King's Research Portal

DOI:

[10.1016/j.media.2014.05.013](https://doi.org/10.1016/j.media.2014.05.013)

Document Version

Early version, also known as pre-print

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

Citation for published version (APA):

Peressutti, D., Penney, G. P., Kolbitsch, C., & King, A. P. (2014). Personalising population-based respiratory motion models of the heart using neighbourhood approximation based on learnt anatomical features. *Medical Image Analysis*, 18(7), 1015-1025. <https://doi.org/10.1016/j.media.2014.05.013>

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.

Personalising population-based respiratory motion models of the heart using neighbourhood approximation based on learnt anatomical features

Devis Peressutti^{a,*}, Graeme P. Penney^a, Christoph Kolbitsch^a, Andrew P. King^a

^aDivision of Imaging Sciences and Biomedical Engineering, King's College London, King's Health Partners, St. Thomas' Hospital, London, SE1 7EH, United Kingdom.

Abstract

Respiratory motion models have been proposed for the estimation and compensation of respiratory motion during image acquisition and image-guided interventions on organs in the chest and abdomen. However, such techniques are not commonly used in the clinic. Subject-specific motion models require a dynamic calibration scan that interrupts the clinical workflow and is often impractical to acquire, while population-based motion models are not as accurate as subject-specific motion models. To address this lack of accuracy, we propose a novel personalisation framework for population-based respiratory motion models and demonstrate its application to respiratory motion of the heart. The proposed method selects a subset of the population sample which is more likely to represent the cardiac respiratory motion of an unseen subject, thus providing a more accurate motion model. The selection is based only on anatomical features of the heart extracted from a static image. The features used are learnt using a neighbourhood approximation technique from a set of training datasets for which respiratory motion estimates are available. Results on a population sample of 28 adult healthy volunteers show average improvements in estimation accuracy of 20% compared to a standard population-based motion model, with an average value for the 50th and 95th quantiles of the estimation error of 1.6mm and 4.7mm respectively. Furthermore, the anatomical features of the heart most strongly correlated to respiratory motion are investigated for the first time, showing

*Corresponding author. devis.peressutti@kcl.ac.uk.

the features on the apex in proximity to the diaphragm and the rib cage, on the left ventricle and interventricular septum to be good predictors of the similarity in cardiac respiratory motion.

Keywords: Respiratory motion, population-based modelling, neighbourhood approximation, anatomical features, MRI

1. Introduction

Despite recent advances in medical imaging, respiratory motion remains a major limiting factor for image acquisition and image-guided interventions on organs in the chest and abdomen. In image acquisition, respiratory motion causes artefacts such as blurring, affecting the validity of subsequent diagnosis and treatment planning decisions. Image degradation for different organs (Keall et al., 2006; Scott et al., 2009) and imaging modalities (Nehrke et al., 2001; Blackall et al., 2006; Nehmeh and Erdi, 2008) has been extensively reported in the literature. In image-guided interventions, respiratory motion causes misalignments between the static images used for guidance and the moving anatomy, making the guidance information misleading and potentially dangerous (Hawkes et al., 2005; Linte et al., 2010). For instance, respiratory motion during radiotherapy treatment can cause the tumour to move, resulting in healthy tissue being irradiated (Keall et al., 2006).

Over the last decade, respiratory motion models have been proposed as a viable solution for respiratory motion estimation and compensation (McClelland et al., 2013). Respiratory motion models aim to model the relationship between the respiratory motion of the internal organs (which typically cannot be imaged in real-time) and some measurable surrogate data (e.g. diaphragm or skin surface displacement). This relationship is learnt in a calibration phase, during which the motion of the organs is estimated from a dynamic calibration scan acquired simultaneously with the surrogate data. When applying the model, the learnt relationship is used to estimate the motion based only on the acquisition of the surrogate data.

Respiratory motion models can be *subject-specific* or *population-based*. When building subject-specific models, the respiratory motion is estimated from a dynamic calibration scan of the same subject that the model will be applied to. The calibration scan is typically acquired using Computed Tomography (CT) or Magnetic Resonance Imaging (MRI). However, it is often impractical, or even impossible, to acquire this scan due to interruptions

to the clinical workflow, subject considerations such as bariatric patients or patients with MRI-incompatible implants, dose issues or high cost. Furthermore, the image quality of the calibration scan might not be sufficient for reliable motion estimation.

To overcome the need for a subject-specific calibration scan, population-based models have been proposed (Fayad et al., 2010; He et al., 2010; Ehrhardt et al., 2011; Klinder and Lorenz, 2012; Preiswerk et al., 2012; Samei et al., 2012). For such models, the motion in the calibration phase is estimated from calibration scans acquired previously from a sample of the population of subjects, and the model is subsequently applied to a subject not belonging to the sample. Respiratory motion can vary significantly between subjects (Keall et al., 2006), so current population-based models are generally less accurate than subject-specific models, as they average out inter-subject variations in motion (McClelland et al., 2013).

Typically, population-based models are personalised to individuals by registering a template image representing the average anatomy of the population sample to a corresponding image of the new subject, and transforming the average motion model from the population coordinate system to that of the new subject. Most population-based models proposed to date use all subjects in the population sample to build the average motion model which is then applied to any out-of-sample subject. In this paper, we propose a technique for selecting a subset of the population sample to build the model, with the aim of improving its accuracy. A related work was presented in Samei et al. (2012). The authors proposed a technique for making more accurate population-based motion estimates by weighting the contributions of the subjects in the population sample according to similarities of the surrogate. The technique was devised to estimate the respiratory motion and drift of the liver during radiotherapy treatment, but proved effective to estimate the respiratory drift only. As opposed to the personalisation proposed by Samei et al. (2012), our proposed method is applied in the model formation phase instead of the model application phase, and considers static anatomical features only for the personalisation of the motion model, without any information on the respiratory motion or respiratory surrogates of the out-of-sample subject. Hence, to date no work has investigated the use of anatomical features learnt from static images to personalise population-based models in order to achieve more accurate motion estimates.

In our framework, a subset of the population sample that is more likely to represent the cardiac respiratory motion of the out-of-sample subject is

selected to personalise the model. The selection is based on anatomical features of the heart derived from a static image which are learnt using a neighbourhood approximation technique. Our underlying hypothesis is that the respiratory motion of the heart and its morphology and position are in some way correlated. To the authors’ knowledge, this is the first work investigating such a hypothesis. Work has been performed to investigate such a correlation for cardiac cycle motion (Metz et al., 2012), but not, so far, for respiratory motion. The intended application for the proposed framework is cardiac respiratory motion estimation in image-guided interventions (Shechter et al., 2004; King et al., 2009).

