DOI:
10.1109/TMI.2009.2028022

Document Version
Peer reviewed version

An Adaptive and Predictive Respiratory Motion Model for Image-Guided Interventions: Theory and First Clinical Application

A. P. King, K. S. Rhode, R. Razavi and T. Schaeffter

***Abstract*—This paper describes a predictive and adaptive single parameter motion model for updating roadmaps to correct for respiratory motion in image-guided interventions. The model can adapt its motion estimates to respond to changes in breathing pattern, such as deep or fast breathing, which normally would result in a decrease in the accuracy of the motion estimates. The adaptation is made possible by interpolating between the motion estimates of multiple sub-models, each of which describes the motion of the target organ during cycles of different amplitudes. We describe a predictive technique which can predict the amplitude of a breathing cycle before it has finished. The predicted amplitude is used to interpolate between the motion estimates of the sub-models to tune the adaptive model to the current breathing pattern. The proposed technique is validated on affine motion models formed from cardiac magnetic resonance imaging (MRI) datasets acquired from seven volunteers and one patient. The amplitude prediction technique showed errors of 1.9mm-6.5mm. The combined predictive and adaptive technique showed 3-D motion prediction errors of 1.0mm-2.8mm, which represents an improvement in modelling performance of up to 40% over a standard non-adaptive single parameter motion model. We also applied the combined technique in a clinical setting to test the feasibility of using it for respiratory motion correction of roadmaps in image-guided cardiac catheterisations. In this clinical case we show that 2-D registration errors due to respiratory motion are reduced from 7.7mm to 2.8mm using the proposed technique.**

***Index Terms*—motion modelling, image-guided interventions, adaptive, predictive, cardiac, MRI**

I. Introduction

In recent years the use of medical images to guide interventions has been growing in popularity. Such image guidance systems typically make registered preprocedure images available during interventions. The images are static, so the organ's position and shape is also assumed to be static. For this reason, high accuracy was initially only achieved in rigid, bony regions such as the head. Organs in the chest and abdomen, such as the heart, lungs and liver, move significantly during respiration. Therefore the accuracy of guidance information during image-guided interventions on these organs is reduced. To overcome this problem, motion models have been proposed that can predict and correct for breathing motion. However, our previous work has suggested that the accuracy of such models can be reduced during non-standard breathing patterns, such as fast or deep breathing [1]. In this paper we propose a novel predictive and adaptive motion model that will address this limitation.

There are two underlying physiological causes of respiration: movement of the rib cage caused by the rib cage muscles (e.g. intercostals) and contraction of the diaphragm leading to an increase in intrathoracic volume. These two actions cause a reduction in intrapleural pressure and a consequent inhalation of air into the lungs [2][3]. A distinction is often made between *abdominal* and *rib cage* breathing. Although these do not match exactly with the actions of the diaphragm and rib cage muscles it can be a useful simplification to make. The relative contributions of rib cage and abdominal breathing during respiration are highly variable. They can vary greatly from subject to subject [4], and are also dependent on the position of the subject and the rate and depth of breathing being performed. Abdominal breathing is more dominant when the subject is supine, whereas rib cage breathing typically becomes relatively more important when the subject is upright, when tidal volume increases (e.g. the subject breathes deeply) or during fast breathing [5]. Generally, it is clear that changes in breathing rate and volume can lead to changes in the type of respiration (i.e. abdominal or rib cage) that occurs.

Many respiratory motion models in the literature are based on a single respiratory parameter, or surrogate. Previous examples of surrogates used in such models include the 1-D superio-inferior translation of the diaphragm [6][7][8], a respiratory bellows and the motion of surface points, such as the chest or abdomen. Our recent work has suggested that such single parameter models can have lower accuracy in the presence of changes in breathing pattern, such as deep or fast breathing [1]. We postulate that this lower accuracy is the result of a change in the relative contributions of the underlying physiological causes of respiration (i.e. abdominal or rib-cage breathing), and hence in the relationship between the single respiratory parameter and the motion. The inability to adapt to such changes is a fundamental limitation of existing single parameter models. One solution to this problem is to introduce extra parameters into the model. For example, physically realistic models of lung motion have been proposed in [9][10], in which the motion is determined by

A. P. King, K. S. Rhode, R. Razavi and T. Schaeffter are with the Division of Imaging Sciences, King's College London, U.K. This work was co-funded by the Technology Strategy Board's Collaborative Research and Development programme, following an open competition (grant number 17352), Philips Healthcare, Best, the Netherlands, and EPSRC grant EP/D061474/1. 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.

setting multiple boundary conditions as surrogates. Although this produces realistic deformations, it is not always easy to acquire such parameters during an intervention. In [11] a model of cardiac respiratory motion was built from magnetic resonance imaging (MRI) data based on three surrogate values. The surrogates used were MRI navigators positioned on the dome of the right hemidiaphragm, the chest wall and the right margin of the heart. This resulted in improved motion predictions. The application in [11] was prospective motion correction for coronary MR angiography, so it was relatively straightforward to acquire the extra parameter values when forming and applying the model. However, when using motion models during image-guided interventions it is often difficult to acquire these extra parameters to apply the model. In [12] the use of extra parameters derived from the first parameter (amplitude and first derivative) was proposed, but no method to estimate the amplitude in real-time was described.

Our adaptive model is based on physiological research [5] suggesting that fast and deep breathing are different in terms of their underlying physical causes. These breathing types tend to occur frequently during interventions in which the patient is sedated. For example, when the patient is nervous they may breathe quickly and when they are in discomfort or pain they may take deep breaths. In our proposed technique we do not explicitly decouple the underlying physical causes of breathing. Rather, we implicitly decouple them by noting that fast breathing cycles tend to have a smaller amplitude, whereas deep breathing cycles tend to have a larger amplitude. Therefore, for each subject, we form multiple motion models that represent the motion of the target organ during cycles of different amplitudes. The final motion estimate is determined by interpolating between the motion estimates of these models based on the amplitude of the current cycle. We also present a technique that can be used to predict the respiratory amplitude in advance of the completion of each cycle. We demonstrate that the combination of the predictive technique and the adaptive model results in a more accurate respiratory motion correction technique that is suitable for use in a range of image-guided interventions.

Preliminary results of this work have previously been described in [13]. Here we present some refinements to the adaptive model formation process, more extensive validation and analysis, and describe for the first time the application of the complete predictive and adaptive technique in a clinical setting. Although our technique is general and has a wide range of potential applications, in this paper we focus on building cardiac respiratory motion models from MRI data, using a pencil beam navigator applied on the diaphragm as the respiratory surrogate (this signal has been previously shown to have a good correlation with heart motion [6][7][8]). We demonstrate application of the model for image-guided cardiac catheterisations on a clinical dataset.

II. Method and Materials

A. Method

Figure 1 illustrates the processes involved in forming and applying the adaptive model. The following description refers to the labels (A)-(G) in the figure.

