DOI:
10.1016/j.media.2009.01.003

Document Version
Early version, also known as pre-print

A subject-specific technique for respiratory motion correction in image-guided cardiac catheterisation procedures

A. P. King [a,*] R. Boubertakh [a] K. S. Rhode [a] Y. L. Ma [a]
P. Chinchapatnam [b] G. Gao [b] T. Tangcharoen [a] M. Ginks [a]
M. Cooklin [c] J. S. Gill [c] D. J. Hawkes [b] R. S. Razavi [a]
T. Schaeffter [a]

[a] *Interdisciplinary Medical Imaging Group, Division of Imaging Sciences, 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.*

[b] *Centre for Medical Image Computing, University College, Gower Street, London, WC1E 6BT, U.K.*

[c] *Department of Cardiology, Guy's and St Thomas' NHS Foundation Trust, St Thomas' Hospital, London, SE1 7EH, U.K.*

Abstract

We describe a system for respiratory motion correction of MRI-derived roadmaps for use in X-ray guided cardiac catheterisation procedures. The technique uses a subject-specific affine motion model that is quickly constructed from a short pre-procedure MRI scan. We test a dynamic MRI sequence that acquires a small number of high resolution slices, rather than a single low resolution volume. Additionally, we use prior knowledge of the nature of cardiac respiratory motion by constraining the model to use only the dominant modes of motion. During the procedure the motion of the diaphragm is tracked in X-ray fluoroscopy images, allowing the roadmap to be updated using the motion model. X-ray image acquisition is cardiac gated. Validation is performed on four volunteer datasets and three patient datasets. The accuracy of the model in 3-D was within 5mm in 97.6% of volunteer validations. For the patients, 2-D accuracy was improved from 5-13mm before applying the model to 2-4mm afterwards. For the dynamic MRI sequence comparison, the highest errors were found when using the low resolution volume sequence with an unconstrained model.

Key words: cardiac, respiratory motion, catheterization, augmented reality

1 Introduction

Medical imaging has a very important role to play in the diagnosis of cardiovascular disease and is also increasingly used to guide minimally-invasive treatment. X-ray fluoroscopy is currently the modality of choice for procedure guidance due to its high spatial and temporal resolution allowing good visualisation of interventional devices such as guide-wires and catheters. However, X-ray imaging offers very little soft tissue contrast resulting in poor visualisation of the cardiovascular anatomy. In fact, iodine-based contrast agents are routinely and repeatedly used to highlight target structures during procedures. Furthermore, one of the major disadvantages of X-ray imaging is that it uses ionising radiation, a potential hazard to both patients and staff. X-ray radiation dose is of particular importance in very long procedures such as electrophysiology (EP) studies. To overcome the lack of soft tissue contrast, we have previously described the development of a system that allows structures derived from magnetic resonance imaging (MRI) to be overlaid in real-time onto X-ray fluoroscopy images (Rhode et al., 2003, 2005). An MRI-derived roadmap is aligned to the X-ray image coordinate system using a combination of system calibration and real-time tracking. However, misalignment occurs if there is patient motion between MRI imaging and the X-ray catheterisation procedure, for example movement of the heart due to the cardiac cycle, bulk patient movement and respiratory motion. In this paper we address the problem of respiratory motion and deformation for image-guided cardiac catheterisations in adult and paediatric patients. In most adult cases the procedure is carried out with the patient sedated and free breathing. However, paediatric cases are normally performed with the patient under a general anaesthetic and with their breathing controlled by a ventilator.

The effect of respiration on the motion of the heart has been studied by a number of researchers. Manke et al. (2002) constructed subject-specific affine models of cardiac respiratory motion with the purpose of motion correction in coronary MR angiography. They noted that for most subjects an affine model was sufficient, but some subjects had significant non-linear local deformations. More recently, the same group have proposed a technique that employs multiple navigators to capture variability in the breathing cycle (Manke et al., 2003). Shechter et al. (2004) used free breathing coronary angiograms to analyse the motion of the coronary artery during respiration. They noted that the majority of the translational motion is in the superio-inferior (S-I) direction, although small anterio-posterior (A-P) and very small medio-lateral (M-L) translations are also sometimes present. All subjects had some rotational component about the M-L axis, and some also had small rotations about the

* Corresponding author. Tel.: +44-2071888376, Fax: +44-2071885442
Email address: andrew.king@kcl.ac.uk (A. P. King).

A-P axis and the S-I axis. In McLeish et al. (2002), we formed subject-specific statistical shape models of cardiac respiratory motion and deformation. Significant variation was observed between individuals in both the rigid-body and deformation components of the motion and some subjects had local deformations of a single chamber of the heart. These models were formed using images acquired at breath-hold. However, Nehrke et al. (2001) and Keegan et al. (2002) have both reported the issue of *hysteresis* in cardiac respiratory motion: the motion and deformation of the heart is not the same in the inspiration and expiration phases. All of these studies have indicated that respiratory motion is subject-specific.

So far, simple motion models have been commonly employed in MRI image acquisition. Most MRI scanners will use a pencil-beam navigator on the diaphragm to track respiratory phase, and acquire data only within a small gating window of navigator values. Inside this window small movements are compensated for by using a simple translational model (Wang et al., 1995). More recently, affine (Manke et al., 2002; Nehrke and Bornert, 2005) and even nonrigid (Ablitt et al., 2004) models have been proposed. However, during image-guided catheterisations the patient is moved out of the MRI scanner before the procedure begins. Therefore we can no longer use the MRI navigator to measure the motion of the diaphragm during the procedure. A different means of measuring respiratory phase must be devised.

Previous work on respiratory motion correction in X-ray images has mostly been 2-D image-based techniques, in which motion is extracted purely from the X-ray images which are then processed to correct for the motion (Meijering et al., 1999). A 3-D model-based approach was described in Shechter et al. (2005), in which a model of cardiac and respiratory motion of the coronary arteries was constructed from biplane contrast-enhanced X-ray image sequences. The model was applied by tracking the motion of the diaphragm in subsequent (non-enhanced) X-ray images. However, forming the model from X-ray images under contrast injection means that it will be constructed from quite a limited amount of data. Finally, in Timinger et al. (2004) a technique for respiratory motion correction of static roadmaps was demonstrated on a dynamic heart phantom.

In this paper we propose and validate an MRI-based technique for whole heart respiratory motion compensation for use in X-ray guided cardiac catheterisation procedures. Although our main application area is EP procedures, the technique has potential application to other procedures. The technique employs a subject-specific affine motion model that is easily formed from a short preprocedure MRI scan. We investigate two different MRI sequences to form this model. We also compare two different registration techniques: a full affine registration and a constrained affine registration. The constrained registration incorporates prior knowledge of the likely nature of cardiac respiratory motion

by limiting the motion parameters to those that are known to be dominant. During the procedure respiratory phase is tracked using real-time X-ray images. The X-ray images are synchronised with the electrocardiagram (ECG) signal and cardiac-gated to eliminate motion due to the cardiac cycle. The model formation process is validated on volunteer data, and the complete system (including motion corrected overlays on real-time X-rays) is demonstrated on three patients.

2 Method and Materials

Figure 1 gives an overview of the processes involved in the proposed respiratory motion correction technique. Motion correction is applied to an MRI-derived roadmap based on a motion model derived from MRI and respiratory phase information derived from X-ray images. The motion-corrected roadmap is combined with the X-ray images to provide a fused display. The following sections deal with each of the subsystems in turn. Section 2.2 describes the formation of the motion model, including details of the MRI sequences and the model itself, and Section 2.3 outlines the application of the model, including respiratory phase estimation and the techniques used to motion-correct the roadmap. Section 2.4 provides details of our validation procedures, including a system simulation and validation on clinical cases.

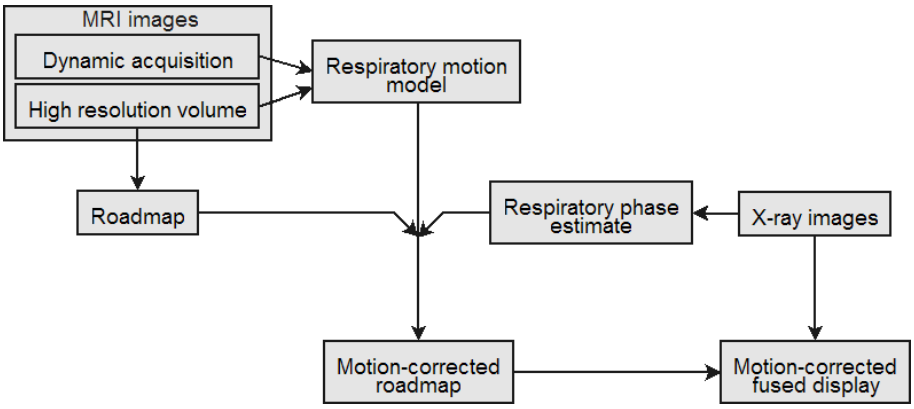

Fig. 1. An illustration of the processes involved in the respiratory motion correction technique. A roadmap is derived from MRI, and MRI data is also used to form a subject-specific model of respiratory motion. Respiratory phase is tracked in fluoroscopic X-ray images, and this phase information is combined with the motion model and the roadmap and fused with the X-ray images to provide motion-corrected image guidance.