Preliminary work was presented in Peressutti et al. (2013b). In this initial work only affine respiratory deformations could be modelled. Furthermore, the correlation analysis was limited to differences in bulk orientation and shape of the heart. In this paper, we extend the framework to model non-rigid respiratory deformations, and make use of non-rigid descriptors of the orientation and morphology of the heart, allowing generalisation of the framework to any other organ affected by respiratory motion. Furthermore, improvements to the methodology are presented, as well as the use of a larger sample of subjects. An extensive investigation of the correlation hypothesis is presented for the first time in Section 3.2.

2. Methods and materials

2.1. Overview

An overview of the proposed framework is shown in Figure 1. The relationship between the respiratory motion of the heart and its anatomical features is learnt in a prior population model calibration phase. Provided with a static high resolution image of an out-of-sample subject as input, the respiratory motion model is personalised based on its anatomical features.

Each of the N datasets in the population sample consists of a high resolution image of the heart, a dynamic free-breathing calibration scan capturing the cardiac respiratory motion and some surrogate data (see Section 2.2). Similar to Ehrhardt et al. (2011) and Klinder and Lorenz (2012), an average shape atlas of the anatomy is built using the N high resolution images (Section 2.2). Respiratory motion estimates are derived from the N dynamic calibration scans and transformed to the atlas’s coordinate system to produce a motion atlas (Section 2.3.2). Unlike the population-based models proposed so far (McClelland et al., 2013; Ehrhardt et al., 2011; Klinder and Lorenz,

Figure 1: Overview of the proposed framework for formation of the population-based model and personalisation to out-of-sample subjects.

2012; Preiswerk et al., 2012) in which an average motion model of the N datasets is used for any out-of-sample subject, the proposed personalisation selects only $K \ll N$ subjects from the population sample which are more likely to represent the respiratory motion of the out-of-sample subject. These are selected using a neighbourhood approximation technique that will be described in Section 2.3.4. The inter-subject variability in motion is therefore exploited to obtain motion estimates that are more accurate than standard population-based model estimates.

We now describe each of these stages in more detail.

2.2. Materials

We imaged 28 healthy adult volunteers. All datasets were acquired using a 1.5T Philips Achieva MRI scanner (Philips Healthcare, Best, The Netherlands). Each dataset consisted of a static high resolution 3-D MRI volume and a free-breathing dynamic 3-D MRI calibration scan:

- **High resolution 3-D:** 3-D balanced TFE, cardiac triggered at late diastole, respiratory gated at end-exhale, 5mm navigator window, typically 120 sagittal slices, $TR = 4.4ms$, $TE = 2.2ms$, flip angle = 90° , acquired voxel size $2.19 \times 2.19 \times 2.74mm^3$, reconstructed voxel size $1.37 \times 1.37 \times 1.37mm^3$, the acquisition window was optimised for each volunteer and was on average 100ms, scan time approximately 5 minutes.
- **Dynamic 3-D calibration scan:** 3-D TFEPI, ECG-triggered at late diastole, typically 20 slices, $TR = 10ms$, $TE = 4.9ms$, flip angle = 20° , acquired voxel size $2.7 \times 3.6 \times 8.0mm^3$, reconstructed voxel size $2.22 \times 2.22 \times 4.0mm^3$, TFE factor 26, EPI factor 13, TFE acquisition time 267.9ms.

The high resolution 3-D image is part of many routine clinical protocols and provides high spatial resolution information about the anatomy and pathology of the heart. The dynamic calibration scan was ECG-triggered, therefore the images represented the motion of the heart due to respiration only. The scan acquired 40 3-D images while the volunteer was breathing normally. Given the short acquisition time of this scan, the field of view of the dynamic images was limited, covering most, but not all, of the four cardiac chambers. The superior-inferior displacement of the left hemi-diaphragm was measured in each image and this was used as the respiratory surrogate s for model formation.

Hereafter, we denote by $I_n, n = 1, \dots, N$ the 3-D high resolution images, while D_{np} denotes the 3-D dynamic calibration image p of subject n , $p = 1, \dots, P$, $n = 1, \dots, N$, where $N = 28$ and $P = 40$. For any subject n , $D_{n_{ref}}$ denotes the reference end-exhale dynamic image with the highest corresponding value of the surrogate s . The high resolution and dynamic scans were performed consecutively and the DICOM header information was used to transform all images into the same coordinate system. A subsequent rigid intensity based registration was performed between the I_n and the corresponding reference dynamic image $D_{n_{ref}}$ to correct for any possible patient motion between scans.

2.3. Methods

The constituent parts of the proposed framework can be summarised as:

1. building of an average shape atlas in its natural coordinate system;
2. transformation of respiratory motion estimates from each subject to the atlas's coordinate system;
3. analysis of respiratory motion similarities between subjects in the population sample;
4. neighbourhood approximation using respiratory motion similarity as distance metric and anatomical features as predictors;
5. personalisation to out-of-sample subject based on a static image.

We describe each of these parts in the following sections.

2.3.1. Average shape atlas

A typical initial step in the formation of any population-based motion model is the creation of an average shape atlas of the organ investigated, providing a common coordinate system in which the motions of the different subjects can be represented. An average shape atlas of the heart in its *natural coordinate system* (Frangi et al., 2002) was built by iteratively registering and averaging together the high resolution 3-D images, as we describe below.

To remove positional differences from the subsequent non-rigid registrations, the high resolution images (see Section 2.2) were firstly aligned according to the centre of mass of the heart. The centres of mass were computed over a manually drawn binary mask covering the whole heart and the ascending aorta. See Figure 2 for an example of a binary mask. The same masks were later used for motion estimation and evaluation purposes.

After alignment, one image was randomly selected as a reference I_{ref} and the remaining images were registered to it using a hierarchical non-rigid registration algorithm (Buerger et al., 2011). The $N - 1$ images (i.e. all apart from I_{ref}) were warped using the registration results. The intensities of the warped images and I_{ref} were averaged to obtain an average intensity image I_{avg_0} which was used as reference I_{ref} in the next iteration. All I_n were registered and warped to the new reference I_{ref} and the intensities of the warped images averaged together to obtain I_{avg_1} . This non-rigid registration and averaging process was repeated until the Normalised Cross Correlation (NCC) similarity measure between I_{avg_t} and $I_{avg_{t-1}}$ was higher than a predefined threshold of 0.99.