(A) Preprocedure imaging data is acquired together with a single respiratory surrogate signal. During imaging the subject performs different breathing manoeuvres such as normal, quick and deep breathing to vary the amplitudes of the breathing cycles.
(B) The surrogate signal is processed to classify each image into one of a number of categories, based on the amplitude of the breathing cycle during which it was acquired.
(C) Separate sub-models are formed for each category based on the image data, each representing the motion of the organ during cycles of different amplitude ranges. The sub-models describe the motion of the target organ as a function of the respiratory surrogate value.
(D) These sub-models comprise the adaptive model.
(E) During the image-guided procedure, the respiratory surrogate signal is acquired in order to apply the model.
(F) The adaptive model can produce an estimate of the motion of the target organ given knowledge of the current surrogate value and the amplitude of the current breathing cycle. Obviously the true amplitude is not known before the cycle has finished. Therefore we use an amplitude estimation technique to predict the current cycle's amplitude before it has completed.
(G) Each sub-model produces a motion estimate based on the current surrogate value. The final motion estimate of the adaptive model is made by interpolating between the motion estimates of the sub-models based on the value of the estimated amplitude.

In the following sections we describe the different stages involved in forming and applying an adaptive affine motion model of the heart from MRI data. To apply this general technique to a different application only the nature of the imaging data, the basic sub-models and the surrogate value need be changed (i.e. labels A, C and E in Figure 1). All other details of the adaptive model and the predictive technique would remain the same.

1) Preprocedure Imaging Data: Two cardiac gated MRI sequences are required to form the basic sub-models: a high resolution end-expiration scan to acquire the anatomy and a dynamic scan to acquire the respiratory motion.

The high resolution scan is a respiratory gated free-breathing scan covering the four chambers and major vessels of the heart (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, scan time approximately 5 minutes). This scan is also used to form the roadmap for the intervention.

The dynamic scan is a free-breathing scan obtaining a number of near real-time free-breathing acquisitions that cover a range of respiratory positions (3-D TFEPI, cardiac triggered and gated at late diastole, 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).

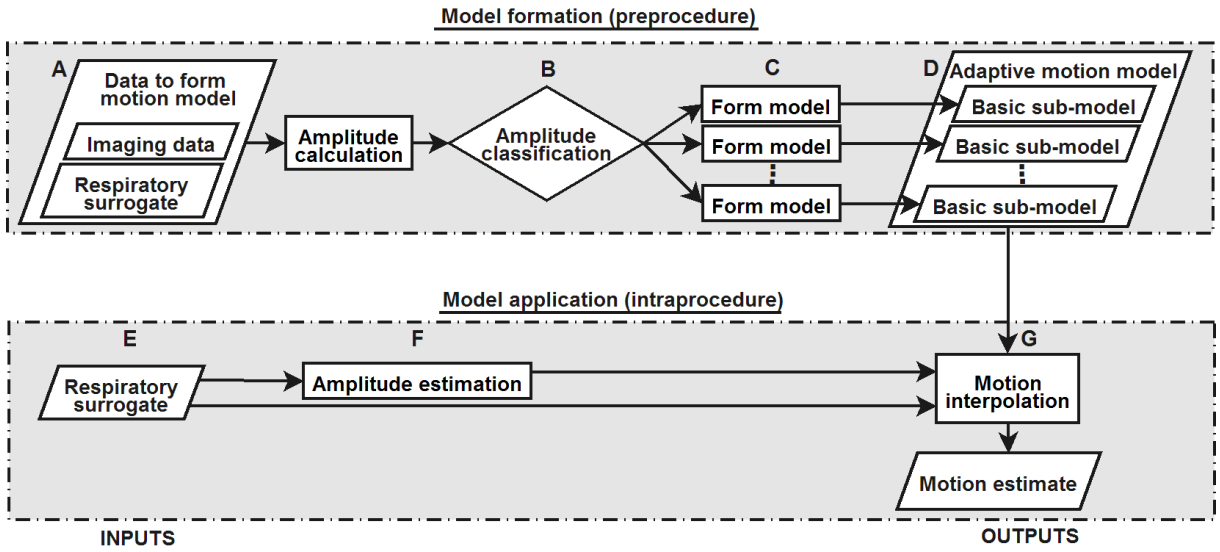

Fig. 1. An overview of the formation and application of the adaptive motion model: (A) Preprocedure imaging data is acquired together with a single respiratory surrogate signal; (B) The surrogate signal is processed to classify each image into one of a number of amplitude categories; (C) Separate sub-models are formed for each category based on the image data; (D) These sub-models comprise the adaptive model; (E) During the image-guided procedure, the respiratory surrogate signal is acquired in order to apply the model; (F) An amplitude estimation technique is used to predict the current cycle's amplitude before it has completed; (G) The final motion estimate is made by interpolating between the estimates of the sub-models based on the value of the estimated amplitude.

This sequence is similar to that used previously to form models for motion-corrected MRI image acquisition [6]. For the data used in this paper we used 120 dynamics to form each adaptive model, resulting in a typical scan time of 2 minutes. During the dynamic scan a pencil beam navigator is applied on the dome of the right hemi-diaphragm to estimate its 1-D superior-inferior translation immediately before and after each dynamic acquisition. The average of these lead and trail navigators is used as the respiratory surrogate to form the models. Since the affine motion model we use for the basic sub-model models the inspiration and expiration phases separately each acquisition is classified as either inspiration or expiration by comparing its navigator value with that of its predecessor.

Respiratory navigators are widely used for respiratory gating of MRI scans, and are commonly positioned on the dome of the right hemi-diaphragm because of the high contrast between the lung and the liver in this area. Such navigators have also previously been used to form respiratory motion models in [6][14][8].

To ensure that the resulting adaptive model captures the full range of breathing motions, a breathing protocol is used each time the dynamic MRI scan is acquired. In our volunteer experiments, we split the 120 dynamics into 3 sets of 40, during which the following breathing instructions were given to the subjects:

- first 40 dynamics: no special breathing instructions (i.e. normal breathing)
- next 40 dynamics: instructed to breathe quickly (i.e. small amplitude)
- final 40 dynamics: instructed to take deep breaths (i.e. large amplitude)

The intention of using this protocol was to acquire data for enough different cycles to form three sub-models representing the three different breathing patterns. For the patient dataset, we used only normal and deep breathing instructions (for 60 dynamics each), and only two sub-models were formed.

2) Amplitude Calculation and Classification: The sub-models are formed by classifying each dynamic into one of N categories according to the respiratory amplitude of the breathing cycle it is a part of. We calculate amplitude for each *half-cycle*, i.e. we compute the differences between the respiratory surrogate values of adjacent extreme (either end-expiration or end-inspiration) positions. Extreme positions are identified as those where the sign of the first time derivative of the surrogate signal changes. For model formation, all dynamics within the half-cycle are assigned the same amplitude value. We compute amplitude for each half-cycle rather each full cycle to allow the model to adapt quicker to changes in amplitude.