2.1 Materials

All volunteer experiments were carried out in the XMR suite at Guy's Hospital, London. This consists of a 1.5 Tesla cylindrical bore MRI scanner (Philips Achieva I/T) and a single plane mobile cardiac X-ray set (Philips BV Pulsera). Patients can be moved between the MRI scanner and the X-ray set on a sliding bed in less than 60 seconds. Image guidance is provided by overlaying a roadmap derived from MRI images onto real-time X-ray images, aligned with an accuracy of 2mm (Rhode et al., 2003, 2005). All patient data was acquired in an identical XMR suite at Evelina Children's Hospital in St Thomas' Hospital, London.

2.2 Motion Model Formation

2.2.1 MRI Image Acquisition

For both patients and volunteers, a number of different MRI images were acquired to construct and validate the respiratory motion models. First, to form the models, a series of dynamic acquisitions, D_i, was acquired at different respiratory positions. Next, to form the roadmap, a high resolution end-expiration anatomy scan, H_1, was acquired. Finally, for validation purposes on volunteers only, three additional high resolution anatomy scans, H_2, H_3 and H_4, were acquired at different respiratory positions. Below we detail the sequences used for each of these acquisitions.

Dynamic Acquisitions

To form each respiratory model a dynamic scan was applied obtaining a number of near real-time free-breathing acquisitions that cover a range of respiratory positions. For the volunteer data we used 40 dynamics to form each model, and for patient data we used 100. However, in our experience 40 dynamics is sufficient to form an accurate enough model. We tested two different MRI sequences for these dynamic acquisitions and compared their results. Each sequence was cardiac triggered and gated at late diastole, resulting in one acquisition per heartbeat. Each dynamic acquisition takes approximately 1 second to acquire. Whilst acquiring these dynamic sequences a pencil-beam navigator was applied on the dome of the right hemi-diaphragm immediately before and after dynamic acquisition. These real-valued navigator values represent the 1-D translation of the diaphragm in the S-I direction. The average of these lead and trail navigators, V_i, was used in forming the model. Each dynamic was classified as either inspiration, expiration or end-cycle based on its navigator value and those of its predecessor/successor. Dynamic i was clas-

sified as inspiration if $V_{i-1} > V_i > V_{i+1}$, expiration if $V_{i-1} < V_i < V_{i+1}$, or end-cycle if either $V_{i-1} < V_i > V_{i+1}$ or $V_{i-1} > V_i < V_{i+1}$ (see Figure 2).

The different MRI sequences we tested for the dynamic acquisitions were:

(1) *Low resolution 3-D*: 3-D TFEPI, typically, 20 slices, TR = 10ms, TE = 4.9ms, flip angle = 20^o, acquired voxel size $2.7 \times 3.6 \times 8.0\text{mm}^3$, acquired matrix size 128×77, reconstructed voxel size $2.22 \times 2.22 \times 4.0\text{mm}^3$, reconstructed matrix size 144×144, TFE factor 26, EPI factor 13, TFE acquisition time 267.9ms. This sequence is similar to that used previously to form models for motion-corrected MRI image acquisition (Manke et al., 2002).

(2) *High resolution 2-D*: Multislice balanced TFE, typically, 2 sagittal slices, TR = 2.7ms, TE = 1.37ms, flip angle = 60^o, acquired voxel size $1.78 \times 1.8 \times 8.0\text{mm}^3$, acquired matrix size 196×142, reconstructed voxel size $1.37 \times 1.37 \times 8.0\text{mm}^3$, reconstructed matrix size 256×256, TFE factor 47, TFE acquisition time 128.7ms.

The acquisition windows of these scans (267.9ms and 2×128.7ms) are small compared with the length of a respiratory cycle, so we believe that errors introduced into the model by temporal blurring will be minor. Both sequences are cardiac gated at the diastolic rest period to minimise cardiac cycle motion.

To test the reproducibility of the model formation process on the volunteer datasets, for each of the two MRI sequences we acquired enough data to form the model three times. 40 dynamic acquisitions were acquired for each model, making a total of 120 acquisitions for each of the sequences listed above.

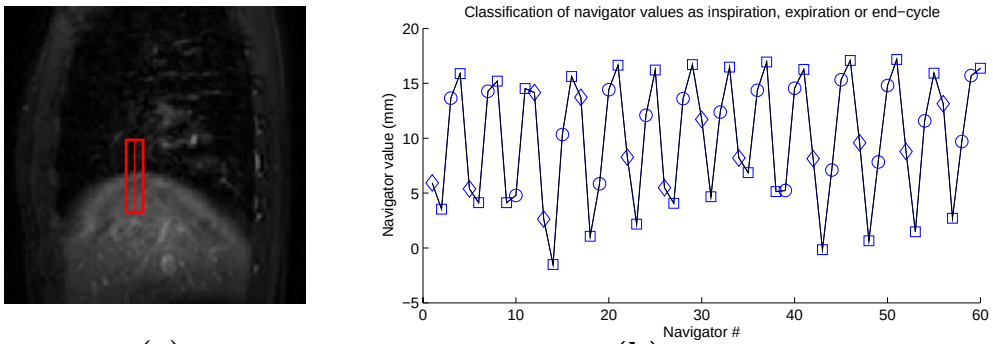

(a) **(b)**

Fig. 2. Classifying navigator values according to respiratory direction: (a) Positioning the MRI scanner navigator on the diaphragm; (b) Classification of navigator values as inspiration, expiration or end-cycle based on the neighbouring values.

High Resolution Volume

As part of the standard imaging protocol for the XMR image guidance system a high resolution free-breathing whole heart scan is acquired (3D balanced TFE, respiratory gated at end-expiration, cardiac triggered and gated at late diastole, typically, 120 sagittal slices, TR=4.4ms, TE=2.2ms, flip angle=90^o, acquired voxel size $2.19 \times 2.19 \times 2.74\text{mm}^3$, acquired matrix size 160×120, reconstructed voxel size $1.37 \times 1.37 \times 1.37\text{mm}^3$, reconstructed matrix size 256×256, acquisition window ≈ 100ms, navigator window 5mm, scan time approximately 5 minutes). This scan is acquired during the diastolic rest period, minimising the effect of cardiac cycle motion. Some blurring due to respiratory motion will occur because of the size of the navigator window, but this is unavoidable as decreasing this window would result in an unacceptably long scan time. The roadmap is formed by segmenting structures of interest, such as the four chambers and/or major vessels, from the high resolution scan using the ANALYZE software package (Biomedical Imaging Resource, Mayo Foundation, Rochester, MN, USA). The marching cubes algorithm implemented in VTK (Schroeder et al., 1997) is used to extract a surface from the binary images. The high resolution image was also used in the model formation process.

For validation purposes on the volunteer datasets, three additional high resolution images covering the same field of view were acquired during free-breathing. The four volumes (i.e. $H_1 \ldots H_4$) were acquired simultaneously using a modified respiratory navigated sequence that allows multiple gating windows (Kolmogorov et al., 2003). The sizes of the navigator gating windows were the same as for the end-expiration volume (5mm). Therefore the validation results computed using these images should be interpreted with this figure in mind. Four respiratory positions were selected: end expiration, end-inspiration, one inspiration image in mid-cycle, and one expiration image in mid-cycle. The total acquisition time for all four volumes was 20-30 minutes. During acquisition, the navigator values were saved for each measurement as well as the gating status (accepted or rejected) and the respiratory position being acquired. The average navigator values, $U_1 \ldots U_4$, of each respiratory position over the whole acquisition duration was computed off-line from the corresponding saved navigator values. The three additional images were not used in forming the model.

2.2.2 The Respiratory Model

Figure 3 illustrates the roles of each of the MRI images in forming the respiratory model. The algorithm can be summarised as follows (King et al., 2008):

(1) One of the dynamic acquisitions $D_1 \ldots D_N$ is chosen as a reference. This

is the one with the highest navigator value, which represents the end-expiration position - we refer to this as D_{ref}.

(2) The reference acquisition D_{ref} is rigidly registered to the high resolution end-expiration volume H_1 using an intensity based registration technique that seeks to maximise the normalised mutual information (Studholme et al., 1999) between the images. The registration algorithm uses a gradient descent technique to maximise the similarity measure, which is computed over the whole area of intersection between the high resolution volume and the dynamic acquisition. We use a rigid registration over the entire field of view of the volume at this stage because we assume that H_1 and D_{ref} are acquired at the same respiratory position. We denote the result of this rigid registration by R_{ref}.

(3) Next, H_1 is registered to every other dynamic acquisition (apart from D_{ref}) using an affine registration. This time we use R_{ref} to transform the dynamic acquisition to the coordinate system of the high resolution volume before the registration commences. The similarity measure is computed only within a manually delimited elliptical region of interest centred on the heart. Therefore this registration computes the transformations required to align H_1 from end-expiration to the respiratory positions of each dynamic acquisition. Two different registration techniques were tested: a full 12 parameter affine registration and a constrained 5 parameter affine registration. This constrained registration allowed variation in only the three translations, the M-L axis rotation and the S-I scaling.

