King's Research Portal

DOI:

[10.1016/j.media.2011.08.003](https://doi.org/10.1016/j.media.2011.08.003)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

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

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Thoracic Respiratory Motion Estimation from MRI Using a Statistical Model and a 2-D Image Navigator

A. P. King^{a,*}, C. Buerger^a, C. Tsoumpas^a, P. Marsden^a, T. Schaeffter^a

^aDivision of Imaging Sciences and Biomedical Engineering, King's College, 4th Floor Lambeth Wing, St Thomas' Hospital, London, SE1 7EH, U.K, and NIHR Biomedical Research Centre at Guy's & St Thomas' NHS Foundation Trust and King's College London, U.K.

Abstract

Respiratory motion models have potential application for estimating and correcting the effects of motion in a wide range of applications, for example in PET-MR imaging. Given that motion cycles caused by breathing are only approximately repeatable, an important quality of such models is their ability to capture and estimate the intra- and inter-cycle variability of the motion. In this paper we propose and describe a technique for free-form nonrigid respiratory motion correction in the thorax. Our model is based on a principal component analysis of the motion states encountered during different breathing patterns, and is formed from motion estimates made from dynamic 3-D MRI data. We apply our model using a data-driven technique based on a 2-D MRI image navigator. Unlike most previously reported work in the literature, our approach is able to capture both intra- and inter-cycle motion variability. In addition, the 2-D image navigator can be used to estimate how applicable the current motion model is, and hence report when more imaging data is required to update the model. We also use the motion model to decide on the best positioning for the image navigator. We validate our approach using MRI data acquired from 10 volunteers and demonstrate improvements of up to 40.5% over other reported motion modelling approaches, which corresponds to 61% of the overall respiratory motion present. Finally we demonstrate one potential application of our technique: MRI-based motion correction of real-time PET data for simultaneous PET-MRI acquisition.

*Corresponding author. andrew.king@kcl.ac.uk, Tel.: +44-2071888259, Fax: +44-2071885442

Keywords: Respiratory motion, modelling, statistical model, PCA, MRI, PET-MR

1. Introduction

Motion caused by respiration can cause difficulties in a wide range of image acquisition applications and image-guided interventions. For example, in lung radiotherapy it can cause healthy tissue to be irradiated (Seppenwoolde et al., 2002); in image-guided cardiac interventions it can cause misalignment between the guidance information and the anatomy (King et al., 2009a); and in magnetic resonance imaging (MRI) it can result in artefacts that degrade image quality (Manke et al., 2002). Our intended application is motion-corrected positron emission tomography (PET) reconstruction in the thorax using a simultaneous PET-MRI system (Tsoumpas et al., 2010). In PET imaging, respiratory motion can cause artefacts due to motion of the target structures and also due to errors in attenuation correction information. In a simultaneous PET-MRI system, it is not feasible to continuously acquire MRI data for motion estimation due to the requirement to perform other MRI imaging for diagnostic purposes. Therefore the aim of our project is to form a subject-specific motion model from MRI data and use a fast two-dimensional (2-D) image navigator to apply the model to motion correct the simultaneously acquired PET data. Extra MRI imaging can be acquired in addition to the 2-D image navigator.

Motion models are typically formed from imaging data acquired at different points in the breathing cycle. These images are coregistered to estimate motion fields, which are used to form a model that can subsequently interpolate the motion at any point in the cycle. The conventional way to apply motion models is to acquire one or more one-dimensional (1-D) breathing signals, or *surrogates* of the true motion, which are used as inputs to estimate motion using a model (Manke et al., 2002; Shechter et al., 2005; McClelland et al., 2006; King et al., 2009a). Examples of common surrogate signals are the translation of the diaphragm (Manke et al., 2002; Shechter et al., 2005; King et al., 2009a) or the motion of surface markers on the abdomen (McClelland et al., 2006). Alternatively, a data-driven approach can be applied, in which the motion model is used to constrain a registration between a reference image (usually acquired at end-exhale) and a dynamically acquired 2-D (Blackall et al., 2005) or 3-D (King et al., 2010) target image (in these cases ultrasound images) covering all or part of the region of interest. The

value(s) of the surrogate signal(s) is optimised to maximise the similarity between the warped reference image and the target image. The key contribution of these works is that the information in the target image allows more accurate motion estimates than the use of simple 1-D input signals, but that the motion model will constrain the registrations to physiologically more plausible types of motion.

A common assumption of most previous work on respiratory motion modelling and estimation is that the motion caused by respiration is *repeatable*, i.e. that the motion paths are the same from cycle to cycle (Manke et al., 2002; Blackall et al., 2005; McClelland et al., 2006; Ehrhardt et al., 2009; King et al., 2009a) or between the expiration and inspiration phases of a single cycle (Manke et al., 2002; Blackall et al., 2005; Ehrhardt et al., 2009). However, this is only approximately true, and a number of papers have reported significant variation in breathing motion paths (Sonke et al., 2008; Blackall et al., 2006; von Siebenthal et al., 2007; Mutaf et al., 2010; McClelland et al., 2011). Breathing motion variability can be classified as either *intra-cycle* variability (i.e. the *hysteresis* effect (Nehrke et al., 2001), in which the motion observed during the inspiration phase is different to the motion observed during expiration) or *inter-cycle* variability (in which the motion paths are different from one breathing cycle to the next). Given the need for increased accuracy of motion estimates in many applications (such as PET-MRI), it is desirable to construct motion models that can capture and estimate these variabilities to some degree. A number of groups have proposed solutions to the problem of intra-cycle variability by classifying data as inspiration or expiration and modelling each separately (Shechter et al., 2005; McClelland et al., 2006; Zhang et al., 2007; King et al., 2009a). Proposals to deal with inter-cycle variability have been less widespread, although some do exist. For example, in image-guided cardiac interventions, we proposed a MRI-derived motion model that could adapt to motion changes caused by changes in breathing depth (King et al., 2009b). In lung motion modelling, Low et al. (2005) and Yang et al. (2008) proposed motion models derived from 4-D computed tomography (4DCT) data that could estimate breathing variability using a derivative of the input signal to the model, as well as the signal itself. However, in these works the 4DCT images that the models were formed from were retrospectively constructed from segments acquired during multiple breathing cycles, and relatively small numbers of images were used (less than 20). Therefore, although the models could, in principle, estimate inter-cycle variability, the data used meant that only intra-cycle variability

was captured.

Some groups have addressed the issue of breathing variability through the use of statistical techniques such as principal component analysis (PCA) to determine the statistical modes of variation of the motion fields during respiration. For example, in Klinder et al. (2008) PCA models of respiratory motion were formed (from CT), but only inhale and exhale pairs of images were used to form the models, which limited their ability to describe breathing variation at intermediate breathing states. Manke et al. (2003) proposed a technique for constructing affine cardiac respiratory motion models from MRI data in which PCA was used to capture breathing variation and also to establish the relationship between the PCA modes of variation and the input signals to the model (the current MRI navigator and a precursory navigator). This approach was extended to thoracic free-form deformations by Zhang et al. (2007). However, Zhang et al. (2007) employed 4DCT data to form the model, and so suffered from the same limitations as Low et al. (2005) and Yang et al. (2008): there were a small number of images used to form the model (10 in this case), each formed from segments acquired during multiple breathing cycles. Therefore, the model captured only intra-cycle variability and not inter-cycle variability. Klinder et al. (2010) also recently used a statistical approach for respiratory motion modelling of the lungs, again based on 4DCT data.