Since the pose and size of I_{avg_t} was still biased towards the initial choice of I_{ref} , we employed the method proposed in Frangi et al. (2002) and Rueckert et al. (2003) to remove any remaining bias, as follows. The final average image I_{avg_t} was non-rigidly registered to each $I_n, n = 1, \dots, N$ and the resulting N deformation fields were averaged together and used to warp I_{avg_t} . The result of this warping is the final atlas image I_{atlas} in its *natural coordinate system*, which requires the least residual deformation to explain the anatomical variability across all individuals. An example of I_{atlas} formed from the 28 healthy adult volunteers is shown in Figure 2, overlaid with the binary mask used for registration. In the following, T_n denotes the diffeomorphic non-rigid transformation that maps each high resolution image I_n to I_{atlas} .

2.3.2. Motion estimates in atlas space

To compare the respiratory motion of the different subjects in the population sample, the motion was first estimated from each of the dynamic calibration scans D_{n_p} (see Section 2.2). These motion estimates were transformed to the atlas’s natural coordinate system, where they could be directly compared. We describe this process below.

To estimate the respiratory motion of the heart, the dynamic image with the highest corresponding value of the surrogate s (see Section 2.2) was selected as the reference end-exhale image $D_{n_{ref}}$ and all remaining images were registered to it. As reported in the literature (Shechter et al., 2004; King et al., 2009), cardiac respiratory motion can be well described by affine deformations, therefore we employed a 12 parameter affine registration algorithm to align $D_{n_{ref}}$ to all other D_{n_p} . This resulted in a set of $P - 1$ affine transformations R_{n_p} that map $D_{n_{ref}}$ to D_{n_p} for each subject n . An identity transformation was also included for the reference dynamic $D_{n_{ref}}$, making P

Figure 2: Axial (a), coronal (b) and sagittal (c) views of the average shape atlas of the heart in its natural coordinate system. This atlas was derived from 28 healthy volunteers. Despite the blurring effect of the averaging process, all major cardiac features are well represented and discernible. The outline of the binary mask used to derive the anatomical features (see Section 2.3.4) is overlaid in yellow. This mask was computed as the average of all subjects’ binary masks warped to the atlas’s natural coordinate system.

affine transformations in total for each subject. To localise the registration to the heart only, the dynamic images were masked using the same binary masks employed to build the atlas.

The transformations¹ R_{n_p} were transformed to the atlas’s coordinate system using the method proposed in Rao et al. (2002). Denoting by $\mathbf{r}_{n_p}$ the respiratory deformation field corresponding to R_{n_p} , the transformation was computed as follows

$$\mathbf{r}_{n_p}^{atlas} = J_n \mathbf{r}_{n_p}, \quad (1)$$

where J_n is the Jacobian matrix of T_n (see Section 2.3.1) which is analytically determined and $\mathbf{r}_{n_p}^{atlas}$ the deformation field transformed to the atlas’s coordinate system. In our original version of the framework (Peressutti et al.,

¹For clarity’s sake, in the following, we denote by a capital letter the transformation mapping one image to another, while a small bold letter denotes the deformation field corresponding to the transformation.

2013b), the composition of the transformations

$$R_{n_p}^{atlas} = T_n \circ R_{n_p} \circ T_n^{-1}, \quad (2)$$

was used instead. This method assumes T_n to be diffeomorphic and requires the estimation of both T_n and T_n^{-1} . Because the applied registration algorithm is not symmetric and does not compute T_n and T_n^{-1} simultaneously, the previous transformation method has been replaced in order to limit the effects of inaccuracies in the independent computation of T_n and T_n^{-1} .

2.3.3. Quantifying motion similarity

The previous two steps are typical of many population-based motion models and allow the creation of a common coordinate system where the subject-specific respiratory motion estimates can be represented. The end result is a set of $N \times P$ respiratory deformation fields $\mathbf{r}_{n_p}^{atlas}$ in the atlas's natural coordinate system.

The aim of the next step was to quantify the similarities between the respiratory motions of the different subjects. The motivation for this was to learn the relationship between the anatomical features and respiratory motion. Quantifying motion similarities is a non-trivial task due to the lack of sampling correspondence between subjects, i.e. the surrogate values at which motions are estimated are different.

To overcome the lack of sampling correspondence between subjects, we derived a similarity measure from the pairwise comparison of subject-specific respiratory motion models. We denote by $r_{n_{px}}^{atlas}$, $r_{n_{py}}^{atlas}$ and $r_{n_{pz}}^{atlas}$ the x , y and z components of the respiratory deformation field $\mathbf{r}_{n_p}^{atlas}$ corresponding to the respiratory state p of subject n . In order to compare the respiratory positions of the different subjects, the respiratory surrogate s of each subject was normalised so that the minimum and maximum values were $[-1, 0]$. The surrogate s was the supero-inferior diaphragm displacement as measured from the dynamic images (see Section 2.2). The subjects were given no particular breathing instructions during the dynamic scan acquisition, so we assume that the acquired sequence is an unbiased sample of their normal breathing pattern. The amplitude range of the respiratory surrogate s computed over all 28 subjects was $23.8 \pm 8.5mm$.

Similar to King et al. (2011), a subject-specific polynomial motion model $\mathcal{M}_n(s)$ was formed for each subject n by fitting a polynomial function to the P values of $r_{n_{px}}^{atlas}$, $r_{n_{py}}^{atlas}$ and $r_{n_{pz}}^{atlas}$ as a function of the surrogate s for every

control point of a grid covering the whole imaging volume. For details of the order of the polynomial function and the spacing of the grid see Section 3. Provided with a value of the surrogate s , the model estimates the respiratory motion field $\mathbf{r}_{n,s}^{atlas}$ at each control point of the grid by using the surrogate value as input to the polynomial function. Although the original motion estimates were affine, due to numerical approximations and registration inaccuracies, the deformation fields $\mathbf{r}_{n_p}^{atlas}$ are not strictly affine after transforming them using Eq. (1). Compared to our preliminary version (Peressutti et al., 2013b) which used affine motion, the modelling of the non-rigid deformation field allows extension of the framework to any non-linear motion/deformation.