(4) The next stage is to fit curves to the variation of each parameter with navigator value. We do this separately for inspiration and expiration. The classification of dynamics into inspiration, expiration or end-cycle as described in Section 2.2.1 is used to determine which data to use for each curve. End-cycle data is used in estimating both inspiration and expiration curves. Erroneous cycles that are due to small fluctuations in the navigator value are eliminated by rejecting any cycle for which the difference between end-expiration and end-inspiration navigator values is less than 5% of the maximum difference over all cycles. Each curve is modelled as a second order polynomial function of navigator value. The polynomial coefficients are determined using a linear least-squares fit that is further constrained to ensure that the inspiration and expiration curves meet at the minimum and maximum navigator values. This was done by incorporating an equality constraint into the fit using the 'lsqlin' function in Matlab (The Mathworks, Natick, MA).

Each full affine model consists of 24 sets of polynomial coefficients: 12 for inspiration and 12 for expiration. Using the constrained registration technique, the models consist of 10 sets of polynomial coefficients (5 each for inspiration and expiration).

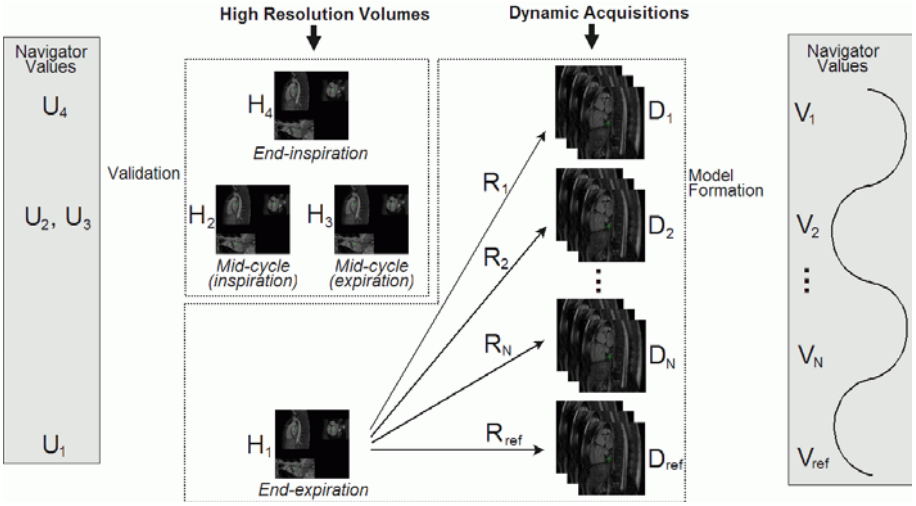

Fig. 3. The MRI images used in constructing and validating the motion model. The high resolution end-expiration volume, H_1, is registered to each dynamic acquisition, $D_1 \ldots D_N$. The resulting registrations $R_1 \ldots R_N$, together with the navigator values for the dynamic acquisitions, $V_1 \ldots V_N$ and the high resolution volume, U_1, are used to form the motion model. The three additional high resolution volumes, $H_2 \ldots H_4$ and their navigator values, $U_2 \ldots U_4$, are used for validation purposes only.

2.3 Motion Model Application

2.3.1 Intraprocedure Diaphragm Tracking

Movement of the diaphragm during the procedure was tracked using live fluoroscopic X-ray images. First, a view was selected which contained both the region of interest (i.e. the roadmap to be overlaid onto the X-ray image) and the diaphragm. Next, a rectangular strip of the image was manually delimited in which the motion of the diaphragm was visible and no other features were present (see Figure 4). One X-ray image was chosen as a reference. The diaphragm motion of subsequent X-ray images was determined by computing the 1-D translation (along the long axis of the rectangle) that minimised the mean sum of squared differences between the intensities in the current image and the reference image within the rectangular strip.

2.3.2 Cardiac Gating

The respiratory motion model is constructed based on information from two MRI sequences: the high resolution end-expiration anatomy scan, H_1, and the dynamic scan. Both of these scans are cardiac triggered and gated at diastole. Therefore the model is only valid at diastole. However, real-time X-ray images are acquired at arbitrary cardiac phases. Therefore to preserve the integrity of the model we must apply it only on X-ray images acquired during the diastolic

Fig. 4. A manually delimited region of interest for tracking diaphragm motion in live fluoroscopic X-ray images.

phase. We performed cardiac gating by synchronising X-ray acquisition with the digital output of the electrocardiogram (ECG) signal of the patient. A time window after the initial R-wave was defined to be the same window used for MRI data acquisition. Images not acquired within this time window were rejected.

2.3.3 Updating the Roadmap

The final stage in the respiratory motion correction technique is to use the diaphragm motion tracked as described in Section 2.3.1 to update the position and shape of the roadmap formed from H_1 on cardiac gated X-ray images. Recall that the motion model is based on the value of a pencil-beam navigator applied during acquisition of the dynamic scan. This represents the 1-D movement of the diaphragm relative to some arbitrary initial position. The values resulting from our X-ray diaphragm tracking technique represent 1-D movements relative to a different arbitrary initial position, possibly of a different part of the diaphragm. Therefore we must relate these two sets of values. We do this by recording the values of the extreme respiratory positions (end-expiration and end-inspiration) using both techniques and computing the linear transform that relates the two sets of values.

For sedated patients there are potential difficulties with this approach: the end-inspiration position is known to be highly variable during normal free breathing, so if the end-inspiration positions during MRI image acquisition and X-ray image acquisition were different this would lead to errors in computing the linear transform. This problem is less significant for general anaesthetic patients. These patients have their breathing controlled by a ventilator so it is possible to ensure that it is regular. We overcome this difficulty by noting that the end-inspiration position during deep breathing is much more repeatable than during normal breathing. Therefore for clinical cases where the patient was free breathing all images used in computing this conversion (MRI and X-ray) were acquired during deep breathing. For the MRI scanner navigator

values we used the maximum and minimum navigator values over the entire dynamic scan acquisition (during which the patient was asked to take some deep breaths), and for the X-ray navigator values we used the maximum and minimum values found during a short period of deep breathing. In the rest of this paper, we refer to the end-inspiration position during deep breathing as full-inspiration.

2.4 Validation

In this section we describe the experiments we performed to validate the respiratory motion correction technique. To validate the motion model, we performed MRI experiments to assess the accuracy and reproducibility of the model formation process on four volunteer datasets. The four volunteers were all male and aged between 23-33 years. The sizes of the hearts of volunteers A, B and D were average. Volunteer C's heart was slightly smaller. To validate the overall technique we used a numerical simulation and X-ray data from clinical cases.

2.4.1 Validation of the Motion Model

Landmarks

The validation described in the proceeding sections is based on the locations of ten anatomical landmarks, which were selected for their relevance to EP procedures. These landmarks were manually localised in the high resolution volumes, $H_1, \ldots, H_4$. They were the left lower/upper pulmonary vein, the right lower/upper pulmonary vein, the junction of the right atrium with the inferior and superior vena cava, the tricuspic valve and three points on the interventricular septum. To test the reproducibility of landmark localisation, five separate localisations of all 10 points on the same image were performed. The RMS error of the repeated landmark localisations (i.e. the intra-observer variability) was found to be less than 1mm.

Reproducibility

For each dynamic sequence we acquired three sets of 40 dynamics for each volunteer. From these three sets we formed motion models three times. These were used to assess the reproducibility of the model formation process. For each set of three models, one hundred different respiratory positions were defined (50 in inspiration and 50 in expiration) by taking evenly spaced navigator values in a range covered by all three models. For each respiratory position, each of the models was used to predict the locations of the ten anatomical

landmarks from the high resolution volume, H_1. From these three sets of predicted locations, a mean predicted location was computed for each landmark. The root mean square (RMS) error was computed away from these mean locations for each landmark for each of the three models over the whole respiratory cycle.

Note that the three models were formed only to test reproducibility. Any single model could be used for motion correction. The reproducibility test was only performed on the volunteer datasets. For the clinical cases described in Section 2.4.3 only a single model was formed.

Validation Technique 1 - Target Registration Error

The accuracy of the model formation process was assessed using two techniques. First, the additional high resolution volumes mentioned in Section 2.2.1 and illustrated in Figure 3 were used to test the predictive power of the models. The average navigator value of H_2, H_3 and H_4, relative to the average navigator value for H_1 was used to predict affine transformations for the model to be tested. These transformations were used to transform the anatomical landmarks from H_1 and predict their positions in H_2, H_3 and H_4. These predicted positions were compared with the manually localised positions. The RMS error was computed for each model. Formally, the RMS error for model M tested against high resolution image j can be written as

$$E_1(M,j) = \sqrt{\frac{1}{10}\sum_{l=1}^{10}\left|(\vec{L}_l(j) - M(U_j - U_1) * \vec{L}_l(1))\right|^2}, j = 2,3,4 \qquad (1)$$

where $\vec{L}_1(j) \ldots \vec{L}_{10}(j)$ are vectors representing the coordinates of the ten landmarks in H_j and U_j denotes the average navigator value for H_j. The values of the dynamic sequence navigators V_i and the high resolution image navigators U_j were related by adding an offset to the high resolution navigators to ensure that their maximum (i.e. end-expiration) values were identical. $M(\cdot)$ denotes the affine transformation matrix produced by applying a navigator value to model M.

This first validation technique assesses the overall target registration error (in 3-D) of the respiratory modelling technique. This includes errors in the registrations of the dynamic acquisitions to H_1 and also errors in modelling the registration parameters. However, it only assesses the errors at three respiratory positions (those of the three additional high resolution volumes H_2, H_3 and H_4).