Statistical approaches such as PCA have also been used to form inter-subject respiratory motion models (Ehrhardt et al., 2009, 2011). However, in this work the model, once adapted to individual subjects, represented the motion during an average breathing cycle and inter-cycle motion variability was not addressed.

Most of the work described above has made the motion estimates to form the models from CT data. Related work on MRI-based techniques in the lung includes von Siebenthal et al. (2007), who made respiratory motion estimates using retrospective stacking of MRI images. They also investigated breathing variability using these estimates. In Blackall et al. (2006), MRI-derived affine motion models of the lung were used to investigate inter- and intra-cycle motion variability. MRI-based techniques have been more common in other organs such as the liver (Blackall et al., 2005) and the heart (Manke et al., 2002, 2003; King et al., 2009a,b).

In summary, a wide range of approaches have been used for thoracic motion modelling and estimation. The vast majority of these have been based on the use of 1-D input signals, or *surrogates*, to drive the model. Very

few data-driven approaches have been described. This has been due mainly to the difficulty in acquiring real-time CT or MRI data to be used in data-driven registrations. However, developments in MRI scanner technology such as fast gradients and multiple receive coils for parallel imaging now make such approaches feasible in MRI. In addition, in most previous work the models have been formed from 4DCT data which has limited their ability to capture inter-cycle motion variability. In this paper we present a technique for thoracic respiratory motion modelling and estimation in which subject-specific motion models are formed from dynamic 3-D MRI data and applied using a data-driven approach based on a 2-D MRI image navigator. We employ a PCA-based motion model that can capture both intra- and inter-cycle breathing variability. Our technique has several advantages that we investigate in this paper:

- the use of dynamic 3-D MRI data to form the model means that a large amount of imaging data, acquired during different breathing cycles, is available - this means that the PCA model will be able to capture and estimate inter-cycle as well as intra-cycle breathing motion variability;
- the detailed anatomical information available in the 2-D MRI image navigator can produce better motion estimates than those achieved using a simple 1-D surrogate signal, particularly when considering breathing motion variability;
- the registration between the motion model and the 2-D image navigator can be used to indicate whether or not a particular model can make sufficiently accurate motion estimates using the observed navigator data.

This last point is significant. A large amount of breathing motion variability occurs during changes in breathing pattern, such as the depth or speed of respiration (King et al., 2009b; von Siebenthal et al., 2007; Blackall et al., 2006). Therefore, if a model is formed during, for example, normal breathing, we need to establish if the accuracy of motion estimates is degraded when the breathing pattern changes to, for example, deep breathing. To the authors' knowledge, no motion modelling technique has yet exploited this potential.

2. Method and materials

Before describing our algorithm in detail, we first give an overview of our proposed workflow (see Figure 1). The first stage is the formation of the PCA-based model: this is based on an initial acquisition of 3-D dynamic

MRI data (detailed in Section 2.1), which is coregistered to estimate the motion fields for each dynamic image (see Section 2.2). The motion model is formed from these motion estimates (Section 2.3). The next stage is the application of the model based on a 2-D MRI image navigator. First, we choose the positioning of the navigator to optimise the motion estimation performance of the model. This is done by using the model to analyse the estimated motion at different potential navigator locations, as described in Section 2.4. Next, the 2-D image navigator data is continuously acquired and used as input to a data-driven application of the motion model (Section 2.5). After each application of the model a decision is made as to whether the motion estimates being made are good enough, i.e. it is checked whether the current model is still applicable to the observed motion. The basis for this decision is outlined in Section 2.6. If the model is not applicable then more 3-D dynamic MRI data is acquired and the model updated. Otherwise the model application continues using the current model.

We now describe each of these processes in more detail.

Figure 1: Illustration of proposed workflow. The section numbers (2.1, 2.2, etc.) refer to the corresponding section numbers in the text.

2.1. Dynamic 3D MRI acquisition

The 3-D dynamic MRI data were acquired on a 1.5 Tesla cylindrical bore MRI scanner (Philips Achieva I/T) using a 32 channel coil. The sequence details are:

- *Dynamic 3-D MRI*: T1-weighted FFE sequence, ECG triggered SENSE protocol with SENSE-factor 8, SENSE-factor of 2 in anterior-posterior

(AP) and SENSE-factor of 4 in right-left (RL). Typically $TR/TE = 3.3\text{ms}/0.9\text{ms}$, flip angle = 10° , field of view $500 \times 450 \times 245\text{mm}^3$, acquired image resolution $1.5 \times 4.1 \times 5\text{mm}^3$ (head-foot (HF),RL,AP), reconstructed image resolution $1.5 \times 1.5 \times 5\text{mm}^3$ (HF,RL,AP). The acquisition time for each 3-D dynamic image was $\sim 0.7\text{s}$, and acquisition was triggered to acquire one image every heart beat. Note that we have used a higher resolution in the HF direction as the majority of respiratory motion is in this direction.

Figure 2(a) shows a sample dynamic 3-D image acquired using this sequence. A 1-D navigator signal positioned on the dome of the right hemi-diaphragm was generated by postprocessing the images using a cross-correlation technique. Details of this navigator generation technique can be found in Savill et al. (2011). For the workflow we describe here, this 1-D navigator signal was only used to select the registration reference image (see Section 2.2) and not to form the PCA model. The 1-D navigator signal was also used in forming other types of models for comparison purposes (see Section 3.1).

Figure 2: Sample MRI images acquired from volunteer H: (a) A coronal slice through a sample 3-D dynamic MRI volume. The cardiac-triggered sequence acquires one such volume every heart beat; (b) A sample 2-D MRI image navigator image acquired at the same position as the coronal slice in (a).

MRI datasets were acquired from 10 volunteers, A-J, aged 22-37, 8 male and 2 female. In order to maximise the amount of inter-cycle motion vari-

ability, for each dataset, three subsets of 3-D dynamic MRI images were acquired for model formation: one during normal breathing, one for which the subject was instructed to breathe quickly and one for which the subject was instructed to breathe deeply. Our motivation for choosing these three breathing patterns was that they would maximise the variation in breathing amplitude and frequency, thereby increasing the likelihood that motion paths would be different from cycle to cycle. In addition, they are likely to occur in real subjects due to the long scanning time of PET imaging. 35 dynamic images were acquired in each subset, resulting in 105 dynamic 3-D MRI images in total for each volunteer. The 105 images typically took 70-100s to acquire, so the data covered approximately 30-50 breathing cycles. These images were used for model formation and also to simulate the 2-D image navigators for the experiments described in Sections 3.1 and 3.2.

2.2. Image registration

To estimate motion fields from the dynamic 3-D MRI images we used the hierarchical local affine registration algorithm we recently proposed in Buerger et al. (2010, 2011). Whereas many free-form registration algorithms work by optimising the displacements of a grid of control points (e.g. Rueckert et al. (1999)), this algorithm achieves efficiency improvements by successively breaking down the problem into simpler affine registrations, subsequently combining them into a single free-form deformation. The result of this algorithm is an $x/y/z$ displacement for each voxel in the image.

