tumour is).

By relating the marker position to an externally monitored surrogate, which can be measured frequently and easily, continuous X-ray imaging can be avoided throughout treatment. Given a surrogate synchronised with the acquisition, motion models Cho et al. [12] have built from this data and used to track marker position (inferring tumour position) in real time. Parameters of a simple linear

motion model are found via a least squares estimation approach. The parameters were found via an iterative approach and manually set to 5 iterations. Although mentioned that it can be extended to do so, the simple linear model detailed by Cho et al. cannot account for hysteresis. By utilising a surrogate, it is possible to easily take a new surrogate measurement and estimate the 3D tumour position using the motion model. Interestingly, Cho et al. also discuss how to update the model should the relationship between surrogate and tumour motion change, as well as extending the method to other RT treatment machines (e.g. Cyberknife's dual X-ray setup).

Krauss et al. [32] compare four approaches to respiratory motion prediction of a lung tumour, based linear regression, neural networks, kernel density estimation and support vector regression. An implanted gold marker's position is assessed with stereoscopic x-ray flouroscopy. The 3D marker positions are then fed into each of the approaches for motion prediction estimation. Based on a training data-set, each of the approaches essentially provides an estimate of the tumour position a set time after the last sample (the 'look-ahead' length), to be used to position estimation. The sampling rate could also be varied, making the approach applicable for different uses, such as real-time position management systems (e.g. RPM system, Varian Medical, Palo Alto, CA, requiring a sampling rate of around 30 Hz) or gantry mounted flat-panel imagers (sampling rate of around 5 Hz). The authors report an accuracy loss of up to 60% for the kernel density estimation and support vector regression approaches, on the patient level. Neural network approaches reported a maximum accuracy loss of up to 40% and linear regression reported the least accuracy loss of around 30%. On average these errors reduced to 7% for kernel density estimation and support vector regression,

6% for neural network and nearly 3% for linear regression. The authors conclude that an appropriate choice of the patient independent parameters is supposedly more important than the choice of actual prediction method itself.

The linear regression motion model approach in particular has been integrated into both imaging and treatment applications. Fast et al. [18] integrated the linear regression motion modelling approach directly into CBCT acquisition. A motion model, using either implanted electro-magnetic markers or an external surrogate as an indicator of tumour position, was used to modify the CBCT acquisition timing such that all phases of the respiratory cycle are sampled with a more even distribution. For traditional phase-based binning in 4D-CBCT, for example, this allows improved reconstruction for phases which are traditionally undersampled. Although demonstrating a useful application of tumour-oriented motion models, this approach increases the time taken to acquire the scan; unless a sacrifice in the number of reconstructed phases is made. Krauss et al. [33] also used the linear regression approach to actively adapt the apeture of an multi-leaf collimator to track a tumour undergoing respiratory motion. The approach is able to account for motion perpendicular to the treatment beam, which the authors argue features the directions usually having steep dose gradients.

However, Cho et al's, and the methods presented by Krauss et al. all crucially still rely on a marker to identify tumour position in each projection. Implantation of markers, however, is a surgical procedure associated with a risk of pneumothorax [23]. It is also possible for marker migration to occur [49]. Therefore methods which do not rely on implanted markers are desirable.

Lewis et al. [36] proposed a markerless method to identify the tumour, using simulated projections of the 4DCT planning scan. A region encompassing the tu-

mour from every phase of the 4DCT planning scan is used to segment the tumour regions. After a registration (rigid, bony-anatomy-based registration between phase of the 4DCT and CBCT reconstruction), simulated CBCT projections are then taken of every segmented tumour volume (one simulated CBCT will result for each phase of the 4DCT) and saved as templates for comparison. Projections of the actual CBCT are then compared to the simulated projections at the same angle and the closest match used to determine the position of the tumour centre in the actual projection. The position is determined via a 2D translation-only registration between the two best-matching templates and the actual projection, from which the final position is determined from a Gaussian-weighted average of the two results. Figure 2.5 demonstrates this matching process. The tumour trajectory is then calculated in a similar way to Becker et al. [4], by first determining the average phase-based trajectory and then determining the CBCT-specific full trajectory by determining the closest point on the line (from tumour in projection to source position) to the average position (using the phase associated with the projection). As the method does not use markers, there is more error in determining the tumour centres. The results inducing the most error in estimating the average position for each phase were iteratively discarded until agreement was within a pre-defined threshold. Figure 2.6 shows the additional errors when using a markerless method to determine the average phase-binned trajectory. In a subsequent piece of work [35], Lewis et al. are able to use this average trajectory to produce a motion compensated reconstruction of the tumour region.

The method by Lewis et al. [36] involves matching CBCT projections with simulated projections of the 4DCT planning scan, the latter a modality with different scatter, imaging energy and geometrical properties to CBCT. Changes

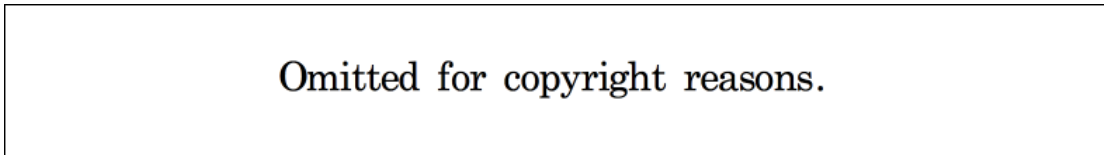

Figure 2.5: From Lewis et al. [36]. A cone-beam projection (left) and generated DRR [projection] (right) for one patient. The ROI [tumour region] is the brighter, highlighted region.

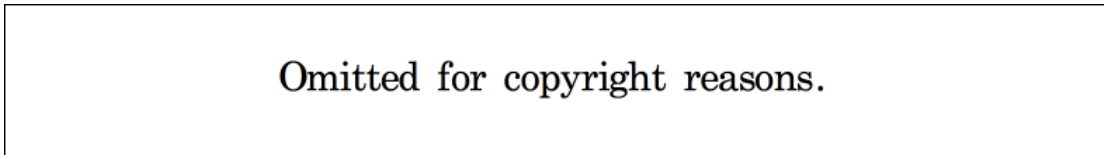

Figure 2.6: From Lewis et al. [36]. Examples of phase-binned tracking results for full exhalation (upper left) mid inhalation (upper right) full inhalation (lower left) mid exhalation (lower right). Line segments connect the source position to the position that the tumor appears at in the imager for each projection.

in tumour size and motion over the course of treatment are not accounted for when using simulated projections of the 4DCT. Outlier rejection techniques were also needed to mitigate the effect of bad matches in projections where the tumour is not visible, or where other high-intensity structures overlap the region to be matched. Although the authors use the tumour motion information to deblur the motion corrupted CBCT reconstruction, this benefit is just confined to the tumour region. It would be more beneficial to have a deblurred reconstruction of the entire patient if possible. As with the implanted marker methods not utilising a motion model [52, 4], Becker et al.'s approach is retrospective and cannot be used to determine tumour position in real-time without continous X-ray imaging and modifications to the method.

Fassi et al. [17] also utilise a prior planning 4DCT scan to build a patient-specific motion model of respiratory motion, which is then adjusted to the day of treatment via a CBCT scan. Respiratory motion is represented via three param-

eters: baseline, amplitude and phase. Registrations between different phases of the 4DCT are initially used to build the model, relevant at the time of the 4DCT. The model is then adjusted to the day of treatment. The CBCT is used to determine the baseline, or mean, tumour position. An external surrogate is then used to monitor the respiratory state, which can be used to calculate amplitude and phase during radiotherapy treatment. The approach was tested on seven patients. The actual tumour centres were manually marked on 100-200 projections of the CBCT scan and compared to that estimated from the motion model, obtaining a median tracking accuracy of 1.5 mm in the vertical and horizontal directions on the CBCT projections. Note that this approach was published after the tumour modelling approach published by Martin et al. [28], which forms the basis of Chapter 3 in this thesis.

Many of the cited approaches use a motion model of some form. It is worth mentioning that this is by no means necessary to effectively monitor tumour position. Gendrin et al. [22] present an approach based on 2D/3D registration. An intensity-based registration of simulated projections of a prior CT scan are compared with the actual CBCT projections. An iterative optimisation yields a rigid-body transformation from which tumour position could be calculated. The method utilises a GPU to allow the registrations to occur in around 0.5 seconds, which allows near real-time determination of tumour position. However, this approach will still need to take this latency into account. The approach was tested on a simple respiratory phantom, and could follow the motion with a mean RMS error of 2.6 mm. Due to the nature of the 2D/3D registration, the approach cannot resolve tumour motion in the imaging beam axis. This was fixed in the authors' approach, which may contribute to errors. Indeed, utilising a motion

model is one way to overcome this issue, as motion from other projections can view the associated tumour displacement from a different perspective. However, using a motion model will inevitably impose a relationship between the tumour motion and some surrogate, which may not always be the case.

In addition to tumour position, knowing the location and motion of the OAR is also desirable to help ensure that any dosimetric constraints are not violated when tracking the tumour motion or modifying the plan in other ways. It is therefore advantageous to have a whole-patient motion model to determine positions of other key patient anatomy during treatment.

Rit et al. [57] use a series of registrations between the phases of the 4DCT planning scan to generate a series of deformation vector fields (DVFs). After choosing the mean position as the reference volume, the DVFs describe how to deform the reference volume to match each of the phases of the 4DCT. Figure 2.7 shows how a point maps to the other phases after registration, using two different reference locations. During treatment, a rigid shift is frequently made to align the CBCT reconstruction with the planning scan. A rigid shift is used clinically as it is only used to correct for the baseline shift of the tumour. This shift was applied to the DVFs to move them into the space of the CBCT scan, from which the motion seen during planning was assumed to be identical to that seen during treatment. By applying this shift to the whole patient, it is possible for incorrect alignment to occur. In addition, assuming that respiratory motion remains identical between planning and treatment fraction is essentially ignoring inter-fraction variations that have occurred, which may introduce further errors. For each CBCT projection, the phase was calculated and used to determine the estimated deformation to deform the patient state at the time of the projection

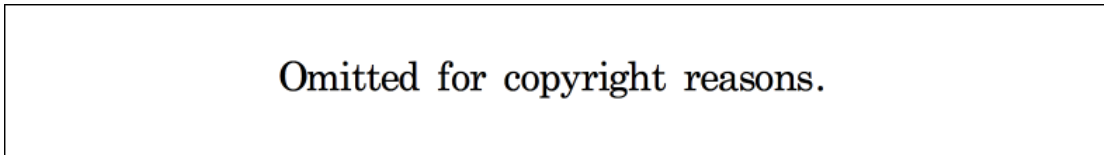

Figure 2.7: From Rit et al. [57]. 2D illustration of the trajectory of one voxel over the ten frames of the 4D planning CT image with (a) the fifth frame as a reference and (b) the mean position as a reference.

to the reference position. Linear interpolation of the DVFs between the two neighbouring phases was used if the phase lied between those calculated from the 4DCT. Although it can account for hysteresis effects and seems consistent with that assumed for the 4DCT, the phase-based motion model assumes that each breath cycle is identical with the last. For patients exhibiting varying breathing depths this model may not be allow suitable flexibility to describe the full respiratory variations seen in these patients.

Li et al. [37] present an interesting approach to determine tumour position directly from a single X-ray projection, based on a motion model built of the entire patient. First each phase of the 4DCT planning scan is registered to a reference phase, to obtain a set of DVFs. Principal component analysis is then used to determine the main variations in the DVFs, with the three largest principal components retained for the model. The DVFs are then moved from 4DCT space to that of the CBCT via a rigid registration on the bony anatomy. A sum of squared differences-based cost function is used to match a projection of the deformed reference image with the desired X-ray projection. Parameters varied to determine the match are those to vary the magnitude of the PCA components, and parameters to facilitate a linear approach to rescaling the intensities (4DCT will often be of different intensity levels to that seen in the CBCT, hence the intensities seen in the projections will be different). As with Rit et al. [57],

using a rigid alignment based on registering bony anatomy risks misalignment, hence may introduce errors. Beyond rescaling each principal component, the rigid alignment also means that the inter-fraction variation that can be accounted for is very restricted. Li et al. attempt to compensate for matching projections of a volume of a different modality to the actual projections with a linear rescaling. In practice, this relationship is non-linear, possibly leading to inaccuracies. This approach can be used in nearly real-time, as the tumour position can be calculated in under 0.5 seconds, but requires a new X-ray projection be taken. Therefore for continual tumour monitoring, this approach will require continuous X-ray projections, increasing the dose delivered to the patient.

Zeng et al. [75] builds a model of the respiratory motion from CBCT projection data. Parameters of a time-varying deformable model are optimised over, improving similarity between simulated and CBCT projections. The cost function included an aperiodicity penalty term, favouring similar deformations at the same point in the breathing cycle, plus deviations from this based on spatio-temporal motion smoothness. However, they do not relate the motion to a surrogate signal, so the model is retrospective and cannot be used to estimate the exact motion that occurs during treatment. Real-time monitoring of patient deformations is not possible as an entire CBCT scan is first needed prior to the optimisation scheme commencing. Due to the large number of parameters, it is both time consuming to fit the model and is susceptible to local minima, the latter therefore not being representative of the true motion.

Zhang et al. [76] propose a method to build a motion model of the whole patient from 4D-CBCT scans of the thorax region. As a motion model is built, this approach can be used to provide real-time estimates of the entire patient, based

on a new surrogate measurement. A standard 4D-CBCT set of reconstructions are produced, which are registered together. This generates a set of deformation fields describing how to deform the reference image to match the reconstruction. 4D-CBCT reconstructions are of particularly low quality, hence registrations between each of the phases may introduce errors at this stage. Principal component analysis is used to reduce the effects of these errors. All but the first two principal components are then disregarded. The model is then calibrated by determining the components of a matrix such that the motion observed is best correlated with the measured surrogate at each phase and that measured one-third of a respiratory cycle prior. Zhang et al. minimise an L2-norm-based distance measure to determine the best correlation. This motion model can account for hysteresis and variable depths of breath, but the choice of surrogate signal (dome of diaphragm extracted from projections) may lead to errors, as extracting the surrogate from projections may contain innaccuracies in the source-to-detector direction. As an alternative to registrations between phases of the 4D-CBCT, Zhang et al. also propose building the model from registrations between phases of respiratory-correlated CT data. However, to move this to the day of treatment they propose a rigid registration on the bony anatomy, which suffers from the same issues (incorrect alignments and ignores inter-fraction variation) as that described with Rit et al.'s approach.

Other authors have proposed methods of building a motion model from other modalities, for example 4DCT and MRI.

Odille et al. [51] proposes an iterative scheme to determine parameters of a motion model, updating the MCR with each iteration. The authors name this method the Generalised Reconstruction of Inversely Coupled Systems (GRICS).

GRICS consists of a generalised reconstruction and motion model optimisation steps, which are iterated over using a multi-resolution approach. A linear motion model is proposed parameterise the motion during the MRI scan. The approach is applied to moving phantom and free breathing patient data. In both sets of data, the MCRs showed improved contrast and less motion blurring. The approach offers a promising method to coping with respiratory motion during a scan. However, it should be ensured that the motion model is physically realistic. During the motion model optimisation step, Odille at al. optimise a cost-function which minimises a quadratic error term, which attempts to make the updates account for the current differences between the actual scan (with full motion) and estimated motion, and also an additional smoothness constraint. The smoothness contraint on the parameters of the motion model. The smoothness contraint used may not hold in regions where sliding motion occurs, such as between the chest wall and lung. In its current form, the model takes over 40 minutes to fit the model. The authors intend to reduce this dramatically, as the approach is amenable to parallelisation.

Hinkle et al. [26], proposes a maximum a posteriori (MAP) algorithm for tracking organ motion. The approach estimates reconstructions of the motion-free image and deformations in the anatomy. A cost function describing the log-data likelihood of observing the data is constructed, consisting of a motion-free motion free image estimate and velocity field estimate. This likelihood is then maximised to obtain the most likely velocity field and motion-free image, given the parameterisation of the velocity field and method of image reconstruction. The maximisation is done via a Euler-Lagrange method, which provides updates to alternately apply to current estimate of the motion-free image and velocity field.

These alternate updates are then repeated until convergence of the algorithm. The method was applied to 4D respiratory-correlated CT (4D-RCCT) phantom data, increasing the signal-to-noise ratio (SNR). A select (central) point animated with the fitted velocity field also showed excellent correlation (0.9988) in the SI-direction to the displacement of the respiratory signal from the 4D-RCCT. The approach was also used on a single patient dataset, again showing an improvement in SNR and reducing slice mismatch artefacts (compared to using a binning-based 4DCT approach to reconstruction). This approach is an interesting approach, however the time taken to obtain convergence was not mentioned, and may be clinically unrealistic. In addition, the approach has been applied on a slice-by-slice basis, with the assumption that the motion occurring during the acquisition of a single slice (0.5 s) is minimal. This assumption may not be completely accurate, as within 0.5 s respiratory motion would be expected in various parts of the breathing cycle. If this were to be applied to CBCT, a similar analogue to motion free slabs is not present. One would have to extend the approach to deal with the individual x-ray projections, which could be assumed to be motion free.

Ehrdhardt et al. [16] proposed an optical-flow-based approach to determining respiratory motion during a cine-CT scan. To create a reconstruction at a desired respiratory phase, a non-linear registration method is used to estimate the optical flow between neighbouring respiratory phases. An interpolated image can then be created. Each series of slices of the cine-CT had an associated respiratory phase via monitoring of a surrogate. To account for hysteresis, the inhale and exhale parts of the respiratory cycle were treated seperately. Gaussian smoothing is performed to limit the possible differences in magnitude and direction of neighbouring velocity vectors. The method was applied to four lung cancer pa-

tient datasets. More than 2500 slices were acquired per patient, taking around 10 minutes. A decrease in motion artefacts was seen in the reconstructions and reconstructions at specific respiratory phases were possible, including visualising a range of snapshots to better assess respiratory motion. However, the approach currently requires an extended scan of the patient, possibly exposing them to additional imaging dose. The authors mention that the approach is computationally expensive, taking over 30 hours. With cine-CT slices, full 3D reconstructions are available for a limited set of slices, allowing a volume to volume registration to be performed. In attempting to move this technique to other modalities, there may also be issues for data of different dimensionality. For example CBCT, one is dealing with 2D X-ray projections, as opposed to 3D volumes.

Table 2.1 lists some of the methods described in a more easily comparable manner.

Table 2.1: Key features of discussed approaches to accounting for respiratory motion during RT treatment. *Note that this approach was published after the tumour modelling approach published by Martin et al. [28], which forms the basis of Chapter 3 in this thesis. In the table, complexity of variation modelled can be either amplitude, hysteresis of full inter-cycle variation.

Paper (author)	Modalities	Target region	Extra dose for tracking?	Markerless?	Can drive daily treatment?	Complexity of variation modelled
Poulsen et al.[52]	CBCT	Tumour	Yes	Markers	Yes	Full
Becker et al.[4]	4DCT,CBCT	Tumour	Yes	Markerless	Yes	Full
Cho et al.[12]	CBCT,mono and dual x-ray	Tumour	Yes	Markers	Yes	Amplitude
Fast et al.[18]	CBCT	Tumour	No	Merkerless	No	Hysteresis
Krauss et al.[32]	CBCT	Tumour	Yes	Markers	Yes	Hysteresis
Lewis et al.[36]	4DCT,CBCT	Tumour	Yes	Markerless	Yes	Full
Fassi et al.*[17]	4DCT,CBCT	Tumour	No	Markerless	Yes	Full
Gendrin et al.[22]	4DCT,CBCT	Tumour	Yes	Markerless	Yes	Full
Rit et al. [57]	4DCT,CBCT	Whole patient	No	Markerless	Yes	Hysteresis
Li et al.[37]	4DCT,CBCT	Whole patient	Yes	Markerless	Yes	Full
Zeng et al.[75]	CBCT	Whole patient	N/A	Markerless	No	Full
Zhang et al.[76]	CBCT	Whole patient	Yes	Markerless	Yes	Full
Odille et al.[51]	MRI	Whole patient	N/A	Markerless	N/A	Full
Hinkle et al.[26]	4D-RCCT	Whole patient	N/A	Markerless	N/A	Full
Ehrdhardt et al.[16]	Cine-CT	Whole patient	N/A	Markerless	N/A	Full

Chapter 3

Tumour-only motion model

The work presented in this chapter was initially inspired by Odille et al's [51] work in MRI. From these conceptual beginnings, I built up the mathematical underpinnings and overall approach to building the motion model from CBCT and surrogate data from scratch. A large proportion of my time was also spent acquiring the data and performing data analysis. Important components of the approach adapted from other authors include tumour enhancement in the CBCT projections [43], the CBCT reconstruction package [56] and form of the motion model [38]. This work has led to first author journal, conference and workshop publications, detailed in chapter 7.

3.1 Introduction

The most critical region to monitor during the delivery of RT treatment is tumour position. For clinical RT treatments (e.g. IGRT), this region is scanned for shifts between fractions and corrected for if necessary. Even with cutting-edge tracked treatment systems, such as Cyberknife, it is only the tumour region which is

exclusively tracked during treatment. With this in mind, a logical starting point would be a method to monitor tumour position throughout the fraction. With just one region to track, one would hope the problem is simplified significantly. However, due to the difficulty of determining 3D tumour position in the CBCT projections, it still remains a difficult challenge.

In this chapter, a method is derived to build a motion model of the tumour region. The approach has multiple advantages over other methods, such as being markerless, having minimal additional scanning requirements and not delivering additional dose to the patient whilst monitoring tumour position. Apart from importing the PTV (tumour region plus a margin) from the planning scan, the model is exclusively built from a standard CBCT scan. As CBCT scans are taken immediately prior to, and during, every fraction of treatment, the approach can account for inter-fraction variations. The form of the motion model used can account for complex respiratory motion, such as hysteresis and variations in length and depth of breathing cycle. Although complex motion can be modelled, only an external indicator of the breathing state, a surrogate, needs to be monitored. The motion model is related to this surrogate, so it is required that during model building from a CBCT scan, a synchronised surrogate trace be available. After building the model, only a new surrogate trace is needed to predict tumour position throughout treatment.

What it intended is a completely markerless approach, where the surrogate is continuously monitored and used to determine respiratory motion of the tumour region. Many forms of surrogate can be used with this approach. In this work, an optical-based stereo-camera system is used. The presented approach could therefore be integrated into a larger system, which could determine tumour posi-

tion based on a motion model built from a CBCT scan, plus additionally monitor the patient continuously for large shifts in their position. If such a shift occured, treatment could be stopped and the new positioning CBCT scan used to re-build the motion model.

The proposed approach could be applicable to many of the available radiotherapy systems with on-board CBCT. In the first instance, the proposed treatment system could be a modified standard SABR or arc-therapy treatment system. The main addition would be a method to acquire a surrogate signal during treatment, which includes a method to synchronise this to the CBCT projections if required. The additional cost of a surrogate is much less than buying a completely bespoke system especially for tracked treatment. Therefore this approach may enable a more cost-effective way to achieve a tracked treatment capacity, via an upgrade of existing equipment (as opposed to a complete replacement). After describing the general method, results are presented on phantom, simulated and clinical data.