Classification of each half-cycle is performed by comparing it's amplitude with N predetermined per-category mean amplitude values. Each half-cycle is assigned to the category whose mean amplitude value is closest to it's amplitude. The per-category mean amplitude values are determined from the dynamic scan surrogate values using a clustering technique: the precomputed amplitudes of all half-cycles in the dynamic scan are used as input to the k-means clustering algorithm [15], which computes N mean amplitude values.

In principle, any number of categories can be used. However, for the volunteer experiments presented in this paper, we used three categories, approximately representing fast breaths, normal breaths and deep breaths. For the patient dataset we used two categories. Half-cycles with an amplitude less than 5% of the maximum were rejected as they were probably due to noise in the surrogate signal.

As an illustration, Figure 2 shows the precomputed amplitudes for one volunteer, together with the results of the

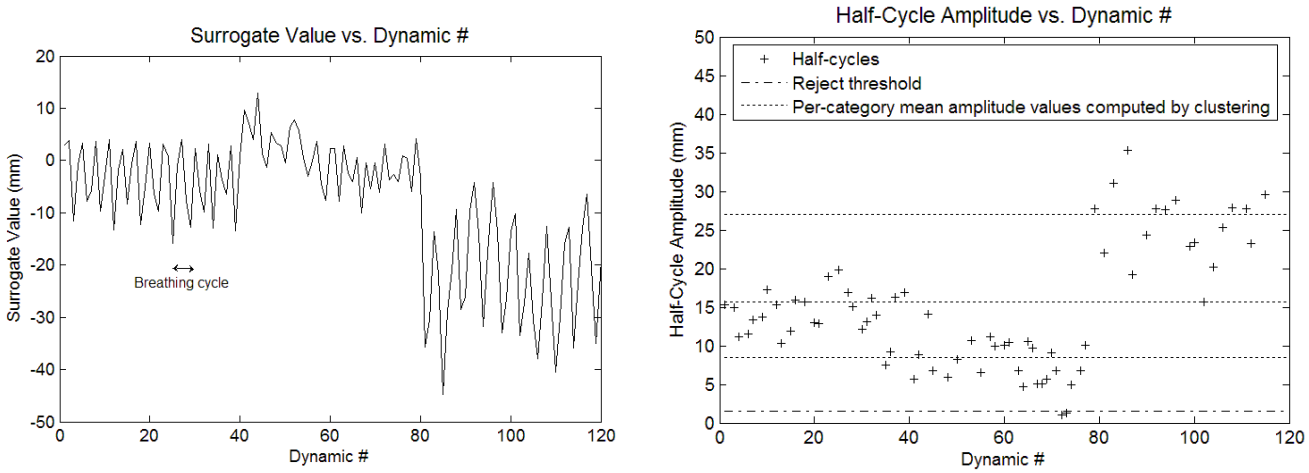

(a) (b)

Fig. 2. (a) A surrogate signal from the dynamic scan for a volunteer, showing the three different breathing patterns: normal breathing (dynamics 1-40), fast breathing (dynamics 41-80) and deep breathing (dynamics 81-120); (b) Clustering the amplitudes of half-cycles. The '+' symbols represent the computed amplitudes of all half-cycles in the dynamic scan, plotted against the dynamic number. The dotted lines represent the three mean amplitude values computed using the k-means clustering algorithm. These mean amplitudes were used to classify each half-cycle into one of three categories for forming the adaptive model. The dash-dotted line represents the rejection threshold for very small half-cycles. Two half-cycles around dynamic 70 were rejected because their amplitudes were less than this threshold.

clustering.

3) Basic Respiratory Model of the Heart: Separate basic sub-models are formed for each category based on the classified dynamic scans and associated navigator values. In the basic respiratory models, the motion and deformation of the heart are represented by a set of second-order polynomial functions of navigator value. The motion parameters used to estimate the functions are found by performing an affine intensity-based registration between each dynamic acquisition and the high resolution end-expiration image. The registration maximises the normalised mutual information between the two images in the area of overlap between them. Each of the twelve affine motion parameters is modelled as two separate polynomial functions: one for inspiration and one for expiration. The classification of dynamic acquisitions into inspiration and expiration is used to determine which dynamics to use for these two functions (the dynamics with the global minimum and maximum navigator values are used for both inspiration and expiration). Modelling inspiration and expiration separately enables hysteretic effects [16][14] to be captured by the model. Once the model has been formed, it can produce an estimate of an affine transformation given a navigator value and a breathing direction (i.e. inspiration or expiration). The transformation represents the motion and deformation of the heart from end-expiration to the current respiratory position. This model can predict cardiac respiratory motion to within 2-4mm [8].

4) Adaptive Model: Finally, an adaptive model, M^a, is formed. This consists of a number of basic sub-models, $M_i, i = 1 \ldots N$.

Our technique for applying the adaptive model consists of two steps. First, we estimate the amplitude of the current half-cycle before it has finished. Next, we use this estimate to interpolate between the motion estimates produced by each of the sub-models. In the next section we describe the amplitude prediction technique, which is expressed in a probabilistic framework. We describe the motion interpolation in the following section.

5) Predicting Amplitude:

Posterior Probability: The single respiratory surrogate is considered to be a continuous signal sampled at a number of discrete time points, k. At each point we make an estimate of what the next extreme position of the surrogate signal will be. We define $z_k = (v_k, \bigtriangledown v_k)$ to be a tuple consisting of the surrogate value and its first time derivative at point k. We denote the actual (unknown) next extreme position of the surrogate by x_k and its current estimate by $\hat{x}_k$ (see Figure 3). The technique for estimating the next extreme position is based on recursive Bayesian estimation [17]. We attempt to find the estimate $\hat{x}_k$ of the next extreme position that maximises the posterior probability of x_k given all previous values of z_k. The posterior probability of x_k is defined as

$$p(x_k|Z_k) = \frac{p(z_k|x_k) \cdot p(x_k|Z_{k-1})}{p(z_k|Z_{k-1})} \quad (1)$$

where Z_k is the set of all surrogate values/derivatives up to and including position k (i.e. $z_1 \ldots z_k$). Recursive Bayesian estimation is a technique for estimating a sequence of unknown states from noisy measurements. In our case, the unknown states, x_k, represent the next extreme position within a half-cycle, and the measurements are the navigator values and their derivatives at each sample point.

Prior Probability: In (1), $p(x_k|Z_{k-1})$ is the prior probability of the surrogate value of the next extreme position. The form of the prior probability expresses knowledge about likely respiratory amplitudes before the half-cycle begins. Here, we assume that the amplitude will remain approximately constant over a short period of time. The prior probability is modelled as a weighted sum of the last three extreme positions of the same type (i.e. end-expiration or end-inspiration). The form of

the prior probability is a Gaussian distribution centred on the value of this weighted sum,

$$p(x_k|Z_{k-1}) = exp(-\frac{\|x_k - x_{prior}\|^2}{2\sigma_p^2}) \quad (2)$$

where

$$x_{prior} = \frac{\sum_{n=1}^{3} \frac{1}{n} \cdot x_{prev,n}}{\sum_{n=1}^{3} \frac{1}{n}}$$

is the weighted sum of the previous three extreme positions of the same type, $x_{prev,n}$. The standard deviation of the Gaussian distribution in (2), σ_p, is an empirically determined constant (we used a value of 10). In principle more than three previous extreme values could be used but for the experiments presented in this paper we used only the last three. If less than three end-cycle positions were available then only those known were used in the weighted sum.