To quantify motion similarities across subjects, the difference in respiratory deformation estimation between each pair of subject-specific motion models was computed as a motion similarity metric. Given 40 values of the surrogate s linearly distributed in $[-1, 0]$, Target Registration Errors (TREs) between each pair of 40 motion model estimates were computed using all $H \approx 30,000$ voxels in the binary mask covering the heart of the atlas as target points (see Figure 2). For any given pair of subjects, i and j , the TRE was a vector of length $40 \times H$, with the element for voxel h and surrogate s defined as

$$TRE_{i,j,s,h} = \|\mathcal{M}_i(s)_h - \mathcal{M}_j(s)_h\|_2 = \|\mathbf{r}_{i,s,h}^{atlas} - \mathbf{r}_{j,s,h}^{atlas}\|_2, \quad (3)$$

where $\mathcal{M}_i(s)_h$ and $\mathbf{r}_{i,s,h}^{atlas}$ denote the model deformation estimation for subject i for the given surrogate values s at the h^{th} voxel location. The 95th quantiles of these pairwise TREs computed over the 40 values of s and all voxels $h = 1, \dots, H$ were employed as the motion similarity metric, $\rho(i, j) = 95^{th}quantile\{TRE_{i,j,s,h}\}$, $\forall h, s, i, j = 1, \dots, N$ and a $N \times N$ distance matrix $\mathbb{P} = \{\rho(i, j)\}$ was built. The statistical distribution of TREs was highly asymmetric and non-Gaussian, thus quantiles represented a robust statistic of the distribution. In particular, 95th quantiles provided a measure of the worst-case error in respiratory motion estimation.

2.3.4. Neighbourhood approximation

The main novelty of this work is the investigation of the hypothesis that the respiratory motion of the heart is in some way correlated with the cardiac morphology and position. Considering the mechanics of respiration, the contraction/relaxation of the diaphragm and the movement of the rib cage are the driving forces of motion (Sharp et al., 1975). Therefore, it is reasonable

to hypothesise that the morphology of the heart and its pose with respect to the rib cage and diaphragm will affect its respiratory motion. Preliminary results in our original description of the framework (Peressutti et al., 2013b) supported this hypothesis.

In our previous framework, a two step procedure (i.e. clustering and classification) was used, and each required careful optimisation of the parameters. In this paper, we address the limitations of our previous method by employing the Neighbourhood Approximation Forests (NAF) technique recently proposed by Konukoglu et al. (2013). The aim of our application of this technique is to determine the neighbourhood of an out-of-sample subject by learning the underlying manifold implied by the respiratory motion similarities.

The training database consisted of the anatomical deformation fields $\mathbf{t}_{atlas \rightarrow n}$, corresponding to T_n^{-1} , that map the atlas I_{atlas} in its natural coordinate system to each subject I_n , and the pairwise distance matrix $\mathbb{P}$. An example of $\mathbf{t}_{atlas \rightarrow n}$ is shown in Figure 3. In the learning phase, NAF learns the subset of anatomical deformations that are most strongly correlated with the neighbourhood structure implied by the respiratory motion similarities $\mathbb{P}$. Compared to our previous method, the anatomical features used as predictors are the non-rigid deformations $\mathbf{t}_{atlas \rightarrow n}$ instead of the affine anatomical parameters, and the two step clustering/classification process is embedded in a single and simpler to optimise approach.

Figure 3: Coronal views of (a) the atlas I_{atlas} in its natural coordinate system and (b) one subject I_n of the population. The deformation field $\mathbf{t}_{atlas \rightarrow n}$ mapping the anatomical transformation T_n^{-1} is overlaid onto the atlas and shown in (c). The scale of the colormap is mm .

An overview of the training phase of the NAF is shown in Figure 4.

Since we are interested in the morphology of the heart only, the subset of H

Figure 4: Toy example of binary tree training of the NAF. The $N = 8$ anatomical deformation fields $\mathbf{t}_{atlas \rightarrow n}^H$ are the input vector features. At each node, the training dataset is partitioned into smaller balanced sets. The partitioning optimises over the anatomical vector features able to describe the neighbourhood structure implied by the respiratory motion similarity distances $\rho(i, j)$, here represented as similarities in colour of the frame.

anatomical deformations $\mathbf{t}_{atlas \rightarrow n}^H$ defined over the binary mask covering the heart and major vessels of the atlas were considered as features, instead of the deformations $\mathbf{t}_{atlas \rightarrow n}$ defined over the entire imaging volume. Note that parts of the rib cage and the diaphragm are included in the binary mask (see Figure 2). In order to remove translational offsets from the feature analysis, the mean value of the deformation vectors $\mathbf{t}_{atlas \rightarrow n}^H$ was subtracted. Since our predictors are 3-D vectors, the original NAF technique (Konukoglu et al., 2013) was extended to deal with vector predictors rather than scalars. In particular, the optimisation at each node of the tree is performed for a vector valued threshold rather than a scalar threshold. The set of features selected at each node for each tree in the forest was used to infer correlation between anatomical features and cardiac respiratory motion (see Figure 7). The parameters to be optimised were the number of trees in the forest, the minimum number of samples per node Δ and the number of random features h selected at each node q . These parameters were determined by means of cross-validation and the values are reported in Section 3.

2.3.5. Personalisation

The personalisation of the proposed population-based model is carried out before model application. The input is a static high resolution image I_{oos} of the heart for an out-of-sample subject, while the output is a respiratory motion model personalised for the out-of-sample subject.

The anatomical deformation field $\mathbf{t}_{atlas \rightarrow oos}$ was determined by registering I_{atlas} to I_{oos} . After the subtraction of the mean value, the anatomical deformations $\mathbf{t}_{atlas \rightarrow oos}^H$ over the binary mask were given as input to the trained NAF. The out-of-sample anatomical deformations were propagated through each node in every tree in the forest by testing the features learnt in the previous step (see Section 2.3.4). The individual tree predictions were combined by computing an affinity measure $A(I_{oos}, I_n) = \sum \mathbf{1}(I_n), \forall I_n$ for every tree in the forest, where $\mathbf{1}(\cdot)$ is the indicator function, i.e. it is 1 if I_n is in the same terminal node as I_{oos} , and 0 if not. The affinity measure therefore counted the occurrences of I_n in the same terminal node as I_{oos} for all trees in the forest. The K subjects used to build the motion model for the out-of-sample subject were the K subjects with the highest affinity measure, representing the K closest neighbours on the manifold implied by the respiratory motion distances $\rho(\cdot, \cdot)$.

The respiratory motion estimates of the K neighbours were transformed to the coordinate system of the out-of-sample subject using equation (1), but employing the inverse Jacobian matrix corresponding to T_n^{-1} . A polynomial motion model, similar to those used to determine respiratory motion similarities (see Section 2.3.3), was then built using the transformed respiratory deformations. Given the respiratory surrogate values for the out-of-sample subject, respiratory motion model estimates were produced.

3. Experiments

The aims of the evaluation were two-fold. First, the estimation accuracy of the proposed personalised model was investigated by comparing it to other estimation techniques, namely no respiratory motion estimation, a standard average population-based motion model (Fayad et al., 2010; He et al., 2010; Ehrhardt et al., 2011; Klinder and Lorenz, 2012; Preiswerk et al., 2012), our previously proposed version (Peressutti et al., 2013b) and a subject-specific motion model (King et al., 2011). The subject-specific motion model accuracy represented the best target accuracy achievable. Details of this experiment are presented in Section 3.1.