Validation Technique 2 - Comparing Predicted/Actual Registrations

The second validation technique tested predicted transformations against registration results for each dynamic acquisition. For this we used a leave-one-out test. From each dataset, models were formed by leaving out each of the dynamic acquisitions $D_1 \ldots D_N$ in turn. We denote these models $M_{leave_1} \ldots M_{leave_N}$ and the navigator values of the N dynamic acquisitions as $V_1 \ldots V_N$. These models, together with the navigator values, were used to predict transformations for each left-out dynamic acquisition. These predicted transformations were compared with the actual registration results for the dynamic acquisitions. The predicted and actual registration results were compared by using the two registrations to transform the ten landmarks and finding the distance between them. Formally, the RMS error for model M can be written as

$$E_2(M) = \sqrt{\frac{\sum_{i=1}^{N}\sum_{l=1}^{10}\left|(R_i * \vec{L}_l(1) - M_{leave_i}(V_i) * \vec{L}_l(1))\right|^2}{10N}} \tag{2}$$

where R_i is the actual registration result for dynamic acquisition i.

This second validation technique assesses the residual error of the modelling process. It does not include any errors in the dynamic acquisition registrations used to form the model. However, it does have the advantage that we can perform the validation at many different respiratory positions.

2.4.2 System Simulation

To assess the overall accuracy of the respiratory motion correction technique we produced a system simulation. The error sources that contribute to the overall accuracy are:

- The 3-D to 2-D registration error of the XMR image guidance system, $\vec{e}_{XMR}$.
- Tracking errors in estimating the diaphragm motion in X-ray images, e_{track}.
- Scaling errors when relating the extreme respiratory positions (end-expiration and full inspiration) between the MRI navigator and the X-ray image (as described in Section 2.3.3), e_{scale}.

In addition there will be errors in the motion model. The 3-D to 2-D registration error is due to errors in calibrating and tracking the components of the XMR suite. The tracking error is due to noise, distortion and calibration errors in the X-ray system. The scaling error will occur if, for example, the subject does not breath in to the same full inspiration position during MRI scanning and X-ray image acquisition.

The numerical simulation assessed the overall system error based on typical values for these error sources. Figure 5 illustrates the transformations and error sources involved in the simulation. There are four transformations, which we have labelled A-D in Figure 5:

(A) For each additional high resolution image, H_j, the navigator value $U_j - U_1$ was modified by the scaling and tracking errors, e_{scale} and e_{track}. The anatomical landmarks $\vec{L}_l(1)$ were transformed by using the motion model and this noisy navigator value to compute an affine transformation.
(B) Next, these motion-corrected landmarks were transformed using a typical 3-D to 2-D registration matrix, R.
(C) Zero mean Gaussian noise was added to each landmark location to simulate the XMR registration error.
(D) The gold standard landmark locations were computed by transforming the manually localised landmarks in image j by R, without any error sources.

Finally, these 'noisy' landmarks were compared with the gold standard landmark locations to compute the overall system error. Formally, we compute the overall system simulation error E_{sim} as

$$E_{sim} = \sqrt{\frac{\sum_{l=1}^{10} \left| (R * (M(U_{j,noisy}) * \vec{L}_l(1)) + \vec{e}_{xmr}) - R * \vec{L}_l(j) \right|^2}{10}}, \quad j = 2,3,4 \tag{3}$$

where $U_{j,noisy} = (1 + e_{scale})(U_j - U_1 + e_{track})$.

Note that because the respiratory model is involved in the simulation and is a source of error itself we run the simulation separately for each model.

2.4.3 Clinical Validation

2D and 3-D validation of the complete respiratory motion correction technique was performed in three clinical cases. Patient A was a 63 year old adult who underwent a catheter ablation procedure to treat atrial fibrillation (Baszko et al., 2002). This patient was sedated and breathing freely throughout the procedure. Patients B and C were both paediatric cases, aged 2 and 4 years, who were diagnosed with ventricular septal defect. They underwent pulmonary vascular resistance (PVR) studies (Muthurangu et al., 2004) whilst under general anaesthetic. Their breathing was controlled by a ventilator. For all three patients the low resolution 3-D dynamic sequence was used with the constrained registration technique. The field of view of the dynamic volumes was positioned to cover the structures that were to be segmented for the

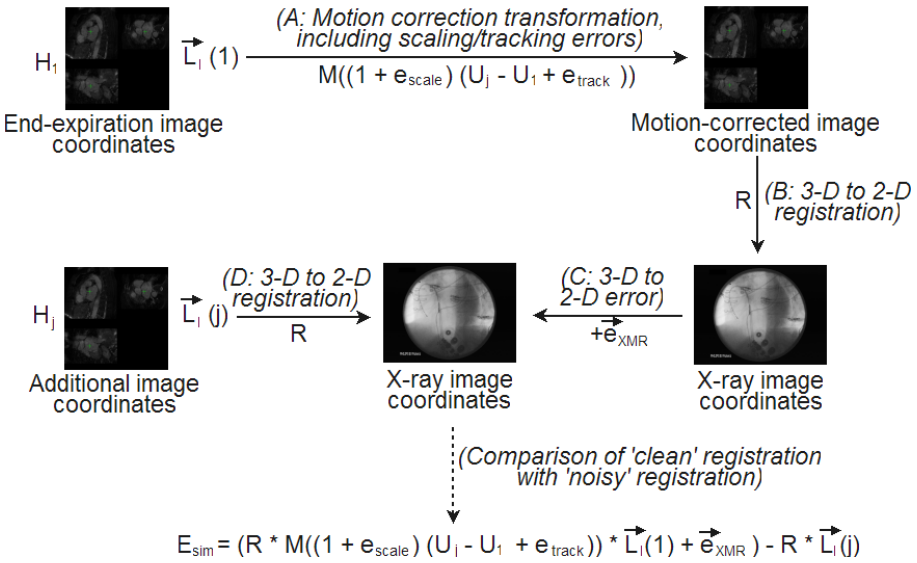

Fig. 5. Illustration of the numerical simulation of the overall respiratory motion correction system. Anatomical landmarks $\vec{L}_l(1)$ are localised in the end-expiration image H_1, and motion-corrected using the model M. The navigator value U_j used as input to the model is modified by a scaling error e_{scale} and a tracking error e_{track}. These motion corrected landmarks are projected onto the X-ray image using a typical 3-D to 2-D registration matrix R, and zero mean Gaussian noise with a standard deviation of 2mm is added to each landmark to represent the errors in the XMR image guidance system. These noisy landmarks are compared with the landmarks localised from the three additional high resolution volumes $H_j(j = 2, 3, 4)$ and transformed by R.

roadmap.

2-D Validation

We performed 2-D validation by retrospectively overlaying a rendering of a vessel onto X-ray images that showed a catheter positioned inside the vessel. Visual assessment of registration error at a number of points along the catheter was carried out. 2-D validation was performed on all three clinical procedures. For Patient A a rendering of the coronary sinus was overlaid onto X-ray images showing a catheter that was known to lie inside the coronary sinus. For Patients B and C a catheter known to lie inside the pulmonary artery was used. The 2-D registration error between the overlaid structure and the catheter was estimated by an observer, both before and after application of the motion correction technique. This was done by manually localising a number of points (between 8 and 13) on the catheter in a single image. The image was manually selected to be at full inspiration and late diastole. For each point, the closest corresponding point on the rendering was also localised. RMS and maximum errors between the two point sets were computed.

3-D Validation

3-D validation was performed on patients B and C by using biplane X-ray views to reconstruct the 3-D contour of a segment of the catheter. (No biplane X-ray data was available for patient A.) The two X-ray images were manually selected to be as closely matched as possible with regard to cardiac and respiratory phase. Piecewise-linear contours were manually defined in both X-ray views. 3-D contours were formed from these using automated software that used the epipolar constraint to reconstruct 3-D point sets from the 2-D contours. The 3-D point sets typically contained around 100 points. This process was performed using X-ray image pairs at both end-expiration and full inspiration. The error before respiratory motion correction was computed between these two point sets. For each point in the end-expiration point set the nearest point in the full inspiration point set was identified. A RMS error was computed between all pairs of corresponding points. The end-expiration point set was then transformed using the motion model to full inspiration. The error after motion correction was computed as the distance between this transformed point set and the actual point set reconstructed at full inspiration.

3 Results

3.1 Validation of the Motion Model

The validation described in Section 2.4.1 was performed for all four volunteers. For each volunteer, three models each were tested for each of four different combinations of dynamic MRI sequence (low resolution 3-D and high resolution 2-D) and registration technique (full affine and constrained). An example of a high resolution volume (for volunteer A) is shown in Figure 6.

The range of diaphragm displacements found during normal tidal breathing in the volunteers ranged between 10-20mm, which is similar to results in the literature (Nehrke et al., 2001; Manke et al., 2002). In general, when using the full affine registration, the dominant parameters were the S-I and A-P translations, the S-I scaling, and the M-L axis rotation. The models for some subjects also featured a small M-L translation. Hysteresis was observed mostly in the A-P translation, the M-L axis rotation, and sometimes the S-I scaling. It was rarely observed in the S-I translation. This confirms the findings reported elsewhere in the literature (Shechter et al., 2004). There was significant inter-subject variability, not only in the ratio of diaphragm motion to S-I translation, but also in the nature and importance of the other parameters. For example, Figures 7(a) and (b) show plots of S-I translation against diaphragm motion for volunteers B and D. Significantly different slopes are found after fitting

curves to the two datasets. Figures 7(c) and (d) compare one of the other dominant parameters (the S-I scaling) for volunteers A and C. Again, there is significant inter-subject variation.