Likelihood: $p(z_k|x_k)$ is the likelihood. We can view this as a model of how measurements are made from the unknown states. In our case, it is a model of how a half-cycle of a given amplitude can produce measurements of the navigator value and its first derivative. For example, if the navigator's first derivative is low then it is likely that the next extreme position is either very far (i.e. the half-cycle has only just started) or very near (it has almost finished). This is illustrated in Figure 3.

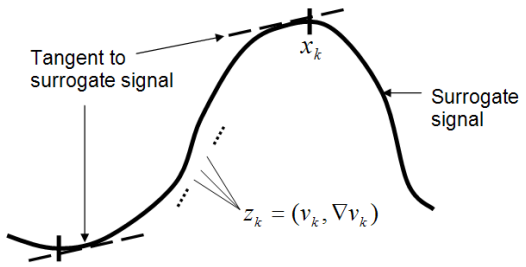

Fig. 3. A sequence of sample points k within a half-cycle of a respiratory surrogate signal. The half-cycle ranges between the two true extreme positions of the surrogate signal (end-expiration/end-inspiration) marked by the vertical lines. At each sample point the surrogate value, v_k, and its first derivative, ∇v_k, are recorded, and an estimate, $\hat{x}_k$, is made of the surrogate value of the next extreme position, x_k. Note that low derivative values such as those illustrated by the two tangents to the curve can indicate that the next extreme position is either near or far from the current point.

The likelihood estimate is based on previously acquired surrogate data analysed to determine the relationship between the current surrogate first derivative value and the distance to the next extreme position. We model this relationship as a sine wave of a known amplitude and wavelength. The sine wave allows us to predict what the derivative might be given a possible extreme surrogate value. The form of the likelihood for a possible extreme value x_k is defined by

$$p(z_k|x_k) = exp(-\frac{\| \bigtriangledown v_k - \bigtriangledown \hat{v}_k\|^2}{2\sigma_l^2}) \quad (3)$$

where $\bigtriangledown v_k$ is the actual derivative at the current point, and $\bigtriangledown \hat{v}_k = A.sin(\frac{(x_k - v_k)\pi}{\lambda})$ is the estimated derivative given that the next extreme position will be x_k. A and λ are the amplitude and wavelength of the sine wave. The values of A and λ represent knowledge of what derivative values and breathing amplitudes respectively are likely in the surrogate signal. These can be chosen empirically for any given type of surrogate signal. For all of the experiments in this paper, the value of λ was assigned to be the maximum of 15mm and the previous breathing cycle's amplitude. The value of A was dependent on the surrogate type and details will be provided in Sections II-B and III-C. The standard deviation of the Gaussian function in (3), σ_l, indicates how much tolerance there is in the derivative estimates. This was empirically determined, and we used a value of 5 for all experiments. Note that $p(z_k|x_k)$ will have two peaks, one at each of the two intersections of the current derivative value with the sine wave. These correspond to the near and far possibilities of the next extreme position illustrated in Figure 3. A mask was used to exclude any peaks beyond these two, and also to exclude the possibility that the next peak is before the current surrogate value.

Finally, the divisor in (1), $p(z_k|Z_{k-1})$, is the prior probability of the measurements. We assumed this to be uniform.

Based on this formulation, we make an estimate of the next extreme position at time point k,

$$\hat{x}_k = \underset{x_k}{\operatorname{argmax}}(p(x_k|Z_k)) \quad (4)$$

Because of the limited range of possible values for x_k, a simple linear search is employed to find $\hat{x}_k$.

6) Model Interpolation: Each of the N sub-models in the adaptive model will normally give a different estimate for each of the motion parameters for any given surrogate value. The adaptive model produces an estimate for each parameter by interpolating between the motion estimates of its sub-models based on the current amplitude estimate. We denote by $T_j(M_i, v_k)$ the value of the j^{th} motion parameter for the i^{th} sub-model using surrogate value v_k. We define the adaptive estimate of each motion parameter j by a model M^a based on current surrogate value v_k and current amplitude estimate $\hat{x}_k$ as

$$\hat{T}_j(M^a, v_k, \hat{x}_k) = \phi(\tilde{T}_j, \hat{x}_k), j = 1 \ldots 12 \quad (5)$$

where $\tilde{T}_j = (T_j(M_1, v_k), \ldots T_j(M_N, v_k))$, a vector containing the estimates for the j^{th} parameter for each of the N sub-models. The function ϕ is an interpolant. For the experiments presented in this paper, ϕ was a linear interpolation function. More complex interpolants could also be used.

B. Experiments

We validate each component of the predictive/adaptive model separately using data acquired from seven volunteers and one patient. First, we test the effectiveness of the amplitude prediction technique alone using the MRI navigator values as the respiratory surrogate. We also test the frequency sensitivity of the amplitude prediction technique using diaphragm tracking data acquired from X-ray images on the patient dataset. Next, we validate the performance of the combined predictive and adaptive technique using MRI image data and surrogate values. Finally, we illustrate how the overall system (i.e. formation of an adaptive model and application of the model using the amplitude prediction technique) could

be applied in a clinical setting on the patient dataset. Our clinical application is motion-correction of roadmaps in image-guided cardiac catheterisations [18][8] using a hybrid X-ray/MR (XMR) catheter laboratory.

1) Materials: The predictive technique and adaptive model volunteer experiments were performed using data acquired from a 1.5T Philips Achieva MRI scanner. The clinical data was acquired in an XMR catheter laboratory featuring a 1.5T Philips Achieva MRI scanner and a Philips BV Pulsera mobile cardiac X-ray system. The patient could be easily moved between the two devices on a sliding bed.

Seven volunteer datasets and one patient dataset were acquired. Volunteers A-G consisted of six males and one female who were aged between 21 and 33 years. Patient A was male, aged 55 years, and underwent a pulmonary vein ablation procedure to treat atrial fibrillation. During preprocedure imaging, the volunteers and patient did not undergo any treatment and were only asked to follow the breathing protocol. Prior to the procedure the patient was sedated and remained so throughout the procedure. All subjects gave informed consent.

For the patient the breathing protocol mentioned in Section II-A3 was modified slightly so that only two breathing patterns were performed: normal and deep breathing. For all volunteers three patterns were performed: normal, fast and deep breathing.