Secondly, the relationship between the anatomical features and the respiratory motion of the heart was analysed in the set of experiments detailed in Section 3.2.

For experiments (II), (III) and (IV) in Section 3.1 and both experiments in Section 3.2, a leave-one-out cross-validation was employed. This means that each of the 28 volunteer datasets was left out in turn and the remaining 27 datasets were used as the population sample (i.e. $N = 27$). The left-out dataset was used as the out-of-sample subject I_{oos} .

For a thorough accuracy evaluation, we non-rigidly registered (Buerger et al., 2011) the dynamic images D_{n_p} of each dataset n to the dynamic end-exhale reference image $D_{n_{ref}}$. This process resulted in $P - 1$ gold-standard non-rigid motion fields which were employed to evaluate the accuracy of all techniques compared in each experiment. By warping $D_{n_{ref}}$ using the gold-standard non-rigid motion fields, we obtained a set of artificial images with known, realistic motion fields which were employed instead of the original images D_{n_p} for motion estimation (Section 2.3.2), for motion classification (Section 2.3.3) and for the evaluation described in this Section.

Target Registration Errors between the non-rigid gold-standard deformations and the respiratory deformations estimated by the compared techniques were computed over the H voxels in the binary mask covering the major chambers and vessels of the heart for all respiratory positions $p = 1, \dots, P - 1$. Median and 95th quantile of the $H \times (P - 1)$ values of the TREs are the statistics considered as accuracy measures. Results for each experiment are reported in Section 4. Mean and standard deviation of the N values for the median and quantile are reported for compactness of representation.

3.1. Evaluation of estimation accuracy

The aim of this experiment was to evaluate the estimation accuracy of the proposed personalisation framework compared to state-of-the-art techniques. To this end, we compared:

- (I) *NOEST*: no respiratory motion estimation;
- (II) *AVPOP*: standard population-based motion model (Fayad et al., 2010; He et al., 2010; Ehrhardt et al., 2011; Klinder and Lorenz, 2012; Preiswerk et al., 2012). All N subjects were used to build an average motion model which was applied to the out-of-sample subject I_{oos} . To build such an average population-based model we transformed all respiratory

motion deformations $\mathbf{r}_{n_p}^{atlas}$, $\forall n, \forall p$ to the coordinate system of I_{oos} using equation (1). Given the normalised ranges $[-1, 0]$ of the surrogate values, we formed an average motion model of the $N \times P$ respiratory deformations at all control points of a grid covering the heart of I_{oos} as a function of the normalised surrogates;

- (III) *PERCLUST*: our previously proposed framework (Peressutti et al., 2013b), which employed clustering/classification using affine anatomical features as predictors. A spectral clustering technique (von Luxburg, 2007) and a random forest classifier were employed. The inputs to the spectral clustering were the respiratory motion similarities $\mathbb{P}$, the number of clusters C and the definition of the similarity graph, while the inputs to the random forest were the clusters and the affine anatomical features. Given the affine anatomical features of the out-of-sample subject, the K subjects in the cluster predicted by the classifier were employed to personalise the model. See Peressutti et al. (2013b) for more details;
- (IV) *PERNAF*: our proposed personalised motion model. As in (II) - *AVPOP*, but only the K closest neighbours as predicted by the NAF were used to form the model for the out-of-sample subject I_{oos} ;
- (V) *SUBSPEC*: subject-specific motion model (King et al., 2011). The respiratory motion model was formed using the respiratory deformations derived directly from the dynamic calibration scan of the out-of-sample subject. The same approach employed to build the motion models for the quantification of motion similarity (see Section 2.3.3) was used to form the subject-specific model. A leave-one-out validation was used to form the subject-specific motion model, meaning that, to estimate the respiratory position p , the corresponding respiratory deformations were not included in the model formation. The accuracy of the subject-specific model estimates represent a limit to the accuracy achievable by population-based motion models.

For all motion models compared (II, III, IV, V), a grid spacing of $5mm$ was employed and a 2^{nd} order polynomial function was fitted to the x, y, z components of the respiratory deformation fields at all control points of the grid (King et al., 2011). For methods (III) *PERCLUST* and (IV) *PERNAF*, the parameters were optimised using a subsample of 15 subjects from the

population sample. For technique (III) *PERCLUST* the optimal values were $C = 6$ clusters, 3-nearest neighbour graph for the spectral clustering and 500 trees, minimum number of samples per leaf $\Delta = 3$, entropy splitting criteria and $h = \sqrt{H}$ random features for the random forest classifier. The optimal NAF parameters for technique (IV) *PERNAF* were 500 trees, minimum number of samples $\Delta = 3$ and $h = \sqrt{H}$ randomly selected features.

3.2. Evaluation of correlation hypothesis

The aim of this experiment was to provide more insights into the relationship between the anatomical features of the heart and its respiratory motion in healthy adult volunteers. Since we hypothesised that some cardiac anatomical features can be used to predict similarities in respiratory motion of the heart, we compared our proposed personalised model to a technique that employed the anatomical similarities only, i.e. without knowledge of respiratory motion similarities, to select the K subjects used to personalise the motion model.

The techniques compared were:

- (i) *PERNAFANAT*: NAF using anatomical similarity. Our proposed framework (IV) *PERNAF* was used, but rather than using the respiratory motion distance $\rho(\cdot, \cdot)$ defined in Section 2.3.3, the mean anatomical deformation difference

$$\rho_{anat}(oos, n) = \frac{1}{H} \sum_{w=1}^H \|\mathbf{t}_{atlas \rightarrow oos}^{H_w} - \mathbf{t}_{atlas \rightarrow n}^{H_w}\|_2, \quad (4)$$

was used instead. In Eq. (4), H_w denotes the w^{th} voxel of the H voxels in the binary mask. The more similar the anatomies of I_{oos} and I_n were, the lower the value of $\rho_{anat}(oos, n)$. This technique corresponds to unsupervised clustering with respect to the entire set of anatomical features. The best anatomical features selected by the NAF would therefore correlate with the neighbourhood structure implied by the anatomical similarities rather than the respiratory motion similarities;

- (ii) *PERNAF*: our proposed model.

By comparing techniques (i) *PERNAFANAT* and (ii) *PERNAF* we investigated the difference between the subset of features most strongly correlated with the anatomical similarity and those most strongly correlated with

the similarity in respiratory motion. The same optimal NAF parameters used in the experiments detailed in Section 3.1 were used for techniques (i) and (ii).

4. Results

Results of the experiments on the accuracy are reported in Section 4.1, while Section 4.2 presents results for the experiments on the correlation hypothesis.