3.2 Methods

3.2.1 Modelling tumour motion

As it is only the tumour motion that is of interest, the motion model is confined to a relatively small region of the patient. In this chapter, the motion of the tumour is approximated as the motion of a solid, non-deforming mass. With this approximation, a rigid translation-only motion model is able to capture the required range of motion (hysteresis, variable depth and length of breathing cycle) whilst keeping the smallest possible number of parameter to optimise.

During the CBCT scan, it is assumed that the tumour motion can be described

by a series of translations, F_n, of a motion-free image of the tumour region, $V_{ref}(x) : x \in \Omega_{TR}$. Ω_{TR} is the tumour region, imaged during the CBCT scan and $n = 0, 1, ..., (N-1)$ is an index corresponding to the time of each of the N projections. The transformed tumour region, V_n, is then:

$$V_n(x) = V_{ref}(x + F_n(x)). \tag{3.1}$$

The translation-only motion model can be expressed as:

$$F_n(x) = s_n \mu_1 + \dot{s}_n \mu_2. \tag{3.2}$$

μ_1 and μ_2 are the motion model parameters, each being three-element vectors. s_n is the surrogate signal, which is a scalar quantity. $\dot{s}_n$ is the temporal derivative of s_n. The motion model parameters determine the superior-inferior (SI), anterior-posterior (AP) and left-right (LR) motion of the tumour region as the surrogate varies. μ_1 determines the contribution to the overall tumour region motion for changes in s_n. Similarly, μ_2 determines the contribution for changes in $\dot{s}_n$.

The motion model presented in (3.2) has a contribution from s_n and $\dot{s}_n$. Figure 3.1 shows an example patient respiratory trace, plus a variety of tumour trajectories for various values of the motion model parameters. To briefly speculate upon the physical justification for the repsiratory properties that the surrogate signals are correlated to; imagine that s_n is a spirometry signal. s_n therefore indicates the volume of air in the lungs and should correspond to the depth of breathing. $\dot{s}_n$ is then the airflow to the lungs, which gives an indication of the which part of the breathing cycle the patient is in [38], for example inhale or exhale. $\dot{s}_n$ can therefore be used to measure the hysteresis effect. Here, instead

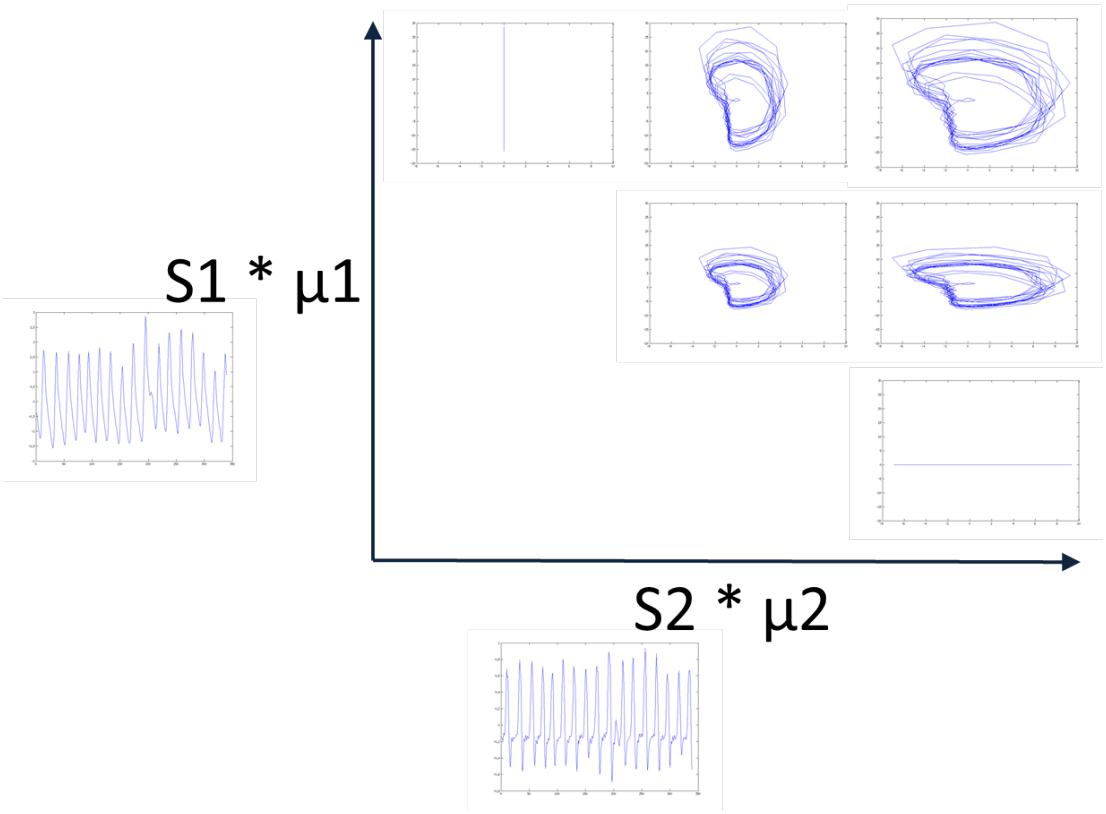

Figure 3.1: An example patient respiratory trace (S1) and temporal derivative (S2), plus description of resulting trajectory from changes in the motion model parameters, μ_1 and μ_2. $*$ is a multiplication. As the figure is in 2D, only the SI component of μ_1 and AP component of μ_2 are being varied, with the resulting motion shown on a sagittal plane.

of using a spirometry signal, we use essentially the average height of the patient chest above the treatment couch. This form of surrogate signal has been shown [27] to correlate well to a spirometry signal. Chapter 3.2.3 describes how this surrogate signal is acquired in more detail.

To fully fit the motion model, six parameters need to be found. This small number of parameters should provide a good balance between modelling complex aspects of respiratory motion and not over-fitting the CBCT data. As CBCT is of a poor quality (reconstruction artefacts, limited field-of-view, scatter) a model which can be robust to these problems is important.

3.2.2 Updating the motion model

The motion model parameters are what determine the motion of the tumour with the surrogate signal. In this approach, given an estimate of the motion model parameters, a method is presented to update them based on comparison with the actual motion seen in the CBCT projections. With an estimate of the motion model parameters, V_n is now an animated tumour volume. By comparing projections of a simulated CBCT of this animated tumour volume with the actual CBCT projections, differences seen can be described in terms of updates to the parameters. For what follows, it is assumed that the motion-free tumour region, V_{ref}, is available.

Introducing an error term for each projection, ε_n:

$$p_n = P_n \left(V_n + \varepsilon_n \right) \tag{3.3}$$

where p_n are the actual projections and P_n is the projection operator. If the simulated projections are not too different from the actual projections, ε_n can be rewritten using a first order Taylor expansion.

$$\varepsilon_n \approx \nabla V_n \cdot \delta F_n \tag{3.4}$$

where δF_n is the first order correction to F_n to make the estimated projection more like the actual projection. Substituting (3.4) into (3.3) and rearranging:

$$R_n \approx P_n \left(\nabla V_n \cdot \delta F_n \right), \tag{3.5}$$

where the residues $R_n = p_n - P_n \left(V_n \right)$ is the actual minus the estimated projection,

at time n. Using an SSD-like cost function, residues at all projection angles can be incorporated into a single update of μ, $\delta\mu$:

$$\delta\mu = \underset{\delta\mu}{argmin} \left[\sum_{n} \sum_{pixels} \left(R_n - P_n \left(\nabla V_n \cdot \delta F \right) \right)^2 \right] \tag{3.6}$$

Given a rigid translation-only motion model (3.2), δF can be re-expressed:

$$\delta F = s_n \delta\mu_1 + \dot{s}_n \delta\mu_2. \tag{3.7}$$

Substituting (3.7) into (3.6) and algebraically expanding (see Appendix A for a detailed derivation), the cost function can be shown to simplify to the following form:

$$\delta\mu_1, \delta\mu_2 = \underset{\delta\mu_1, \delta\mu_2}{argmin} \left[\sum_{\theta,\phi=0}^{6} \lambda_\theta \lambda_\phi \sum_{n} \sum_{pixels} C_\theta C_\phi \right], \tag{3.8}$$

where

$$\begin{aligned}
&\lambda_0 = 1; && C_0 = R_n; \\
&\lambda_1 = \delta\mu_{1,x}; && C_1 = -s_n P_n \left(\nabla_x V_n \right); \\
&\lambda_2 = \delta\mu_{1,y}; && C_2 = -s_n P_n \left(\nabla_y V_n \right); \\
&\lambda_3 = \delta\mu_{1,z}; && C_3 = -s_n P_n \left(\nabla_z V_n \right); \\
&\lambda_4 = \delta\mu_{2,x}; && C_4 = -\dot{s}_n P_n \left(\nabla_x V_n \right); \\
&\lambda_5 = \delta\mu_{2,y}; && C_5 = -\dot{s}_n P_n \left(\nabla_y V_n \right); \\
&\lambda_6 = \delta\mu_{2,z}; && C_6 = -\dot{s}_n P_n \left(\nabla_z V_n \right).
\end{aligned} \tag{3.9}$$

∇_x, ∇_y and ∇_z are the partial derivatives in LR, AP and SI directions, respectively. The sum over pixels occurs over all pixels in the projection. Because of the highly constrained nature of the motion model used, further regularisation (for

example Tikhonov regularisation [68]) associated with more complex deformation motion models are not necessary. Note that the motion model parameter updates, λ_θ, have been moved outside of the sums over projection and pixel. This speeds up optimisation considerably, as sums over C_θ need only be calculated once, and not for every trial value of λ_θ. The parameter space was searched for the minimum, giving the motion model parameter updates $\delta\mu$. The cost function is of a positive definite quadratic form, hence is convex. This global minimum can be found quickly and simply. A BFGS Quasi-Newton method with a cubic line search procedure [11] was used to search the parameter space.

After the updates are calculated, the motion model parameters can be updated:

$$\mu_{new} = \mu + \delta\mu. \tag{3.10}$$

As mentioned earlier, these updates rely on knowledge of the motion-free reconstruction, V_{ref}. Unfortunately, this is not available and must also be determined in parallel with the fitting of the model model parameters. Described in 3.2.5, this is done in an iterative manner, with a motion-free reconstruction being calculated with a modified FDK reconstruction algorithm.

3.2.3 Align RT and extracting the surrogate signal

Align RT (Vision RT Limited, London, UK) was used to produce the surrogate signals. Their bespoke equipment consists of 3 pods which are mounted on the ceiling around the radiotherapy system. Each pod projects a speckled pattern on the surface of the patient and images it from two camera angles. Each part of the

pseudo-random, speckled pattern can be uniquely identified, and from two views, each point's 3D position can therefore be determined. Three pods are used to account for the gantry occlusion that occers from certain views as it rotates.

The system produces 3D surfaces at a rate of approximately 15 Hz. Currently, this frame rate is only available in the research implementation of their software. The raw bitmap images (with the speckled pattern) need to be post-processed after the treatment to produce the 3D surfaces. Although this is currently not real-time, it could be made to be with appropriate resource applied, based on prioritisation on offering this feature to Vision RT's customers.

Vision RT's bespoke software was used to process the bitmap images. These take into account the geometry of the pods, lighting conditions in the room and desired resolution required from the surface. The surfaces are written in ".obj" format, which is easily read by other third party software. In this work, the surfaces were then loaded into MatLab (MathWorks, Massachusetts, USA) to produce the 1D surrogate trace from the 3D surfaces.

Surrogate traces were extracted from the skin surface data via the use of a manually defined 5-point bounding box. The bounding boxes were manually drawn for each CBCT acquisition. Figure 3.2 shows an example patient chest surface with bounding box overlaid. The mean height of the surface above the treatment couch, within the enclosed region, was then used as the respiratory trace. The raw surrogate signal was then normalised (mean subtracted; divided by standard deviation).

An X-ray detector (Black Cat Systems, Westminster, USA) was used to synchronise the skin surface data with the CBCT projections. A simple program, running on the Align RT host computer, sampled the photon count every 8ms.

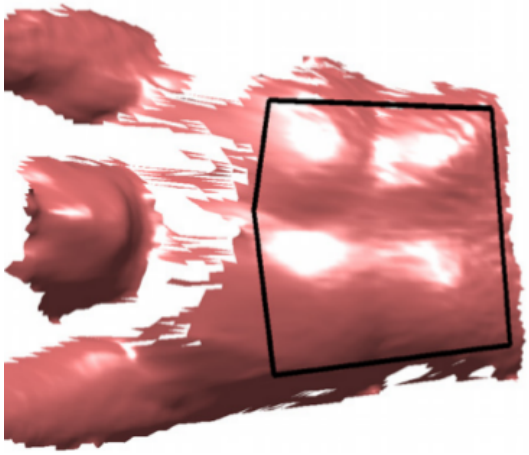

Figure 3.2: Example Vision RT surface with 5-point bounding box.

Peaks corresponding to individual projections were identified and the system clock used to put projections and skin surfaces in the same time reference. Once in same time reference, the surrogate and projections are now synchronised and can be used for model building and MCRs.

3.2.4 Enhancing the tumour region in the CBCT projections

The motion model introduced above (3.2) is only valid for projections of the tumour region. Therefore, prior to fitting a motion model to the data, the projections are first pre-processed to enhance the tumour region. Previously described in [44], the tumour enhancement method attempts to improve tumour contrast in the projections by subtracting anatomy outside the tumour region from the projections.

A standard reconstruction is performed, from which the tumour region is delineated. The tumour region is chosen such that it is sufficiently large to include all of the tumour's motion. For this work, the tumour region used was based on

the PTV from planning. A simple program was constructed in MatLab (MathWorks, Massachusetts, USA), which allows the PTV to be imported and then stretched in SI, AP and LR directions. Importing the PTV was done for convenience, and could be drawn manually instead. A mask is created of the tumour region and used to create two volumes. One is of the tumour region with voxels outside set to the same intensity as air (0). The other is of non-tumour region with the tumour region voxels set to zero. Simulated projections of both the tumour region and the non-tumour region are generated using the same geometry as the real CBCT acquisition. The non-tumour region projections are then subtracted from the original CBCT projections.

If the non-tumour anatomy was stationary during the CBCT acquisition, only the tumour region would be left in the subtracted images (for an ideal reconstruction suffering from no artefacts or CBCT-related problems, such as incomplete field of view). However, as the other anatomy is moving, artefacts are generally present in the subtracted images. The effect of these artefacts is minimised by masking the projections according to the projected outer edges of the tumour region, giving the enhanced tumour projections. Figure 3.3 illustrates the tumour enhancement procedure. Also shown are part of an example projection before, during and after application of tumour enhancement.

3.2.5 Iterative approach

Given some estimate on the motion, the motion model updates could be accurately determined if the motion-free image were known. However, to determine the motion-free image one must already know the motion correction to apply during calculation of the MCR, which would mean knowledge of the true motion.

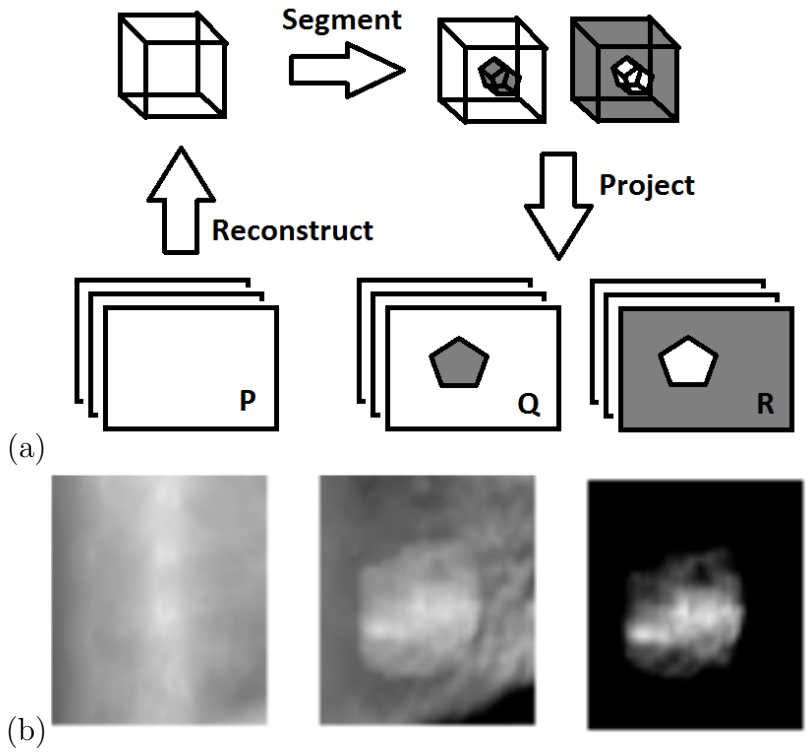

Figure 3.3: (a) Visualisation of tumour enhancement procedure. P are the original projections. Q and R are projections of the non-tumour and tumour regions, respectively. (b) Example projection after each stage of tumour enhancement process. Original (P), original minus non-tumour region (P-Q) and tumour enhanced (P-Q masked by R) projections, from left to right respectively.

In this work, this circular problem is solved with an iterative approach.

An MCR is first calculated assuming there is no motion (i.e. standard reconstruction). The MCR is then set as the reference volume and the motion model parameter updates calculated. Because the MCR will have motion-related artefacts present, the updates may not be completely accurate. However, some of the motion would be expected to be recovered, resulting in parameter updates closer to estimating the seen motion than the no motion assumption. Equation (3.10) can then be used to update μ_1 and μ_2. The updated motion model is then used to recalculate an improved MCR. This MCR is now set as the reference volume and the procedure repeated until convergence conditions are met. For this work, the algorithm was terminated if the maximum change in the motion, over the

entire CBCT acquisition, was less than one voxel in SI, AP or LR directions. The volumes all had 1 mm isotropic voxels.

3.3 Results

3.3.1 Phantom experiment

A modified Modus QUASAR (Modus Medical Devices, London, Canada) phantom was used to test the algorithm on clinical phantom data. An Elekta Synergy (Elekta, Crawley, UK), at Guy's and St. Thomas' Hospital (London, UK), was used to image the phantom. Emulating current protocol, approximately 350 projections are taken with the gantry turning one complete rotation, at a frame-rate of 5.5Hz.

Via a signal generator, parts of the phantom can be set to oscillate in two orthogonal directions in a linearly related way. Initially, an approach was taken which avoided the tumour enhancement technique. This was done to independently test the general iterative approach and see if it had merit. An 'M4' metal nut was attached to an 'internal' moving point in the phantom, moving in the SI direction. An analogous patient situation could be a metal marker implanted in the site of the tumour. An 'external' surface of the phantom was monitored with Align RT, which was moving in the AP direction. As with the simulated data, the surface data was converted into a surrogate signal to be used for fitting the motion model. In this case, precise knowledge of the position of the metal nut was known, as was the fact that a linear relationship existed between the SI position of the metal nut and the AP position of the tracked surface (surrogate). The phantom was placed upon the treatment couch, with the metal nut motion

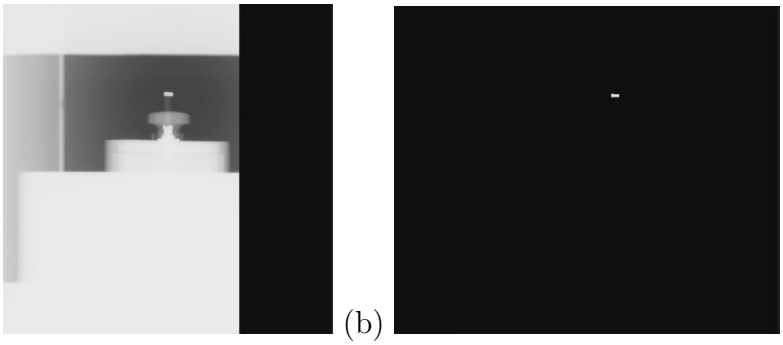

Figure 3.4: Example projection of phantom: (a) raw (with offset detector); (b) thresholded and masked.

approximately about the centre of rotation of the CBCT scanner. The metal nut was set to approximately a 4 second oscillation period.

For this example, it was possible to avoid using the tumour enhancement method. Simple thresholding and masking steps were applied to the CBCT projections to extract only the parts of the images of the metal nut. Figure 3.4 shows an example projection before and after application of these pre-processing steps. After running the proposed approach on the CBCT data, the algorithm terminated after 3 iterations. Improvements were seen in the edge contrast after each iteration. Figure 3.5 shows slice from the reconstruction of the nut region along with select edge profiles. Measurements of the nut were taken from the MCR, all of which agreed with the actual dimensions to within 1 mm. As 1 mm^3 voxels were used for the reconstruction this is a reasonable agreement.

3.3.2 Simulated data

The approach was applied to simulated dataset, based on actual patient data. This also included application of the tumour enhancement technique on the tumour region of the patient volume. A CBCT acquisition was simulated using a deformable registration-based motion model built from real patient data [46, 45]. An

end-exhale 4DCT volume was non-rigidly registered to a series of free-breathing CT volumes (taken in cine-CT mode). Unlike 4DCT, this approach allows the motion to be simulated for any surrogate signal value. The motion model uses the surrogate signal value and its gradient, which ensured the simulated motion included some complex features of respiratory motion, such as hysteresis and breath-to-breath variations. An Align RT surrogate trace, synchronised with the patient data, was used to animate the patient volume. A simplified CBCT geometry and projection angle sequence (of similar properties to the actual CBCT system) were used. To approximate set-up prior to radiotherapy treatment, the mean tumour position over the simulated acquisition was placed at the centre of rotation of the simulated CBCT.

Relative to the time-weighted mean position of the tumour centre of mass, the mean error was reduced from 4.1 to 1.0 mm, without and with motion correction respectively. The maximum error was reduced from 16.5 to 4.6 mm. Figure 3.6 shows a lateral view through the tumour region before and after motion corrections, as well as the ground truth. There appears to be sharper edge contrast of the tumour region after motion compensation, which is also a closer representation of the ground truth. With each iteration, the reconstruction got noticeably sharper. The greatest corrections to the motion occurred at the earliest iterations. Parts of the anatomy not moving as the tumour does, such as the ribs, appear more blurry. Five iterations were needed for convergence.

3.3.3 Clinical data

To determine whether the proposed approach can be of use in actual clinical treatment, the method was tested on actual clinical data. Ethics approval was given

Table 3.1: Patient and tumour characteristics.