2) Amplitude Prediction: To validate the predictive technique the navigator values from the dynamic MRI scans were processed using the proposed technique and end-cycle values were predicted for each sample point. The sequence of navigator values was subsequently processed to determine the true end-cycle values. These were used as the gold standard. The results of the predictive technique were compared to this gold standard. To show the benefit of using the Bayesian approach we tested the use of only the prior probability and also the full Bayesian technique in this way.

The value of the constant A in the likelihood calculation in (3) was set to 20 for all experiments involving the MRI navigator as the surrogate signal.

Obviously the performance of the amplitude prediction technique will be dependent on the frequency of the breathing cycles in the surrogate signal (this will affect the accuracy of estimating the derivative of the surrogate). We can simulate a change in the frequency of a signal by altering the sampling rate whilst keeping the time step between adjacent samples fixed. Therefore to test the sensitivity of the technique to changes in frequency we ran the amplitude prediction technique on a real diaphragm tracking sequence from a clinical case for patient A (a subset of the complete sequence for the procedure). The diaphragm was tracked from fluoroscopic X-ray images (see Section II-B4 for details). The algorithm was run on the original signal, and also versions of the signal that had been subsampled by different factors. The accuracy results for the different signals were compared to the result obtained by using the prior probability alone (i.e. that the current half cycle's amplitude will be a weighted sum of the previous three half cycle's, as described in Section II-A5).

3) Adaptive Model: To validate the formation of the adaptive model we used a leave-one-out test. Recall that normally the adaptive model would be formed from the results of registering 120 dynamic acquisitions to a high resolution volume. In the leave-one-out test, we formed 120 separate adaptive models by leaving out each of the registration results in turn. Therefore each of these 120 adaptive models was formed from 119 registration results. Each of these models was used to predict the registration result for its left-out dynamic based on its MRI navigator value. The predicted registrations were compared with the actual registration results. Errors were computed at ten clinically relevant anatomical landmarks. The landmarks were the centres of the four pulmonary vein ostia (left upper, left lower, right upper and right lower pulmonary veins), the centres of the junctions of the right atrium with the inferior and superior vena cava, the tricuspic valve and three points on the interventricular septum.

Using this validation, we tested three different techniques for forming and applying the motion model. For each left-out dynamic, the models for the three techniques were formed from exactly the same data. Only the technique for forming and applying the model differed. The three techniques were:

1) A single basic model formed as described in Section II-A3 from all 120 dynamic scans. The motion prediction was made using the MRI navigator alone.
2) An adaptive model with perfect knowledge of the cycle amplitudes. The adaptive model was formed based on 'true' amplitude values computed and classified by postprocessing the MRI navigator signal. The motion prediction was made by the adaptive model using the same 'true' amplitude values together with the MRI navigator value.
3) An adaptive model using the amplitude prediction technique. The adaptive model was formed using the 'true' amplitudes as above, but the motion predictions were made using amplitude values computed using the amplitude prediction technique.

Note that technique 3 is the clinically realistic case. It would normally not be possible to use the 'true' amplitude value when applying the model as technique 2 does. Therefore the real test of the proposed predictive and adaptive method is to compare technique 3 with technique 1. We introduce technique 2 in order to be able to test the impact of imperfect knowledge of the amplitudes on the overall accuracy of the adaptive model. This can be done by comparing techniques 2 and 3. Finally, by comparing techniques 1 and 2 we test the effectiveness of forming an adaptive model rather than a basic model. By giving the model perfect knowledge of the amplitudes when making motion predictions we exclude the impact of imperfect knowledge of the amplitudes from this validation.

4) Clinical Application: We illustrate how the overall system (i.e. formation and application of the predictive and adaptive model) can be applied in a clinical setting by using it for respiratory motion correction of roadmaps in an image-guided cardiac catheterisation in an XMR catheter laboratory [18][8]. This image guidance technology enables a roadmap derived from MRI data to be overlaid onto real-time X-ray fluoroscopy images. To perform respiratory motion correction of the roadmap a motion model is formed from

MRI images and applied during the procedure by tracking the superio-inferior motion of the diaphragm in the X-ray images [8]. Diaphragm tracking is performed by defining a rectangular region of interest in the X-ray images, and finding the 1-D translation that minimises the mean sum of squared intensity differences within this region between each X-ray image and a reference image. Since the motion model is based on gated MRI data acquired during the late diastolic rest period of the heart, the X-ray images were cardiac gated to late diastole. Therefore no cardiac motion correction is performed.

We formed an adaptive motion model for the patient from MRI data acquired before the procedure. Referring back to Figure 1, label A corresponds to the MRI data and labels B-D represent the formation of the adaptive model from this data. During the procedure X-ray image sequences were acquired showing the motion of the patients' diaphragm over several breathing cycles. The diaphragm motion was tracked as described above (label E in Figure 1). This tracking information was used as the input surrogate signal for the amplitude prediction technique (label F). The surrogate values and the estimated amplitudes were used as inputs to the adaptive model (label G). The transformations produced by the adaptive model were applied to the roadmap. When applying the predictive technique on this X-ray derived surrogate signal the value of the constant A in the likelihood calculation in Equation (3) was set to 10. Recall that this value reflects knowledge of likely values of the surrogate signal's first derivative. Therefore it's value will be dependent on the sampling rate of the surrogate signal.

Accuracy assessment was performed by overlaying a roadmap showing the coronary sinus onto X-ray images that showed a coronary sinus catheter positioned inside this vessel. Since the catheter shows up well in X-ray images the error between the overlay and the underlying anatomy can be assessed. Errors were assessed before and after motion correction using the adaptive model.

III. Results

A. Amplitude Prediction

Table I shows the mean prediction errors for the seven volunteers and the patient dataset. Every subject shows lower errors for the Bayesian approach compared to using the prior only approach. This shows the benefit of adapting amplitude predictions based on the current surrogate derivative value. When the amplitude changes by a large amount from one half-cycle to the next the prior only amplitude prediction has significant errors whereas the Bayesian technique is able to detect the change in derivative value and adapt its amplitude estimate accordingly. Mean errors using the Bayesian approach vary from 1.9mm to 6.5mm. The errors for volunteer D show a high standard deviation. This was due to a sudden change from quick breathing to very deep breathing that caused high errors in a small part of the sequence. However, it is clear that the Bayesian approach was able to adapt better to this change than the prior only approach.

Figure 4b shows the results of the frequency response experiment. We can see that on the original signal the Bayesian technique outperforms the prior only approach, but it degrades as the subsample factor reaches 3 or higher. We can see from Figure 4a that this is when the sampling rate approaches the number of samples in an average half-cycle in the original signal. This emphasises the dependence of the amplitude prediction technique on acquiring a surrogate signal with a high enough sampling rate (there should be several samples per half-cycle) and a good enough estimate of the surrogate derivative. Note that the Bayesian technique results never degrade completely because they include the prior probability which constrains the result even when the likelihood gives a poor prediction.