4.1. Estimation accuracy

Results of the leave-one-out cross-validation for the estimation accuracy experiments are reported in Table 1. Mean and standard deviation of the N values for the median and 95th quantile of the TREs corresponding to the estimation accuracy experiments are shown. 95th quantiles of the TREs for each subject n are shown in Figure 5. Results for no motion estimation (I) *NOEST* are not reported in the Figure for clarity of representation.

Different values for the number of neighbours K (i.e. $K=3, 5, 7$) used for the personalisation of our proposed model (IV) *PERNAF* were considered, with no significant statistical difference between the accuracy results over all subjects. However, the value of K did affect the individual estimation accuracies. For instance, a small K generated more accurate motion estimations for the subjects showing respiratory motion significantly different from the average motion of the population sample, while poorer accuracy was achieved for subjects with an average respiratory motion. On the other hand, $K \rightarrow N$ generated better estimation accuracies for subjects showing an average respiratory motion, but poorer accuracy for the other subjects. Therefore, to better highlight the performance of our method, hereafter we report the results using the closest $K = 3$ neighbours. The proposed personalised model estimates (IV) *PERNAF* achieved an average improvement in estimation accuracy of 20% for both median and 95th quantile compared to an average population-based motion model (II) *AVPOP*. A paired t-test showed a statistically significant difference with $p < 0.01$. The highest improvements of the 95th quantile were achieved for subjects 27 (59.2%), 18 (52.8%) and 9 (48.8%). The accuracy improvement was higher for subjects with a respiratory motion differing significantly from the average motion of the population sample. The poor accuracy results achieved for subjects 28 (-40.8%), 15 (-21.0%) and 22 (-11.7%) might be due to sub-optimal parameters

of the NAF and the number of neighbours K . In fact, these subjects show an average respiratory motion and accuracy results can be improved by increasing K . Furthermore, increasing the number of subjects in the population sample would likely lead to more robust performance improvements, given the increased respiratory motion variability represented by the population sample.

Compared to our previously proposed technique (III) *PERCLUST*, average accuracy improvements of 11% and 7.8% in median and 95th quantile were achieved, respectively ($p < 0.10$). In addition to accuracy improvements, the framework proposed in this paper requires less parameters to be optimised and allows extension to non-rigid anatomical deformations, which are necessary to describe anatomical variations for many organs affected by respiratory motion (e.g. lungs, liver).

	Mean/standard deviation (mm)					Impr. in mean (%)	
	<i>NOEST</i> (I)	<i>AVPOP</i> (II)	<i>PERCLUST</i> (III)	<i>PERNAF</i> (IV)	<i>SUBSPEC</i> (V)	wrt (II)	wrt (III)
Median	3.5/2.0	2.0/1.1	1.8/0.7	1.6/0.6	0.9/0.3	20.0	11.1
95 th quant	11.3/4.4	5.9/2.4	5.1/1.7	4.7/1.5	2.5/0.8	20.3	7.8

Table 1: Leave-one-out estimation accuracy results. Mean and standard deviation of the median and 95th quantile of TREs over the N subjects are reported. The last two columns show the accuracy improvement in the mean value achieved by the proposed technique (IV) *PERNAF* against the average population-based model (II) *AVPOP*, and our previously proposed technique (III) *PERCLUST*.

4.2. Correlation hypothesis

Results of the leave-one-out cross-validation for the experiment on the correlation hypothesis are reported in Table 2 and Figure 6.

The closest $K = 3$ neighbours were considered for the personalisation of the motion models for both techniques (i) *PERNAFANAT* and (ii) *PERNAF*. Accuracy results in respiratory motion estimation achieved by technique (i) *PERNAFANAT* highlight the fact that the similarity in the anatomy of the heart is not a good predictor of respiratory motion. In fact, estimation accuracy for technique (i) *PERNAFANAT* was slightly worse than an average population-based model (II) *AVPOP*.

Figure 7 shows the frequency maps of the anatomical features that were selected as best features by the NAF for methods (i) *PERNAFANAT* and (ii) *PERNAF*. These frequency maps show the occurrence of the best features selected during the training of the trees in the forest normalised by the

Figure 5: 95th quantiles of the TREs for each left-out subject for the estimation accuracy experiment. Results for the average population-based motion model (II) *AVPOP*, our previously proposed method (III) *PERCLUST*, our proposed personalised motion model (IV) *PERNAF* and subject-specific motion model (V) *SUBSPEC* are shown.

overall number of nodes. For clarity of representation, the frequency maps shown are scaled to the maximum value, which was different for method (i) *PERNAFANAT* and (ii) *PERNAF* (see colourmap ranges). It can clearly be seen that the most relevant features implied by the anatomical similarities are different from those implied by the similarity in respiratory motion. In particular, referring to Figure 7(b), anatomical features most strongly correlated to respiratory motion are found on the apex of the heart in proximity to the diaphragm and the rib cage, on the left ventricle and interventricular septum. This seems plausible, considering that the main drivers of respiratory motion are the diaphragm and the rib cage movement. On the contrary, the set of features most strongly correlated to anatomical similarity only (Figure 7(a)) are less localised, especially in the inner regions of the heart. This experiment emphasises the importance of the respiratory motion manifold in constraining the selection of the anatomical features relevant for respiratory motion estimation.

The anatomical features selected by the NAF could be used to derive local high-order descriptors of the cardiac shape which might be meaningful from a clinical perspective. For instance, the average shape of the apex for

the different subsets in the population sample could be extracted and used as a clinical parameter for motion classification. However, the extraction of these parameters requires further evaluation and goes beyond the scope of this paper.

	Mean/standard deviation (mm)		Impr. in mean (%)
	<i>PERNAFANAT</i> (i)	<i>PERNAF</i> (ii)	
Median	2.1/1.0	1.6/0.6	23.8
95 th quant	6.4/2.6	4.7/1.5	26.6

Table 2: Leave-one-out results for the correlation hypothesis experiment. Mean and standard deviation of the median and 95th quantile of TREs over the N subjects are reported. The last column shows the accuracy improvement in the mean value achieved by technique (ii) *PERNAF* with respect to (i) *PERNAFANAT*.

Figure 6: 95th quantiles of the TREs for each left-out subject for the correlation hypothesis experiment. Results for the NAF on anatomical similarities (i) *PERNAFANAT* and the proposed model (ii) *PERNAF* are shown.