Patient ref.	Histology	Tumour location	TNM
1	Squamous cell carcinoma	Left upper lobe	T1bN0M0
2	Squamous cell carcinoma	Right lower lobe	T1N0M0
3	Adenocarcinoma	Right lower lobe	T1bN0M0

to access a series of datasets of lung cancer patients undergoing RT treatment at Guy's and St. Thomas' NHS Foundation Trust (GSTT). At GSTT, patients are planned on a combination of free-breathing CT and 4DCT scans. Organs at risk are delineated on the free-breathing CT by a clinician. The 4DCT scans, together with the free-breathing CT scan, are used to define a target volume (the internal target volume, or ITV) which encompasses the tumour at all breathing phases. CBCT is used to confirm alignment of the tumour with the ITV prior to each fraction of treatment. If the CBCT shows that a shift is required, another CBCT will be acquired to confirm alignment after the shift. An Elekta Synergy (Elekta Limited, Crawley, London, UK) treatment machine, with orthogonally mounted CBCT scanner to the treatment beam, was used for CBCT scans. For all CBCT scans, synchronous 3D skin surface data was acquired from Align RT, synchronised to the CBCT and had a surrogate trace extracted via the method described earlier.

Six SABR datasets, from three patients, were used. All the scans were "fast" scans, acquiring almost 350 projections over approximately 1.5 minutes. For each patient, two scans were used, corresponding to the first and second scans of a fraction (i.e. before and after the patient has been shifted to correct for the basline shifts in tumour position). See Table 3.1 for patient details and their tumour characteristics.

Since the ground truth of the 3D tumour location at the time of each projection

was not available, the tumour location was identified in a subset of projections where the tumour was visible. To assist the clinician in identifying the tumour location, a bespoke piece of software was made to allow the clinician to manually align projections of the MCR with the tumour. To reduce the workload of the clinician, every fifth projection was chosen for a reduced dataset (approximately 70 per scan). Errors were calculated as the 2D distance between the manually identified tumour location and: 1) the projection of the machine isocentre (centre of rotation of the CBCT), 2) the projection of the mean tumour position (as determined from the model fitting), and 3) the projection of the model estimated tumour location. Projections where the tumour was not visible could be skipped. To make the errors more meaningful, the units were scaled from pixels on the projections to mm at the isocentre (i.e. multiplied by pixel dimension and scaled from projection to isocentre position). Errors are referred to as either in the SI or transverse directions, corresponding to the vertical and horizontal directions in the projections, respectively. Figure 3.7 shows a projection after a clinician aligned projections of the MCR.

To understand the observer variation, the delineations were repeated a second time by the clinician. The results of the repeated delineation are shown in Table 3.3. The mean absolute variation of the observer delineation was under 2 mm in SI and transverse directions, with maxima of nearly 8 mm. The clinician was unable to delineate the tumour to a good accuracy in approximately two thirds of the provided projections, which were skipped. Reasons for this include the poor quality of the projections, similarity of the tumour edges to the background and high contrast anatomy passing behind the tumour. The majority of the skipped projections were from a lateral view, meaning the transverse errors mostly

correspond to left-right errors.

Averaged over all cases, the use of a motion model reduced mean absolute errors to under 2.5 mm, in either SI or transverse directions. For each individual case, all bar one of the mean absolute errors are within the maximum observer absolute error for that case. The mean absolute error was only 0.21 mm larger than the maximum observer variation for the exception. Maximum absolute errors with motion model tracking are under 7.5 and 7 mm in the SI and transverse directions, respectively. Fitting the motion model also provides an estimate of the mean position of the tumour. Using the projected mean position produced mean absolute errors of under 3.5 mm in either SI or transverse directions, averaged over all cases. Excluding one scan (3A), where the patient inhaled sharply, maximum absolute errors under 8.5 mm were seen with the projected mean position. For the case with the sharp inhalation, a maximum absolute error of 22 mm was seen in the SI direction. Motion model tracking reduced this maximum absolute error to 7.5 mm. Convergence was achieved in under 5 iterations for all cases. Table 3.2 show the full results for the patients.

Figure 3.8 shows the apparent improvement in edge contrast of the tumour region with motion correction. A clinician compared the MCR to the planning scans and confirmed that the motion corrected images provided better visualisation of the tumour features (e.g. spicules) and nearby vessels.

3.4 Discussion and conclusions

Initial experiments on the phantom has shown that the method can be applied to a clinical CBCT scanner. For phantom data, where the surrogate-driven mo-

Table 3.2: SI and transverse absolute errors, given in mm at the machine isocentre (centre of CBCT rotation). SI absolute errors are displayed first, followed by transverse errors. Mean and maximum absolute errors given, with maximum absolute errors in brackets.

Scan ref.	Error relative to isocentre	Error relative to calculated mean	Error with motion model tracking
1A	2.03 (5.42), 2.88 (6.67)	1.58 (4.71), 2.20 (4.88)	1.61 (4.51), 2.20 (5.00)
1B	1.72 (3.96), 3.53 (7.92)	1.70 (3.77), 2.70 (6.51)	1.69 (3.88), 2.70 (6.61)
2A	3.60 (9.48), 1.56 (3.96)	3.04 (7.34), 1.66 (4.90)	1.95 (4.87), 2.17 (4.43)
2B	2.96 (9.06), 1.42 (4.66)	2.44 (6.88), 1.41 (5.63)	1.18 (3.58), 1.19 (5.21)
3A	17.4 (40.4), 2.17 (6.23)	6.48 (22.0), 2.22 (5.05)	3.99 (7.50), 2.48 (4.95)
3B	7.93 (16.0), 2.10 (5.53)	4.06 (8.16), 1.99 (5.84)	3.18 (6.83), 2.04 (5.61)

Table 3.3: Observer variation evaluated at the isocentre. Tumour cannot be delineated indicates proportion of projections clincian was unable to successfully delineate the tumour.

Scan ref.	Observer variation mean (maximum)	Tumour not sufficiently visible to be delineated
1A	1.88 (4.65), 0.984 (2.75)	67%
1B	1.27 (3.87), 1.15 (3.36)	66%
2A	0.760 (2.40), 1.21 (3.79)	62%
2B	0.698 (2.06), 1.56 (7.55)	57%
3A	1.45 (3.78), 1.22 (2.75)	72%
3B	1.36 (3.78), 0.96 (2.75)	72%

tion model is an accurate reflection of the actual motion, the approach is able to improve the MCR and recover the motion model parameters. Although the intention is for this method to be used to allow for markerless tracked RT treatments, this method could offer benefits in marker-based treatments. Markers are much simpler to enhance in the projections, only requiring thresholding in the case of gold markers. This means, assuming the tumour position is accurately reflected by the marker position, artefacts present from using the tumour enhancement approach are avoided. By comparing the known measurements of the marker with those taken from the MCR, a measure could be derived as to how well the actual motion of the marker can be modelled by both the motion model and fitted parameters.

The simulated experiment allowed the method to be evaluated on data unaffected by CBCT-specific artefacts. However, the simulation was based on actual patient data and contained fully deformable motion within the tumour region. Since the ground truth was known, 3D errors could be calculated. The tumour position could be determined to within 1 mm on average, with maximum errors brought under 5 mm. These errors are of a clinically useful accuracy (under 5 mm) for motion management in radiotherapy [30].

The method was tested on a number of patient datasets exhibiting tumours in different areas within the lung, moving to varying degrees. The 3D position of the tumour at the time of each projection was not available for a ground truth comparison. Indeed this is a reason markers have proven popular for tracking tumours [61, 50]. To give some assessment of the ability of the proposed method to correctly model the tumour position and motion, the tumour position was manually identified in the projections by a clinical expert. Projections of the

MCR were manually positioned by the clinician to best overlap the tumour on the original projections. Using projections of the MCR allows the tumour position estimated by the model and manually located tumour position to be directly compared, providing an indication of the relative misalignment in the projections. The misalignment between the tumour positions on the projections was scaled so as to give the corresponding misalignment at the isocentre. This gives a more meaningful estimate of the error as the tumour is likely to be located near the isocentre.

Figure 3.8 shows the improvement in edge contrast of the tumour region after motion correction. Qualitatively, this suggests that the fitted motion is reflective of the actual tumour motion. The clinical delineations can be used to quantitatively assess the motion estimates made by the model. In the SI and transverse directions, all of the maximum errors were within a clinically useful threshold of 5 mm, if the maximum observer variations of each scan were taken into account. This also includes one case where the patient took a sharp breath, possibly corresponding to a cough.

From previous studies of respiratory motion during SABR, it was hypothesised that gated treatments may only be necessary for cases where the tumour motion is larger than 10-15 mm [25]. It was deemed that treatment at the time-weighted mean tumour position was adequate for the majority of cases. The presented technique could be used to determine the time-weighted tumour position, requiring a shorter acquisition sequence than that needed for 4D-CBCT. The extent of tumour motion could then be assessed and used to determine if treatment at the mean position, or if gated or tracked treatment is most suitable. The observed motion could also be used to calculate the maximum possible duty cycle [41] to

reduce the error in tumour position to acceptable tolerances.

The clinician was asked to repeat the alignment to understand the observer variation. In all of the patient cases, error-reductions (relative to manually identified tumour positions) support the assumption that the tumour positions could be estimated by the model well, given the observer variation. However, it is acknowledged that the presented delineations were only repeated once, and all were completed by the same clinician, on markerless data. This was the best possible analysis given constraints on available data and clinician delineation time. The mean and maximum errors in manually locating the tumour of up to 2 and 8 mm, respectively, can be explained by the difficulty of identifying the tumour edges. In the majority of projections, the clinician was unable to delineate the tumour. Given the difficulties with identifying the tumour in the projections, the reported inaccuracies do not seem unreasonable.

For sequential scans of the same patient, tumour motion does remain reasonably consistent. This would be expected, given the patient breathing should not change dramatically between subsequent scans in the majority of cases. The three patients exhibited different types of tumour motion and locations, demonstrating the ability of the surrogate-driven motion model to be used in a variety of clinical situations. For the patient with an upper lobe tumour, the method predicted small surrogate-dependent motion, which was verified on the planning 4DCT.

For patient 2, tracking in the transverse direction provided an interesting insight into the effect on tumour position of a shift applied after the first scan. The shift was done to better align the tumour with the target volume prior to treatment. For this patient, this shift had the effect of better aligning the tumour with the machine isocentre, in the LR and AP directions (see Figure 3.9). If just

the manually located tumour positions are used, it would not be possible to draw this conclusion.

The method presented uses an optical tracking system to provide a surrogate trace. It should be noted that a surrogate obtained by other means could also be used, for example extracted directly from the projections themselves [77, 29]. The optical tracking system was adopted as it provided a consistent 3D reference surface from which to produce a surrogate. For surrogates obtained via projections, it is not trivial to determine the 3D position of a projected feature, hence extracting a surrogate unaffected by projection angle from them is difficult. The optical tracking system also utilises multiple pods, each producing a surface which is stitched together. This is beneficial as a stable surrogate signal can be obtained even as the gantry rotates, obscuring the view of select pods.

The limited field of view of the CBCT led to incomplete reconstructions of transverse slices of the patient. This, in turn, introduced artefacts when generating the enhanced tumour projections. Because the scan is centred on the tumour, the missing anatomy is on the side of the patient furthest from the isocentre, which included parts of the chest, ribs, and lung. Most of the artefacts were introduced in lateral projections of the CBCT. To improve the quality of data available, it is recommended that the patient be scanned so as to allow complete transverse slices. This would include offsetting the detector further to increase the field of view of the reconstruction, and aligning the centre of the patient with the isocentre (rather than the tumour). However, even with artefacts in the enhanced tumour projections, a stable estimate of the motion model parameters was attained. To check convergence, extra iterations were run after the stopping criteria had been reached and the motion model parameters remained stable. The

authors attribute the convergence properties to the form of the motion model (rigid translation-only) and global matching (using all projections to update the parameters) technique. The form of the motion model is to be extended to more complex deformations, such as non-rigid motion. This would allow the rest of the patient, including the organs at risk, to be monitored during gated and tracked treatments.

The limitations of using a rigid-translation only motion model are acknowledged. Although this is suitable for modelling the tumour motion, it cannot model the motion of other parts of the anatomy that move differently to the tumour. For example the ribs in Figures 3.6 and 3.8c appear more artefact afflicted in the MCRs as they have had the tumour's motion imposed on them. In the following chapter, this methodology is significantly extended to incorporate non-rigid deformations. This not only allows for potentially non-rigid motion within the tumour region, but allows the motion of the whole region imaged by the patient to be modelled.

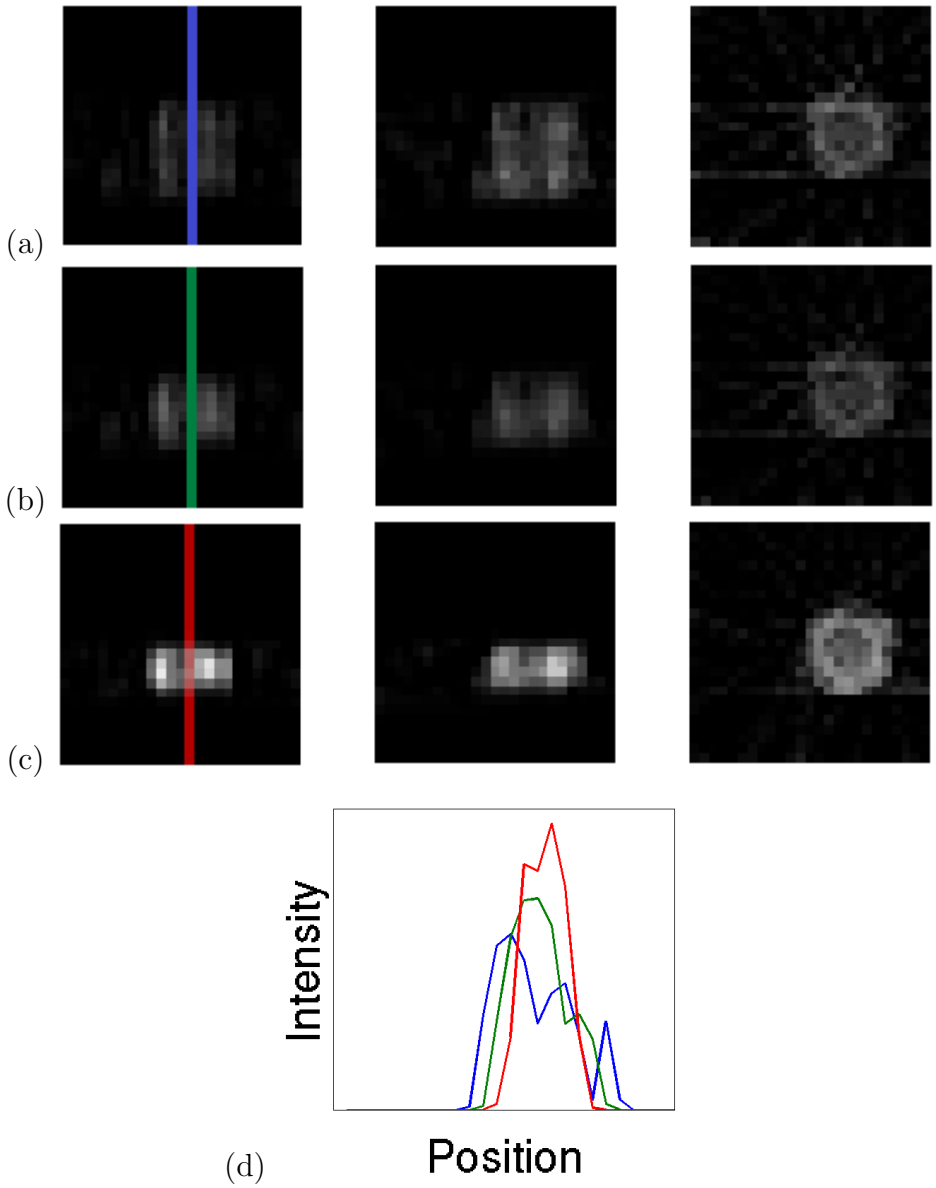

Figure 3.5: Slices though MCRs for phantom data with (a) no motion compensation (b) 1 iteration of the algorithm and (c) 3 iterations. (d) Edge profiles along indicated lines.

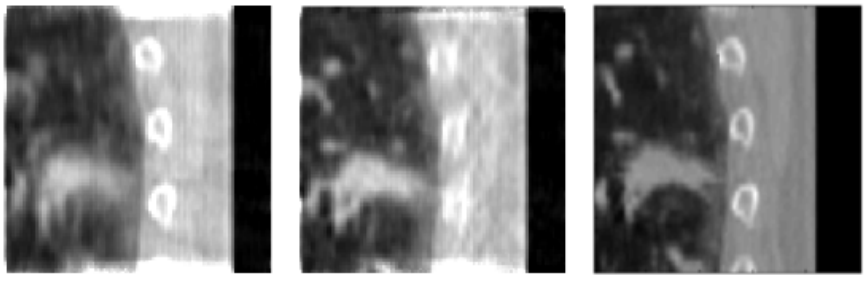

Figure 3.6: Sagittal slices of reconstructions around tumour region before (left) and after (middle) motion correction. For comparison, the ground truth used in the simulation is shown (right).

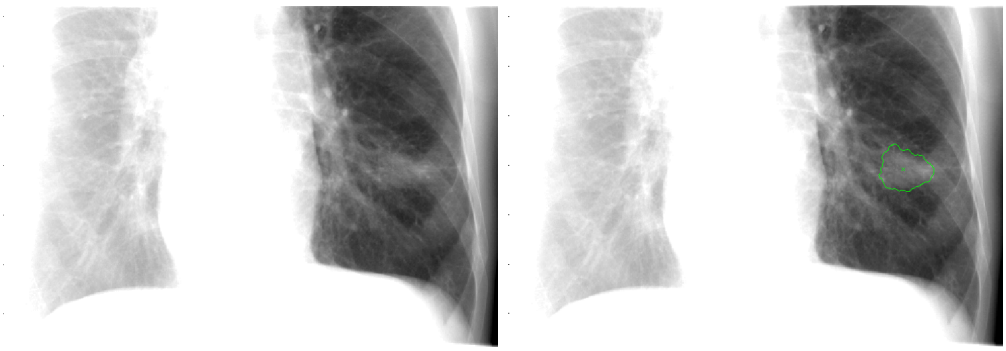

Figure 3.7: Example CBCT projection before and after alignment of projected MCR region by clinician.

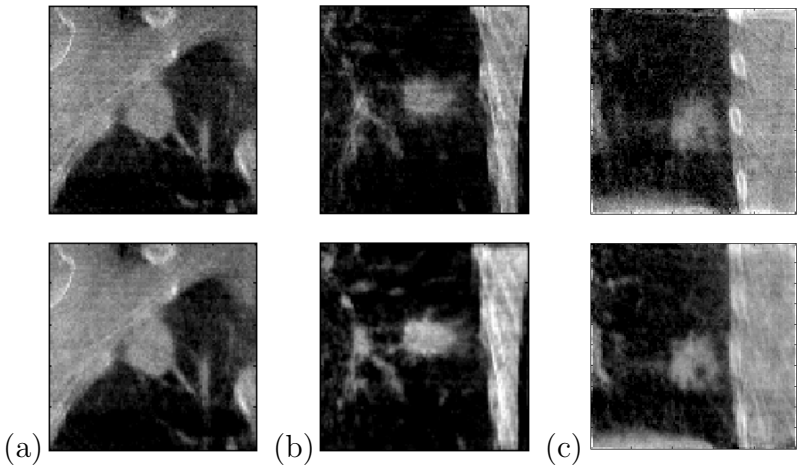

Figure 3.8: Coronal slices of tumour region reconstructions, using original projections with (bottom rows) and without (top rows) motion compensation, for select scans. Patient 1 (a) had a left upper lobe tumour with little respiratory motion. Patients 2 (b) and 3 (c) both had right lower lobe tumours.

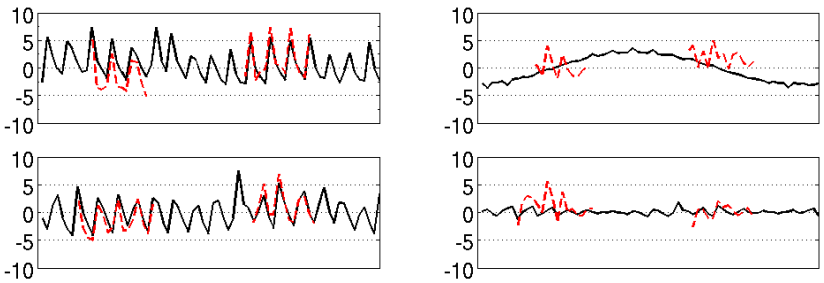

Figure 3.9: Motion estimated (solid black) and manually identified (dotted red) tumour positions in SI (left) and transverse (right) directions, for first (top) and second (bottom) scans of patient 2. The patient was shifted for better tumour alignment between the scans. Units are in mm, scaled from projection to a parallel plane passing through machine isocentre.

Chapter 4

Moving to a whole patient motion model

Upon achieving promising results, it was natural to consider how to extend the model beyond that of simply the tumour region. I completely developed the presented approach to building a fully deformable motion model, which attempts to balance the complexities required from a fully deformable motion model with the limitations of CBCT in the clinical environment. This includes understanding that one only has a finite time with which to optimise the parameters of the motion model (on the day of treatment) and, again due to restrictions on scanning time, often needs to work with a limited number of CBCT projections. The results of this work has led to first author conference and workshop publications, detailed in chapter 7. Future work has also been influenced by the developments made in this chapter.

4.1 Introduction

Although knowledge of tumour motion is the most important feature to understand, other regions of the patient can offer critical insight during RT treatment. One of the key regions which is also important to track during RT treatment are the OAR. These regions have been identified during planning and their exposure limits were a key aspect in determining the most effective dose plan to prescribe to the patient. Unlike the tumour region, motion of the OAR cannot be encapsulated with a rigid-translation only motion model. This is because regions such as the lungs cannot be described as a solid, non-deforming mass.

In this chapter, the previously described translation-only model is modified and extended to a fully deformable one. Previously only applicable to the tumour region, this new approach can now be used to track the tumour region and OAR. This approach can now retain many of the advantages offered by the previous approach, but now accommodate for the demands of a fully deformable motion model of the whole patient. With a fully deformable motion model, rich detail on patient respiratory motion can be derived, and for example with tracked treatments, it can be ensured that beam modifications to track the tumour region do not over-expose the OAR.

A surrogate-driven motion model is used to overcome the difficulties associated with real-time imaging of the internal anatomy. Once the model has been built, only a new surrogate measurement is needed to estimate the patient state. The surrogate can be monitored easily throughout treatment, avoiding implanting markers or taking additional scans. The motion model is posed in a way which exploits the relative strengths of the datasets most often available during

RT treatment: The planning 4DCT and positioning CBCT. The higher quality planning 4DCT is used to construct what will be termed weighting arrays. The weighting arrays are essentially deformation fields within the model, which describe how the anatomy of the patient could vary with the surrogate signals. The motion model parameters (after being weighted by the weighting arrays) are then fitted to the CBCT, which can account for the inter-fraction variation that occurs between planning and the day of treatment. The motion model parameters are global in nature (same over the whole volume). This means the number of model parameters to fit on the day of treatment is low and they can be quickly found. The motion model is also robust against many of the previously mentioned problems affecting CBCT. In addition, delineations from planning can be imported onto the model, providing tumour and OAR boundaries based on those used to construct the plan.