B. Adaptive Model

The leave-one-out test results for the adaptive models formed for the seven volunteers and the patient dataset are shown in Table II. Recall that the 120 dynamics were divided into three sections, during which the subject breathed normally, quickly and deeply. We present the mean overall error and also the error for the section that showed the maximum improvement of technique 3 over technique 1. For the patient dataset only two breathing types were performed (normal and deep) so the 120 dynamics were split into two sections of 60 dynamics each. Comparing the first and second techniques in Table II shows that the adaptive model has the potential to significantly improve the accuracy of motion models, if it has knowledge of the true amplitude of each half-cycle. Comparing the second and third techniques shows that the amplitude prediction technique is performing well, as the errors are very similar for the two approaches. In all eight cases the predictive-adaptive technique gave statistically significant improvements over the single model ($p < 0.01$ in a one-tailed, paired student's t-test), with a maximum improvement of 40%. The maximum improvements were typically found during the non-standard breathing patterns (i.e. quick and deep) suggesting that the single basic model does not adequately capture the organ motion during these patterns, whereas the predictive and adaptive technique does.

A sample parameter plot from the adaptive model for volunteer D is shown in Figure 5a. This shows the relationship between the anterior-posterior translation and the surrogate value for the three sub-models. It is clear that there are significant differences between the motion patterns for cycles of different amplitudes. In particular the sub-model for cycles of large amplitude has a strong hysteresis effect (i.e. there is a large difference between the inspiration and expiration curves). The other two sub-models approximately follow the inspiration curve for the large cycle sub-model. Using a normal single parameter motion model could lead to errors of up to 4mm in this parameter alone. Figure 5b shows a plot of the superio-inferior translation from the adaptive model for volunteer F. In this case little hysteresis is present in any of the sub-models. However, there is a significant difference in the slopes of the curves, this time between the sub-model for cycles of small amplitude and the other two sub-models. Errors of up to 3mm could result from using the wrong sub-model. In both cases the predictive and adaptive technique would have much improved

Subject	Overall breathing depth (mm)	Mean +/- 1 s.d. prediction errors (mm)		% improvement of Bayesian over prior only
		Prior only	Bayesian	
Vol. A	42.5	3.4 +/- 3.0	2.6 +/- 3.1	23.5
Vol. B	54.7	4.2 +/- 4.3	3.0 +/- 3.8	28.6
Vol. C	40.4	3.8 +/- 5.1	2.8 +/- 4.4	26.3
Vol. D	60.7	10.2 +/- 17.8	6.5 +/- 13.4	36.3
Vol. E	54.2	2.4 +/- 3.3	1.9 +/- 2.1	20.8
Vol. F	57.6	3.5 +/- 3.5	2.6 +/- 3.3	25.7
Vol. G	70.4	5.7 +/- 7.2	4.0 +/- 5.9	29.8
Pat. A	19.4	2.4 +/- 3.0	1.7 +/- 2.5	29.2

TABLE I

ERRORS IN COMPUTING THE NEXT END-CYCLE POSITION USING THE PROPOSED AMPLITUDE PREDICTION TECHNIQUE. THE SURROGATE DATA PROCESSED USING THE TECHNIQUE WAS A SEQUENCE OF MRI NAVIGATOR VALUES. THE PREDICTIONS WERE COMPARED WITH THE TRUE END-CYCLE POSITIONS COMPUTED BY POSTPROCESSING. ERRORS ARE SHOWN USING ONLY THE PRIOR PROBABILITY AND USING THE BAYESIAN APPROACH. THE MAXIMUM BREATHING DEPTH IS ALSO SHOWN FOR REFERENCE (I.E. THE DIFFERENCE BETWEEN THE MAXIMUM NAVIGATOR AND MINIMUM NAVIGATOR VALUES)

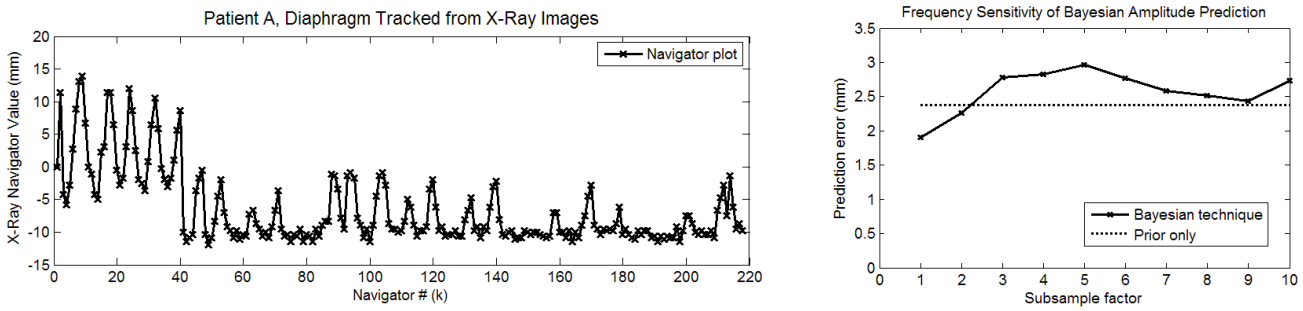

(a) (b)

Fig. 4. Analysing the frequency response of the Bayesian amplitude prediction technique. We simulate different signal frequencies by subsampling the signal whilst keeping the time step between adjacent samples fixed. (a) A diaphragm tracking sequence acquired from a fluroscopic X-ray sequence for patient A. (b) The amplitude prediction errors for the original signal (subsample factor = 1) and for subsampled versions of the signal. The dotted line shows the error for using the prior probability only. The Bayesian technique degrades (i.e. performs worse than the prior only technique) at a subsample factor of 3 or greater in this case. It can be seen from (a) that this is when the subsample factor approaches the number of samples in a typical half-cycle of the original signal.

performance, as with a reasonable amplitude estimate it would apply a motion estimate from the appropriate sub-model(s).

Table III shows a quantification of the variation of the parameter values for the different sub-models. Only three of the parameter values are shown (the anterio-posterior translation, the superio-inferior translation and the medio-lateral axis rotation) because we have previously found them to be the dominant modes of motion in cardiac respiratory motion [8]. Because some parameters commonly have an almost linear relationship with the surrogate value (e.g. the superio-inferior translation), we present the mean slope of the parameter curves. To quantify the hysteresis effect we present the maximum difference between the inspiration and expiration curves at any point in the breathing cycle. Both the large hysteresis effect in the third sub-model for volunteer D (illustrated in Figure 5a) and the differing slopes in the superio-inferior translation for volunteer F (illustrated in Figure 5b) can be seen in this table.

C. Clinical Application

The validation results for the predictive technique and the adaptive model for the clinical case (patient A) were presented in Tables I and II. The adaptive model showed a maximum improvement of 7.7% over the normal motion model.

During the catheterisation we applied the model by tracking the diaphragm in a sequence of X-ray images. We mentioned in Section II-B that the value of the constant A in the likelihood calculation was set to 10 when applying the model in the clinical application. Figure 6 illustrates how this number was arrived at. The data points were acquired from sample X-ray surrogate data processed to determine the true relationship between the surrogate derivative and the distance to the next extreme position. The sample X-ray sequence used was from a different clinical case and was not used in the validation.