First, one of the images from each dataset, acquired at the end-exhale position, was chosen as a reference. The image with the maximum 1-D navigator value (computed as described in the previous section) was selected for this purpose. Next, every other image in the dataset was registered to the reference image. The result was a set of 3-D motion fields, one for each dynamic image apart from the reference image. We denote these motion fields by $\mathbf{u}_n(\mathbf{s})$, $n \in [1 \dots N]$, where $\mathbf{s}$ is a coordinate location in the images and n is the index of the dynamic image.

2.3. Model formation

Our subject-specific respiratory motion model is similar in concept to the PCA-based model described in Zhang et al. (2007), with the main difference being that our proposed approach uses a data-driven technique (using a 2-D image navigator) to apply the model, so we do not need to incorporate surrogate signals(s) into the PCA formulation.

Formally, we first manually define a regular $I \times J \times K$ 3-D grid of modelling control points $\phi_{i,j,k}$ across a region of interest in the dynamic MRI images. The rectangular region of interest used covered as much of the thorax as was visible in the field of view of the dynamic MRI images: typically extending beyond the thoracic cage into the arms in the RL direction, and between the bottom of the liver and the shoulders in the HF direction. The input data to the model consists of a set of 3-D motion estimates at the control points $\phi_{i,j,k}$ for each of the N dynamic images. These motion estimates were the result of the image registrations described in the previous section. The 3-D motion estimate at each control point for dynamic image n is denoted by $\mathbf{u}_n(\phi_{i,j,k}), n \in [1 \dots N]$.

To construct the PCA model, first, we form a vector, $\mathbf{p}_n$, for each dynamic image,

$$\mathbf{p}_n = (\mathbf{u}_n^x(\phi_{1,1,1}), \mathbf{u}_n^y(\phi_{1,1,1}), \mathbf{u}_n^z(\phi_{1,1,1}), \dots, \mathbf{u}_n^x(\phi_{I,J,K}), \mathbf{u}_n^y(\phi_{I,J,K}), \mathbf{u}_n^z(\phi_{I,J,K}))^T \quad (1)$$

These vectors represent the respiratory motion of the control points in each dynamic image, as estimated by the registration algorithm. Next, we define the mean motion state over all dynamic images:

$$\bar{\mathbf{p}} = \frac{1}{N} \sum_{n=1}^N \mathbf{p}_n \quad (2)$$

The motion of each vector, $\mathbf{p}_n$, away from this mean state is defined as,

$$\tilde{\mathbf{p}}_n = \mathbf{p}_n - \bar{\mathbf{p}} \quad (3)$$

Next, we construct a matrix consisting of the vectors $\tilde{\mathbf{p}}_n$ as columns,

$$\mathbf{P} = [\tilde{\mathbf{p}}_1, \tilde{\mathbf{p}}_2, \dots, \tilde{\mathbf{p}}_N] \quad (4)$$

At this point, the traditional approach to PCA-based shape models involves computing the eigenvalues and eigenvectors of the covariance matrix $\mathbf{P}\mathbf{P}^T$. However, the size of this covariance matrix is $(IJK) \times (IJK)$, which can be large for any reasonably dense control point spacing. Therefore, the computational complexity of the eigenvector computation places a limit on the density of the control point spacing. However, Zhang et al. (2007) described a technique for reducing the complexity of this computation, allowing for

denser control point spacings. They pointed out that there are only $N - 1$ non-zero eigenvalues of the covariance matrix. Furthermore the non-zero eigenvalues of $\mathbf{P}^T \mathbf{P}$ are the same as those of $\mathbf{P} \mathbf{P}^T$, and the eigenvectors of $\mathbf{P}^T \mathbf{P}$ are equal to the eigenvectors of $\mathbf{P} \mathbf{P}^T$ premultiplied by $\mathbf{P}$ (Zhang et al., 2007). Therefore, the problem can be reformulated to compute the eigenvalues and eigenvectors of $\mathbf{P}^T \mathbf{P}$, which is of size $N \times N$ (N is the number of dynamic images), and then premultiplying all resulting eigenvectors by $\mathbf{P}$. Details of this derivation can be found in Zhang et al. (2007), but note that the analytical result is the same as the traditional approach of finding the eigenvectors of $\mathbf{P} \mathbf{P}^T$. We denote the resulting eigenvectors by $\mathbf{E}_e$ and the eigenvalues by λ_e .

To apply the PCA-based model we need to specify a number of weights, one for each eigenvector being used. For a vector of f weights, $\mathbf{b} = (b_1, \dots, b_f)^T$, the control point positions can be computed as,

$$\hat{\mathbf{p}} = \bar{\mathbf{p}} + \mathbf{E} \mathbf{b} \quad (5)$$

where $\hat{\mathbf{p}}$ is an instance of the motion vector (i.e. the control point positions) and $\mathbf{E} = (\mathbf{E}_1, \dots, \mathbf{E}_f)$ is a matrix comprising the first f eigenvectors as columns. The length of the vector $\mathbf{b}$ represents the number of degrees of freedom of the PCA-based model.

Once the 3-D motions at each of the control points $\phi_{i,j,k}$ (from $\hat{\mathbf{p}}$) are known, the overall dense motion field is computed using B-spline interpolation (Rueckert et al., 1999). We denote this dense motion field by $\mathbf{v}(\cdot)$.

2.4. MRI 2-D image navigator selection and acquisition

The next stage is to apply the motion model based on acquired 2-D MRI image navigators. These image navigators are single 2-D coronal slices. The sequence details are:

- *2-D MRI image navigator*: T1-weighted FFE sequence, ECG triggered SENSE protocol with SENSE-factor of 4 in RL, typically TR/TE = 2.6ms/0.88ms, flip angle = 10° , field of view $450 \times 497 \times 10\text{mm}^3$, acquired image resolution $1.5 \times 5.5 \times 10\text{mm}^3$ (HF,RL,AP), reconstructed image resolution $1.5 \times 1.5 \times 10\text{mm}^3$ (HF,RL,AP), acquisition time 45.4ms.

A sample 2-D image navigator is shown in Figure 2(b).

The performance of the data-driven model application will depend to some degree on the spatial positioning of the 2-D image navigators. We

propose to select the navigator position automatically based on information extracted from the motion model. The basic principle we apply is that a potential navigator position is ‘useful’ if it is likely to contain a lot of information about the subject’s thoracic breathing motion and motion variability. Therefore, if the expected motion within a potential 2-D slice is high, then a navigator positioned there is likely to be more useful than one positioned in a region where the expected motion is small. In the following we will describe how this is accomplished.