4.2 Methods

4.2.1 Modifying the model to account for non-rigid motion

During the CBCT scan, it is assumed that the patient motion can be described by a series of deformations, F_n, of a motion-free image of the patient, $V_{ref}(x) : x \in \Omega_{CBCT}$. Ω_{CBCT} is the region imaged during the CBCT scan and $n = 0, 1, ..., (N-1)$ is an index corresponding to the time of each of the N projections. The deformed patient volume, V_n, is then:

$$V_n(x) = V_{ref}(x + F_n(x)). \tag{4.1}$$

A variety of non-rigid motion models [47] have been proposed in the literature. For this work, a non-rigid motion model is required which can account for hysteresis and changes in depth and length of breathing cycle. A non-rigid motion model, similar in form to that used to for the tumour region-only motion model, is chosen:

$$F_n(x) = s_n D_1(x) \circ \mu_1 + \dot{s}_n D_2(x) \circ \mu_2, \tag{4.2}$$

where $\circ$ is an element-wise multiplication (Hadamard product). μ_1 and μ_2 are the motion model parameters (3D vectors). s_n and $\dot{s}_n$ are the scalar surrogate signal associated with the n^{th} projection, and its rate of change, respectively. D_1 and D_2 are weighting arrays for μ_1 and μ_2 respectively. The weighting arrays are deformation fields, returning a 3D vector at a point x in the patient volume. The use of weighting arrays allows for spatially-dependent variations within the model, which are essential for a fully deformable motion model of the patient. The weighting arrays are pre-built from the planning 4DCT and remain fixed during fitting of the motion model parameters. It is the final fit of the motion model parameters that gives the day-specific estimate of the motion. The motion model parameters are independent of position within the patient volume, hence are global in nature.

One could think of (4.2) as a set of global, motion model parameters which are spatially weighted by the weighting arrays. The motion model parameters describe how key features of the respiratory variation (e.g. hysteresis, depth of breathing) of a whole patient correlate with the surrogate signals, as seen in the

CBCT projections. The weighting arrays then use this information to generate a fully deformable motion model, which is based on anatomical variation seen with respiration in the planning 4DCT.

4.2.2 Global/local approach to describing respiratory deformation

As far as the author is aware, decomposing the non-rigid motion model into local and global components as in (4.2), in each spatial dimension, is conceptually new. Because of this it is worth discussing this model in more detail.

The weighting arrays control the relative motion between different patient anatomy, ultimately realising the non-rigid motion of the patient with changes in the surrogate signals. The weighting arrays achieve this by specifying the distance each part of the anatomy moves relative to the other, in each spatial dimension. Therefore, the weighting arrays can be thought of as providing a local description of the motion as they are dependent on spatial position.

The quality of the planning 4DCT is relatively high, making this scan a good candidate to build the local description of anatomical variation with the surrogate signals empirically. There are two weighting arrays, one to describe the local variation with each surrogate signal. The magnitude of motion has not been completely detached in their current implementation, although using normalised weighting arrays would also be possible. The magnitude information is present as the weighting arrays are built from a series of registrations between phases of the planning 4DCT, and still contain physically relevant deformation information. However, to account as much as possible for some of the inter-fraction variations

(specifically baseline shifts, anatomical changes and some expected changes in respiratory motion because of these changes) a registration between the average 4DCT and CBCT is used to move the weighting arrays to CBCT space.

The motion model parameters are fitted to the CBCT, and allow changes of the respiratory motion, beyond that described in the local-based weighting arrays, to be made. The changes to the respiratory motion that these parameters impose are global in nature, as each parameter has the same value throughout the patient volume. They act to essentially rescale, in each individual direction, the weighting arrays over the whole volume simultaneously. These global parameters may then give an indication of some more general features of the respiration seen on the day of treatment. This could include the magnitude of the hysteresis effect, or the depth of breathing, as the surrogate varies. Figure 4.1 shows an example case of the effects of varying the motion model parameters and weighting arrays.

Although providing up-to-date information describing the daily motion and anatomy of the patient, CBCT data poses a variety of problems. One key issue is only projection data being available. Unlike phases of the 4DCT, whole volumes are not available, with instead a long series of reduced-dimension (2D) X-ray projections. For each projection, this effectively removes the information available in the source-to-detector (isocentre) direction. For each projection, it would be very difficult to have a well-enough constrained, local approach to building the motion model empirically. For this reason, it makes sense to fit parameters to the CBCT projection data which are well-enough constrained. The global nature of the motion model parameters suitably constrain the fitting. All the projections can therefore provide an insight into the respiratory motion seen in each dimension, avoiding problems caused by having projections instead of complete

volumes. This approach also adds robustness to the fit. Over the whole volume, the sensitivity of the method to other issues affecting CBCT reconstructions, such as artefacts (e.g. cupping, streaking), scatter and incomplete field-of-view, is reduced. On the day of treatment, this approach also allows for a low number of parameters to be fitted. This greatly reduces optimisation time as the algorithm used only needs to search a much reduced parameter space. The resulting cost function to optimise is convex - it is of a positive-definite quadratic form with a global, unique minimum. This removes the possibility of getting stuck in local minima.

Having a fast-to-optimise cost function which can be robustly fitted to CBCT potentially offers a promising way to meet clinical requirements. The clinical implementation of a next generation, tracked RT treatment system will need to meet a variety of demands, such as quick model fitting on the day of treatment, and if possible, a stepped approach to increasing the ability of current treatments. Easy integration to current clinical protocol will also increase the target audience, as it is more favourable to upgrade existing systems and keep changes to existing protocols minimal. Another point to note is that ideally one would simply upgrade the CBCT to a different modality, for example 4DCT. However, clinical limitations (e.g. cost, space available) prevent such a solution. The CBCT modality is well integrated into, and is a critical element of, current RT treatments. This will continue to be the case for the foreseeable future due to its popularity, practicality and relatively low cost.

4.2.3 New form of the cost function

A similar approach to that used to determine the parameters of the tumour only motion model can be taken with this newly modified motion model. Explained in the next section, the weighting arrays have been built and moved into the space of the CBCT scan to then fit the motion model parameters.

As with the tumour only motion model, for an estimate of the motion model parameters, a method is needed to update them. Similar to the derivation of the cost function used to update the parameters of the tumour-only motion model, equations (3.3,3.4,3.5) remain valid. Note, however, that the motion model is now applicable to the whole of the patient. Therefore V_{ref} is now referring to the whole patient volume imaged during the CBCT scan rather than just the tumour region. V_n is then the animated patient volume, using the model described by (4.1). This changes the equation used to update the parameters to:

$$\delta\mu = \underset{\delta\mu}{argmin}\left[\sum_{n}\sum_{pixels}\left(R_n - P_n\left(\nabla V_n \cdot \delta F_n\right)\right)^2\right], \tag{4.3}$$

where δF_n now takes the following form:

$$\delta F_n\left(x\right) = s_n D_1\left(x\right) \circ \delta\mu_1 + \dot{s}_n D_2\left(x\right) \circ \delta\mu_2. \tag{4.4}$$

Substituting (4.4) into (4.3) and algebraically expanding (see Appendix A for a detailed derivation), the cost function now becomes:

$$\mu_1, \mu_2 = \underset{\mu_1,\mu_2}{argmin}\left[\sum_{\theta,\phi=0}^{6}\lambda_\theta\lambda_\phi\sum_{n}\sum_{pixels}C_\theta C_\phi\right], \tag{4.5}$$

where

$$\lambda_0 = 1, \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix} = \mu_1, \begin{pmatrix} \lambda_4 \\ \lambda_5 \\ \lambda_6 \end{pmatrix} = \mu_2. \tag{4.6}$$

and

$$C_0 = p_n - P_n(V_n),$$

$$\begin{pmatrix} C_1 \\ C_2 \\ C_3 \end{pmatrix} = -s_n \begin{pmatrix} P_n\left(\nabla_x V_n \circ D_{1,x}\right) \\ P_n\left(\nabla_y V_n \circ D_{1,y}\right) \\ P_n\left(\nabla_z V_n \circ D_{1,z}\right) \end{pmatrix}, \tag{4.7}$$

$$\begin{pmatrix} C_4 \\ C_5 \\ C_6 \end{pmatrix} = -\dot{s}_n \begin{pmatrix} P_n\left(\nabla_x V_n \circ D_{2,x}\right) \\ P_n\left(\nabla_y V_n \circ D_{2,y}\right) \\ P_n\left(\nabla_z V_n \circ D_{2,z}\right) \end{pmatrix}.$$

The x, y, z subscripts of the weighting matrices indicate the relevant component of the vector field.

To determine the various parts of the cost function requires multiple volume deformations. Section 4.2.5 explores determining the updates to the undeformed MCR as an option to greatly reduce the frequency of these calculations. A faster algorithm is useful to make this approach more clinically acceptable.

4.2.4 4DCT, CBCT and moving weighting arrays to CBCT space

4DCT scans are routinely used during planning to outline the tumour region (including respiratory motion seen between phases) and OAR. In this work, 4DCT

is used to build the weighting arrays via a series of registrations. The phases of the 4DCT are referred to as I_{pq}, where p is a percentage and q "In" (inhale) or "Ex" (exhale). For example I_{100In} is the 100% inhale phase of the 4DCT. Note that in this notation I_{0In} is equivalent to I_{100Ex} and I_{100In} to I_{0Ex}.

Since D_1 describes the spatially-dependent weightings of the motion model deformation due to the surrogate signal s, it can be built from registrations between two phases of the 4DCT where $\dot{s}$ is assumed to have the same value. Vice-versa for building D_2, which describes the spatially-dependent weightings due to $\dot{s}$. Here, D_1^{0In} is built from a registration between I_{0In} and I_{100In}. D_2^{50Ex} is built from registrations between I_{50Ex} and I_{50In}. The superscripts indicate that the weighting arrays are in different spaces to what is needed for the motion model.

The weighting arrays need to be moved into CBCT space Ω_{CBCT} prior to fitting the model. This is achieved by first moving the weighting arrays into a common average 4DCT space, Ω_{ave}, followed by finally moving into CBCT space. Note that one could think of Ω_{ave} and Ω_{CBCT} as the coordinate systems of the 4DCT and CBCT scanners, respectively. Assuming we have a homeomorphism $\Psi_{a\to b} : \Omega_a \to \Omega_b$ between the coordinate systems Ω_a and Ω_b, the transformation (each weighting array can be regarded as a transformation) $T_a : \Omega_a \to \Omega_a$ in the coordinate system Ω_a has a topological conjugate $T_b : \Omega_b \to \Omega_b$ if [16]:

$$T_b = \Psi_{a\to b} * T_a * \Psi_{a\to b}^{-1}, \tag{4.8}$$

where $*$ is a function composition. First the homeomorphism $\Psi_{0In\to ave}$, is determined. This is calculated by averaging the deformation fields from registrations of 0In to all the other phases of the 4DCT. D_1^{0In}, currently in the space of Ω_{0In}

is moved into the space of Ω_{ave} via:

$$D_1^{ave} = \Psi_{0In \to ave} * D_1^{0In} * \Psi_{0In \to ave}^{-1}. \tag{4.9}$$

D_2^{50Ex}, currently in the space of Ω_{50Ex} can also be moved:

$$D_2^{ave} = \Psi_{0In \to ave} * \Psi_{50Ex \to 0In} * D_2^{50Ex} * \Psi_{50Ex \to 0In}^{-1} * \Psi_{0In \to ave}^{-1}. \tag{4.10}$$

The final stage of moving D_i^{ave}, $i = 1, 2$ into the space of Ω_{CBCT} is performed on the day of treatment. The average 4DCT image, I_{ave}, is registered to a standard reconstruction from the CBCT (with an initial rigid registration) to determine $\Psi_{ave \to CBCT}$, with

$$D_i^{CBCT} = \Psi_{ave \to CBCT} * D_i^{ave} * \Psi_{ave \to CBCT}^{-1}. \tag{4.11}$$

Figure 4.2 provides a diagrammatical representation of the various registration steps. For the final registration step, the standard CBCT reconstruction is assumed to be an approximation to the average image [35].

Registrations are critical to building the weighting arrays and moving them to the space of the CBCT scanner. Therefore, description of the non-rigid local motion hinge on robust registrations. Intra-modality are required between different phases of the 4DCT. An inter-modality registration is also required, namely between the average 4DCT scan and the CBCT reconstruction.

An open-source software package [48], was used for the registrations and to calculate inverses of the deformation fields. The registrations are based on B-

splines [58] with a control point spacing of 5 mm. Bending energy was used as a constraint and normalised mutual information as the similarity measure. All the resulting deformation fields had a positive Jacobian determinant throughout and were assumed to represent a homeomorphism. To avoid errors being introduced by artefacts and inaccurate registration results near the edges of the reconstruction, a mask was applied to points in the weighting arrays within 10 voxels from the edges of the reconstructed volume.

The approach could also be used with a variety of measures other than normalised mutual information. For inter-modality registration, this measure may be suitable, but for intra-modality registrations, an intensity-based measure may offer a useful alternative. Another interesting alternative would be a registration based on biomechanical constraints (e.g. finite element driven registrations).

4.2.5 MCR-based update to the parameters, form of cost function

The full cost function (4.5,4.6,4.7) presented in chapter 4.2.3 requires deformations be taken for the time of every projection. These constant deformations increase the computational workload. In order to greatly speed up the time taken to solve the cost function, using a comparison with projections of the MCR is proposed. Taking this approach means the whole model parameters will be calculated at every iteration, as opposed to updates to the previous estimate of the parameters.

By taking projections through the unanimated MCR, the cost function simplifies to the following:

$$\mu_1, \mu_2 = \underset{\mu_1,\mu_2}{argmin} \left[\sum_n \sum_{pixels} \left((p_n - P_n(V_{ref})) - P_n(\nabla V_{ref} \cdot F_n) \right)^2 \right], \tag{4.12}$$

where P_n is the forward projection operator (n^{th} projection), given the CBCT geometry, and p_n the actual CBCT projections. Algebraic expansion of (4.12) allows the parameters to be factored out of the projection operator and summations. The cost function remains of a positive definite quadratic form and all the sums over projections can be calculated prior to the optimisation step.

$$\mu_1, \mu_2 = \underset{\mu_1,\mu_2}{argmin} \left[\sum_{\theta,\phi=0}^{6} \lambda_\theta \lambda_\phi \sum_n \sum_{pixels} C_\theta C_\phi \right], \tag{4.13}$$

where

$$\lambda_0 = 1, \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix} = \mu_1, \begin{pmatrix} \lambda_4 \\ \lambda_5 \\ \lambda_6 \end{pmatrix} = \mu_2. \tag{4.14}$$

and

$$C_0 = p_n - P_n(V_{ref}),$$

$$\begin{pmatrix} C_1 \\ C_2 \\ C_3 \end{pmatrix} = -s_n \begin{pmatrix} P_n(\nabla_x V_{ref} \circ D_{1,x}) \\ P_n(\nabla_y V_{ref} \circ D_{1,y}) \\ P_n(\nabla_z V_{ref} \circ D_{1,z}) \end{pmatrix}, \tag{4.15}$$

$$\begin{pmatrix} C_4 \\ C_5 \\ C_6 \end{pmatrix} = -\dot{s}_n \begin{pmatrix} P_n(\nabla_x V_{ref} \circ D_{2,x}) \\ P_n(\nabla_y V_{ref} \circ D_{2,y}) \\ P_n(\nabla_z V_{ref} \circ D_{2,z}) \end{pmatrix}.$$

The x, y, z subscripts of the weighting matrices indicate the relevant component of the vector field. The same software package used to calculate the MCR was modified for the forward projections.

With this form, it is possible to determine the various aspects of the cost function quickly. For example, projections of the undeformed MCR are taken, completely removing the need to deform the patient volume at the time of each projection. In the previous version of the cost function, projections of the gradient of the deformed volumes, pre-multiplied by particular components of the deformed weighting arrays (composed with the deformation), were needed. These now become projections of the gradient of the MCR pre-multiplied by undeformed weighting arrays. As it is essentially one, unchanging, 3D array which is being projected, approaches tailored to quickly simulating a CBCT of one volume can be modified for calculating various parts of this cost function. For this work, a GPU-accelerated forward projector originally tailored to simulating CBCTs of a patient volume was modified to determine the projections required for this cost function.

WIth the GPU accelerated version, one CBCT acquisition consisting of around 350 projections can be simulated in under 15 seconds. If the patient volume needs to be deformed prior to each projection, this time is increased to 3-4 minutes.

4.2.6 Iterative approach

As with the tumour only approach, an iterative method is taken to simultaneously determine the MCR and motion model parameters. The MCR is set as the approximation to V_{ref}, which should become more accurate as the parameters produce an MCR which is more representative of V_{ref}. An open-source, FDK-

based software package (www.openrtk.org) was modified to calculate the MCR. MCRs are performed by back-projecting each projection through a deformed volume. The deformation to use should be the inverse of the forward transformation from the motion model (i.e. F_n^{-1} from (4.2)).

An initial MCR is calculated assuming there is no motion (i.e. a standard reconstruction is performed). This is then set as V_{ref} and new estimates for μ_1 and μ_2 are calculated. The new estimates of the motion model parameters are then used to recalculate the MCR, giving an improved estimate of V_{ref}. The process is then iterated until the convergence criteria are met. The algorithm is terminated if the absolute change over all the parameter components are less than 5% of their previous values, or had an absolute change of less than 0.005. With reasonable assumptions (motion model parameters are all of the order of 0.5, maximum respiratory motion of around 20 mm), the termination conditions correspond to changes less than 0.5 mm. The absolute threshold is used to avoid changes in smaller parameter values (which will have a larger percentage change) unnecessarily postponing convergence.

4.2.7 Masking the weighting arrays

In this approach, to avoid artefacts, or inaccurate registration results, near the edges of the reconstructed volume overly effecting results, a masking step was used. Rather than mask the volume itself, the weighting arrays were instead masked. This is essentially a switch to control which regions in the reconstructions are motion corrected, and which are not. (i.e. by masking the weighting arrays, the anatomy in the masked positions will not be motion corrected when calculating the MCR.) However, when the projections are taken, the anatomy will

still be present, albeit with motion blurring. When determining the parameter updates, differences between estimated and actual projections are taken. With the anatomy still being present in the estimated projections, the subtracted projections will have a closer match than if it was not present. Therefore with this approach, without the benefit of motion correction, it is still possible to remove this masked anatomy from the projections, which should improve the appearance of differences in non-masked regions (which is the region of the patient that is of interest for fitting the motion model parameters).

Although in the ideal situation one would want motion correction in this region, this approach stikes a balance between how much the artefacts near the edges of the reconstruction will affect the fit of the motion parameters, versus the improvement from suitable motion correction in this region. Figure 4.3 shows the effects of motion correction with masked weighting arrays. Reconstructions with and without motion correction are shown for an example sagittal slice through the reconstruction, with the used masked delineated in red.

There are two interesting insights from using this masking approach: (i) If the mask region was reduced to just the tumour region, and all the elements of the weighting arrays set to one (rigid motion only), this approach simplifies to the tumour-only case. This approach could therefore be considered as a more theoretically complete version of the rigid approach, which can account for fully deformable motion. (ii) With the current implementation, the motion model parameters describe some of the global properties of the patient respiratory motion. One approach to make the parameters specific to a particular region, say OAR, would be to set the mask to delineate that OAR. Therefore the motion model parameters may be able to describe how that particular OAR deforms with the

surrogate signal, which may be more physically realistic.

4.3 Results

4.3.1 Simulated results

For this work, simulated data was used to assess the error reduction possible after model fitting, as the true motion was known. CBCT acquisitions were simulated using motion models built from actual patient 4DCT datasets. I_{0In} was transformed to the space of Ω_{ave}. The weighting arrays were in the space of Ω_{ave}. The motion model parameters were then manually set to 0.5, inducing respiratory motion approximately equal to that seen in the 4DCT. The I_{0In} 4DCT volume (moved into average 4DCT space) is used as the reference volume, and is deformed by the motion model to simulate the motion seen during a CBCT acquisition. A CBCT was then simulated from this, with the simulated motion being driven by a real patient surrogate trace.

Simulated data of four patients were produced. There were variations over the patients, including size, breathing pattern and tumour/OAR positions.

Three iterations of the algorithm were required to fit the models. Visible reductions in blurring were seen in the reconstructions after motion compensation. Figure 4.4 shows reconstructions using the fitted motions. Table 4.1 shows the mean and maximum L2 norm errors before and after fitting the model for the simulated cases. The errors are relative to the mean positions of the anatomy. Tumour, lower lung and sternum regions were manually identified in the reconstructions, for more localised analysis. The lower lung region is the lower part of the ipsilateral lung. Over all simulated cases, the tumour motion was found

Table 4.1: Mean (maximum) L2 norm errors for simulated datasets, in mm, for various regions in the patient. Results are given with (MC) and without (-) motion correction. For no motion correction, the anatomy is assumed to remain at its average position. The various regions were manually identified in the reconstructions. Lower lung refers to the lower part of the ipsilateral lung.

Patient	Tumour	Lower lung	Sternum	Whole volume
1 -	2.10 (6.21)	2.94 (18.2)	0.645 (2.28)	1.39 (29.9)
1 MC	0.401 (1.37)	1.32 (7.29)	0.216 (0.730)	0.479 (12.8)
2 -	4.15 (12.9)	3.99 (16.7)	2.51 (11.6)	1.41 (24.0)
2 MC	1.09 (3.50)	1.02 (4.61)	0.711 (3.22)	0.371 (6.58)
3 -	4.48 (14.8)	3.79 (11.9)	0.369 (1.95)	0.570 (21.8)
3 MC	1.37 (4.86)	0.866 (3.96)	0.101 (0.646)	0.178 (7.11)
4 -	0.444 (5.02)	1.84 (12.9)	0.404 (2.56)	0.534 (22.5)
4 MC	0.147 (1.78)	0.667 (4.76)	0.108 (0.678)	0.172 (8.38)

with a mean (max) L2 error over the entire CBCT of 0.752 (4.86) mm. Without motion correction the L2 errors were 2.79 (14.8) mm. For the lower lung regions, motion correction reduced the L2 errors from 3.14 (18.2) mm to 0.968 (7.29) mm. For the sternum, L2 errors were reduced from 0.982 (11.6) mm to 0.284 (3.22) mm.

4.3.2 Clinical results

4DCT planning scans and CBCT scans were collected from four patients undergoing SABR treatment at Guy's and St. Thomas' NHS Foundation Trust. An Elekta Synergy (Elekta Limited, Crawley, London, UK) treatment machine, with a CBCT scanner orthogonally mounted to the treatment beam, was used for the SABR treatment. During the CBCT scans, synchronised 3D skin surface data was acquired from Align RT, a stereo camera system (Vision RT, London, UK). These surfaces were then converted to the raw 1D surrogate signal by taking the average height between the couch and skin surfaces, within a manually defined bounding box covering the chest and abdomen. The raw surrogate is smoothed and normalised (mean subtracted; divided by standard deviation), yielding s_n.