5. Discussion and conclusions

We have proposed a novel framework for the personalisation of population-based respiratory motion models of the heart. The proposed method selects a subset of the population sample which is more likely to represent the cardiac

(a)

(b)

Figure 7: Frequency maps. The frequency of the best features selected during the training of the trees in the forest normalised by the overall number of nodes is overlaid onto the atlas. Coronal views in the antero-posterior direction show (a) features more correlated to similarities in anatomy (method (i) *PERNAFANAT*), and (b) features more correlated to similarities in respiratory motion (method (ii) *PERNAF*). Note the scaling difference of the two colour maps.

respiratory motion of an out-of-sample subject based on learnt anatomical features only. These learnt anatomical features are able to predict the neighbourhood structure implied by the similarity in respiratory motion.

Results on the cardiac respiratory motion of adult healthy volunteers derived from MRI images were reported. The proposed framework showed average improvements in respiratory motion estimation accuracy of 20% compared to state-of-the-art population-based motion models. Results reported in this paper support the hypothesis of correlation between the anatomical appearance of the heart and its motion due to respiration. To the authors' knowledge, this is the first work investigating such a hypothesis. In particular, features of the heart on the apex in proximity to the diaphragm and rib cage, on the left ventricle and on the interventricular septum were found to be good predictors of the similarity in respiratory motion of the heart.

There is a conceptual link between the framework described here and prior work on multi-atlases. For example, in Aljabar et al. (2009) and Wolz et al. (2010) more accurate atlas-based segmentations were achieved by exploiting similarities between an out-of-sample subject and a subset of a database of atlas images. This type of approach can be viewed as defining a manifold of atlas images and then determining the neighbours of a new image in this manifold. Atlas segmentations were then propagated from the neighbours. In our framework, we define an implicit manifold of motion and use NAF to find the neighbours of an out-of-sample subject in the manifold. Motion models are then formed from the motion estimates of the neighbours.

To date, there is only one example of clinical translation of a motion model-based technique: the Cyberknife Synchrony system (Schweikard et al., 2000). This technique uses a subject-specific respiratory motion model to synchronise dose delivery in radiotherapy. Because of the imaging technology and clinical protocol used in this application, it is more feasible to acquire the subject-specific calibration scan without significant interruptions to the workflow. For most other situations, however, the need for a subject-specific calibration scan and the relatively poor estimation accuracy of state-of-the-art population-based motion models represent significant impediments to the clinical translation of respiratory motion models. The proposed framework addresses these limitations, providing more accurate motion estimation without the need for a subject-specific calibration scan, thus opening up new possibilities for the use of respiratory motion models in routine clinical protocols. The intended application of our framework is respiratory motion estimation during cardiac image-guided interventions (Shechter et al., 2004;

King et al., 2009), although the personalised model could also be applied for motion estimation during cardiac MRI image acquisition (Manke et al., 2003). For some cardiac image-guided interventions, such as ablation procedures or coronary artery bypass grafting, an accuracy of $\sim 5mm$ represents a clinically adequate goal (Linte et al., 2010), which is on average achieved by our proposed framework. A static high resolution image of the organ is often already available in the patient’s record so, ideally, our personalisation framework can be applied without the need for any extra imaging, reducing scanning time and cost, and patient discomfort. As a further example, the implementation of the proposed framework in our previously proposed Bayesian respiratory motion model (Peressutti et al., 2013a) may lead to robust respiratory motion estimates for image-guided interventions fully based on echography images.

To evaluate the proposed framework, a set of gold standard motion fields and corresponding dynamic images was generated. This enabled us to perform a thorough quantitative analysis of performance whilst using realistic images and motions. All registration results were visually inspected to ensure their accuracy and to assess any possible bias towards the registration technique used. In this work, a small-deformation model was employed as opposed to a large-deformation model (Ashburner, 2007). Free-Form Deformations were used to characterise the variation in anatomy and the method proposed by Rao et al. (2002) was employed to transport the respiratory transformations to the atlas’s coordinate system. More theoretically sound but complex methods could be used instead to derive and compute statistics on the transformations (Pennec, 2009).

To quantify the respiratory motion similarities, the range of the respiratory surrogate was normalised and a motion model was built subsequently for each subject. Therefore, to apply the personalised motion model, the respiratory surrogate signal of the out-of-sample subject needs to be normalised as well. This might require a preparatory phase where the surrogate is acquired to establish minimum and maximum values, so that the surrogate can be normalised in real time. More sophisticated techniques for the quantification of motion similarities might be the focus of future work, as well as the extension to other respiratory patterns, such as shallow or deep breathing.

Further work is also required to investigate the optimal settings for the NAF parameters and the number of neighbours on a larger sample of subjects. It is possible that the relatively poor performance of our method for some subjects is due to non-optimal NAF parameters. A neighbourhood ap-

proximation technique was employed in this work, but other methods could be used to provide an explicit formulation of the manifold implied by the respiratory motion. Moreover, further investigations are required to analyse the effect of pathologies and abnormalities of the heart on the cardiac respiratory motion. In the case of abnormal anatomies, difficulties in registering patient images to an atlas derived from healthy subjects might occur. To overcome these possible issues, inclusion of abnormal anatomies in the atlas formation or the use of multiple atlases could be adopted. Furthermore, more robust techniques to derive the anatomical features or the use of non image-based descriptors might be required in these cases. Future investigations could also analyse the links between the cardiac cycle motion and the respiratory motion of the heart since some pathologies (e.g. ischaemia) may cause changes to the cardiac cycle motion which might impact on the respiratory motion.

In future work we plan to investigate the validity of our assumption (correlation between organ anatomy and respiratory motion of the organ) to other organs affected by respiratory motion, such as the lungs or the liver. Such investigation could lead to a personalisation framework for the respiratory motion of the whole thorax, provided with the subject-specific anatomy only. Furthermore, previous studies have shown differences in respiratory motion of the lungs due to lung diseases and tumours (Ehrhardt et al., 2009). It is therefore possible that an explicit formulation of the manifold implied by respiratory motion could be used as a diagnostic tool for disease detection or staging. Such an approach could have application in treatment/detection of lung cancer or diseases such as chronic obstructive pulmonary disease.

Acknowledgements

This work was funded by EPSRC grant EP/H046410/1. The authors acknowledge financial support from the Department of Health via the National Institute for Health Research (NIHR) comprehensive Biomedical Research Centre award to Guy’s & St Thomas’ NHS Foundation Trust in partnership with King’s College London and King’s College Hospital NHS Foundation Trust.

References

- Aljabar, P., Heckemann, R.A., Hammers, A., Hajnal, J.V., Rueckert, D., 2009. Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy. *NeuroImage* 46, 726 – 738.