The 2-D image navigator is a coronal slice, so we need to choose a coordinate value, y , along the AP direction for the navigator to be positioned at. We first define a slice selection term, $\xi(y)$, for each potential navigator position that expresses the amount of motion that would be expected if the navigator were positioned there,

$$\xi(y) = \left[\sum_{e=1}^f \sum_{b_e=-3\sqrt{\lambda_e}}^{+3\sqrt{\lambda_e}} \sum_{\mathbf{s} \in \Omega_y} (\mathbf{v}_{\mathbf{b}}(\mathbf{s}) \cdot \mathbf{HF})^2 \right]^{\frac{1}{2}} \quad (6)$$

where $\mathbf{v}_{\mathbf{b}}(\cdot)$ is the dense 3-D motion field generated using the set of PCA weights $\mathbf{b}$ (as described in Section 2.3), $\mathbf{HF}$ is a unit vector in the head-foot direction and $\mathbf{s} \in \Omega_y$ are the 3-D voxel coordinates within the domain of the 2-D coronal slice defined by y . The weights vector $\mathbf{b}$ consists of only zeros apart from the current mode, e , which varies between $-3\sqrt{\lambda_e}$ and $+3\sqrt{\lambda_e}$ in small steps. λ_e is the eigenvalue for mode e and represents the variance of the data along the corresponding eigenvector. Therefore, the weight for each mode varies within 3 standard deviations of the mean, within which most of the variation could be expected to lie. In practice, to speed up the computation of $\xi(y)$, we used a regularly sampled subset of the complete image domain, Ω_y (subsampled by a factor of 10 along the x/z axes).

An intuitive explanation of Eq. (6) is that, for a given slice position, we compute the root-mean-square of the motion vector components in the HF direction over all modes and over the entire 2-D image navigator. The root-mean-square was used as it provides a good measure of the expected motion, and we want a 2-D image navigator containing significant HF motion. To select the best navigator positioning, we pick the coordinate in the AP direction, $\hat{y}$, that maximises this term,

$$\hat{y} = \underset{y}{\operatorname{argmax}} \xi(y) \quad (7)$$

2.5. Registration of the model to the 2-D navigator image

We apply our PCA-based motion model using a data-driven approach based on a 2-D MRI image navigator as a target image. This approach is similar in concept to the use of ultrasound data as target images proposed by Blackall et al. (2005) and King et al. (2010) (see Section 1), with the exception that the signals we optimise are the weights of our PCA model.

We denote our 2-D MRI target image by I_{target} . To apply the model we estimate the set of PCA weights that yield the best match to the target image,

$$\hat{\mathbf{b}} = \underset{\mathbf{b}}{\operatorname{argmax}} \operatorname{Sim} \left(I_{\mathbf{b}}, I_{target} \right), \quad (8)$$

where $I_{\mathbf{b}}$ is the reference image warped using the motion field computed using the set of PCA weights, $\mathbf{b}$ (as described in Section 2.3) and Sim is the similarity measure. For all experiments in this paper we used normalised cross correlation (NCC) for the similarity measure. Note that $I_{\mathbf{b}}$ is a 3-D image whereas I_{target} is a 2-D image. Therefore the similarity measure is computed only in the area of overlap between the two images. We used a gradient descent scheme to optimise $\hat{\mathbf{b}}$.

2.6. Model applicability test

As well as providing the target image for the data-driven model application, the 2-D MRI image navigator is used to provide feedback on the applicability of the current motion model. We test the applicability of the current model by comparing the final similarity measure of an image navigator after registration (i.e. $\operatorname{Sim} \left(I_{\hat{\mathbf{b}}}, I_{target} \right)$, see Eq. (8)) to threshold values. To compute the thresholds we examine the final similarity measures of all 2-D image navigators acquired up to that point. Data is flagged as not applicable to the current model if:

- the similarities of the last three image navigators have been less than the overall mean similarity up to that point minus one standard deviation; or
- the current similarity is less than the overall mean similarity minus three standard deviations.

Assuming a normal distribution of similarities, approximately 16% of samples will be less than 1 standard deviation below the mean. Therefore, there would only be a 0.4% chance of three consecutive similarities falling below this threshold if the distribution of the similarities was unchanged, i.e the

breathing pattern was the same. There is a 0.1% chance of any sample being less than three standard deviations below the mean. Therefore, the chance of ‘false positives’ in this model applicability test should be small.

3. Experiments

We performed three experiments to validate the three main hypotheses of this work, namely that:

1. data-driven model application of a PCA-based motion model using a 2-D image navigator estimates motion more accurately than the standard surrogate based approach (i.e. a 1-D navigator);
2. the PCA-based motion model can be used to automatically select the best position for the 2-D image navigator;
3. the use of a 2-D image navigator provides feedback as to how applicable the model is to the current motion.

In addition we performed a fourth experiment to demonstrate the application of our proposed technique to correct PET data for the effects of respiratory motion.

Validation in experiments 1 and 2 was carried out using a leave-one-out test (see Figure 3):

- for each dynamic image, a motion model was formed using all dynamic images apart from the current one;
- a 2-D image navigator was simulated by extracting a single coronal slice from the left-out dynamic image;
- this simulated image navigator was used for registration with the motion model;
- the result of the registration was used to warp a reference 3-D image;
- finally, the warped image was compared to the entire 3-D left-out dynamic image.

3.1. Experiment 1 - comparison of model formation/application approaches

To test hypothesis 1 we compared four different approaches to model formation and application. All four techniques were based on the same motion estimates, computed at a regular grid of modelling control points, $\phi_{i,j,k}$, with a control point spacing of $10 \times 10 \times 10\text{mm}^3$. These motion estimates were determined using the hierarchical local affine registration algorithm (Buerger et al., 2010, 2011), as described in Section 2.2.

Figure 3: Validation using a leave-one-out test. Each dynamic 3-D image is left out in turn. A motion model is formed from all other dynamic 3-D images, and a 2-D image navigator is simulated by extracting a coronal slice from the left-out 3-D dynamic. Registration between the reference dynamic 3-D image and the 2-D image navigator using the motion model results in a warped 3-D image. This is compared with the original left-out 3-D dynamic image.

1. *Average cycle, surrogate driven.* This model is similar in concept to a number of average cycle motion models that have been described in the literature (Manke et al., 2002; Blackall et al., 2005; McClelland et al., 2006; Ehrhardt et al., 2009; King et al., 2009a). It represents the 3-D motion of the region of interest in an average breathing cycle. The 1-D MRI navigator signals described in Section 2.1 were used as the input signals for the model. The motion model consists of a set of second-order polynomial fits to the variation of the $x/y/z$ displacements of the control points as functions of the MRI navigator value. Data are classified into inspiration and expiration phases and polynomial fitting is performed for each phase separately with the constraint that the inspiration and expiration curves should meet at the extreme navigator positions (King et al., 2009a). Therefore each model is defined by a vector of 3 polynomial coefficients, α , for each control point location (i, j, k) , each coordinate direction $(x, y \text{ or } z)$ and each breathing phase,

d . To apply the model, the $x/y/z$ displacements for all control points are estimated based on a single input navigator value, γ , and a breathing phase, d (inspiration or expiration). The motion estimate at control point $\phi_{i,j,k}$ is defined by

$$\hat{\mathbf{u}}(\phi_{i,j,k}) = \begin{pmatrix} \Psi(\alpha_{i,j,k,x,d}, \gamma) \\ \Psi(\alpha_{i,j,k,y,d}, \gamma) \\ \Psi(\alpha_{i,j,k,z,d}, \gamma) \end{pmatrix} \quad (9)$$

where $\Psi(\alpha, \gamma)$ represents the application of a second-order polynomial function using the coefficients α and input variable γ . The breathing phase d is determined by comparing the current navigator value with its predecessor. Once the 3-D motions at each of the control points are determined the overall dense motion field is computed using B-spline interpolation (Rueckert et al., 1999). Note that this model captures intra-cycle variability but does not capture any inter-cycle variability in breathing motion.