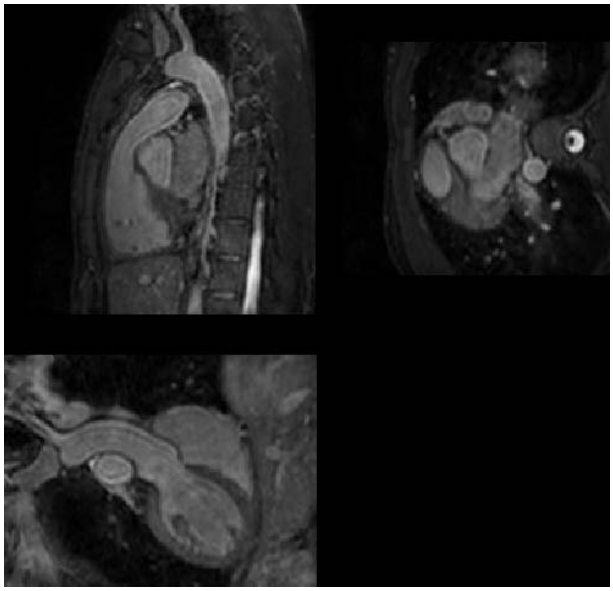

Fig. 6. Three orthogonal slices through a sample high resolution MRI volume.

3.1.1 Reproducibility

Table 1 shows the RMS and maximum reproducibility errors for the four volunteers. Figures are presented for both dynamic sequences (high resolution 2-D and low resolution 3-D) and both registration techniques (full affine and constrained affine). The RMS error between the three models is less than 1mm in every case, and the highest maximum error is 2.1mm. In most EP procedures, the smallest structures encountered are the pulmonary veins, which are more than 10mm in diameter for most adults. Therefore we estimate that the clinical accuracy requirement for EP procedures is 5mm for adults and slightly less for children. These results are within this requirement.

3.1.2 Validation Technique 1 - Target Registration Error

Table 2 shows the results for the first validation technique for all four volunteers. The target registration errors (TRE) for each additional high resolution image are provided separately - we do not include a single overall TRE because some images were out of the range of some of the models (i.e. the navigator value for the high resolution image was outside the range of navigator values for the dynamic scan from which the model was formed). In these cases the results are marked 'N/A'. It can be seen that TREs are typically around 1-4mm, with the highest errors occurring for the low resolution 3-D sequence using the

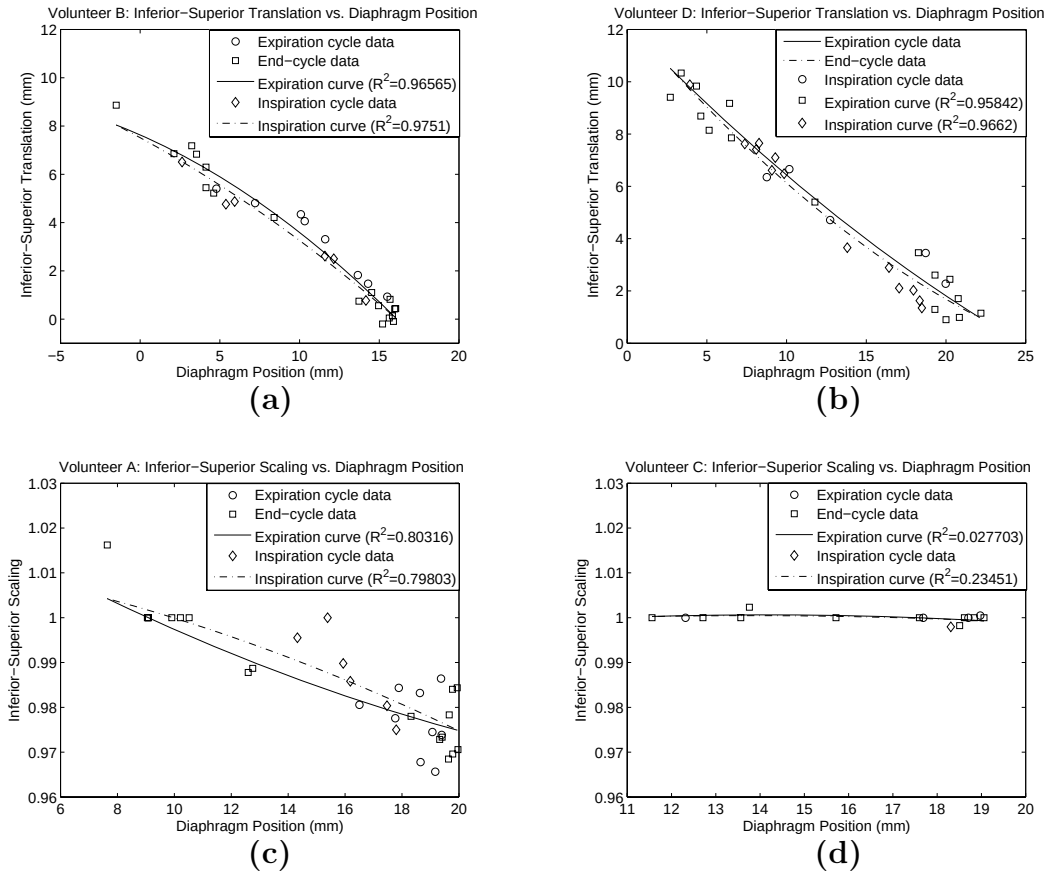

Fig. 7. (a), (b): Plots of S-I translation against diaphragm motion for (a) volunteer B; and (b) volunteer D. The ratio of S-I translation to diaphragm motion is approximately 0.6 for volunteer B and 0.5 for volunteer D. The figure used by most MRI scanners to correct heart motion is 0.6 (Wang et al., 1995). (c), (d): Plots of S-I scaling against diaphragm motion for (c) volunteer A; and (d) volunteer C. The magnitudes of the S-I scaling for volunteer A is significantly more variable with diaphragm position than that of volunteer C. All four plots are for the high resolution 2-D sequence and used a full affine registration. R^2 denotes the *coefficients of determination* , a measure of the proportion of the variability in the data sets that is accounted for by the model (maximum 1.0).

unconstrained registration. TREs without applying the motion model (i.e. the distance between the target points in the different high resolution images) were approximately 3-8mm. Our clinical accuracy requirement for EP procedures is 5mm for adults, so the model achieves a RMS error within this requirement in every case for this volunteer. Over all volunteers, 123 of the 126 high resolution image validations (97.6%) were within this clinical accuracy requirement. We also include the percentage of the overall motion recovered by the model, defined as the magnitude of the recovered motion expressed as a percentage of the gold standard motion. This is approximately 30-50% in most cases (with a maximum of 56.8%). Some lower figures are found, particularly for volunteer

Subject	Reproducibility errors (mm) for models formed using different MRI sequences: RMS/Max. using full affine registration [RMS/Max. using constrained registration]	
	High resolution 2-D	Low resolution 3-D
Volunteer A	0.5/1.2 [0.3/0.8]	0.4/1.2 [0.2/1.1]
Volunteer B	0.3/0.8 [0.4/1.0]	0.5/1.6 [0.3/0.6]
Volunteer C	0.3/0.9 [0.3/1.0]	0.4/1.2 [0.3/0.9]
Volunteer D	0.5/1.5 [0.6/1.4]	0.4/2.1 [0.4/1.4]

Table 1
The reproducibility of constructing breathing models for four volunteers. The values are the root mean square (RMS) and maximum errors for 10 target points that are clinically relevant for EP procedures. For each volunteer, one hundred different navigator values were defined (50 in inspiration and 50 in expiration), equally spaced in a range covered by the three models. For each navigator value, a set of mean target locations was computed by predicting their locations using each of the models. The RMS and maximum errors were calculated for the three models away from these mean locations.

C. We believe this is because when the overall magnitude of motion is smaller the effects of blurring artefacts in the additional high resolution images used for computing the TREs (see Section 2.2.1) and also uncertainty in identifying the true locations of the anatomical landmarks become more pronounced. Although observers were able to locate the landmarks reasonably reproducibly, there was still uncertainty that the localised points were the landmarks' true locations. Volunteer C had a slightly smaller heart than the other volunteers so the overall motion was less, making the effect of this uncertainty on the figures more significant.

We also computed the difference in TRE between modelling inspiration and expiration separately and using a single polynomial curve for both respiratory phases. The overall improvements in TRE for each volunteer were -1.62%, -0.43%, +4.47% and -3.03% respectively for volunteers A, B, C and D. This means that only 1 of the four volunteers showed an improvement from modelling inspiration and expiration separately. However, we believe that all of these figures are within the uncertainty introduced by the blurring artefacts in the extra high resolution volumes. Therefore these results are inconclusive and further investigation would be required to quantify the impact of separating the respiratory phases in the model.

3.1.3 Validation Technique 2 - Comparing Predicted/Actual Registrations

Table 3 presents the results of the second validation technique for all four volunteers. The leave-one-out test results are computed using every image in the dynamic acquisition sequence. Recall that this validation technique can be viewed as a residual error of the modelling process. We can see that using a constrained registration technique significantly reduces the residual error. This is to be expected because constraining the registration to use only five of the affine parameters instead of the full twelve reduces the number of parameters that can cause errors in the modelling.