Figure 7 illustrates the use of the predictive and adaptive technique at a single sample point during the X-ray sequence for patient A. The first derivative of the diaphragm tracking data (in the top figure) at the current sample point is low, so the likelihood (the dashed line in the bottom figure) has two peaks representing possible next extreme positions (the near and far possibilities) at approximately -15mm and 35mm. The prior probability (the dotted line in the bottom figure) is a Gaussian distribution centred on the expected next extreme position, i.e. that it will be the same as the previous one, -16.5mm. The posterior probability (the solid line in the bottom figure) is the product of the prior and likelihood. The current peak posterior probability indicates a predicted next extreme position of -14.5mm, as indicated by the dotted line

Subject		Mean +/- 1 s.d. leave-one-out test errors (mm)			Improvement in mean error of technique 3 over technique 1
		1. Single basic model	2. Adaptive model, perfect knowledge of amplitudes	3. Adaptive model, amplitude prediction	
Vol. A:	Mean	2.10 +/- 0.98	1.89 +/- 0.98	1.85 +/- 0.95	11.9%
	Max.	2.14 +/- 1.04	1.85 +/- 0.97	1.75 +/- 0.84	18.2%
Vol. B:	Mean	1.88 +/- 1.11	1.81 +/- 1.15	1.81+/- 1.15	3.9%
	Max.	1.98 +/- 1.20	1.61 +/- 1.07	1.61 +/- 1.07	13.6%
Vol. C:	Mean	1.22 +/- 0.62	1.05 +/- 0.56	1.04 +/- 0.56	14.8%
	Max.	1.50 +/- 0.67	1.21 +/- 0.65	1.21 +/- 0.66	19.3%
Vol. D:	Mean	1.41 +/- 0.84	1.21 +/- 0.78	1.18 +/- 0.73	16.3%
	Max.	1.46 +/- 0.88	1.15 +/- 0.58	1.13 +/- 0.55	22.6%
Vol. E:	Mean	1.80 +/- 0.98	1.61 +/- 0.84	1.59 +/- 0.84	11.7%
	Max.	1.93 +/- 1.10	1.56 +/- 0.74	1.58 +/- 0.74	18.1%
Vol. F:	Mean	1.81 +/- 0.87	1.45 +/- 0.76	1.37 +/- 0.72	24.3%
	Max.	1.75 +/- 0.70	1.17 +/- 0.56	1.05 +/- 0.54	40.0%
Vol. G:	Mean	2.74 +/- 1.56	2.36 +/- 1.36	2.43 +/- 1.41	11.3%
	Max.	3.42 +/- 1.70	2.79 +/- 1.41	2.79 +/- 1.39	18.4%
Pat. A*:	Mean	1.38 +/- 0.70	1.29 +/- 0.69	1.29 +/- 0.69	6.5%
	Max.	1.29 +/- 0.59	1.18 +/- 0.55	1.19 +/- 0.55	7.7%

TABLE II

LEAVE-ONE-OUT TEST RESULTS FOR THREE DIFFERENT MODEL FORMATION TECHNIQUES. THE PERCENTAGE IMPROVEMENT IS SHOWN FOR THE THIRD TECHNIQUE OVER THE FIRST. THE 120 DYNAMICS OF THE VOLUNTEER DATASETS WERE DIVIDED INTO THREE SECTIONS, DURING WHICH THE SUBJECT BREATHED NORMALLY, QUICKLY AND DEEPLY. THE PERCENTAGE IMPROVEMENT WAS COMPUTED OVER THE WHOLE DATASET AND ALSO SEPARATELY FOR EACH SECTION. THE MEAN OVERALL IMPROVEMENT AND THE MAXIMUM IMPROVEMENT FOR ANY SINGLE SECTION ARE GIVEN. IN EVERY CASE AN IMPROVEMENT IS OBSERVED FOR THE PROPOSED PREDICTIVE-ADAPTIVE TECHNIQUE (TECHNIQUE 3) OVER THE SINGLE BASIC MODEL (TECHNIQUE 1). THE LARGE MAXIMUM IMPROVEMENTS WERE TYPICALLY FOUND DURING THE NON-STANDARD BREATHING PATTERNS (I.E. QUICK AND DEEP) SUGGESTING THAT THE SINGLE BASIC MODEL DOES NOT ADEQUATELY CAPTURE THE MOTION DURING THESE PATTERNS.

* FOR THE PATIENT DATASET ONLY TWO BREATHING TYPES WERE USED: NORMAL AND DEEP. THEREFORE THE 120 DYNAMICS WERE SPLIT INTO TWO SECTIONS OF 60 FOR COMPUTING THE MAXIMUM IMPROVEMENT.

in the top figure. This corresponds roughly to the peak in the likelihood representing the near possibility. The amplitude prediction error over the entire X-ray sequence was 2.8mm +/- 3.0mm using the prior probability only, and 2.0mm +/- 3.0 using the Bayesian technique.

Figure 8 shows three sample X-ray frames, all acquired at late diastole and full inspiration. The coronary sinus catheter is indicated by the yellow arrow in Figure 8a. The red overlays represent the roadmap, which is a rendering of the coronary sinus segmented from the preprocedure MRI volume. The left column in Figure 8 shows the overlays before motion correction, and the right column shows the overlays after motion correction. The misalignment between the catheter and the overlay was assessed in each image by manually localising five points on the catheter and the closest points to each on the centre line of the vessel overlay. The root mean square (RMS) and maximum errors before motion correction were estimated to be 7.7mm and 11.9mm respectively. The same figures after motion correction were 2.8mm and 4.1mm.

IV. DISCUSSION

We have described a technique for constructing and applying single parameter motion models that has the capacity to adapt to different breathing patterns. The proposed technique is a combination of an adaptive motion model and an amplitude prediction technique. Validation of the amplitude prediction technique has indicated that the Bayesian approach offers significant improvements over the prior-only approach. This is due to being able to react to changes in the first-derivative of the surrogate value that signal a change in amplitude before the half-cycle has completed. For the adaptive model it is important that a reasonable amplitude estimate is formed as early as possible in the half-cycle, to ensure that the appropriate sub-model(s) can be used to compute the motion estimate. Results have shown that the combined predictive and adaptive technique gives significant improvements in modelling performance over normal single parameter respiratory models. Our results suggest that the magnitude of the improvement achieved seems to be subject dependent. The reason for the improvements can be seen from Figure 5: altering the amplitude of respiration causes a modification to the relationship between the single respiratory parameter and the motion of the organ. The likely underlying cause for this modification is a change in the actions of the different muscles involved in respiration. This change has long been appreciated in the physiology literature [4][5]. More recently, Nehrke et al [16] proposed that a change in the interplay of the respiratory muscles between inspiration and expiration is a likely cause for the hysteresis effect observed in respiratory motion. (Results presented in [16] suggest that the change in motion due to hysteresis can be as large as 5mm in the left ventricle of the heart.) Therefore it seems likely that the changing actions of these muscles during different breathing patterns will also cause changes to the motion function, whether due