2. *Average cycle, data driven.* This technique uses the same average cycle motion model described above. However, a data-driven approach to model application is employed, based on the 2-D MRI navigator image. The input signals to the model (1-D navigator γ and breathing phase d) are optimised to maximise the similarity between the 2-D navigator image I_{target} and the reference image warped using the model,

$$\hat{\gamma}, \hat{d} = \underset{\gamma, d}{\operatorname{argmax}} \operatorname{Sim}(I_{\gamma, d}, I_{target}) \quad (10)$$

where $I_{\gamma, d}$ is the reference image warped using the motion field computed using γ and d according to Eq. (9). This approach is similar to that used with an affine motion model and ultrasound data in Blackall et al. (2005) and King et al. (2010).

3. *PCA-based, surrogate driven.* This model is similar in concept to that described by Zhang et al. (2007). PCA is used to capture the modes of variation of the motion of the control points as well as the relationship between the PCA modes and the input signals. The 1-D navigator signal and a precursory navigator (i.e. the 1-D navigator signal at a previous time step) were used as input signals. A time step is defined to occur at an acquisition of a dynamic image. The images are cardiac-triggered so a time step in this case corresponds to one heart beat, or

between 0.7 second and 1 second. Based on our experiments we found a lag of 2 time steps produced the best results for this technique, hence as inputs we selected the navigator signals at times t and $t - 2$. Since 2 input signals were used, the limitations of the PCA technique employed mean that only the first 2 PCA modes could be employed. Details of this restriction are given in Zhang et al. (2007). This technique has previously been demonstrated to capture intra-cycle variation (Zhang et al., 2007). Here we validate it on data likely to contain inter- as well as intra-cycle variation.

4. *PCA-based, data driven.* This is the model formation and application approach we propose and describe in this paper. This approach represents the main novel contribution of our work. It can capture both intra-cycle and inter-cycle motion variation. Based on observations of the amount of variation captured by the major modes of variation, we tested this technique using the first 5 PCA modes in our experiments. To allow for a fair comparison with technique (3) we also tested it using the first 2 modes.

The data used to compare these different model formation approaches were the 3-D dynamic MRI data described in Section 2.1. The subsets for all three breathing patterns were combined and all 105 images for each volunteer were used to form the models. For each volunteer, the following images were excluded from the leave-one-out test validation:

- the reference image (end-exhale);
- the first two images in the sequence (these would have no known precursory navigator for technique (3));
- the end-inhale image (as it has the smallest 1-D navigator value, this would result in the navigator being out of the range of the leave-one-out model computed using techniques (1) and (2)).

The same coronal slices were used as 2-D image navigators for techniques (2) and (4) (i.e. automatically selected as described in Section 2.4).

Validation of registration accuracy was carried out using three measures:

- The final (registered) value of the similarity measure (i.e. NCC) between the left-out 3-D dynamic image and the warped 3-D dynamic image.
- The distance between the lung segmentations of the same two images. The lung segmentations were automatically determined as described in Buerger et al. (2011). Sample segmentation results are shown in Figure 4. We computed the root-mean-square of the shortest distances

between the points on the surface of the first segmentation and those on the surface of the second.

- The DICE coefficient between the lung segmentations of the same two images.

It was not possible to perform any landmark based validation owing to the poor signal-to-noise ratio of structural information inside the lungs.

Figure 4: Illustration of lung surface distance validation: (a) coronal slice through a left-out 3-D dynamic image, overlaid with contours from its automatic segmentation; (b) coronal slice through the corresponding warped dynamic image to be tested, overlaid with contours from its automatic segmentation; (c) rendering of surface of segmentation from (a), with surface colour indicating shortest distance to surface of segmentation from (b).

3.2. Experiment 2 - validation of automatic 2-D image navigator positioning

To test hypothesis 2 we first formed a PCA-based model and computed our proposed slice selection term as defined in Eq. (6). We then performed separate registrations using all possible image navigator positions (i.e. by extracting all coronal slices from the 3-D dynamic images). Finally, we compared the variation of the slice selection term (and the choice of the best slice) with the corresponding variation of registration accuracy. As for experiment 1, the data for all 3 breathing patterns were combined for use in this experiment. The final value of the NCC similarity measure was used as a measure of registration accuracy. We performed these experiments for all volunteers.

3.3. Experiment 3 - validation of model applicability test

To test hypothesis 3 we investigated the effectiveness of the final similarity measure in indicating changes in breathing motion. We formed a motion model using dynamic 3-D MRI data acquired during a single breathing pattern (normal breathing), and applied it using real 2-D image navigator data acquired during three breathing patterns (normal, fast and deep breathing), recording the value of the final 2-D similarity measure for each. The intention here was to investigate if the final 2-D similarity measure could detect whether a change in the nature of the breathing motion has occurred. We performed this experiment for volunteers B, C, E and H. Note that this experiment used real 2-D image navigators acquired as described in Section 2.4, and not simulated 2-D image navigators as used in experiments 1 and 2.

3.4. Experiment 4 - PET-MRI demonstration

As a final experiment, we provide a demonstration of the use of our proposed model formation/application technique (technique (4)) for our intended application of MRI-based motion correction of PET imaging in the human thorax. We simulated a 35 second motion-corrupted PET acquisition using the approach we have previously described in Tsoumpas et al. (2009). For each frame, one million coincidence events were simulated for the Philips Gemini TF PET scanner including the attenuation effect. All PET images ($128 \times 128 \times 87$ voxels of dimension $4 \times 4 \times 2\text{mm}^3$) were obtained using the STIR package (Thielemans et al., 2007) with the ordered subsets expectation maximisation (OSEM) reconstruction algorithm (23 subsets, 2 iterations). The PET simulations were formed from attenuation and emission maps that were manually segmented from high resolution MRI images (Tsoumpas et al., 2009) acquired from volunteer B. To simulate the effects of respiratory motion, these maps were warped using the motion fields estimated by the 3-D registrations (deep breathing section only, consisting of 35 dynamic images). 35 PET images were simulated from the warped attenuation and emission maps, corresponding to the respiratory positions of the 35 dynamic MRI images. An uncorrected overall PET image was formed by averaging all 35 simulated PET images. The motion corrected overall PET image was formed as follows:

- 2-D image navigators were simulated from the 3-D dynamic MRI images as described above;

- the PCA-based motion model for volunteer B (constructed from all 105 dynamic images from all breathing types) was used in the data-driven approach to model application to estimate a motion field for each corresponding dynamic image;
- these motion fields (which transform from end-exhale to inhale states) were inverted using the iterative technique described in Crum et al. (2007) so that they transformed inhale states to end-exhale;
- the inverted motion fields were applied to motion correct the corresponding simulated PET images;
- finally, all 35 motion corrected PET images were averaged to produce the motion corrected overall PET image.

Note that the original motion fields estimated by the registration algorithm were used for simulating the 35 PET images, but the motion fields estimated by a motion model were used for motion correcting them. Similar to experiments 1 and 2, we performed a leave-one-out test in this experiment, so for each of the 35 images, the motion field used in simulating the image was not used to form the motion model employed to motion correct it.