To give an idea of the magnitude of motion being recovered by the model, we computed the distance between the target points at end-expiration and the same points transformed by each registration result, i.e.

$$\left|\vec{L}_l(1) - R_i * \vec{L}_l(1)\right|, l = 1 \ldots 10, i = 1 \ldots N. \tag{4}$$

The mean distances for the four volunteers varied between 3 and 7mm, with maximum distances at full inhale of up to 16mm. The results in Table 3 should be read with this in mind.

3.1.4 Comparison of Imaging Sequences/Registration Techniques

In this section we compare the validation results for the two dynamic MRI sequences and the two registration approaches. Figure 8 shows sample slices from the two dynamic MRI sequences.

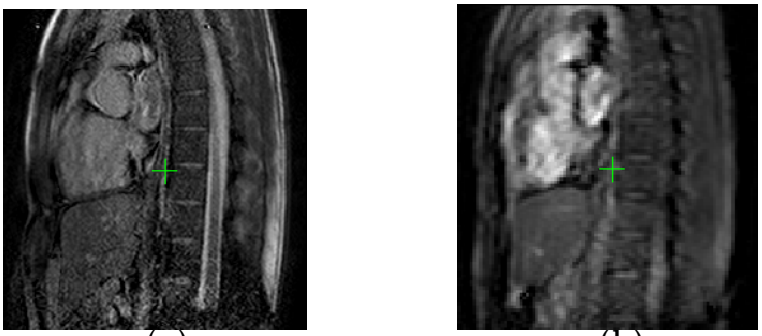

Fig. 8. Sample corresponding slices from the two dynamic MRI sequences: (a) high resolution 2-D sequence; (b) low resolution 3-D sequence.

In order to compare the performance of the models formed using the two dynamic MRI sequences and two registration techniques we produced a single TRE for each model over all volunteer datasets. We computed the mean and standard deviation of the RMS TRE values given in Table 2. Entries marked

'N/A' (i.e. those whose navigator value was outside the range covered by the model) were not included. A total of 28 entries for the high resolution 2-D sequence and 35 for the low resolution 3-D sequence were used. The overall TREs are illustrated in Figure 9(a). We can see that the highest errors occurred for the low resolution 3-D dynamic sequence. However, constraining the registration significantly reduced the TRE for the low resolution 3-D sequence.

We also compare the different model formation approaches using the leave-one-out test results. In this case a mean and standard deviation of the RMS errors from Table 3 was computed (i.e. 12 entries for each approach were used). Figure 9(b) illustrates these errors. The relative magnitudes of the errors are similar to the TRE figures.

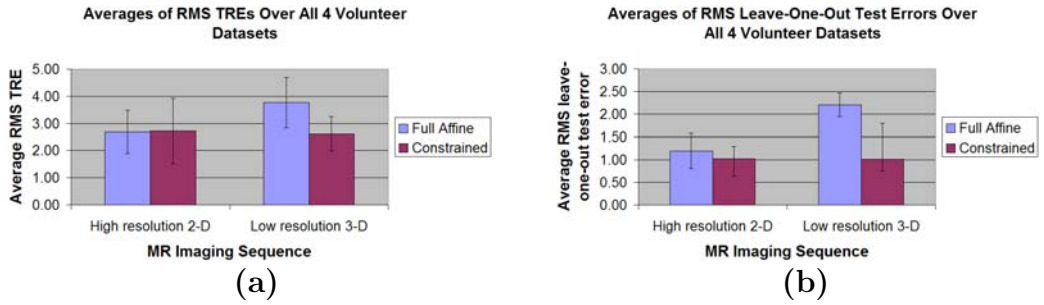

Fig. 9. (a) RMS TREs and (b) RMS leave-one-out test errors for two different MRI sequences and two different registration techniques (full affine and constrained). The mean/standard deviations of the RMS errors were computed over all four volunteer datasets

3.2 System Simulation

The 2-D registration accuracy of the XMR image guidance system is estimated to be 2mm. Therefore we used a value of 2.0mm for $\vec{e}_{XMR}$ in the simulation. Based on analysis of the navigator values of volunteers and patients, we estimate that respiratory range can vary by up to 20% during deep breathing, and by $10-15\%$ for ventilated patients, so a value of 0.2 would be a reasonable top-end estimate for e_{scale}. We estimate that tracking errors will be small, typically a maximum of 1mm. Based on these estimates, the overall system error predicted by the simulation was between 2.5mm and 4mm for all volunteer models formed. This is well within our clinical accuracy requirement of 5mm.

3.3.1 2-D Validation

Table 4 shows the estimated 2-D errors in millimetres for all three patients. For the two PVR cases (Patients B and C) the errors were estimated using a segment of the catheter that was believed to lie along the superior boundary of the pulmonary artery - in these cases the nearest corresponding points on the superior boundary of the rendering were used. For the flutter ablation case (Patient A), the catheter was assumed to lie in the centre of the coronary sinus, so the nearest corresponding point on the centre-line of the rendering were used. We can see that the registration error was reduced from 5-13mm to 2-4mm, with the model recovering up to 78.9% of the overall motion. Even allowing for the slightly smaller size of clinical target structures in the two paediatric cases (patients B and C) these results are very encouraging. Figure 10 shows the images that were used for visual assessment.

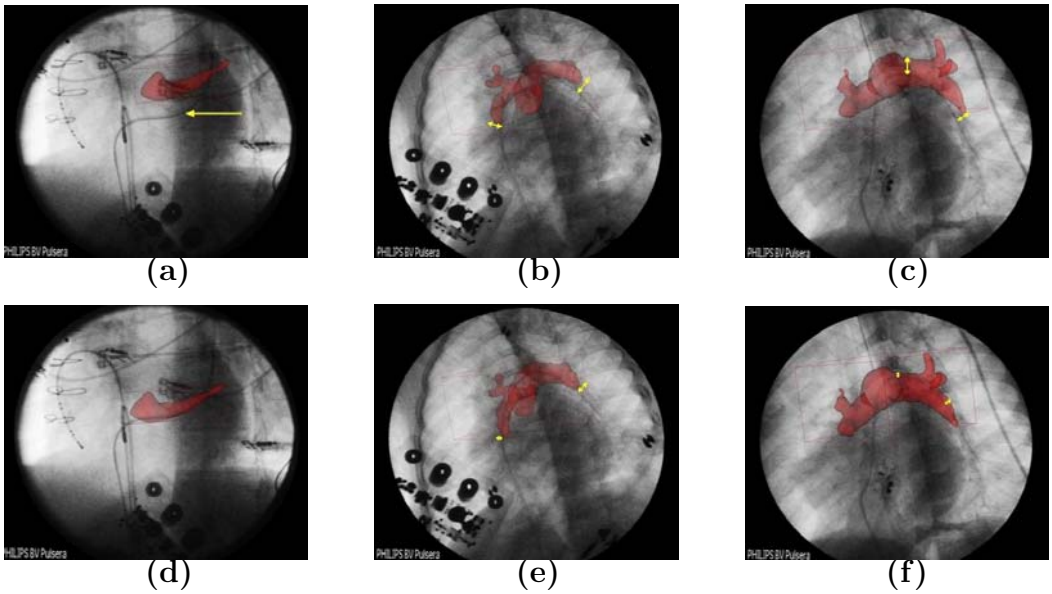

Fig. 10. Images used for visual assessment of registration error: (a) and (d) show an overlay of the coronary sinus for Patient A before and after motion correction. The coronary sinus catheter used for assessment is indicated with an arrow in (a). The mis-registration is clear in (a), whereas the catheter is well positioned inside most of the rendering in (d). The end of the catheter was positioned inside a branch off of the main sinus and so lies outside of the rendering. (b) and (e) show an overlay of the pulmonary artery for Patient B before and after motion correction. The arrows in both images show the estimated mis-registration at either end of the catheter segment used for assessment. (c) and (f) show the same information for Patient C.

3.3.2 3-D Validation

Table 5 shows the RMS and maximum 3-D errors before and after applying the model. Although in both cases the error is reduced by applying the model, Patient B has a smaller reduction than Patient C - this confirms the findings of the 2-D validation shown in Table 4. From analysis of the 3-D contours the mis-registration appears to be in the A-P direction. Renderings showing the three 3-D contours (end-expiration, transformed end-expiration, and full inspiration) for both patients are shown in Figure 11. It should be noted that these figures will be affected by errors in reconstructing 3-D points from biplane fluoroscopy images. These errors are largely due to X-ray calibration errors, which are approximately 1.5mm (Rhode et al., 2005), so the reconstruction error will be similar to this figure.

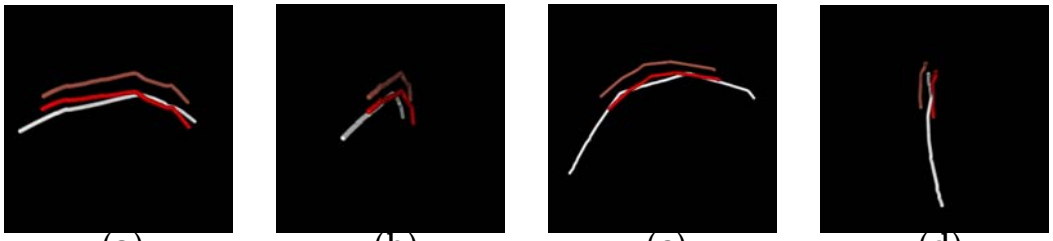

Fig. 11. Images showing the positions of the reconstructed catheter points at end–expiration (pink), full-inspiration (white) and the end-expiration segment transformed by the model to full inspiration (red). Note that the full inspiration segment is longer than the end-expiration and transformed segments - this was done deliberately to ensure that each point on the end-expiration and transformed segments had a corresponding point on the full inspiration segment. (a), (b) volunteer B; (c), (d) volunteer C.