$\dot{s}_n$ is calculated using a finite difference approach, and has been normalised with its own mean and standard deviation.

Data from four patients were used in this study. These were the same patients as those used to generate the simulated data. All CBCT scans took 1.5 minutes, acquiring around 350 projections over a 360 degree rotation. Due to the dimensions of the CBCT scanner, it was not possible to reconstruct complete transverse slices for any of the patients. Therefore all the scans had patient anatomy lying outside the reconstructed volume which was visible in the CBCT projections. One CBCT scan (patient 3) was particularly affected by scatter due to patient size.

Unlike the simulated case, the true motion of the tumour and OAR are not known. Indeed, attempting to understand the full motion is one of the fundamental drivers for this work in the first place. However, one indicator that the motion during the scan has been predicted is achieving a sharper MCR, especially when compared to assuming no motion. Regions of higher contrast, such as tumour or OAR edges should have less motion blurring. Some examples of this are seen in Figures 4.5. Accompanied with this there should be a higher signal to noise ratio. One way to allude to this is to look at regions where one expects a constant value, such as some parts of the tumour, rib, lungs and fatty tissues. Figure 4.6 shows how the signal over the tumour region has a sharper transition at the tumour boundaries, combined with a more constant value within the tumour up to the boundary.

Over all the patients, the algorithm converged within five iterations. Visible improvements in the MCRs were seen, such as increased edge contrast near the diaphragm, chest wall and tumour regions. Figure 4.5 shows examples of these

improvements, including the patient affected by large amounts of scatter.

Figure 4.6 shows surface plots of a central plane through the tumour regions of the two patients which exhibited the greatest motion. After motion correction, both cases exhibit a larger signal within the tumour region, particularly near the edges. Higher edge contrast is also seen after motion correction. Line profiles through the diaphragms of each patient's ipsilateral lung show sharper transitions between the lower lung and diaphragm after motion correction.

4.3.3 Approximation to inverse deformation field of model

By updating the whole motion model parameter from comparisons with projections of the MCR (rather than animated volume), simulations of the projections can be GPU accelerated. However, at this stage, calculation of the MCR is not GPU accelerated, and is the most time-consuming part of each iteration of the algorithm. Much of the time in determining the MCR is used in determing the inverse transform of the motion model at the time of each projection.

One interesting approach to significantly reduce the time taken to calculate the MCR can be made by simply approximating these inverses with the negative of the forward transformation. This, strictly, introduces errors into the reconstruction, especially for large deformations. This is due to the sampling point in the deformation field to determine the actual inverse not being the same as the point in the deformation field at which the inverse deformation is required. The only time these would be coincident is for zero deformation. Figure 4.7 visually demonstrates how assuming a negative forward approximation to the inverse introduces errors, compared to the actual inverse.

Although Figure 4.7 demonstrates how errors can be introduced into the cal-

culation of the MCR with increasing size of deformation, for the actual datasets used, the introduced errors stayed relatively small. Interestingly, for the simulated data, this approximation was able to determine the fitted motion across all cases with an L2 norm maximum error of under 0.836 mm, compared to using the full inverse. The mean error, over all simulated cases with this approximation was 0.0194 mm. For the actual data, using the approximated inverse found the parameters within 1.46 mm, with a mean error of 0.00605 mm.

One possible explanation is that over the whole CBCT, the occurances of extremal deformations are relatively small, compared to the number of smaller deformations. Smaller deformations sample the deformation field close to the actual position needed to determine the true inverse value, hence resulting in a small error with the negative forward transformation approximation. The resulting MCR from deformations of each back-projections using these approximated inverted deformation fields are therefore close to the MCR when using the actual inverse deformation field. The approximated MCR will therefore result in fitted motion model parameters which are close to those fitted with the actual MCR.

One clear benefit of using the negative forward deformation approximation is a significantly reduced MCR calculation time. For a CBCT scan consisting of around 350 projections, and reconstruction size 300x300x150 $1mm^3$ voxels, the MCR can be calculated in around 3 minutes. If the full inverse is to be fully determined for each back-projected volume, this time is significantly increased. Using an ITK-based iterative method to determine the inverse resulted in a calculation time of over 30 minutes. This approach then, may offer a good way to practically reduce the time to calculate an MCR without overly compromising on results. It may also be a good way to quickly build a first guess at the model,

prior to a more accurate (but longer to calculate) update.

4.4 Discussion and conclusions

Key to obtaining an accurate fit is building the weighting arrays and moving them to the space of Ω_{CBCT}. For the final step, a registration between a standard CBCT reconstruction and the average 4DCT is performed. This assumes that the blurring seen in the CBCT is an approximation to the average of all the patient states seen over the CBCT [35]. Since the weighting arrays are derived from sharper images, the registration of two images with blurred edges may result in misalignment of regions of rapid variation in the motion, for example the lung/chest-wall boundary. This issue may be further exacerbated as the MCR becomes more defined with iteration number. After the first iteration, a possible alternative method to that presented could be to derive the final transformation to CBCT space. This would be based on a further registration between a phase of the 4DCT (e.g. I_{0In}), moved to the space of Ω_{ave}, and the MCR. The authors plan to investigate alternative ways to build the weighting arrays and move them into the space of Ω_{CBCT}. Other areas for future research are the use of different modalities to build the weighting arrays and the use of registration techniques that can account for sliding between the lungs and chest wall (e.g. [14]).

Prior to fitting the model, a masking step was used to minimise the effects of artefacts and inaccurate registration results near the edges of the CBCT reconstruction. Instead of masking the MCR itself, the weighting arrays were masked. This produced full CBCT projections whilst masking the derivative-derived projections for the cost function. This approach is advantageous as during the pro-

jection subtraction, regions outside of the mask will have been reconstructed with no motion correction, hence will be included in the simulated CBCT and subtracted from the original projections. This reduces the effects of these regions prior to the model fitting step. Interestingly, if the mask region was reduced to just the tumour region, and all the elements of the weighting arrays set to one (rigid motion only), this approach simplifies to that presented in [28]. This approach could therefore be considered as a more theoretically complete version of the rigid approach, which can account for fully deformable motion.

A common approach for patient set-up is to centre the tumour region on the isocentre of the treatment machine. Due to the limited size of the CBCT detector, this almost always results in incomplete transverse slices of the patient. The effect of incomplete transverse slices is anatomy being visible in the CBCT projections which are outside the field-of-view of the reconstructed CBCT volume. Despite the CBCT having an offset detector for an increased field-of-view, all of our patient datasets suffered from this problem. The presented approach appears to be robust enough to cope with additional anatomy in the projections. This is probably due to the use of the 4DCT data to build the weighting arrays, which has a larger field-of-view than the CBCT reconstruction, and form of the cost function. Better ways to account for this anatomy, for example to centre the CBCT scanner on the patient centre (to obtain images more like that used for the simulated data), or use different reconstruction techniques (e.g. iterative reconstruction [3]) to more accurately account for partially sampled regions, are being considered for future research.

A useful property of the proposed approach is its amenability to fitting a model within a clinically useful timeframe. Due to the form of the motion model, only

the motion model parameters need to be fitted on the day of treatment, which can be done very efficiently (optimising the cost function takes seconds). The most time consuming parts are simulating projections, performing reconstructions and calculating the average 4DCT to CBCT registration. All of these are amenable to speeding up via parallelisation. Work is currently underway to reduce the time taken using GPU technology. With appropriate improvements, the authors expect the model fitting to be completed in under 5 minutes, which is compatible with the current clinical workflow. Depending on the number of iterations and quality required of the MCR, the current implementation can take upwards of 15 minutes to fit the model.

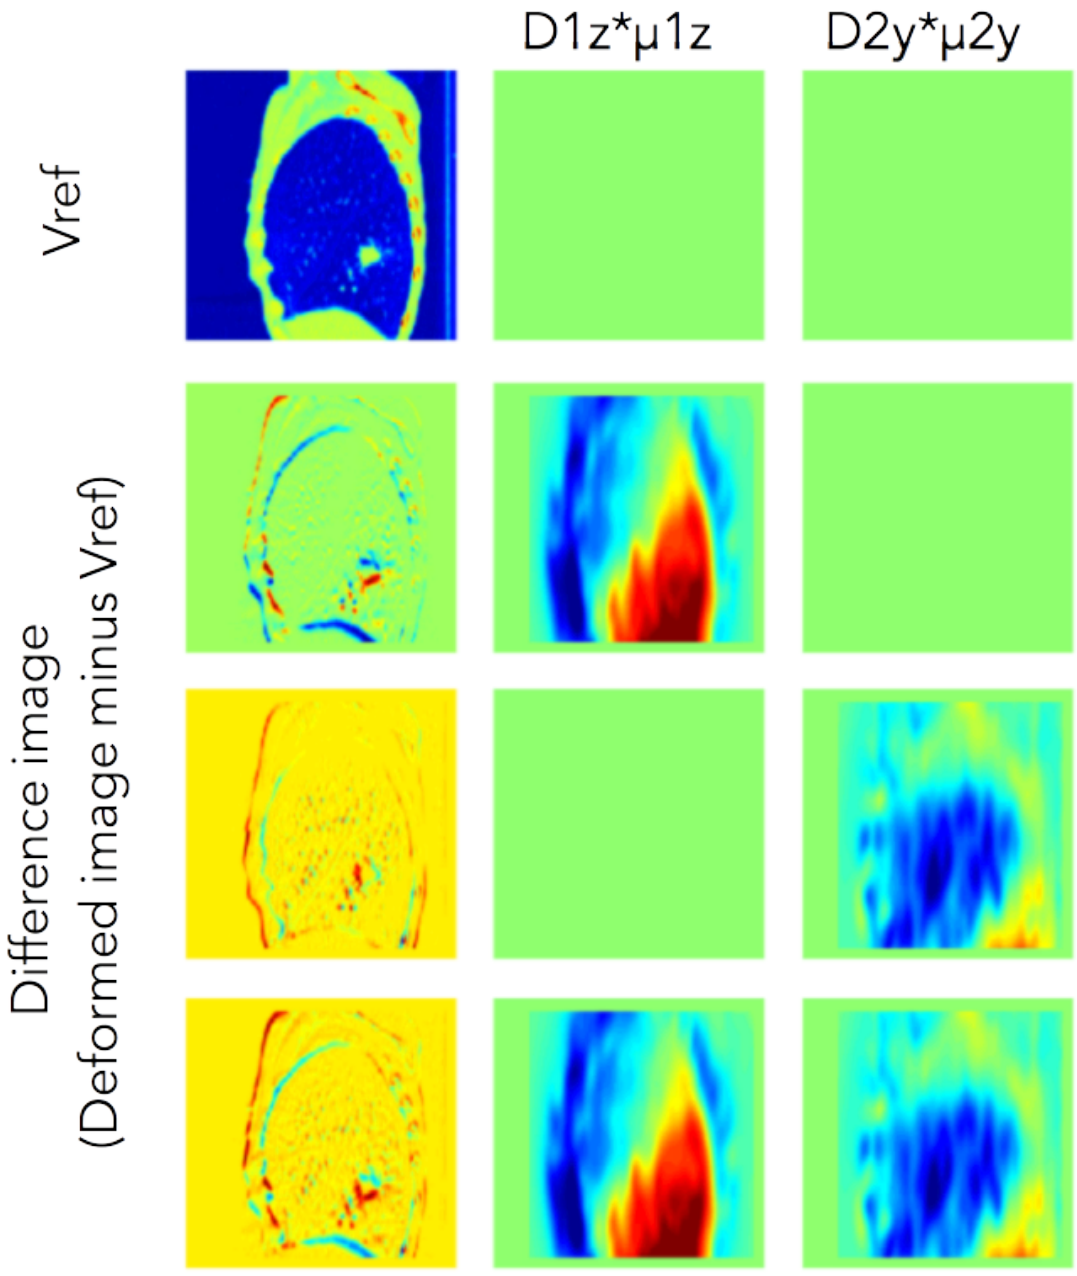

Figure 4.1: Example sagittal slice of patient volume undergoing deformations via weighting array formulation of motion model described in equation (4.2). Reference volume shown in top row, with difference images (between the reference volume) in subsequent 3 rows. In this example, D_1 and D_2 deformations have been limited to the SI and AP directions, respectively.

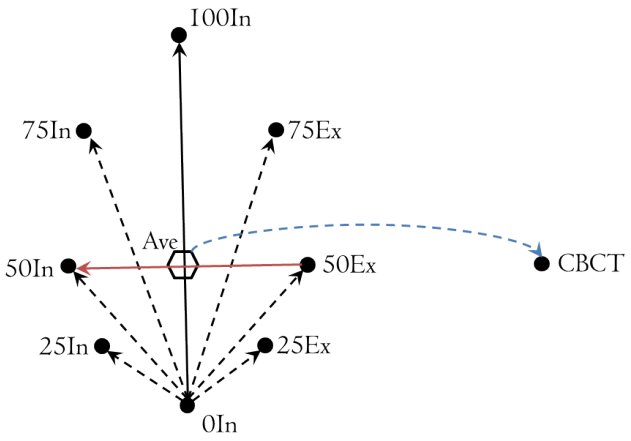

Figure 4.2: Simplified diagram of registrations to build weighting matrices and move them into the space of Ω_{CBCT}. Black arrows (dashed and solid) indicate registrations needed to determine $\Psi_{0In \to ave}$. Solid arrows indicate registrations needed to build D_1^{0In} (black) and D_2^{50Ex} (red). The blue arrow indicates the registration required for $\Psi_{ave \to CBCT}$

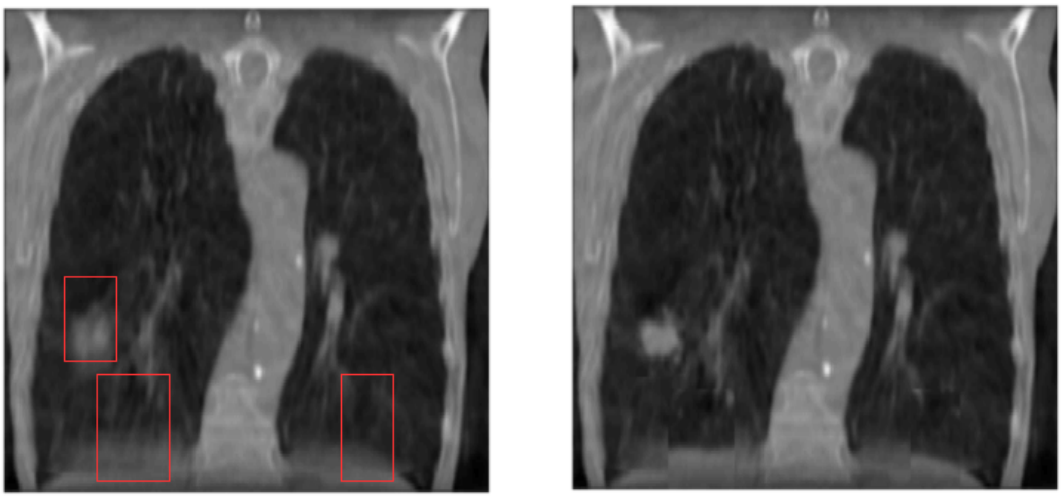

Figure 4.3: Sagittal slice of example reconstruction to demonstrate effects of masking weighting arrays. Reconstruction without (left) and with (right) motion correction, using masked weighting arrays. Mask used shown in red, with field values outside the rectangles set to zero.

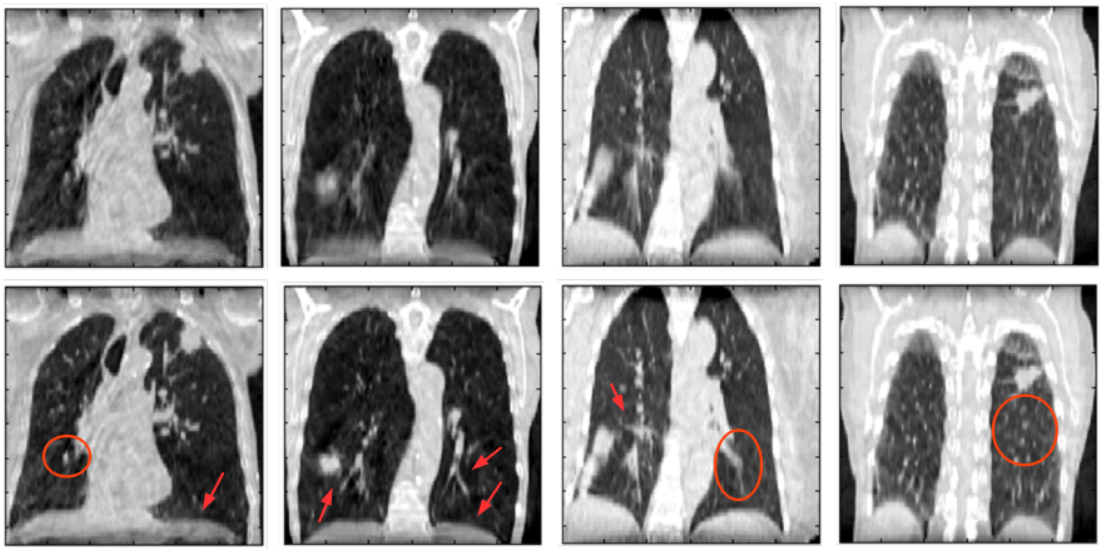

Figure 4.4: Coronal slices of reconstructions before (top) and after (bottom) motion compensation with fitted model, for simulated cases. Slice positions are chosen to intersect with the tumour region. Red arrows and circles indicate areas of improvements with motion correction.

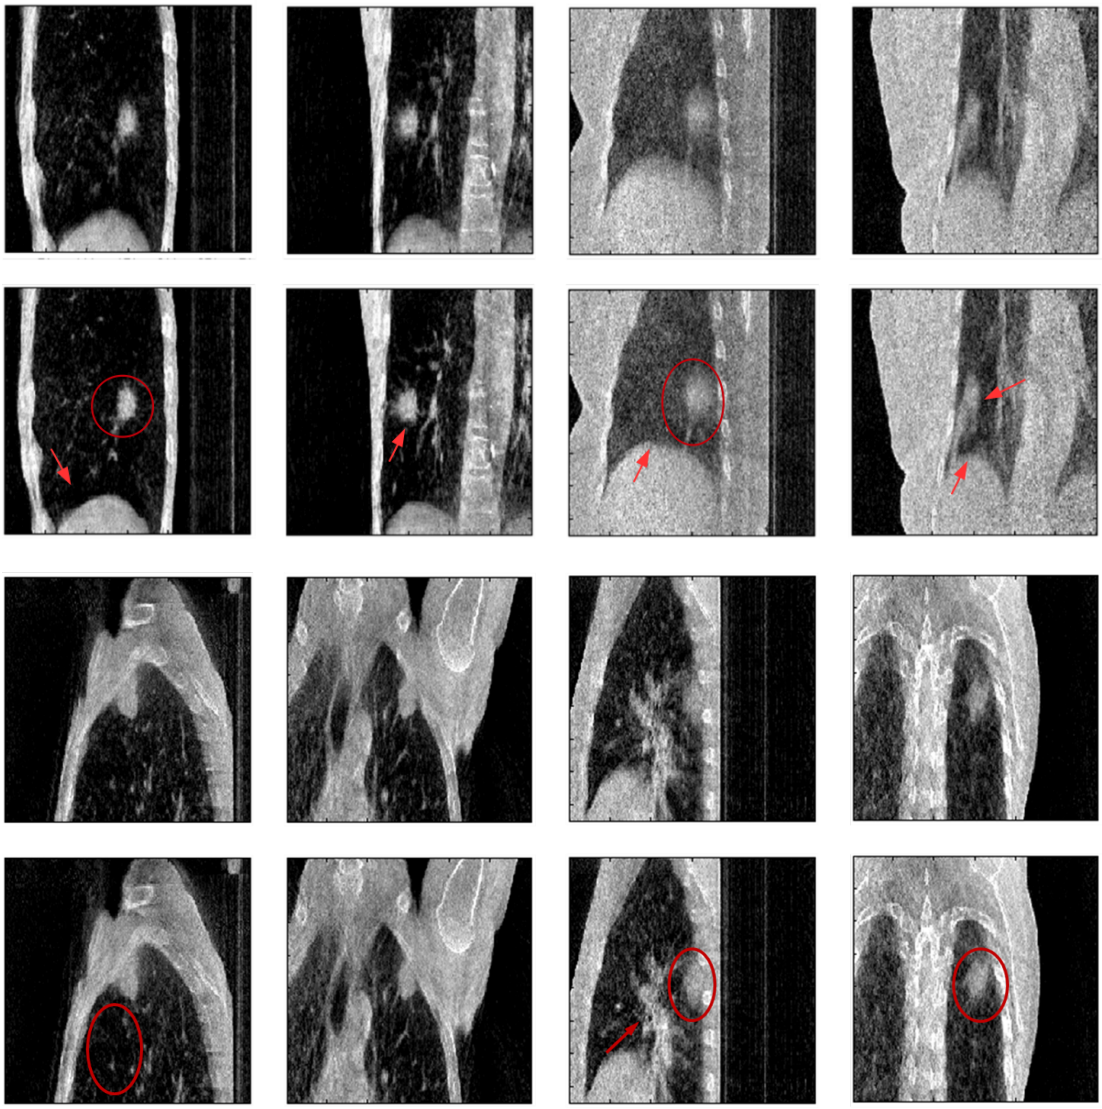

Figure 4.5: Sagittal and coronal slices of patients 1 (bottom left), 2 (top left), 3 (top right) and 4 (bottom right). For each patient, reconstructions without (top row) and with fitted (bottom row) motion correction are given. Patients 2 and 3 exhibited the most motion. Slice position intersects with isocentre of the CBCT scanner, which is approximately aligned with the tumour region. Red arrows and circles indicate areas of improvement with motion correction.