4. Results

4.1. Experiment 1 - comparison of model formation/application approaches

A summary of the registration results using the NCC, lung surface distance and DICE coefficient accuracy measures is shown in Figure 5. Our proposed data-driven PCA-based approach resulted in a NCC measure of 0.964 ± 0.023 , a lung surface distance of $2.75 \pm 0.73\text{mm}$ and a DICE coefficient of 0.951 ± 0.017 . It outperformed all other techniques on all three measures, and statistically significant improvements were found in all cases ($p \leq 0.01$ in a two-tailed paired t-test). Although the improvement in NCC may appear to be modest, this is a global intensity-based measure so improvements at local features may be masked by the similarity in intensities of larger areas such as the insides of the lungs. For the local lung surface distance measure, the percentage improvements of our approach over the surrogate-driven average-cycle, the data-driven average-cycle model and the surrogate-driven PCA-based model were 6.1%, 2.9% and 40.5%. To estimate the motion correction performance of the algorithms, we also computed the accuracy measures before registration. Over all volunteers and all breathing patterns, the NCC was 0.92 ± 0.068 , the lung surface distance was $7.12 \pm 5.21\text{mm}$ and the DICE coefficient was 0.876 ± 0.077 . This means that our data-driven

PCA-based approach recovers 61% of the motion on average based on the lung surface distance measure. Figure 6 shows a comparison of registrations performed using techniques (3) and (4) for a sample image acquired from volunteer H. The difference in NCC between these two registration results was only 0.01, whereas the difference in local error at the lung boundary can be seen to be significant.

For technique (4) (our proposed data-driven PCA-based approach), the results shown are for using the first 5 modes of variation, whereas for technique (3), because of the limitations of the formulation (i.e. the number of PCA modes cannot be more than the number of input surrogates) it was only possible to use the first 2 modes. To make a fairer comparison between the two techniques, we also ran technique (4) using only the first 2 modes. The corresponding mean/standard deviations for this technique using 2 modes were 0.96 ± 0.026 for the NCC, $2.9 \pm 0.82\text{mm}$ for the lung surface distance and 0.947 ± 0.019 for the DICE coefficient. These results are superior to those for techniques (1) and (3), but are similar to those of the data-driven average-cycle model (technique (2)). The reason for technique (4) not being superior to technique (2) when using only 2 modes is that the PCA eigenvectors represent linear variations of the control points from their average positions, whereas the average-cycle model uses polynomials that can represent non-linear motion paths. However, when 5 modes are used technique (4) is clearly superior.

We also computed the percentage of the total variation captured by the PCA models used in technique (4) for each volunteer: the first 2 modes captured 50.8 – 76.6% of the total variation, and the first 5 modes captured 61.2 – 88.7% of the variation.

4.2. Experiment 2 - validation of automatic 2-D image navigator positioning

Figure 7 illustrates a sample result of the automatic image navigator positioning validation. The graph on the left shows the slice selection measure together with the overall registration accuracy achieved by using each individual coronal slice as a 2-D image navigator. The right-hand graph is a zoomed version of the left-hand graph. The dashed/dotted lines in each graph show the slice chosen by our proposed automatic image navigator selection technique (red dashed) and the actual optimum based on running all registrations for each slice as an image navigator (blue dotted). In this case our automatic technique selected an image navigator position with an accuracy close to the optimum: the difference between the registration accuracy

(i.e. mean 3-D NCC value computed over all dynamic images) using the selected slice and that of the overall optimum slice was 0.05%. For the other 9 volunteers the differences were 0.29%, 0.26%, 0.05%, 0.14%, 0.4%, 0.24%, 0.13%, 0.17% and 0.14% respectively. Therefore in all cases the slice chosen resulted in an accuracy very close to the optimum.

4.3. Experiment 3 - validation of model applicability test

Figure 8 shows the results for the model applicability test. The left images show the within-slice NCC values after registration of a normal breathing motion model to 2-D image navigator data from all 3 breathing patterns. The right images show a demonstration of applying the model applicability test. These graphs show the data from which the statistics in the graphs in the left images were derived. The red squares show where the model was flagged as not being applicable to the observed data. The model applicability test was only applied after the first 10 dynamic images to enable reasonable estimates of the mean and standard deviation of the similarity to be made. We can see that for deep breathing the test quickly identifies that the breathing pattern has changed for all volunteers. For fast breathing the model is identified as being not applicable for two of the volunteers (C and H), but not for the other two. This suggests that for some subjects a normal breathing model is sufficiently accurate to estimate fast breathing motion, but not for others. There may have been some inconsistency in how the different volunteers interpreted the instruction to breathe quickly. Some volunteers (notably volunteer B) appear to have taken short breaths near to the end-exhale position, which is known to be quite consistent regardless of breathing pattern. This resulted in high values for the within-slice NCC. Others varied their respiratory position more, resulting in lower values. For normal breathing data the model is always accepted as being applicable for all volunteers. In our proposed workflow, when the model is flagged as being not applicable we would acquire extra 3-D dynamic images in order to ensure the continued high accuracy of the model.

4.4. Experiment 4 - PET-MRI demonstration

Figure 9(c) and (d) show the overall simulated PET images without and with motion correction using technique (4). The tissue boundaries in the uncorrected PET image are significantly more blurred and shifted in the posterior direction (i.e. the principal direction of respiratory motion). In contrast, the motion corrected PET image aligns well with the boundaries

extracted from the (end-exhale) emission image, demonstrating good motion correction performance. Figure 9(a) shows the emission image used for simulation, and Figure 9(b) shows a simulated PET image with no motion transformations applied. Such an image could not be acquired in practice, and we show it only for comparison purposes.

5. Discussion

We have presented a technique for thoracic motion modelling and estimation that can capture and estimate inter-cycle breathing variation as well as intra-cycle variation. The model is formed from MRI data and is based on a PCA of the motion state variability. We apply a data-driven approach to model application based on a 2-D MRI image navigator. The most closely related work is that of Zhang et al. (2007). However, our approach to applying the motion model differs, and is similar in concept to the data-driven approaches described in Blackall et al. (2005) and King et al. (2010): we use the model to constrain a registration to a dynamic 2-D MRI image navigator. This obviates the need to acquire any surrogate signals. Furthermore, this 2-D image navigator is able to provide us with valuable information about how reliable the motion estimates of the model currently are. Our intended application is MRI-based respiratory motion correction of PET imaging of the human thorax in a simultaneous PET-MRI system. However, we believe that our proposed technique has potential application in other areas such as MRI-guided interventions.

The main contributions of our work can be summarised as:

- the use of a large amount of MRI data acquired during different breathing cycles to form a statistical model of respiratory motion to capture inter-cycle motion variability;
- the use of a data-driven approach using a 2-D MRI image navigator to the application of a PCA-based motion model;
- the incorporation of feedback as to how well the model fits the observed motion;
- the use of the motion model to determine details of the image acquisition used for model application (i.e. 2-D image navigator positioning).