- Ashburner, J., 2007. A fast diffeomorphic image registration algorithm. *NeuroImage* 38, 95 – 113.
- 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* 53, 4147–4169.
- 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.
- 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.
- Fayad, H., Pan, T., Roux, C., Visvikis, D., 2010. A generic respiratory motion model for motion correction in PET/CT, in: *Nuclear Science Symposium Conference Record (NSS/MIC)*, 2010 IEEE, pp. 2455 –2458.
- Frangi, A., Rueckert, D., Schnabel, J., Niessen, W., 2002. Automatic construction of multiple-object three-dimensional statistical shape models: application to cardiac modeling. *IEEE Transactions on Medical Imaging* 21, 1151–1166.
- Hawkes, D., Barratt, D., Blackall, J., Chan, C., Edwards, P., Rhode, K., Penney, G., McClelland, J., Hill, D., 2005. Tissue deformation and shape models in image-guided interventions: a discussion paper. *Medical Image Analysis* 9, 163–175.
- He, T., Xue, Z., Xie, W., Wong, S., 2010. Online 4-D CT estimation for patient-specific respiratory motion based on real-time breathing signals, in: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2010*. Springer Berlin Heidelberg. volume 6363, pp. 392–399.

- Keall, P.J., Mageras, G.S., Balter, J.M., Emery, R.S., Forster, K.M., Jiang, S.B., Kapatoes, J.M., Low, D.A., Murphy, M.J., Murray, B.R., Ramsey, C.R., Van Herk, M.B., Vedam, S.S., Wong, J.W., Yorke, E., 2006. The management of respiratory motion in radiation oncology, report of AAPM Task Group 76. *Medical Physics* 33, 3874+.
- King, A., Boubertakh, R., Rhode, K., Ma, Y., Chinchapatnam, P., Gao, G., Tangcharoen, T., Ginks, M., Cooklin, M., Gill, J., Hawkes, D., Razavi, R., Schaeffter, T., 2009. A subject-specific technique for respiratory motion correction in image-guided cardiac catheterisation procedures. *Medical Image Analysis* 13, 419–431.
- King, A., Tsoumpas, C., Buerger, C., Schulz, V., Marsden, P., Schaeffter, T., 2011. Real-time respiratory motion correction for simultaneous PET-MR using an MR-derived motion model, in: *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2011 IEEE, pp. 3589–3594.
- Klinder, T., Lorenz, C., 2012. Respiratory motion compensation for image-guided bronchoscopy using a general motion model, in: *Biomedical Imaging (ISBI)*, 2012 9th IEEE International Symposium on, pp. 960–963.
- Konukoglu, E., Glocker, B., Zikic, D., Criminisi, A., 2013. Neighbourhood approximation using randomized forests. *Medical Image Analysis* 17, 790–804.
- Linte, C., Moore, J., Peters, T., 2010. How accurate is accurate enough? A brief overview on accuracy considerations in image-guided cardiac interventions, in: *Engineering in Medicine and Biology Society (EMBC)*, 2010 Annual International Conference of the IEEE, pp. 2313–2316.
- von Luxburg, U., 2007. A tutorial on spectral clustering. *Statistics and Computing* 17, 395–416.
- 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.
- McClelland, J., Hawkes, D., Schaeffter, T., King, A., 2013. Respiratory motion models: A review. *Medical Image Analysis* 17, 19–42.

- Metz, C., Baka, N., Kirisli, H., Schaap, M., Klein, S., Neefjes, L., Mollet, N., Lelieveldt, B., De Bruijne, M., Niessen, W., Van Walsum, T., 2012. Regression-based cardiac motion prediction from single-phase CTA. *IEEE Transactions on Medical Imaging* 31, 1311–1325.
- Nehmeh, S.A., Erdi, Y.E., 2008. Respiratory motion in Positron Emission Tomography/Computed Tomography: A review. *Seminars in Nuclear Medicine* 38, 167–176.
- 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.
- Pennec, X., 2009. Statistical computing on manifolds: From riemannian geometry to computational anatomy, in: Nielsen, F. (Ed.), *Emerging Trends in Visual Computing*. Springer Berlin Heidelberg. volume 5416 of *Lecture Notes in Computer Science*, pp. 347–386.
- Peressutti, D., Penney, G.P., Housden, R.J., Kolbitsch, C., Gomez, A., Rijckhorst, E.J., Barratt, D.C., Rhode, K.S., King, A.P., 2013a. A novel Bayesian respiratory motion model to estimate and resolve uncertainty in image-guided cardiac interventions. *Medical Image Analysis* 17, 488 – 502.
- Peressutti, D., Penney, G.P., Kolbitsch, C., King, A.P., 2013b. Personalising cross-population respiratory motion models using anatomical features, in: *Proceedings of the 17th Conference on Medical Image Understanding and Analysis*.
- Preiswerk, F., Arnold, P., Fasel, B., Cattin, P., 2012. Robust tumour tracking from 2D imaging using a population-based statistical motion model, in: *Mathematical Methods in Biomedical Image Analysis (MMBIA)*, 2012 IEEE Workshop on, pp. 209 –214.
- Rao, A., Sanchez-Ortiz, G., Chandrashekara, R., Lorenzo-Valds, M., Mohiaddin, R., Rueckert, D., 2002. Comparison of cardiac motion across subjects using non-rigid registration, in: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2002*. Springer Berlin Heidelberg. volume 2488, pp. 722–729.

- Rueckert, D., Frangi, A., Schnabel, J., 2003. Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration. *IEEE Transactions on Medical Imaging* 22, 1014–1025.
- Samei, G., Tanner, C., Székely, G., 2012. Predicting liver motion using exemplar models, in: *Abdominal Imaging. Computational and Clinical Applications*. Springer Berlin Heidelberg. volume 7601, pp. 147–157.
- Schweikard, A., Glosser, G., Bodduluri, M., Murphy, M.J., Adler, J.R., 2000. Robotic motion compensation for respiratory movement during radiosurgery. *Computer Aided Surgery* 5, 263 – 277.
- Scott, A.D., Keegan, J., Firmin, D.N., 2009. Motion in cardiovascular MR imaging. *Radiology* 250, 331–351.
- Sharp, J.T., Goldberg, N.B., Druz, W.S., Danon, J., 1975. Relative contribution of rib cage and abdomen to breathing in normal subjects. *Journal of Applied Physiology* 39, 608–618.
- Shechter, G., Ozturk, C., Resar, J.R., McVeigh, E.R., 2004. Respiratory motion of the heart from free breathing coronary angiograms. *IEEE Transactions on Medical Imaging* 23, 1046–1056.
- Wolz, R., Aljabar, P., Hajnal, J.V., Hammers, A., Rueckert, D., 2010. LEAP: Learning embeddings for atlas propagation. *NeuroImage* 49, 1316 – 1325.