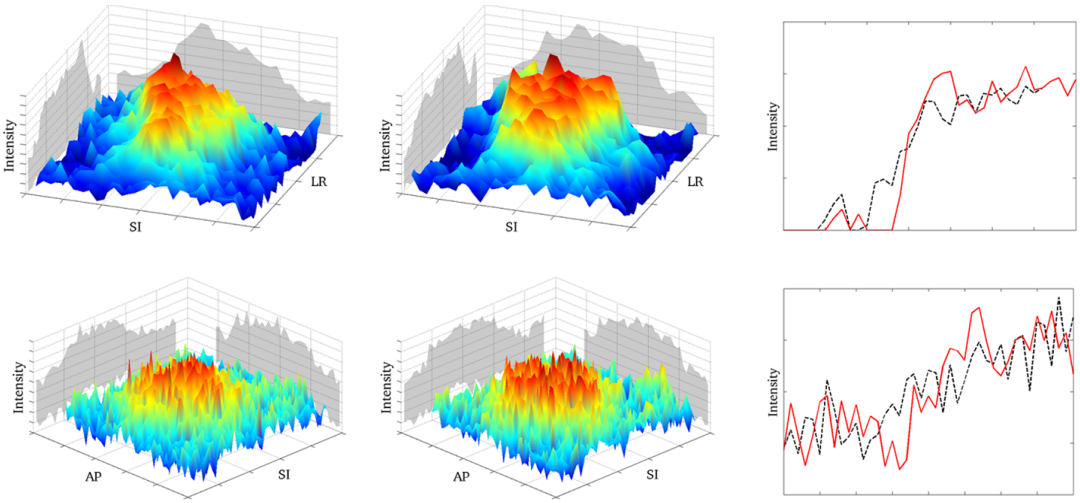

Figure 4.6: Surface plots of a central plane through the tumour region of patients 2 (top) and 3 (bottom), before (left) and after (middle) motion correction. Edge profiles (right) are shown for a vertical profile passing through the diaphragm in the ipsilateral lung. Profiles prior to motion correction are in dashed black, and after motion correction in solid red. Note that patient 3 suffered from large amounts of scatter, resulting in a noisy reconstruction.

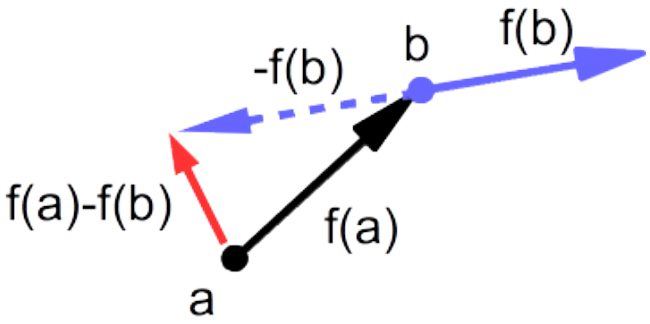

Figure 4.7: Visualisation of error from approximating true inverse with the negative of the forward transformation. For a deformation field f, $f(a)$ is the point to sample to determine the true inverse at point b (i.e. $f(b)^{-1} = -f(a)$). The error $f(a) - f(b)$ is that resulting from approximating the inverse with the negative of the deformation field sampled at b (i.e. $f(b)^{-1} \approx -f(b)$).

Chapter 5

Fully deformable motion model from CBCT

This chapter is the result of an idea I had to try and describe the changes that occurred between two projections, seperated by a small time step. The developments in this chapter are completely my own. The main contribution in this chapter are the concepts behind how to obtain 3D directional information from 2D projections of the CBCT, and computational implementation of the method on patient data. This work has led to first author conference publications detailed in chapter 7, as well as stimulating much future work.

5.1 Introduction

In the previous chapter, a complete motion model of the whole patient is built by splitting the fitted motion into local and global parts. The 4DCT and CBCT are each used to determine the final fitted motion, utilising the various strengths of each dataset.

In this work, we assume that the higher quality, older 4DCT scan is not available. This means the entire non-rigid motion model needs to be fitted to the CBCT. This situation may be useful in situations where large variations are to be expected between planning and the day of treatment. It could also be used to modify the fitted motion using the constrained non-rigid approach in particular regions where large variations may be present, making a fully non-rigid fit to the CBCT data a more suitable approach.

Unlike the previously presented methods, the updates to the motion model are determined in 3D patient space, as opposed to in projection space. A novel method to determine voxel-wise updates to the motion model by extending an optical-flow based method. A modified CBCT reconstruction is proposed, which measures the differences between the actual and current estimate of the motion, and how these differences correlate to the surrogate signals of the motion model. From these reconstructions, the current estimate of the motion can be updated via an analytical expression. It is envisaged that this approach be used to assess respiratory motion immediately prior to treatment, particularly in situations where large inter-fraction variations are expected. These could include weight loss, tumour shrinkage or lung region collapse. Results on a simulated case are presented.

5.2 Methods

5.2.1 Motion model updates

Similar in approach to the constrained non-rigid motion model, it is assumed that the patient motion can be described by a series of deformations, F_n, of a motion-

free image of the patient, $V_{ref}(x) : x \in \Omega_{CBCT}$. Ω_{CBCT} is the region imaged during the CBCT scan and $n = 0, 1, ..., (N-1)$ is an index corresponding to the time of each of the N projections. The deformed patient volume, V_n, is then:

$$V_n(x) = V_{ref}\left(x + F_n(x)\right). \tag{5.1}$$

F_n is parameterised by a motion model which can take into account hysteresis and variations in length and depth of breathing cycle:

$$F_n\left(x\right) = s_n \Psi_1\left(x\right) + \dot{s}_n \Psi_2\left(x\right), \tag{5.2}$$

where s_n and $\dot{s}_n$ are the scalar surrogate signal associated with the n^{th} projection, and its rate of change, respectively. Ψ_1 and Ψ_2 are the motion model deformation fields, which determine a spatially-dependent, linear relationship to s_n and $\dot{s}_n$ respectively. Comparing with the constrained non-rigid motion model (4.2), Ψ_1 and Ψ_2 contain both local and global components of the motion. In this particular case, the motion is to be completely found from the CBCT, making separation of the components into global and local constituents unfavourable.

As with previous work [28], the authors opt to use the Align RT camera system (Vision RT, London, UK) to produce the surrogate signal. The optical stereo-camera system is used to produce high-resolution surface images of the patient's chest. A region of the patient chest (enclosing parts of the thorax and abdomen) is chosen and the average height above the treatment couch within this box is used as the raw surrogate signal. After calculating the derivative trace, both traces are then normalised (mean subtracted; divided by standard deviation) giving the final surrogate traces.

5.2.2 Determining Ψ_1 and Ψ_2

If the motion free image of the patient were known, a Taylor expansion can be used to approximate the first order correction, δF_n^{est}, to an estimated deformation field, F_n^{est}:

$$\begin{aligned} V_n(x) &= V_{ref}\left(x + F_n^{est}(x) + \delta F_n^{est}(x)\right) \\ &\approx V_n^{est}(x) + \nabla V_n^{est}(x) \cdot \delta F_n^{est}(x). \end{aligned} \tag{5.3}$$

In terms of updates to the motion model, the first order correction can be written:

$$\delta F_n^{est} = s_n \delta\Psi_1 + \dot{s}_n \delta\Psi_2. \tag{5.4}$$

Note that $\delta\Psi_1$ and $\delta\Psi_2$ are defined in the space of V_n. Substituting the motion model updates (5.4) into (5.3), and rearranging:

$$V_n - V_n^{est} \approx \nabla V_n^{est} \cdot \left(s_n \delta\Psi_1 + \dot{s}_n \delta\Psi_2\right). \tag{5.5}$$

A demons optical flow approach [67] can then be used to specify the form of the motion model deformation fields Ψ_1 and Ψ_2:

$$s_n \delta\Psi_1 + \dot{s}_n \delta\Psi_2 \approx \frac{\nabla V_n^{est}}{\left(\nabla V_n^{est}\right)^2}\left(V_n - V_n^{est}\right). \tag{5.6}$$

Updates to each motion model deformation field are desired. This is achievable by exploiting the covariance between the surrogate signals. First multiply both sides of (5.6) by s_n, sum over all n and divide by the number of patient states seen over the CBCT scan (i.e. number of projections), N:

$$var(s)\delta\Psi_1 + covar(s,\dot{s})\delta\Psi_2 \approx \frac{1}{N}\sum_n \frac{\nabla V_n^{est}}{(\nabla V_n^{est})^2} s_n (V_n - V_n^{est}), \tag{5.7}$$

where $var(s) = \frac{1}{N}\sum_n s^2(n)$, $covar(s,\dot{s}) = \frac{1}{N}\sum_n s(n)\dot{s}(n)$. Because the surrogate signals have been normalised (mean subtracted; standard deviation divided), *var*

and *covar* are Pearson correlation coefficients and can be simplified: $var(s) = 1$ and $covar(s, \dot{s}) \ll 1$. The latter simplification comes from the approximation of independence between the surrogate signals. The proposed approach could still be used without this approximation, but an additional set of linear equations would need to be solved. Note that satisfying independence could be used as a condition for choosing a suitable set of surrogate signals. The simplified form of (5.7) is:

$$\delta\Psi_1 \approx \frac{1}{N} \sum_n \frac{\nabla V_n^{est}}{(\nabla V_n^{est})^2} s_n \left(V_n - V_n^{est}\right). \tag{5.8}$$

By approximating the sum of all patient states seen over the CBCT by a FDK CBCT reconstruction [19, 35], it is possible to approximate the right side of (5.8) using an FDK reconstruction [57], giving:

$$\delta\Psi_1 \approx \frac{1}{N} \sum_n \frac{\nabla V_n^{est}}{(\nabla V_n^{est})^2} s_n P_n^{\dagger} \left(p_n - P_n(V_n^{est})\right), \tag{5.9}$$

where p_n are the CBCT projections, P_n the projection operator and $P^{\dagger}$ the FDK back-projection operator. A similar equation results for $\delta\Psi_2$ by multiplying (5.6) by $\dot{s}_n$ instead of s_n:

$$\delta\Psi_2 \approx \frac{1}{N} \sum_n \frac{\nabla V_n^{est}}{(\nabla V_n^{est})^2} \dot{s}_n P_n^{\dagger} \left(p_n - P_n(V_n^{est})\right). \tag{5.10}$$

The approximation is essentially using the reconstruction of the differences (in projection space) as an approximation to the average differences in 3D space. The multiplication by the the surrogate is additionally determining which of the differences are correlated with changes in the surrogate signal. As an analogy, if the projections were simply of the animated patient volume, this approximation

would correspond to the assumption that the reconstruction is an estimate of the average of all the states seen over the CBCT in 3D space.

Here the average patient state seen over the entire CBCT is approximated as the standard FDK CBCT reconstruction, as well as the average estimated patient state over the CBCT represented using a standard FDK CBCT reconstruction of the animated CBCT projections. Since the radon transform is additive, these approximations can be written, with the reconstruction of the subtracted projections (original minus estimated projections) being an estimate of the average differences between the actual and animated patient volumes, throughout the CBCT.

5.2.3 Iterative approach

In practice, V_{ref} is not accurately known, as this would require knowledge of the exact respiratory motion from which to perform an MCR. In this work, an iterative approach is taken, with improving estimates of the motion model deformation fields and MCR with each iteration. Starting with a zero motion assumption (i.e. $\Psi_1 = \Psi_2 = \mathbf{0}$), perform an MCR, V^0, and set it to V_{ref}. The updates to the motion model deformation fields $\delta\Psi_1^0$ and $\delta\Psi_1^0$ are then calculated, and the fields updated via:

$$\Psi_i^1 = \Psi_i^0 + \delta\Psi_i^0 \quad \text{for} \quad i = 1, 2. \tag{5.11}$$

The updated motion model deformation fields can then be used to determine the updated MCR and the procedure repeated.

For this work, a fixed number of iterations was set. Ten iterations were run

with the updated motion model deformation fields having Gaussian smoothing applied prior to composition. To reduce calculation time, $\frac{V_n^{est}}{(V_n^{est})^2}$ was approximated by a Gaussian blurred V_{ref}, as $\frac{V_{ref}}{\left(V_{ref}\right)^2}$, allowing it to be moved outside the summation in (5.9,5.10) and applied in one step. The Gaussian blurring applied was the same as that used to smooth the motion model deformation fields. To ensure that the optical flow equation is not unstable at small values of ∇V_n^{est}, updates to $\delta\Psi_1$ and $\delta\Psi_2$ were set to zero if $(\nabla V_{ref})^2 + \left(\frac{1}{N}\sum_n \dot{s}_n P_n^{\dagger}\left(p_n - P_n(V_n^{est})\right)\right)^2 < \varepsilon$, where ε was empically set to 5×10^{-5}.

5.3 Results

5.3.1 Simulated results

The method was tested on a simulated case, built from the 4DCT of an actual patient undergoing radiotherapy treatment. End exhale was registered to each of the seven other phases of the 4DCT and used to determine a transformation to average 4DCT space. Registrations were used to determine transformations from end-inhale to end-exhale, and from mid-exhale to mid-inhale. These transformations were used to determine Ψ_1 and Ψ_2. The end-exhale phase, moved to average 4DCT space was then animated using Ψ_1, Ψ_2 and an actual patient surrogate trace. The accompanying CBCT geometry information for the respiratory trace was used to simulate a CBCT of the animated volume.

After motion correction, improvements in the quality of the MCR were seen. Figure 5.1 shows coronal slices through the tumour region of the patient volume before and after motion correction. The regions of greatest improvement were in the diaphragm and tumour region. The ribs were also seen to have fewer streak

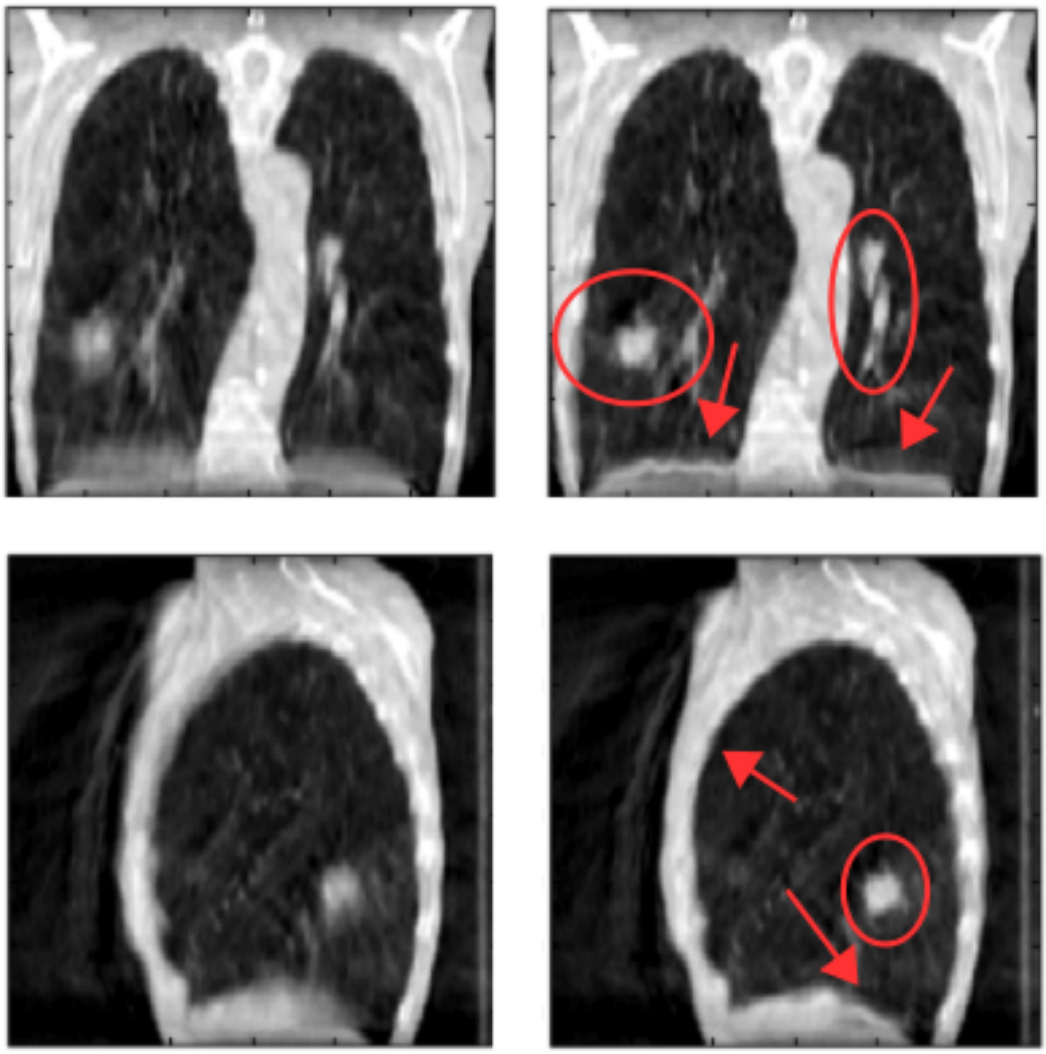

Figure 5.1: Coronal (top) and sagittal (bottom) slices, intersecting the tumour region, of MCRs before (left) and after (middle) fitting the motion model.

artefacts, especially in the SI direction.

CBCT projections of the animated MCR with the fitted motion were seen to be more similar to the actual projections, than projections of the unanimated no motion reconstruction. Over all the projections, the SSD error was reduced by 12%, from 1.6×10^9 to 1.4×10^9. This also corresponds with an improvement in similarity between actual and fitted motion model deformation fields, compared to assumming no motion. Over the whole patient, mean L2-norm errors decreased from 1.4 to 0.46 mm. In the tumour region, the mean L2-norm error was decreased

from 4.2 to 2.6 mm. A movie of the original (left) and simulated (right) CBCT projections with the fitted motion have been included as supplementary material.

5.4 Conclusions and future work

An iterative approach to build a surrogate-driven motion model entirely from CBCT is presented. This approach could be used to reduce respiratory artefacts in CBCT reconstructions, and as an intensity-driven way to estimate patient motion when only the CBCT scan is available.

In its current form, the initial estimate of the motion is zero (no motion). As the whole motion model is built entirely from the CBCT, this was used as no prior knowledge of the motion is assumed to be known. One way to improve upon this estimate would be to use the prior information from the 4DCT planning scan, or perhaps from previous treatments' CBCT scans, as the starting estimate. Although inter-fraction variations may have occurred, this initial estimate is more likely to be closer to the actual motion than assuming no motion. For cases where a prior CBCT-fitted model exists, a registration between the two CBCT reconstructions (without motion correction) should be suitable to move the motion model to the latest day of treatment. For the 4DCT planning scan approach, an approach similar to the previous chapter (Chapter 4) could be used to move the motion model to CBCT space.

In higher constrast areas, there is a visible reduction in motion artefacts in the MCR using the fitted motion. This is because the gradient-based updates rely on changes in intensity within the images to provide an accurate update. However, in regions where this contrast is not available, the updates may not

be as accurate. Smoothing is one way to regularise the updates over the entire patient volume, as well as provide an estimate of the how to update the model in regions of low contrast. One areas of future research, which may improve the accuracy of update in these regions is to combine the presented approach with the constrained non-rigid approach in Chapter 4. For example, the constrained non-rigid approach could be used to provide the first estimate of the patient motion, with the presented approach then being used to update this estimate further. As well as improving upon the accuracy in low contrast regions, the proposed approach may also be able to improve upon the constrained non-rigid approach. This improved accuracy may be in regions of high contrast which show large inter-fraction variation from planning.

It is acknowledged that hard termination conditions were set with the simulated case. For different patients, each exhibiting different motion, a hard constraint such as this may result in different errors on their fitted motion models. For future work, it is intended to add termination conditions, for example on the magnitude of the average update, to determine whether enough iterations have occurred to suitably describe the motion seen throughout the CBCT scan.

With termination conditions set, it is likely that many iterations will be required to suitably describe the motion. To be of clinical use, the number of iterations should be as low as possible. One way to reduce the number of iterations would be to include a more physical basis for updating the motion model, to satisfy the termination conditions in fewer iterations. The authors are planning to investigate a more physically justified basis for the direction-vector updates (i.e. alternatives to $\frac{\nabla V_{ref}}{(\nabla V_{ref})^2}$) to achieve these aims. Using a more physically justified update may also help reduce the effect of the relatively poor quality CBCT scan

on the updates.

Chapter 6

Future work

6.1 Further trials with more patients, better integrated equipment

Three general approaches to modelling motion of a patient undergoing tracked RT treatment has been presented. Each method has its own advantages and disadvantages and therefore potential uses. Two of the methods presented results on actual patient datasets and one on a simulated dataset. In order to better understand the behaviour of the presented methods with actual data, all the approaches should be applied to a larger number of patient cases.

Having results on a larger set of patient data will allow the authors to understand in more detail the reliability, accuracy and general applicability of each method. After these key issues are understood in more detail, it would be possible to make a more concrete conclusion as to which method provides the largest overall benefits during tracked RT treatments. The selected method, then, can be the first approach which is pursued further with regard to an actual clinical

implementation.

Key to a clinical implementation is having a highly optimised, efficient algorithm. Where there is no motion correction, GPU acceleration is currently used to reduce the processing time as much as possible. If possible GPU-accelerating both the simulated projections of the animated volume, and an MCR should be given serious consideration. If both of these processes can be reduced to under 20 seconds per calculation, constrained non-rigid method presented in Chapter 4 can be made clinically useful. Note that for the tumour-only method presented in Chapter 3, if the reconstructed region is centred around the tumour and kept small (e.g. 100 sided-cube of 1 mm^3 voxels) the model can be fitted within a clinically useful timeframe. For the fully deformable method presented in Chapter 5, many iterations will probably be required. Therefore simply using a GPU-accelerated projector of the animated volume and MCR may not be enough given the number of iterations. In addition to this, it would be worth providing a better starting estimate of the motion, described in more detail later in this chapter.

Some of the latest SABR treatment equipment has features which would make the integration of this approach easier than utilising existing equipment. For example the Elekta TrueBeam features access to the live-streamed projections of the CBCT, which would make it much simpler to fit the model during or immediately after the CBCT scan. Currently, the CBCT projections must be transfered to hard disk post-fraction and processed off-site. In order to maximise exposure to this next generation equipment, key partnerships should be considered with groups who will have access to this equipment. Due to the likely timescale for tracked RT treatments to become more accepted, it is likely that this next generation equipment will be standard by that time. The correct choice now of

equipment to develop the clinical implementations of the methods, will result in a smoother transition to the clinic when tracked treatments become more accepted.

6.2 Multiple region motion models within constrained non-rigid approach

For the constrained non-rigid approach, the motion model parameters are global in nature, in each respective dimension. This means there is one assumed relationship with the surrogate signal throughout the patient volume. An extension of this approach results from questioning whether these global parameters should indeed be the same throughout the patient. For example, the lungs and chest wall slide past one another and simply accounting for the majority of the inter-fraction variation by the registration between planning 4DCT and CBCT reconstruction (i.e. by modifying the weighting arrays, or local component of the motion model) may not be enough. In this example, it may be better to split the patient into regions (e.g. lung and non-lung), each with their own regional parameters, as opposed to global parameters over the whole patient. This means each region will now have a distinct relationship to the surrogate signal, rather than simply fitting the global parameters over the whole patient.