To the authors' knowledge, none of these points have previously been reported in the literature. Our technique has been shown to derive accurate thoracic motion estimates that can be acquired with a high temporal resolution, which is an important factor for PET-MRI. Over all volunteers,

the accuracy at the lung boundaries was reduced from $7.12 \pm 5.21\text{mm}$ before registration to $2.75 \pm 0.73\text{mm}$ after registration. Our data-driven PCA-based technique outperformed the three other registration approaches used for comparison, achieving a statistically significant improvement in all cases. We used 5 PCA modes for our statistical model, which captured between $61.2 - 88.7\%$ of the motion variation. More modes could be used, but there would be a trade-off between introducing more modes and the efficiency of the registration to the 2-D image navigator (each extra mode would add an extra degree of freedom to the registration). Although the anatomical plausibility (i.e. avoiding folding artefacts) of the motion fields was not directly evaluated in this paper, the registration algorithm that was employed has previously been shown to produce smooth motion fields (Buerger et al., 2011).

Our results suggest that the data-driven technique of motion model application is more robust than the use of 1-D surrogate signals. This is particularly true for the PCA-based model that uses the current surrogate and a precursory surrogate as inputs (i.e. technique (3)). We believe that the relatively poor performance of this approach was because this model must be able to generalise over a 2-parameter space (i.e. the surrogate and the precursor). Poor registrations seem to occur when the test data is at a point in this 2-parameter space that is far from the training data (i.e. the test data is an unfamiliar surrogate/precursor pairing). This lack of generalisation is more likely to happen in 2 parameter models than in 1 parameter models such as the average-cycle model (technique (1)). Also, the surrogate-driven PCA-based approach contains an implicit assumption that there is a constant relationship between the surrogate/precursor pair and the nature of the motion. This seems not to be the case when there are changes in breathing pattern such as from normal to deep or fast breathing. This makes the technique less suitable for capturing inter-cycle breathing motion variability.

Another important contribution of our work is the formation of a motion model that can capture breathing motion variability from MRI data. The use of MRI data, compared to the more commonly reported modality of 4DCT, allows us to acquire a large amount of dynamic images acquired during different breathing patterns, allowing the PCA to more fully capture inter-cycle variability. Most previously published work (Low et al., 2005; Yang et al., 2008; Zhang et al., 2007) has used 4DCT data consisting of a relatively small number of images, each formed from data acquired during multiple breathing cycles. This limits the resulting models' ability to capture inter-

cycle motion variability. The use of MRI data to form a model capable of capturing breathing motion variability has previously been reported only for affine cardiac motion (Manke et al., 2003), and not for nonrigid free-form deformations as we have used in this work.

However, despite the important advantages of using MRI data to form breathing motion variability motion models, there is also a disadvantage. We can see from the sample dynamic 3-D volume shown in Figure 2(a) that there is very little structural information inside the lungs. This means that the motion fields estimated by the registration algorithm within the lungs are mostly just interpolated from those estimated at the boundaries. For the same reason, it is also difficult to validate the accuracy of the motion estimates inside the lungs. In future work we plan to address this important issue by combining dynamic 3-D MRI data with dynamic 2-D MRI slices that can show improved contrast and structural information inside the lungs due to the in-flow effect of the blood.

The data-driven application of our PCA-based model does not currently work in real time. A single registration using 5 PCA modes takes about 2-3 minutes using our current implementation. PET-MRI motion correction can be performed retrospectively, so the lack of a real-time implementation is not a serious problem for motion correction alone. However, for the model applicability test we propose in Section 2.6 a real-time implementation would be required. We believe that by optimising our code and making use of GPU hardware this could be feasible.

We have also presented a potential technique for automatically positioning the 2-D image navigator slices based on information extracted from the motion model. Our experiments (experiment 2) showed that the slices selected by our technique had registration accuracies between 0.03%-0.4% of that of the overall optimum slice. Although this may seem a small amount, we pointed out in Section 4.1 that small differences in NCC can correspond to large local differences, for example at the lung boundaries. Therefore, although our proposed image navigator positioning technique performs reasonably well, there is still room for improvement, and in future work we will investigate further ways of automatically tuning some of the image acquisition parameters (such as field of view) to optimise the quality/utility of the images acquired.

In this work we have used PCA as a tool to analyse the statistical distribution of motion states. PCA is a linear technique that makes the important assumption that the motion states are normally distributed along the axes of

some reduced dimensionality Euclidian space. For some types of data, this may not be the case, and it remains to be shown whether or not such a technique is optimal for capturing and modelling breathing motion variability. It may be that a non-linear approach (Cayton, 2005) may result in a more efficient reduction in dimensionality, and hence facilitate more accurate motion estimates. In future work we will investigate the applicability of non-linear dimensionality reduction techniques to breathing motion.

Acknowledgements

This work was funded by EPSRC grant EP/H046410/1 and the HYPER-image project which is supported by the European Union under the 7th framework program (201651).

References

- Blackall, J., Penney, G., King, A., Hawkes, D., 2005. Alignment of sparse freehand 3-D ultrasound with preoperative images of the liver using models of respiratory motion. *IEEE Transactions on Medical Imaging* 24, 1405–1416.
- Blackall, J.M., Ahmad, S., Miquel, M.E., McClelland, J.R., Landau, D.B., Hawkes, D.J., 2006. MRI-based measurements of respiratory motion variability and assessment of imaging strategies for radiotherapy planning. *Physics in Medicine and Biology* 51, 4147–4169.
- Buerger, C., Schaeffter, T., King, A.P., 2010. Hierarchical adaptive local affine registration for respiratory motion estimation from 3-D MRI, in: *Proceedings ISBI*, pp. 1237–1240.
- Buerger, C., Schaeffter, T., King, A.P., 2011. Hierarchical adaptive local affine registration for fast and robust respiratory motion estimation. *Medical Image Analysis* 15, 551–564.
- Cayton, L., 2005. *Algorithms for Manifold Learning*. Technical Report CS2008-0923. Univ. of California.
- Crum, W.R., Camara, O., Hawkes, D.J., 2007. Methods for inverting dense displacement fields: Evaluation in brain image registration, in: *Proceedings MICCAI*, pp. 900–907.

- Ehrhardt, J., Werner, R., Schmidt-Richberg, A., Handels, H., 2009. Prediction of respiratory motion using a statistical 4D mean motion model, in: Proceedings MICCAI Workshop on Pulmonary Image Processing, pp. 3–14.
- Ehrhardt, J., Werner, R., Schmidt-Richberg, A., Handels, H., 2011. Statistical modeling of 4D respiratory lung motion using diffeomorphic image registration. *IEEE Transactions on Medical Imaging* 30, 251–265.
- King, A.P., Boubertakh, R., Rhode, K.S., Ma, Y.L., Chinchapatnam, P., Gao, G., Tangcharoen, T., Ginks, M., Cooklin, M., Gill, J.S., Hawkes, D.J., Razavi, R.S., Schaeffter, T., 2009a. A subject-specific technique for respiratory motion correction in image-guided cardiac catheterisation procedures. *Medical Image Analysis* 13, 419–431.
- King, A.P., Jansen, C., Rhode, K.S., Caulfield, D., Razavi, R.S., Penney, G.P., 2010. Respiratory motion correction for image-guided cardiac interventions using 3-D echocardiography. *Medical Image Analysis* 14, 21–29.
- King, A.P., Rhode, K.S., Razavi, R.S., Schaeffter, T., 2009b. An adaptive and predictive respiratory motion model for image-guided interventions: Theory and first clinical application. *IEEE Transactions on Medical Imaging* 28, 2020–2032.
- Klinder, T., Lorenz, C., Ostermann, J., 2008. Respiratory motion modeling and estimation, in: Proceedings MICCAI Workshop on Pulmonary Image Analysis, pp. 53–62.
- Klinder, T., Lorenz, C., Ostermann, J., 2010. Prediction framework for statistical respiratory motion modelling, in: Proceedings MICCAI, pp. 327–334.
- Low, D.A., Parikh, P.J., Lu, W., Dempsey, J.F., Wahab, S.H., Hubenschmidt, J.P., Nystrom, M.M., Handoko, M., Bradley, J.D., 2005. Novel breathing motion model for radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.* 63, 921–929.
- Manke, D., Nehrke, K., Bornert, P., 2003. Novel prospective respiratory motion correction approach for free-breathing coronary MR angiography using a patient-adapted affine motion model. *Magnetic Resonance in Medicine* 50, 122–131.