4 Discussion

We have described a system that can be used to correct for respiratory motion in X-ray image-guided cardiac catheterisations. The results presented in this paper were produced retrospectively, but we are currently in the process of implementing the technique for real-time processing. The system is based on a subject-specific affine model of cardiac respiratory motion and automated tracking of the diaphragm in X-ray images. The method requires minimal manual intervention. Only a region of interest for the image registration needs to be manually specified - the rest of the model construction process is automatic and takes about 45 minutes for a single model on a Pentium 4 2.16GHz PC. This execution time could be reduced significantly by using a faster PC and optimising the code. Preparation of the patient and gaining intravascular

access for a catheterisation typically takes at least 15 minutes so this execution time is reasonable. Once the real-time implementation is available, this model preparation would need to be performed once before each procedure in which motion correction is required. The high resolution scan and heart surface segmentation are routinely performed for guidance purposes at our institution, and becoming increasingly common elsewhere due to the emergence of commercial image-guidance systems from manufacturers such as Phillips and Siemens, and mapping technologies such as CARTOMERGE (Biosense Webster Inc.) and Ensite NavX Fusion. Therefore the only extra scanning requirement is the dynamic MRI scan. Validation on volunteer datasets has demonstrated that the accuracy of the model is 2-4mm in 3-D. A numerical simulation of the overall system predicts a 2-D accuracy of 3-4mm. This is backed up by our experience in three clinical cases, in which 3-D registration accuracy was estimated to be 1.9-3mm, and 2-D accuracy to be 2.2-3.9mm. These errors are very encouraging and demonstrate the feasibility of clinical application of our technique.

Our proposed 2-D MRI sequence to form the motion model uses two sagittal slices separated to cover the area of interest. The reasoning behind using high resolution sagittal slices is that all of the parameters that are known to be dominant in cardiac respiratory motion (i.e. the three translations, the S-I scaling and the M-L axis rotation) are quantifiable from sagittal slices. Therefore acquiring a small number of high resolution images in this plane can provide more accurate and reliable estimates of these parameters. We have also proposed the use of a constrained affine registration approach that allows variation in only the five dominant parameters of cardiac respiratory motion. The new 2-D imaging sequence shows slightly improved accuracy compared to the single low resolution 3-D sequence that has previously been used to construct such motion models. We believe that this improvement is partly due to the extension of the effective spatial coverage of the model. Whereas the high resolution 2-D slice sequence covered the entire region of interest of the heart, the low resolution 3-D sequence did not have such a wide field of view in the M-L direction. Obtaining single volumes with complete coverage of the heart in a single cardiac cycle would have resulted in an unacceptable loss of image quality. But it appears that this loss of coverage can be compensated for by constraining the registration to include only those affine parameters that are known to be dominant. This effect is not so noticeable in volunteer C as this subject had a smaller heart than the other volunteers, meaning that the single volume covered a comparatively larger proportion of the heart. The constrained registration also shows improvements in accuracy in some subjects and could also be used to extend the effective M-L coverage of the model.

However, we should be careful about drawing strong conclusions from direct comparisons between the two dynamic MRI sequences. This is because there are potentially a number of other possible sources of error. For example, the

model formation technique we have proposed assumes that breathing is regular and repeatable, whereas in reality there are sometimes slight inter-cycle variations. Therefore the amount of variation in breathing from cycle to cycle could affect the size of the errors.

The respiratory model we have described differs from most of those previously proposed in that it allows inspiration and expiration motion to be modelled separately. The model proposed by Shechter et al. (2004, 2005) did separate inspiration and expiration, but only covered the coronary arteries and was formed from a relatively short (4-8 seconds) period of contrast injection. Modelling inspiration and expiration separately enables the model to capture the hysteretic effect that has been reported elsewhere in the literature (Nehrke et al., 2001; Keegan et al., 2002). Although our results did not show a conclusive improvement from this modelling approach, our observations of the hysteresis present in different motion models suggest that it can be beneficial. Constraining the curves to meet at the minimum and maximum navigator values ensures that the respiratory cycle produced by the model is cyclic, i.e. when moving to and from either end-expiration or end-inspiration the variation in the transformations produced by the model is smooth. This constraint does, however, contain an implicit assumption that the respiratory cycle is regular over the course of the time taken to acquire the images. Currently our model uses affine transformations, as our experience and the literature suggests that such transformations are adequate to describe most cardiac motion (Manke et al., 2002). However, the same modelling process could be easily extended to model free-form non-rigid deformations by fitting curves to the motion of individual control points during inspiration and expiration.

In order to be able to apply the X-ray diaphragm tracking information it is first necessary to calibrate the extreme respiratory positions (as described in Section 2.3.3). As we have noted, for sedated patients we compute the linear transform that relates the end-expiration and full-inspiration points in the MRI and X-ray images during deep breathing. The deep-breathing protocol during acquisition of the dynamic MRI scan had the added advantage of increasing the range of breathing positions for which the model was applicable. Computing the transform during the procedure involved tracking the diaphragm in X-ray images for only 1 or 2 respiratory cycles. Once the linear transform was computed there was no restriction on the subject's breathing. We should note that we make an assumption here about the nature of the motion of the diaphragm. It is possible that the part of the diaphragm tracked by the MRI navigator will not be visible in the selected X-ray view. Therefore, if we track a different part of the diaphragm during MRI image acquisition and X-ray image acquisition then we assume that there is a linear relationship between the two motions. So long as the relationship is linear, then the (linear) conversion we apply will be valid.

An important point that remains to be addressed is the issue of variation in type of breathing. For X-ray guided catheterisations there can be a significant delay between pre-procedure MRI imaging and the procedure itself. Often imaging is performed one or even several days before the procedure. Also, the patient is physically moved to a different environment between MRI imaging and the procedure. This, combined with the discomfort caused by insertion of the catheter means that the state of relaxation of the patient may be changed. For example, the patient may increase their breathing rate when nervous, or take deep breaths when in discomfort. Patients under general anaesthetic are not aware of such changes and have their breathing controlled by a ventilator so this will not be an issue. However, for sedated patients these changes, together with the effect of sedation, could lead to changes in breathing pattern. Such a change may result in a different type of motion and deformation in the heart. Therefore cardiac respiratory motion/deformation may be different between the time at which data is acquired to form the model and the time at which the model is used in the procedure. This could introduce errors into the system. Although the biggest improvement we found in our clinical validation was for the sedated (adult) patient, this was mostly due to the larger magnitude of respiratory motion in adults compared to children. It is possible that some of the remaining error was due to this variation in breathing pattern. In future work we plan to investigate whether changes in breathing pattern affect the motion and deformation of the heart, and if so, we will investigate strategies to cope with this.

In conclusion, the proposed respiratory motion correction system represents a significant advance in the use of image guidance in cardiac catheterisations. Deriving the model from MRI is quick and involves no additional radiation dose. The technique has the potential to increase the cardiologist's confidence in the accuracy of the MRI-derived roadmap, thus reducing procedure time and X-ray dosage and improving outcome.

Acknowledgements

This work was funded by EPSRC grant EP/D061471/1, DTI technology programme grant 17352 and Philips Healthcare, Best, the Netherlands. 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. We are also grateful to the members of the Imaging Sciences division at St Thomas' hospital, London, particularly Philip Batchelor, for useful discussions and feedback on this work.

References

N. A. Ablitt, J. Gao, J. Keegan, L. Stegger, D. N. Firmin, and G. Z. Yang. Predictive Cardiac Motion Modelling and Correction With Partial Least Squares Regression. *IEEE Transactions on Medical Imaging*, 23(10):1315–1324, October 2004.

A. Baszko, C. A. Rinaldi, R. D. Simon, and J. S. Gill. Atrial Fibrillation Current and Future Treatments: Radiofrequency Ablation and Novel Pacing Techniques. *Int J Clin Pract*, 56(5):370–376, 2002.

J. Keegan, P. Gatehouse, G. Z. Yang, and D. Firmin. Coronary Artery Motion With the Respiratory Cycle During Breath-Holding and Free-Breathing: Implications for Slice-Followed Coronary Artery Imaging. *Magnetic Resonance in Medicine*, 47:476–481, 2002.

A. P. King, R. Boubertakh, K. L. Ng, Y. L. Ma, P. Chinchapatnam, G. Gao, T. Schaeffter, D. J. Hawkes, R. Razavi, and K. S. Rhode. A Technique for Respiratory Motion Correction in Image Guided Cardiac Catheterisation Procedures. In *Proc. SPIE Medical Imaging (in press)*, 2008.

V. N. Kolmogorov, R. Watts, M. R. Prince, R. Zabih, and Y. Wang. Simultaneous Multiple Volume (SMV) Acquisition Algorithm for Real-Time Navigator Gating. *Magnetic Resonance Imaging*, 21:969–975, 2003.