A brief description of the necessary modifications to the cost function follows. For a patient divided into I non-overlapping regions (all of patient belongs to a region), each with an independent surrogate dependence:

$$F(x, s(n), \phi) = \sum_{i=1}^{I} F\left(x, s(n), \phi_i\right). \tag{6.1}$$

Note that the subscripts of ϕ do not correspond to iteration number but region. $F(x, s(n), \phi_i) = s(n)D_1(x) \circ \phi_{1,i} + \dot{s}(n)D_2(x) \circ \phi_{2,i}$. $\phi_{1,i}$ and $\phi_{2,i}$ are the motion model parameters for the i^{th} region (surrogate-dependent and surrogate-derivative-dependent, respectively). The deformation field updates are:

$$\delta F = \sum_{i=1}^{I} F\left(x, s(n), \delta\phi_i\right), \tag{6.2}$$

where $F\left(x, s(n), \delta\phi_i\right) = s(n)D_1(x) \circ \delta\phi_{1,i} + \dot{s}(n)D_2(x) \circ \delta\phi_{2,i}$. Note here we assume each region can deform non-rigidly. It is possible to impose rigid translation-only motion (useful for certain situations, such as modelling tumour region) by filling the weighting matrices with ones. For convenience, let $F[\phi_i] = F\left(x, s(n), \phi_i\right)$ and $F[\delta\phi_i] = F\left(x, s(n), \delta\phi_i\right)$.

The estimated patient volume, at the time of the nth projection, is now:

$$V[n] = V_{ref}\left(x + \sum_{i=1}^{I} F[\phi_i]\right) = \sum_{i=1}^{I}\left(V_{ref}^{i}\left(x + F[\phi_i]\right)\right), \tag{6.3}$$

where V_{ref}^{i} is V_{ref} masked by region i. For notional convenience, let $V_{ref}^{i}\left(x + F[\phi_i]\right) = V_i[n]$, where again the subscript denotes a region and not iteration number. The cost function also requires modification:

$$\delta\phi = \underset{\delta\phi}{argmin}\left(\sum_{n}\sum_{pixels}\left(r(n) - \boldsymbol{P}_n\left(\nabla V[n] \cdot \sum_{i=1}^{I} F[\delta\phi_i]\right)\right)^2\right). \tag{6.4}$$

The sum over regions can be brought outside the projection operator, giving the following expanded form of the cost function:

$$\delta\phi = \underset{\delta\phi}{argmin}\left[\sum_{i=1}^{I}\sum_{\theta,\psi=0}^{6}\lambda_{\theta,i}\lambda_{\psi,i}\sum_{t_n}\sum_{pixels}C_{\theta,i}C_{\psi,i}\gamma_{\psi,\theta}\right], \tag{6.5}$$

$$\lambda_{0,i} = 1, \begin{pmatrix} \lambda_{1,i} \\ \lambda_{2,i} \\ \lambda_{3,i} \end{pmatrix} = \delta\phi_{1,i}, \begin{pmatrix} \lambda_{4,i} \\ \lambda_{5,i} \\ \lambda_{6,i} \end{pmatrix} = \delta\phi_{2,i}. \tag{6.6}$$

and

$$C_{0,i} = r(n),$$

$$\begin{pmatrix} C_{1,i} \\ C_{2,i} \\ C_{3,i} \end{pmatrix} = -s(n) \begin{pmatrix} \boldsymbol{P}_n \left(\nabla_x V_i[n] \circ D_{1,x}\right) \\ \boldsymbol{P}_n \left(\nabla_y V_i[n] \circ D_{1,y}\right) \\ \boldsymbol{P}_n \left(\nabla_z V_i[n] \circ D_{1,z}\right) \end{pmatrix},$$

$$\begin{pmatrix} C_{4,i} \\ C_{5,i} \\ C_{6,i} \end{pmatrix} = -\dot{s}(n) \begin{pmatrix} \boldsymbol{P}_n \left(\nabla_x V_i[n] \circ D_{2,x}\right) \\ \boldsymbol{P}_n \left(\nabla_y V_i[n] \circ D_{2,y}\right) \\ \boldsymbol{P}_n \left(\nabla_z V_i[n] \circ D_{2,z}\right) \end{pmatrix}. \tag{6.7}$$

$$\gamma_{\theta,\phi} = \begin{cases} I^{-1} & \text{if } \theta = \psi = 0 \\ 1 & \text{otherwise} \end{cases}.$$

Note that projections of the ith animated patient region, $V_i[n]$, is needed to bring the spatially independent parameters for that region out of the projection operator.

For this particular form of the cost function, a critical issue is how to deal with the transition between regions. If a simple implementation is made, the regions will effectively result in patient anatomy overlapping if one is not careful with region selection. If the regions correspond to edges of the anatomy, this may not

be the case, as long as the overall deformations remain small enough. However, instead of relying on small deformations, it would be better to impose some form of constraint on the final motion model estimate given the fitted motion. The most simple method to implement would be to apply a smoothing after each fitting of the regional motion model parameters. This would ensure that the changes in estimated motion at the boundaries were not too extreme. However, for the lung/chest-wall boundary, an extreme change may be precisely what is needed. Alternative methods could be to then apply some form of bio-mechanical-based smoothing to the fitted motion. For example, one would expect sliding between the chest wall and lung, so no smoothing would be made at this boundary. For boundaries between regions with expected similar deformations, smoothing would take place to prevent large discrete changes in the estimated motion.

Another interesting approach would be to alter the directions of the region motion model parameters (i.e. they will no longer be in SI, AP and LR directions). [14] describe a b-spline-based approach to describing a patient divided into lung and non-lung regions, which can account for sliding between the lungs and chest wall whilst providing a smooth transformation between the regions. Based on this approach, it is now assumed that the voxel-wise description of the motion model will be generated from a b-spline-based description of the motion model. Registrations between the phases of the 4DCT used to generate the weighting arrays within each region will be done using the approach described in [14]. From these registrations, once moved to the CBCT space, there will now be a regional b-spline-based description for the weighting arrays along with an array which details each region. The voxel-wise description of the global and local components of the constrained motion model can now be deduced from these constituents.

Critical to ensuring smooth transitions between regions are the arrays detailing the various regions. This is because the direction vectors parallel and orthogonal to the region boundaries are needed to ensure the smooth transitions. In order to use this method and still retain the local and global form of the motion model, the global motion model parameters are now directed parallel and orthogonal to the region boundaries (rather than in SI, AP and LR directions). By modifying the cost function to account for this form of the motion model (the key change would be modifying the direction to sample to calculate the gradient-based terms of the cost function), this approach could be a promising way to have multiple regions of anatomical relevance modelled, with smooth transitions between regions modelled whilst still accounting for sliding at the region boundaries. The cost function also retains many of the attractive properties of the constrained non-rigid motion model approach, for example a cost function of positive definite quadratic form, and fitting the motion based on the respective strengths of the available datasets (weighting arrays based on high quality 4DCT; global properties based on day-specific but poor quality CBCT).

6.3 More physically-based updates to motion model in fully non-rigid approach

For the fully deformable approach described in Chapter 5, the updates to the motion model are oriented in the direction of greatest change in the gradient. For some parts of the patient anatomy, for example the ribs, this assumption is incorrect. Heavy smoothing can be used to reduce the impact of these incorrect orientations at each iteration, and a large number of iterations used to eventually

determine the correct motion. However, it may be possible to better constrain the updates to the motion model by using an anatomically-based orientation instead.

Looking at the update equations (5.9,5.10), one can decompose them into two parts, a magnitude and direction parts. For Ψ_1 these are $\frac{1}{N}\sum_n \frac{1}{(\nabla V_n^{est})^2} s_n P_n^\dagger \left(p_n - P_n(V_n^{est})\right)$ and ∇V_n^{est}, respectively. For Ψ_2 the direction component is the same and magnitude is $\frac{1}{N}\sum_n \frac{1}{(\nabla V_n^{est})^2} \dot{s}_n P_n^\dagger \left(p_n - P_n(V_n^{est})\right)$.

It is possible to relate Ψ_1 and Ψ_2 in equation (5.2) to the form of the constrained non-rigid motion model (4.2):

$$\Psi_i\left(x\right) = D_i\left(x\right) \circ \mu_i \quad \text{for} \quad i = 1, 2. \tag{6.8}$$

It is possible to consider fitted constrained non-rigid motion model $D_i \circ \mu_i$ as deformation fields which describe how the anatomy is expected to move. Because of the way they were constructed, the weighting arrays need to be normalised to extract just the orientation information of these deformations.

I propose the following as an initial modification to the update equations, to perform an update based on an anatomical-derived basis, as opposed to completely gradient based:

$$\delta\Psi_1 \approx \frac{1}{N}\sum_n \frac{D_1 \circ \mu_1}{|D_1 \circ \mu_1|} \left|\frac{\nabla V_n^{est}}{(\nabla V_n^{est})^2}\right| s_n P_n^\dagger \left(p_n - P_n(V_n^{est})\right), \tag{6.9}$$

$$\delta\Psi_2 \approx \frac{1}{N}\sum_n \frac{D_2 \circ \mu_2}{|D_2 \circ \mu_2|} \left|\frac{\nabla V_n^{est}}{(\nabla V_n^{est})^2}\right| \dot{s}_n P_n^\dagger \left(p_n - P_n(V_n^{est})\right), \tag{6.10}$$

The reasoning for the modification to the gradient-based term, $\left|\frac{\nabla V_n^{est}}{(\nabla V_n^{est})^2}\right|$, is as follows. Although here we attempt to remove reliance on the gradient-based term, it cannot be completely omitted as it also provides a key part of the update.

The gradient provides information on how the underlying physical deformations relate to intensity changes within the image. Since the new orientations, $\frac{D_i \circ \mu_i}{|D_i \circ \mu_i|}$ are normalised the magnitude of the gradient-based term allows the orientation information to be scaled correctly, based on the CBCT image. Without the gradient term, the orientation may be physically based, but the magnitude of the update may be wrong as the CBCT and 4DCT reconstructions have different intensity-ranges within the images. Alternative ways to account for the different intensity ranges between the 4DCT and CBCT reconstructions are to be left for future development.

Note that since the constrained non-rigid and fully deformable motion models have now been related, it would be possible to first use the constrained non-rigid approach as the starting estimate of the fully deformable approach. This may greatly reduce the number of iterations (compared to assuming zero motion) required, as well as improving the estimate in low contrast regions.

Chapter 7

Contributions

For this thesis, contributions have been made in the development of a variety of methods of building motion models from a CBCT scan, for the purposes of guiding tracked RT treatment.

In Chapter 3, I devised the CBCT-specific underpinnings of an iterative method to build a motion model of the tumour region only. The tumour region is the most critical region to monitor during RT treatment, and the markerless method presented can be used with the standard CBCT scan and protocol used during RT treatment. The only additional requirement is acquiring a synchronised surrogate signal, which can be further monitored to drive tracked treatment. The cost function is of a positive definite quadratic form, with a global minimum which can be found quickly. The method does not require markers, avoiding the risks associated with their implantation. Scans acquired at planning are not needed by the method, so can account for changes of tumour appearance and motion over each fraction of treatment. The motion model is used to perform a motion compensated reconstruction (MCR), which utilises all the projection data so does not suffer from the artefacts seen in 4D-CBCT. The MCR can then be

animated using the motion model to evaluate tumour motion during the CBCT. CBCT scans are routinely acquired immediately prior to and in the middle of each fraction, ensuring the motion model built is appropriate for each fraction of treatment and can be assessed mid-treatment for accuracy. As well as evaluation of tumour motion prior to treatment, real-time surrogate data could be used to predict tumour position during treatment and so guide gated or tracked treatments. In proton therapy, where motion compensation is even more important [8], this method could also be of use [15]. Results on a simulated, phantom and clinical datasets are presented.

In Chapter 4, I extended and generalised this iterative method to a whole-patient motion model. This new approach can now be used to track the tumour region and OAR, whilst retaining many of the attractive features of the tumour-only method. A model of this kind is important for delivering tracked RT treatments, ensuring modifications to the beam do not over-expose the OAR. A surrogate-driven motion model is used to overcome the difficulties associated with real-time imaging of the internal anatomy. Once the model has been built, only a new surrogate measurement is needed to estimate the patient state. The surrogate can be monitored easily throughout treatment, avoiding implanting markers or taking additional scans. The motion model is posed in a way which exploits the relative strengths of the datasets most often available during RT treatment: The planning 4DCT and positioning CBCT. The higher quality planning 4DCT is used to construct what the authors term weighting arrays. The weighting arrays are essentially deformation fields within the model, which describe how the anatomy of the patient could vary with the surrogate signals. The motion model parameters (after being weighted by the weighting arrays) are then

fitted to the CBCT, which can account for the inter-fraction variation that occurs between planning and the day of treatment. The motion model parameters are global in nature (same over the whole volume). This means the number of model parameters to fit on the day of treatment is low and they can be quickly found. The motion model is also robust against many of the previously mentioned problems affecting CBCT. In addition, delineations from planning can be imported onto the model, providing tumour and OAR boundaries based on those used to construct the plan. Results on simulated and clinical datasets are presented.

In some cases, it is advantageous to be able to build a motion model completely from the CBCT scan. In Chapter 5, I presented a novel approach to achieve this. As well as being used to completely build the motion model from scratch, this method can also be used to augment the full patient motion model presented in Chapter 4. This is especially useful where large inter-fraction variations are present. These could include weight loss, tumour shrinkage or lung region collapse. Although this latter approach is left to future work, the connection between the motion models used for Chapters 4 and 5 are described to highlight the connections between the two approaches and how they could compliment each other. By extending an optical-flow-based approach, voxel-wise updates to the motion model are determined. A modified CBCT reconstruction is proposed, which measures the differences between the actual and current estimate of the motion, and how these differences correlate to the surrogate signals of the motion model. From these reconstructions, the current estimate of the motion can be updated via an analytical expression. Results on a simulated case are presented.

Via the work presented in this thesis, it is ultimately intended for a variety of options to be made available to monitor patient respiratory motion, during

tracked RT treatment. Whether it be simply determining tumour motion or a whole patient motion model, the methods presented should satisfy a wide range of situations. With further development, it is hoped that these methods can form a key component of tracked RT treatments, to ultimately improve treatment outcomes for lung cancer patients.

7.1 Publications

The work presented in this thesis have resulted in the following publications:

Journal (1st author)

- Building motion models of lung tumours from cone-beam CT for radiotherapy applications. James Martin, Jamie McClelland, Connie Yip, Christopher Thomas, Clare Hartill, Shahreen Ahmad, Richard OBrien, Ivan Meir, David Landau, David Hawkes. Physics in Medicine and Biology 58 (6), 1809. 2013.

Conference (1st author)

- Determination of tumour-centre motion during a cone-beam CT acquisition - A feasibility study. J Martin, J McClelland, D Hawkes. 1214. Radiotherapy and Oncology. 2011.

- Motion modelling and motion compensated reconstruction of tumours in cone-beam computed tomography. James Martin, Jamie McClelland, Christopher Thomas, Kate Wildermuth, David Landau, Sebastien Ourselin, David Hawkes. Mathematical Methods in Biomedical Image Analysis (MMBIA). 281-286. 2012.

- Markerlessly tracking lung tumours during radiotherapy treatment using Align RT optical surrogate and motion model built from cone-beam CT on day of treatment, J. Martin, J. McClelland, C. Yip, C. Thomas, C. Hartill, S. Ahmad, R. O'Brien, I. Meir, D. Landau, D. J. Hawkes. European Congress on Radiology (ECR) 2013. C-2089. 2013.

- Lung tumour motion models from cone-beam CT. James Martin, Jamie McClelland, Christopher Thomas, Richard OBrien, Shahreen Ahmed, Clare Hartill, Connie Yip, David Landau, Ivan Meir, David Hawkes. Image-Guidance and Multimodal Dose Planning in Radiation Therapy. 50. 2013.

- Fully deformable patient motion models from cone-beam CT for radiotherapy applications. James Martin, Jamie McClelland, Connie Yip, Christopher Thomas, Clare Hartill, Shahreen Ahmed, Ivan Meir, David Landau and David Hawkes. International Conference on the Use of Computers in Radiation Therapy (ICCR) 2013. 2013.

- Building surrogate-driven motion models from cone-beam CT via surrogate-correlated optical flow. James Martin, Jamie McClelland, David Hawkes. Image Processing and Computer Aided Interventions (IPCAI). 61-67. 2014.

Conference (2nd author)

- Lung tumour tracking and motion compensated reconstruction in cone beam CT. J McClelland, J Martin, D Hawkes - Radiotherapy and Oncology. 457. 2011.

Workshop (1st author)

- Lung tumour motion models from CBCT, James Martin, Jamie McClelland, Christopher Thomas, Richard OBrien, Shahreen Ahmad, Clare Hartill, Connie Yip, David Landau, Ivan Meir, David Hawkes. Image-guidance and multimodal dose planning in RT. MICCAI Workshop 2012. 2012.

Appendix A

Derivation of cost function

The cost functions used for the tumour-only (3.2) and constrained non-rigid (4.2) motion models are derived from algebraic expansion and simplification. Here the cost functions are derived in full, starting with the constrained non-rigid motion model from which the simplified form of the tumour-only motion model can be found. The cost function can be written:

$$\delta\mu = \underset{\delta\mu}{argmin}\left[\sum_{n}\sum_{pixels}\left(R_n - P_n\left(\nabla V_n \cdot \delta F_n\right)\right)^2\right], \tag{A.1}$$

where δF_n can take the constrained non-rigid motion model:

$$\delta F_n\left(x\right) = s_n D_1\left(x\right) \circ \delta\mu_1 + \dot{s}_n D_2\left(x\right) \circ \delta\mu_2. \tag{A.2}$$

Substituting (4.4) into (4.3) and algebraically expanding, the right hand side of the cost function (in the square brackets) now becomes:

$$\begin{aligned}\sum_{n}\sum_{pixels}\left(R_n - P_n\left(\nabla V_n \cdot \delta F_n\right)\right)^2 = \sum_{n}\sum_{pixels}&\left(R_n - P_n\left(\nabla_x V_n \delta_x F_n\right)\right.\\ &\left. - P_n\left(\nabla_y V_n \delta_y F_n\right) - P_n\left(\nabla_z V_n \delta_z F_n\right)\right)^2,\end{aligned} \tag{A.3}$$

where the subscripts of ∇ and δ indicate the relevant directional component. Further expanding the squared bracket:

$$\begin{aligned}\sum_{n}\sum_{pixels}&\left(R_nR_n - R_nP_n\left(\nabla_xV_n\delta_xF_n\right) - R_nP_n\left(\nabla_yV_n\delta_yF_n\right)\right.\\&- R_nP_n\left(\nabla_zV_n\delta_zF_n\right) - P_n\left(\nabla_xV_n\delta_xF_n\right)R_n\\&\left.- \ldots + P_n\left(\nabla_zV_n\delta_zF_n\right)P_n\left(\nabla_zV_n\delta_zF_n\right)\right),\end{aligned} \tag{A.4}$$

and writing the full form of each directional component of δF_n:

$$\begin{aligned}\sum_{n}\sum_{pixels}&\left(R_nR_n - R_nP_n\left(\nabla_xV_n\circ\left(s_nD_{1,x}\delta\mu_{1,x} + \dot{s}_nD_{2,x}\delta\mu_{2,x}\right)\right)\right.\\&- R_nP_n\left(\nabla_yV_n\circ\left(s_nD_{1,y}\delta\mu_{1,y} + \dot{s}_nD_{2,y}\delta\mu_{2,y}\right)\right)\\&- R_nP_n\left(\nabla_zV_n\circ\left(s_nD_{1,z}\delta\mu_{1,z} + \dot{s}_nD_{2,z}\delta\mu_{2,z}\right)\right)\\&- P_n\left(\nabla_xV_n\circ\left(s_nD_{1,x}\delta\mu_{1,x} + \dot{s}_nD_{2,x}\delta\mu_{2,x}\right)\right)R_n - \ldots\\&+ P_n\left(\nabla_zV_n\circ\left(s_nD_{1,z}\delta\mu_{1,z} + \dot{s}_nD_{2,z}\delta\mu_{2,z}\right)\right)P_n\left(\nabla_zV_n\right.\\&\left.\left.\circ\left(s_nD_{1,z}\delta\mu_{1,z} + \dot{s}_nD_{2,z}\delta\mu_{2,z}\right)\right)\right),\end{aligned} \tag{A.5}$$

where because the patient volume has been discretised into voxels, the $\circ$ symbol (Hadamard product) is used to indicate a voxel-wise multiplication within the volume, prior to forward projection. The $\circ$ product is unneccesary between the weighting array and motion model updates when only one direction-composition is chosen as is the case here. Further expanding the terms, due to the projection operator being additive:

$$\begin{aligned}\sum_{n}\sum_{pixels}&\left(R_nR_n - R_nP_n\left(\nabla_xV_n\circ\left(s_nD_{1,x}\delta\mu_{1,x}\right)\right)\right.\\&- R_nP_n\left(\nabla_xV_n\circ\left(\dot{s}_nD_{2,x}\delta\mu_{2,x}\right)\right) - R_nP_n\left(\nabla_yV_n\circ\left(s_nD_{1,y}\delta\mu_{1,y}\right)\right)\\&- R_nP_n\left(\nabla_yV_n\circ\left(\dot{s}_nD_{2,y}\delta\mu_{2,y}\right)\right) - R_nP_n\left(\nabla_zV_n\circ\left(s_nD_{1,z}\delta\mu_{1,z}\right)\right)\\&- R_nP_n\left(\nabla_zV_n\circ\left(\dot{s}_nD_{2,z}\delta\mu_{2,z}\right)\right) - P_n\left(\nabla_xV_n\circ\left(s_nD_{1,x}\delta\mu_{1,x}\right)\right)R_n\\&- P_n\left(\nabla_xV_n\circ\left(\dot{s}_nD_{2,x}\delta\mu_{2,x}\right)\right)R_n - \ldots\\&+ P_n\left(\nabla_zV_n\circ\left(s_nD_{1,z}\delta\mu_{1,z}\right)\right)P_n\left(\nabla_zV_n\circ\left(s_nD_{1,z}\delta\mu_{1,z}\right)\right)\\&+ 2P_n\left(\nabla_zV_n\circ\left(s_nD_{1,z}\delta\mu_{1,z}\right)\right)P_n\left(\nabla_zV_n\circ\left(\dot{s}_nD_{2,z}\delta\mu_{2,z}\right)\right)\\&\left.+ P_n\left(\nabla_zV_n\circ\left(\dot{s}_nD_{2,z}\delta\mu_{2,z}\right)\right)P_n\left(\nabla_zV_n\circ\left(\dot{s}_nD_{2,z}\delta\mu_{2,z}\right)\right)\right).\end{aligned} \tag{A.6}$$

$\delta\mu_{i,j}; i = 1, 2, j = x, y, z$ is the same across all voxels for all values of n, hence can be taken outside the Hadamard product, projection operator and sums over

pixel and n. Similarly for s_n, with the exception that it is dependent on n, so can be taken outside the Hadamard product and projection operator (it could also be taken outside the sum over pixel but due to the low computational cost, this is not done here for a more appealing final cost function). By reordering the operations as described, one arrives at the form of the cost function described in Chapter 4.

$$\mu_1, \mu_2 = \underset{\mu_1,\mu_2}{argmin} \left[\sum_{\theta,\phi=0}^{6} \lambda_\theta \lambda_\phi \sum_n \sum_{pixels} C_\theta C_\phi \right], \tag{A.7}$$

where

$$\lambda_0 = 1, \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix} = \mu_1, \begin{pmatrix} \lambda_4 \\ \lambda_5 \\ \lambda_6 \end{pmatrix} = \mu_2. \tag{A.8}$$

and

$$C_0 = p_n - P_n(V_n),$$

$$\begin{pmatrix} C_1 \\ C_2 \\ C_3 \end{pmatrix} = -s_n \begin{pmatrix} P_n\left(\nabla_x V_n \circ D_{1,x}\right) \\ P_n\left(\nabla_y V_n \circ D_{1,y}\right) \\ P_n\left(\nabla_z V_n \circ D_{1,z}\right) \end{pmatrix}, \tag{A.9}$$

$$\begin{pmatrix} C_4 \\ C_5 \\ C_6 \end{pmatrix} = -\dot{s}_n \begin{pmatrix} P_n\left(\nabla_x V_n \circ D_{2,x}\right) \\ P_n\left(\nabla_y V_n \circ D_{2,y}\right) \\ P_n\left(\nabla_z V_n \circ D_{2,z}\right) \end{pmatrix}.$$

This form is advantageous as during the CBCT-based fitting (when determining λ_θ) all the other components ($\sum_n \sum_{pixels} C_\theta C_\phi$) can be determined at the be-

ginning of the optimisation. If this were not done, projections of a continually varying volume would need to be recalculated and summed at each trial value of the motion model parameters, greatly increasing optimisation time.