- Manke, D., Rosch, P., Nehrke, K., Bornert, P., Dossel, O., 2002. Model evaluation and calibration for prospective respiratory motion correction in coronary MR angiography based on 3-D image registration. *IEEE Transactions on Medical Imaging* 21, 1132–1141.
- McClelland, J.R., Blackall, J.M., Tarte, S., Chandler, A.C., Hughes, S., Ahmad, S., Landau, D.B., Hawkes, D.J., 2006. A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy. *Medical Physics* 33, 3348–3358.
- McClelland, J.R., Hughes, S., Modat, M., Qureshi, A., Ahmad, S., Landau, D.B., Ourselin, S., Hawkes, D.J., 2011. Inter-fraction variations in respiratory motion models. *Physics in Medicine and Biology* 56, 251–272.
- Mutaf, Y.D., Scicutella, C.J., Michalski, D., Fallon, K., Brandner, E.D., Bednarz, G., Huq, M.S., 2010. A simulation study of irregular respiratory motion and its dosimetric impact on lung tumors. *Physics in Medicine and Biology* 56, 845–859.
- Nehrke, K., Bornert, P., Manke, D., Bock, J.C., 2001. Free-breathing cardiac MR imaging: Study of implications of respiratory motion - initial results. *Radiology* 220, 810–815.
- Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J., 1999. Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging* 18, 712–721.
- Savill, F., Schaeffter, T., King, A.P., 2011. Assessment of input signal positioning for cardiac respiratory motion models during different breathing patterns, in: *Proceedings ISBI*, pp. 1698–1701.
- Seppenwoolde, Y., Shirato, H., Kitamura, K., Shimizu, S., van Herk, M., Lebesque, J.V., Miyasaka, K., 2002. 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, 822–834.
- Shechter, G., Shechter, B., Resar, J.R., Beyar, R., 2005. Prospective motion correction of X-ray images for coronary interventions. *IEEE Transactions on Medical Imaging* 24, 441–450.

- von Siebenthal, M., Szekely, G., Gamper, U., Boesiger, P., Lomax, A., Cattin, P., 2007. 4D MR imaging of respiratory organ motion and its variability. *Physics in Medicine and Biology* 52, 1547–1564.
- Sonke, J.J., Lebesque, J., van Herk, M., 2008. Variability of four-dimensional computed tomography patient models. *Int. J. Radiat. Oncol. Biol. Phys.* 70, 590–598.
- Thielemans, K., Mustafovic, S., Tsoumpas, C., 2007. STIR: Software for tomographic image reconstruction release 2, in: *IEEE Nuclear Science Symposium Conference Record*, pp. 2174–2176.
- Tsoumpas, C., Buerger, C., King, A.P., Keereman, V., Vandenberghe, S., Schulz, V., Schaeffter, T., Marsden, P.K., 2009. Simulation of dynamic PET data from real MR acquisitions, in: *IEEE Nuclear Science Symposium and Medical Imaging Conference*, pp. 3065–3068.
- Tsoumpas, C., MeckKewn, J., Halsted, P., King, A., Buerger, C., Totman, J., Schaeffter, T., Marsden, P., 2010. Simultaneous PET-MR acquisition and MR-derived motion fields for correction of non-rigid motion in PET. *Annals of Nuclear Medicine* 24, 745–750.
- Yang, D., Lu, W., Low, D.A., Deasy, J.O., Hope, A.J., El Naqa, I., 2008. 4D-CT motion estimation using deformable image registration and 5D respiratory motion modeling. *Medical Physics* 35, 4577–4590.
- Zhang, Q., Pevsner, A., Hertanto, A., Hu, Y.C., 2007. A patient-specific respiratory model of anatomical motion for radiation treatment planning. *Medical Physics* 34, 4772–4781.

Figure 5: Experiment 1: comparison of different model formation/application approaches. Results shown are computed over 10 volunteers, all breathing patterns combined (typically 105 registrations per volunteer): (a) normalised cross-correlation (NCC) accuracy measure, computed between the 3-D reference dynamic image warped using the registration result and the 3-D target dynamic image; (b) root-mean-square of the shortest surface distances between the automatically segmented lung boundaries from the same images. (c) DICE coefficients between the automatically segmented lungs of the same images.

Figure 6: Comparison of sample registrations using techniques (3) and (4) for volunteer H. Left image: target image for registration, overlaid with isocontours indicating the air-tissue boundaries. Centre images: the reference image warped using the registrations computed using technique (3) (top) and technique (4) (bottom). Both images are overlaid with the same isocontours extracted from the target image. Right images: zoomed versions of the centre images, highlighting an area where there is a significant misalignment for technique (3) (indicated with the arrow) but not for technique (4).

Figure 7: Experiment 2: validation of automatic image navigator positioning for volunteer A. The left-hand graph shows the slice selection measure together with the overall registration accuracy achieved by using each coronal slice as a 2-D image navigator. The right-hand graph is a zoomed version of that on the left. In our automatic slice selection approach we choose as a 2-D image navigator the slice with the maximum value for the slice selection measure (shown with a red dashed line). The blue dotted line shows the actual optimal slice in terms of registration accuracy.

Figure 8: Experiment 3: validation of the use of the final 2-D similarity measure (NCC) as a model applicability test. Left column: NCC within navigator slice after registering using the PCA-based approach. The PCA model was formed from a normal breathing dataset and applied to 2-D image navigator data acquired during normal, fast and deep breathing. Right column: Demonstration of model applicability test. The graphs show the final 2-D within-slice NCC for the normal, fast and deep breathing 2-D image navigator data. The red squares show where the test has flagged the model as being not applicable to the data observed. The two thresholds described in Section 2.6 are shown with dotted and dashed lines.

Figure 9: Experiment 4: Demonstration of MRI-based motion correction in PET imaging: (a) Emission image used for PET simulation. (b) Sagittal slice through a simulated PET image with no motion. This was simulated from image (a) with no motion transformations applied. This image is our 'gold standard', and shows what the motion corrected image would ideally look like. Note that in practice such a motion-free image would not be possible, and we show it only for comparison purposes. (c) Sagittal slice through the overall simulated PET image, with motion but without motion correction. This image was simulated from the emission image (a) transformed using the 35 different motion transformations. (d) The same slice through the overall simulated PET image motion corrected using technique (4). In images (b)-(d) the yellow contours show outlines of the lungs and myocardium and were extracted from the emission image (a) used for PET simulation, i.e. they show where these tissue boundaries should be in the absence of motion.