D. Manke, K. Nehrke, and P. Bornert. 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, 2003.

D. Manke, P. Rosch, K. Nehrke, P. Bornert, and O. Dossel. 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(9):1132–1141, September 2002.

K. McLeish, D. L. G. Hill, D. Atkinson, J. M. Blackall, and R. Razavi. A Study of the Motion and Deformation of the Heart Due to Respiration. *IEEE Transactions on Medical Imaging*, 21(9):1142–1150, September 2002.

E. H. Meijering, W. J. Niessen, and M. A. Viergever. Retrospective Motion Correction in Digital Subtraction Angiography: A Review. *IEEE Transactions on Medical Imaging*, 18(1):2–21, 1999.

V. Muthurangu, A. Taylor, R. Adriantsimiavona, S. Hegde, M. E. Miquel, R. Tulloh, E. Baker, D. L. G. Hill, and R.. S. Razavi. Novel Method of Quantifying Pulmonary Vascular Resistance by Use of Simultaneous Invasive Pressure Monitoring and Phase-Contrast Magnetic Resonance Flow. *Circulation*, 110(7):826–834, August 2004.

K. Nehrke and P. Bornert. Prospective Correction of Affine Motion for Arbitrary MR Sequences on a Clinical Scanner. *Magnetic Resonance in Medicine*, 54:1130–1138, 2005.

K. Nehrke, P. Bornert, D. Manke, and J. C Bock. Free-breathing Cardiac MR Imaging: Study of Implications of Respiratory Motion - Initial Results. *Radiology*, 220:810–815, 2001.

K. S. Rhode, D. L. G. Hill, P. J. Edwards, J. Hipwell, D. Rueckert, G. Sanchez-Ortiz, S. Hegde, V. Rahunathan, and R. Razavi. Registration and Tracking to Integrate X-Ray and MR Images in an XMR Facility. *IEEE Transactions on Medical Imaging*, 22(11):1369–1378, Nov 2003.

K. S. Rhode, M. Sermesant, D. Brogan, S. Hegde, J. Hipwell, P. Lambiase, E. Rosenthal, C. Bucknall, S. A. Qureshi, J. S. Gill, R. Rezavi, and D. L. G Hill. A System for Real-Time XMR Guided Cardiovascular Intervention. *IEEE Transactions on Medical Imaging*, 24(11):1428–1440, November 2005.

W. Schroeder, K. Martin, B. Lorensen, L. Avila, R. Avila, and C. Law. *The Visualization Toolkit: An Object-Oriented Approach to 3-D Graphics*. Prentice-Hall, 1997.

G. Shechter, C. Ozturk, J. R. Resar, and E. R. McVeigh. Respiratory Motion of the Heart From Free Breathing Coronary Angiograms. *IEEE Transactions on Medical Imaging*, 23(8):1046–1056, August 2004.

G. Shechter, B. Shechter, J. R. Resar, and R. Beyar. Prospective Motion Correction of X-Ray Images for Coronary Interventions. *IEEE Transactions on Medical Imaging*, 24(4):441–450, April 2005.

C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3-D medical image alignment. *Pattern Recognition*, 32(1):71–86, 1999.

H. Timinger, S. Krueger, J. Borgert, and R. Grewer. Motion Compensation for Interventional Navigation on 3D Static Roadmaps Based on an Affine Model and Gating. *Physics in Medicine and Biology*, 49:719–732, 2004.

Y. Wang, S. J. Riederer, and R. L. Ehman. Respiratory Motion of the Heart: Kinematics and the Implications for the Spatial Resolution in Coronary Imaging. *Magnetic Resonance in Medicine*, 33(5):713–719, November 1995.

Vol.	MRI Sequence	Model No.	RMS target registration errors in mm: using full affine registration [using constrained registration]			% of motion recovered
			Image 2 (Mid-Insp.)	Image 3 (Mid-Exp.)	Image 4 (End-Insp.)	
A	High resolution 2-D	1	3.1 [2.7]	1.7 [1.1]	N/A	26.2 [38.4]
		2	3.0 [2.7]	2.1 [1.2]	2.7 [2.6]	47.8 [54.7]
		3	2.5 [2.5]	1.4 [1.0]	N/A	40.4 [43.4]
	Low resolution 3-D	1	4.3 [2.3]	3.8 [1.4]	5.0 [2.5]	12.5 [56.8]
		2	4.5 [2.4]	4.1 [1.5]	5.0 [2.6]	9.2 [55.1]
		3	4.4 [2.6]	3.8 [1.6]	4.6 [2.6]	14.6 [53.9]
B	High resolution 2-D	1	2.3 [2.0]	2.4 [2.3]	2.2 [2.3]	41.4 [44.0]
		2	2.0 [1.7]	2.2 [2.1]	1.9 [2.2]	48.2 [49.6]
		3	2.0 [1.8]	1.9 [1.9]	1.8 [1.8]	51.6 [53.5]
	Low resolution 3-D	1	3.8 [3.0]	2.9 [2.2]	2.9 [2.4]	18.6 [36.2]
		2	4.2 [3.2]	3.3 [2.4]	2.7 [2.3]	12.4 [33.4]
		3	3.9 [2.7]	3.1 [2.0]	N/A	-11.7 [25.1]
C	High resolution 2-D	1	2.4 [2.6]	1.9 [1.5]	N/A	24.5 [26.6]
		2	N/A	2.8 [2.0]	N/A	-19.9 [14.5]
		3	N/A	N/A	N/A	- [-]
	Low resolution 3-D	1	3.1 [3.3]	2.1 [2.4]	3.7 [3.6]	27.5 [25.1]
		2	3.0 [3.3]	1.9 [2.3]	3.7 [4.0]	29.2 [22.1]
		3	3.3 [3.6]	2.3 [2.6]	3.6 [3.7]	25.1 [20.3]
D	High resolution 2-D	1	3.3 [3.6]	3.4 [4.4]	3.9 [4.7]	38.4 [26.9]
		2	4.9 [3.8]	3.6 [4.4]	5.5 [5.6]	33.4 [19.1]
		3	3.2 [3.7]	3.3 [3.6]	4.0 [4.4]	38.5 [32.2]
	Low resolution 3-D	1	4.2 [3.1]	3.0 [2.8]	6.0 [1.8]	20.2 [54.7]
		2	4.9 [3.2]	3.6 [2.5]	5.5 [1.9]	17.6 [54.7]
		3	4.0 [3.0]	3.2 [2.8]	4.1 [1.9]	34.0 [54.8]

Table 2

The target registration errors of respiratory models for all four volunteers. Results were computed using ten anatomical landmarks manually localised in three high resolution volumes acquired at different respiratory positions. The three additional volumes were not used in model formation. Entries are marked 'N/A' when the test high resolution image was out of range of the model.

Vol.	MRI Sequence	RMS leave-one-out errors in mm: using full affine registration [using constrained registration]		
		Model 1	Model 2	Model 3
A	High resolution 2-D	1.2 [1.0]	1.4 [0.9]	1.4 [1.1]
	Low resolution 3-D	2.5 [1.1]	2.5 [0.8]	2.6 [0.8]
B	High resolution 2-D	1.0 [0.8]	0.9 [0.6]	1.0 [1.1]
	Low resolution 3-D	3.3 [0.9]	3.0 [1.1]	3.2 [0.9]
C	High resolution 2-D	1.3 [1.0]	0.9 [0.6]	0.8 [0.4]
	Low resolution 3-D	1.3 [0.9]	1.2 [0.8]	1.0 [0.6]
D	High resolution 2-D	1.3 [1.5]	1.7 [1.6]	1.3 [1.6]
	Low resolution 3-D	2.0 [1.3]	1.6 [1.2]	2.2 [1.6]

Table 3
The residual error of the curve-fitting process for all four volunteers, three independent models for each. Results were computed using a leave-one-out test: for each dynamic acquisition the registration result was compared with the transformation predicted by the model formed by leaving out the dynamic acquisition in question. Errors were computed at 10 clinically relevant anatomical landmarks.

Patient	Estimated 2-D Registration Error: RMS/Max. (mm)		% of motion recovered
	Before Respiratory Motion Correction	After Respiratory Motion Correction	
A	13.3/14.6	2.8/3.3	78.9
B	5.1/7.6	3.9/5.3	23.5
C	7.5/8.2	2.2/2.8	70.7

Table 4
Estimated 2-D registration errors for three clinical cases. Errors were estimated by an observer on cardiac-gated X-ray images acquired at inhale position. Errors are given in mm - for all three cases 1mm corresponded to approximately 3 image pixels.

Patient	Estimated 3-D Registration Error: RMS/Max. (mm)		% of motion recovered
	Before Respiratory Motion Correction	After Respiratory Motion Correction	
B	4.3/5.0	3.0/4.3	30.2
C	4.0/5.4	1.9/3.1	52.5

Table 5
Estimated 3-D registration errors for two clinical cases. Errors were estimated by reconstructing points on a segment of the catheter positioned inside the pulmonary artery from biplane X-ray images. The segment was reconstructed from image pairs at both end-expiration and full inspiration. The end-expiration reconstruction was transformed using the motion model to full inspiration and compared with the actual full inspiration reconstruction. Only the part of the catheter that lay along the roof of the pulmonary artery was reconstructed because in this area it is reasonable to assume that the catheter will remain in the same position relative to the anatomy as it moves with respiration. Images showing the three point sets are shown in Figure 11.