The steps can be taken with the tumour-only motion model equation to obtain the associated cost function. Alternatively, one can achieve the same result by slightly modifying the constrained non-rigid cost function. Simply set the weighting arrays to 1 within the tumour region (and zero outside). During the calculation of the cost function, the tumour enhanced projections will be automatically calculated within the cost function, hence the tumour enhancement method does not need to be applied with this modified approach.

References

[1] John R Adler Jr, SD Chang, MJ Murphy, J Doty, P Geis, and SL Hancock. The cyberknife: a frameless robotic system for radiosurgery. *Stereotactic and functional neurosurgery*, 69(1-4):124–128, 1997.

[2] Hamideh Alasti, M Peter Petric, Charles N Catton, and Padraig R Warde. Portal imaging for evaluation of daily on-line setup errors and off-line organ motion during conformal irradiation of carcinoma of the prostate. *International Journal of Radiation Oncology* Biology* Physics*, 49(3):869–884, 2001.

[3] A.H. Andersen and A.C. Kak. Simultaneous algebraic reconstruction technique (sart): A superior implementation of the {ART} algorithm. *Ultrasonic Imaging*, 6(1):81 – 94, 1984.

[4] N Becker, W L Smith, S Quirk, and I Kay. Using cone-beam ct projection images to estimate the average and complete trajectory of a fiducial marker moving with respiration. *Physics in Medicine and Biology*, 55(24):7439, 2010.

[5] Arjan Bel, Marcel van Herk, and Joos V Lebesque. Target margins for random geometrical treatment uncertainties in conformal radiotherapy. *Medical physics*, 23:1537, 1996.

[6] Arjan Bel, Pieter H Vos, Patrick TR Rodrigus, Carien L Creutzberg, Andries G Visser, Joep C Stroom, and Joos V Lebesque. High-precision prostate cancer irradiation by clinical application of an offline patient setup verification procedure, using portal imaging. *International Journal of Radiation Oncology* Biology* Physics*, 35(2):321–332, 1996.

[7] Stanley H Benedict, Kamil M Yenice, David Followill, James M Galvin, William Hinson, Brian Kavanagh, Paul Keall, Michael Lovelock, Sanford Meeks, Lech Papiez, et al. Stereotactic body radiation therapy: the report of aapm task group 101. *Medical physics*, 37:4078, 2010.

[8] C Bert and M Durante. Motion in radiotherapy: particle therapy. *Physics in Medicine and Biology*, 56(16):R113, 2011.

[9] Juergen Biederer, Julien Dinkel, Gregor Remmert, Siri Jetter, Simeon Nill, Torsten Moser, Rolf Bendl, Carsten Thierfelder, Michael Fabel, Uwe Oelfke, Michael Bock, Christian Plathow, Hendrik Bolte, Thomas Welzel, Beata Hoffmann, Gnter Hartmann, Wolfgang Schlegel, Jrgen Debus, Martin Heller, and Hans-Ulrich Kauczor. 4d-imaging of the lung: Reproducibility of lesion size and displacement on helical ct, mri, and cone beam ct in a ventilated ex vivo system. *International Journal of Radiation Oncology*Biology*Physics*, 73(3):919 – 926, 2009.

[10] Keith R. Britton, George Starkschall, Susan L. Tucker, Tinsu Pan, Christopher Nelson, Joe Y. Chang, James D. Cox, Radhe Mohan, and Ritsuko Komaki. Assessment of gross tumor volume regression and motion changes during radiotherapy for nonsmall-cell lung cancer as measured by four-dimensional computed tomography. *International Journal of Radiation Oncology*Biology*Physics*, 68(4):1036 – 1046, 2007.

[11] C. G. Broydon. The convergence of a class of double-rank minimization algorithms. *Journal Int. Math Applic.*, 6:76–90, 1970.

[12] Byungchul Cho, Per Rugaard Poulsen, and Paul J Keall. Real-time tumor tracking using sequential kv imaging combined with respiratory monitoring: a general framework applicable to commonly used igrt systems. *Physics in Medicine and Biology*, 55(12):3299, 2010.

[13] Wilfried De Neve, F Van den Heuvel, Marc Coghe, D Verellen, M De Beukeleer, A Roelstraete, P De Roover, L Thon, and G Storme. Interactive use of on-line portal imaging in pelvic radiation. *International Journal of Radiation Oncology* Biology* Physics*, 25(3):517–524, 1993.

[14] V Delmon, S Rit, R Pinho, and D Sarrut. Registration of sliding objects using direction dependent b-splines decomposition. *Physics in Medicine and Biology*, 58(5):1303, 2013.

[15] Claude Dupont. Recent developments making iba proton therapy system more accessible. *Nuclear Medicine Review*, 2011.

[16] Jan Ehrhardt, Tobias Klinder, and Cristian Lorenz. Computational motion phantoms and statistical models of respiratory motion. In Jan Ehrhardt and Cristian Lorenz, editors, *4D Modeling and Estimation of Respiratory Motion for Radiation Therapy*, Biological and Medical Physics, Biomedical Engineering, pages 215–247. Springer Berlin Heidelberg, 2013.

[17] Aurora Fassi, Joël Schaerer, Mathieu Fernandes, Marco Riboldi, David Sarrut, and Guido Baroni. Tumor tracking method based on a deformable 4d ct breathing motion model driven by an external surface surrogate. *International Journal of Radiation Oncology* Biology* Physics*, 88(1):182–188, 2014.

[18] Martin F Fast, Eric Wisotzky, Uwe Oelfke, and Simeon Nill. Actively triggered 4d cone-beam ct acquisition. *Medical physics*, 40(9):091909, 2013.

[19] LA Feldkamp, LC Davis, and JW Kress. Practical cone-beam algorithm. *JOSA A*, 1(6):612–619, 1984.

[20] JD Fenwick, WA Tome, HA Jaradat, SK Hui, JA James, JP Balog, CN DeSouza, DB Lucas, GH Olivera, TR Mackie, et al. Quality assurance of a helical tomotherapy machine. *Physics in medicine and biology*, 49(13):2933, 2004.

[21] Jacques Ferlay, Hai-Rim Shin, Freddie Bray, David Forman, Colin Mathers, and Donald Maxwell Parkin. Estimates of worldwide burden of cancer in 2008: Globocan 2008. *International journal of cancer*, 127(12):2893–2917, 2010.

[22] Christelle Gendrin, Hugo Furtado, Christoph Weber, Christoph Bloch, Michael Figl, Supriyanto Ardjo Pawiro, Helmar Bergmann, Markus Stock, Gabor Fichtinger, Dietmar Georg, et al. Monitoring tumor motion by real time 2d/3d registration during radiotherapy. *Radiotherapy and oncology*, 102(2):274–280, 2012.

[23] Patricia R. Geraghty, Stephen T. Kee, Gillian McFarlane, Mahmood K. Razavi, Daniel Y. Sze, and Michael D. Dake. Ct-guided transthoracic needle aspiration biopsy of pulmonary nodules: Needle size and pneumothorax rate1. *Radiology*, 229(2):475–481, 2003.

[24] J Gildersleve, DP Dearnaley, PM Evans, M Law, C Rawlings, and W Swindell. A randomised trial of patient repositioning during radiotherapy using a megavoltage imaging system. *Radiotherapy and Oncology*, 31(2):161–168, 1994.

[25] Matthias Guckenberger, Anthony Kavanagh, Steve Webb, and Michael Brada. A novel respiratory motion compensation strategy combining gated beam delivery and mean target position concept a compromise between small safety margins and long duty cycles. *Radiotherapy and Oncology*, 98(3):317 – 322, 2011.

[26] Jacob Hinkle, Martin Szegedi, Brian Wang, Bill Salter, and Sarang Joshi. 4d ct image reconstruction with diffeomorphic motion model. *Medical image analysis*, 16(6):1307–1316, 2012.

[27] Simon Hughes, James McClelland, Segolene Tarte, David Lawrence, Shahreen Ahmad, David Hawkes, and David Landau. Assessment of two novel ventilatory surrogates for use in the delivery of gated/tracked radiotherapy for non-small cell lung cancer. *Radiotherapy and Oncology*, 91(3):336–341, 2009.

[28] J. McClelland J. Martin and D. Hawkes. Building motion models of lung tumours from cone-beam ct. *Physics in Medicine and Biology*, 2013.

[29] Anthony Kavanagh, Philip M Evans, Vibeke N Hansen, and Steve Webb. Obtaining breathing patterns from any sequential thoracic x-ray image set. *Physics in Medicine and Biology*, 54(16):4879, 2009.

[30] Paul J. Keall, Gig S. Mageras, James M. Balter, Richard S. Emery, Kenneth M. Forster, Steve B. Jiang, Jeffrey M. Kapatoes, Daniel A. Low, Martin J. Murphy, Brad R. Murray, Chester R. Ramsey, Marcel B. Van Herk, S. Sastry Vedam, John W. Wong, and Ellen Yorke. The management of respiratory motion in radiation oncology report of aapm task group 76. *Medical Physics*, 33(10):3874–3900, 2006.

[31] Stine S Korreman, Anders N Pedersen, Trine Jakobi Nøttrup, Lena Specht, and Håkan Nyström. Breathing adapted radiotherapy for breast cancer: comparison of free breathing gating with the breath-hold technique. *Radiotherapy and oncology*, 76(3):311–318, 2005.

[32] A Krauss, S Nill, and U Oelfke. The comparative performance of four respiratory motion predictors for real-time tumour tracking. *Physics in medicine and biology*, 56(16):5303, 2011.

[33] Andreas Krauss, Martin F Fast, Simeon Nill, and Uwe Oelfke. Multileaf collimator tracking integrated with a novel x-ray imaging system and external surrogate monitoring. *Physics in Medicine and Biology*, 57(8):2425, 2012.

[34] Hideo D Kubo and Bruce C Hill. Respiration gated radiotherapy treatment: a technical study. *Physics in Medicine and Biology*, 41(1):83, 1996.

[35] John H Lewis, Ruijiang Li, Xun Jia, W Tyler Watkins, Yifei Lou, William Y Song, and Steve B Jiang. Mitigation of motion artifacts in cbct of lung tumors based on tracked tumor motion during cbct acquisition. *Physics in Medicine and Biology*, 56(17):5485, 2011.

[36] John H Lewis, Ruijiang Li, W Tyler Watkins, Joshua D Lawson, W Paul Segars, Laura I Cervio, William Y Song, and Steve B Jiang. Markerless lung tumor tracking and trajectory reconstruction using rotational cone-beam projections: a feasibility study. *Physics in Medicine and Biology*, 55(9):2505, 2010.

[37] Ruijiang Li, John H Lewis, Xun Jia, Xuejun Gu, Michael Folkerts, Chunhua Men, William Y Song, and Steve B Jiang. 3d tumor localization through real-time volumetric x-ray imaging for lung cancer radiotherapy. *Medical Physics*, 38:2783, 2011.

[38] Daniel A. Low, Parag J. Parikh, Wei Lu, James F. Dempsey, Sasha H. Wahab, James P. Hubenschmidt, Michelle M. Nystrom, Maureen Handoko, and Jeffrey D. Bradley. Novel breathing motion model for radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, 63(3):921 – 929, 2005.

[39] T Rock Mackie, Timothy Holmes, Stuart Swerdloff, Paul Reckwerdt, Joseph O Deasy, James Yang, Bhudatt Paliwal, and Timothy Kinsella. Tomotherapy: a new concept for the delivery of dynamic conformal radiotherapy. *Medical physics*, 20:1709, 1993.

[40] Thomas Rockwell Mackie, Jeff Kapatoes, Ken Ruchala, Weiguo Lu, Chuan Wu, Gustavo Olivera, Lisa Forrest, Wolfgang Tome, Jim Welsh, Robert Jeraj, et al. Image guidance for precise conformal radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*, 56(1):89–105, 2003.

[41] Gikas S Mageras and Ellen Yorke. Deep inspiration breath hold and respiratory gating strategies for reducing organ motion in radiation treatment. *Seminars in Radiation Oncology*, 14(1):65 – 75, 2004.

[42] Robert Manzke. Cardiac cone-beam ct. *Medical Physics*, 32:3227, 2005.

[43] J. McClelland, J. Martin, and D. Hawkes. 457 poster lung tumour tracking and motion compensated reconstruction in cone beam ct. *Radiotherapy and Oncology*, 99, Supplement 1(0):S184 –, 2011.

[44] J McClelland, J Martin, and D Hawkes. 457 poster lung tumour tracking and motion compensated reconstruction in cone beam ct. *Radiotherapy and Oncology*, 99:S184, 2011.

[45] J R McClelland, S Hughes, M Modat, A Qureshi, S Ahmad, D B Landau, S Ourselin, and D J Hawkes. Inter-fraction variations in respiratory motion models. *Physics in Medicine and Biology*, 56(1):251, 2011.

[46] Jamie R. McClelland, Jane M. Blackall, Sègolène Tarte, Adam C. Chandler, Simon Hughes, Shahreen Ahmad, David B. Landau, and David J. Hawkes. A continuous 4d motion model from multiple respiratory cycles for use in lung radiotherapy. *Medical Physics*, 33(9):3348–3358, 2006.

[47] J.R. McClelland, D.J. Hawkes, T. Schaeffter, and A.P. King. Respiratory motion models: A review. *Medical Image Analysis*, 17(1):19 – 42, 2013.

[48] Marc Modat, Gerard R. Ridgway, Zeike A. Taylor, Manja Lehmann, Josephine Barnes, David J. Hawkes, Nick C. Fox, and Sbastien Ourselin. Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine*, 98(3):278 – 284, 2010.

[49] Christopher Nelson, George Starkschall, Peter Balter, Rodolfo C. Morice, Craig W. Stevens, and Joe Y. Chang. Assessment of lung tumor motion and setup uncertainties using implanted fiducials. *International Journal of Radiation Oncology*Biology*Physics*, 67(3):915 – 923, 2007.

[50] J. J. Nuyttens, J.-B. Prvost, J. Praag, M. Hoogeman, R. J. Van Klaveren, P. C. Levendag, and P. M. T. Pattynama. Lung tumor tracking during stereotactic radiotherapy treatment with the cyberknife: Marker placement and early results. *Acta Oncologica*, 45(7):961–965, 2006.

[51] Freddy Odille, Pierre-Andr Vuissoz, Pierre-Yves Marie, and Jacques Felblinger. Generalized reconstruction by inversion of coupled systems (grics) applied to free-breathing mri. *Magnetic Resonance in Medicine*, 60(1):146–157, 2008.

[52] Per Rugaard Poulsen, Byungchul Cho, and Paul J. Keall. A method to estimate mean position, motion magnitude, motion correlation, and trajectory of a tumor from cone-beam ct projections for image-guided radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, 72(5):1587 – 1596, 2008.

[53] Min Rao, Wensha Yang, Fan Chen, Ke Sheng, Jinsong Ye, Vivek Mehta, David Shepard, and Daliang Cao. Comparison of elekta vmat with helical tomotherapy and fixed field imrt: plan quality, delivery efficiency and accuracy. *Medical physics*, 37:1350, 2010.

[54] Kristin J. Redmond, Danny Y. Song, Jana L. Fox, Jessica Zhou, C. Nicole Rosenzweig, and Eric Ford. Respiratory motion changes of lung tumors over the course of radiation therapy based on respiration-correlated four-dimensional computed tomography scans. *International Journal of Radiation Oncology*Biology*Physics*, 75(5):1605 – 1612, 2009.

[55] Vincent M Remouchamps, Nicola Letts, Frank A Vicini, Michael B Sharpe, Larry L Kestin, Peter Y Chen, Alvaro A Martinez, and John W Wong. Initial clinical experience with moderate deep-inspiration breath hold using an active breathing control device in the treatment of patients with left-sided breast cancer using external beam radiation therapy. *International Journal of Radiation Oncology* Biology* Physics*, 56(3):704–715, 2003.

[56] S Rit, M Vila Oliva, S Brousmiche, R Labarbe, D Sarrut, GC Sharp, et al. The reconstruction toolkit (rtk), an open-source cone-beam ct reconstruction toolkit based on the insight toolkit (itk). In *International Conference on the Use of Computers in Radiation Therapy (ICCR) 2013*, 2013.

[57] Simon Rit, Jochem W. H. Wolthaus, Marcel van Herk, and Jan-Jakob Sonke. On-the-fly motion-compensated cone-beam ct using an a priori model of the respiratory motion. *Medical Physics*, 36(6):2283–2296, 2009.

[58] D. Rueckert, L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, and D.J. Hawkes. Nonrigid registration using free-form deformations: application to breast mr images. *Medical Imaging, IEEE Transactions on*, 18(8):712–721, aug. 1999.

[59] Achim Schweikard, Hiroya Shiomi, and John Adler. Respiration tracking in radiosurgery. *Medical Physics*, 31(10):2738–2741, 2004.

[60] Yvette Seppenwoolde, Hiroki Shirato, Kei Kitamura, Shinichi Shimizu, Marcel van Herk, Joos V. Lebesque, and Kazuo Miyasaka. Precise and real-time measurement of 3d tumor motion in lung due to breathing and heartbeat, measured during radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, 53(4):822 – 834, 2002.

[61] Shinichi Shimizu, Hiroki Shirato, Shigeaki Ogura, Hirotoshi Akita-Dosaka, Kei Kitamura, Takeshi Nishioka, Kenji Kagei, Masaji Nishimura, and Kazuo Miyasaka. Detection of lung tumor movement in real-time tumor-tracking radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, 51(2):304 – 310, 2001.

[62] Hiroki Shirato, Shinichi Shimizu, Tatsuya Kunieda, Kei Kitamura, Marcel van Herk, Kenji Kagei, Takeshi Nishioka, Seiko Hashimoto, Katsuhisa Fujita, Hidefumi Aoyama, et al. Physical aspects of a real-time tumor-tracking system for gated radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*, 48(4):1187–1195, 2000.

[63] Bruce D Smith. Image reconstruction from cone-beam projections: necessary and sufficient conditions and reconstruction methods. *Medical Imaging, IEEE Transactions on*, 4(1):14–25, 1985.

[64] Jan-Jakob Sonke, Joos Lebesque, and Marcel van Herk. Variability of four-dimensional computed tomography patient models. *International Journal of Radiation Oncology*Biology*Physics*, 70(2):590 – 598, 2008.

[65] Jan-Jakob Sonke, Maddalena Rossi, Jochem Wolthaus, Marcel van Herk, Eugene Damen, and Jose Belderbos. Frameless stereotactic body radiotherapy for lung cancer using four-dimensional cone beam ct guidance. *International Journal of Radiation Oncology*Biology*Physics*, 74(2):567 – 574, 2009.

[66] Jan-Jakob Sonke, Lambert Zijp, Peter Remeijer, and Marcel van Herk. Respiratory correlated cone beam ct. *Medical Physics*, 32(4):1176–1186, 2005.

[67] J-P Thirion. Image matching as a diffusion process: an analogy with maxwell's demons. *Medical image analysis*, 2(3):243–260, 1998.

[68] A. N. Tikhonov. *Solutions of Ill Posed Problems*. Winston, Washington and New York, 1977.

[69] Henrik Turbell. *Cone-beam reconstruction using filtered backprojection*. PhD thesis, Linköping, 2001.

[70] Marcel van Herk, Allison Bruce, AP Guus Kroes, Tarek Shouman, Adriaan Touw, and Joos V Lebesque. Quantification of organ motion during conformal radiotherapy of the prostate by three dimensional image registration. *International Journal of Radiation Oncology* Biology* Physics*, 33(5):1311–1320, 1995.

[71] Dirk Verellen, Mark De Riddeer, Nadine Linthout, Koen Tournel, Guy Soete, and Guy Storme. Innovations in image-guided radiotherapy. *Nature Review*, 2007.

[72] Dirk Wolff, Florian Stieler, Grit Welzel, Friedlieb Lorenz, Yasser Abo-Madyan, Sabine Mai, Carsten Herskind, Martin Polednik, Volker Steil, Frederik Wenz, et al. Volumetric modulated arc therapy (vmat) vs. serial tomotherapy, step-and-shoot imrt and 3d-conformal rt for treatment of prostate cancer. *Radiotherapy and Oncology*, 93(2):226–233, 2009.

[73] Di Yan, John Wong, Frank Vicini, Jeff Michalski, Cheng Pan, Arthur Frazier, Eric Horwitz, and Alvaro Martinez. Adaptive modification of treatment

planning to minimize the deleterious effects of treatment setup errors. *International Journal of Radiation Oncology* Biology* Physics*, 38(1):197–206, 1997.

[74] Di Yan, John W Wong, Gary Gustafson, and Alvaro Martinez. A new model for accept or reject strategies in off-line and on-line megavoltage treatment evaluation. *International Journal of Radiation Oncology* Biology* Physics*, 31(4):943–952, 1995.

[75] Rongping Zeng, J.A. Fessler, and J.M. Balter. Estimating 3-d respiratory motion from orbiting views by tomographic image registration. *Medical Imaging, IEEE Transactions on*, 26(2):153 –163, feb. 2007.

[76] Qinghui Zhang, Yu-Chi Hu, Fenghong Liu, Karyn Goodman, Kenneth E Rosenzweig, and Gig S Mageras. Correction of motion artifacts in cone-beam ct using a patient-specific respiratory motion model. *Medical physics*, 37:2901, 2010.

[77] Lambert Zijp, Jan-Jacob Sonke, and Marcel van Herk. Extraction of the respiratory signal from sequential thorax cone-beam x-ray images. In *International Conference on the Use of Computers in Computational Radiotherapy 2004*, 2004.