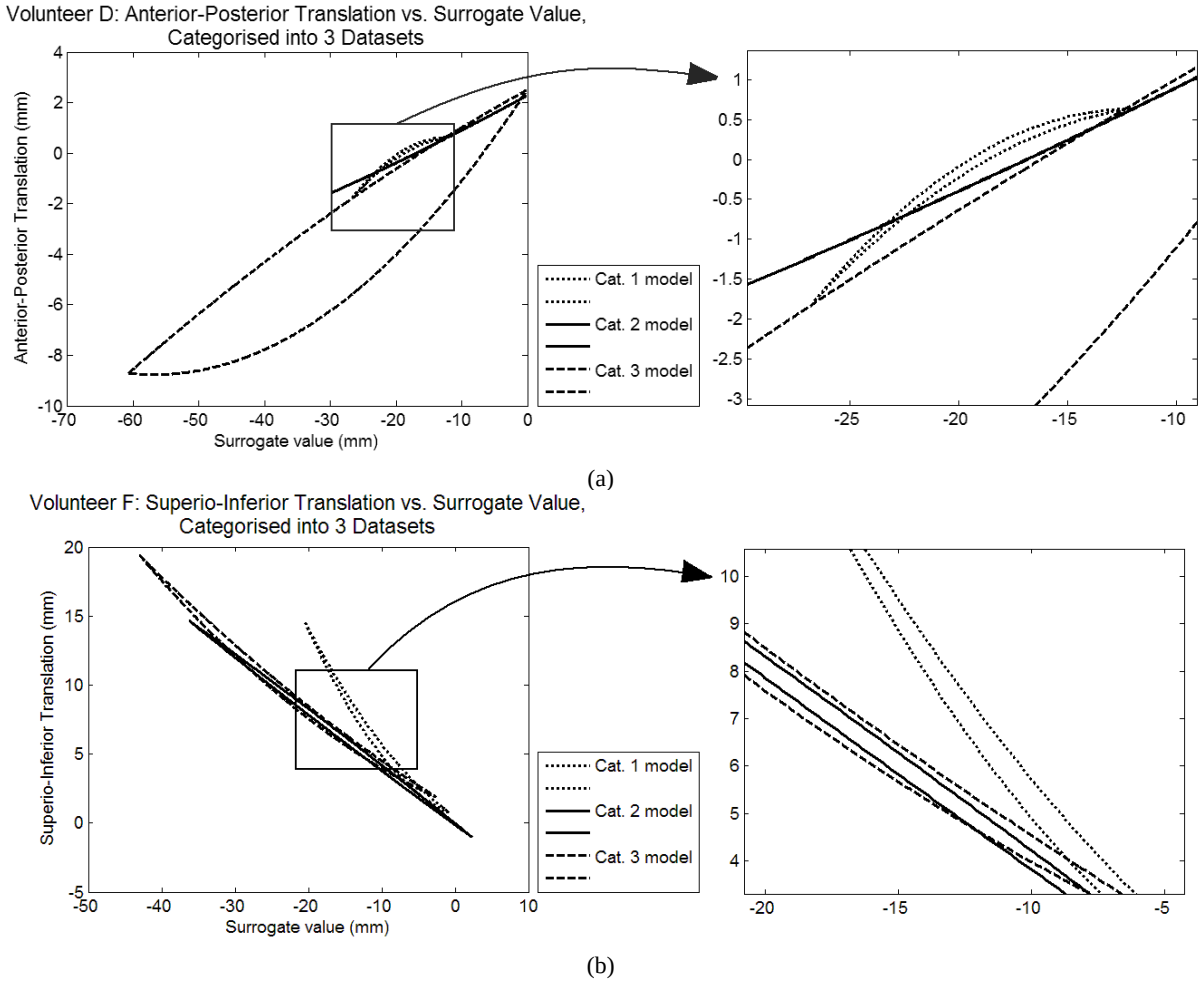

Fig. 5. Sample plots to illustrate the adaptive model. The dotted, solid and dashed curves represent the sub-models for small, medium and large amplitudes respectively. There are two curves for each sub-model: one for inspiration and one for expiration [8]. (a) Plot of the anterio-posterior translation from the adaptive model formed for vol. D. Note the strong hysteresis effect for cycles of large amplitude, i.e. there is a large difference between the inspiration and expiration curves. Both inspiration and expiration curves for cycles of small and medium amplitude approximately follow the inspiration curve for large amplitude cycles. An error of up to 4mm could result in this parameter alone if a sub-model for the wrong amplitude was applied for this volunteer. (b) Plot of the superio-inferior translation from the adaptive model formed for vol. F. In this case there is little hysteresis in any of the sub-models but the curve for cycles of small amplitudes is significantly different to those for medium and large amplitude. In this case a potential error of up to 3mm would result from selecting the wrong sub-model.

to variation in the amount of hysteresis (as in Figure 5a) or because of other changes in the nature of the motion (as in Figure 5b). Note that we do not explicitly decouple the effects of rib-cage and abdominal breathing. Rather, the decoupling is performed implicitly by separating the motion models into different amplitude categories.

Validation of the adaptive model was performed using a leave-one-out test on the registration results of the free-breathing dynamic scan. This effectively excludes any registration errors from the final error assessment - we are evaluating the technique's ability to predict the correct registration given a navigator value and an estimated amplitude. Therefore we cannot interpret the figures in Table II as true target registration errors. However they are useful for comparing the performances of the different modelling techniques. Errors were between 1.0mm-2.8mm (see Table II). Validation of the amplitude prediction technique alone showed errors of 1.7-6.5mm (see Table I). It should be noted that amplitude prediction errors are also included in the adaptive model errors given in Table II. Therefore it seems that the impact of the amplitude prediction errors is reasonably small. The likely reason for this is that the amplitude prediction errors are small compared to the overall breathing depth (see Table I).

We have also demonstrated how the proposed predictive and adaptive technique can be used in a clinical setting, applying it to motion correction of roadmaps in image guided cardiac catheterisations. Validation on this case showed that the model reduced 2-D registration errors from 7.7mm RMS to 2.8mm RMS. We did not perform a direct comparison of 2-D registration errors between the adaptive motion model and the basic motion model as described in [8] because other sources of error and uncertainty (e.g. initial registration error, diaphragm tracking error, errors in segmenting the coronary sinus) would make this comparison unreliable on a single dataset.

For the experiments presented in this paper, we used three sub-models to form an adaptive model for the volunteer datasets, and two sub-models for the patient dataset. In prin-

Subject	Parameter	Quantification of parameter variation between sub-models for short/medium/large amplitude cycles	
		Mean slope	Hysteresis
Vol. A:	Anterio-posterior translation (mm)	0.03/0.05/0.02	0.05/0.51/0.34
	Superio-inferior translation (mm)	-0.48/-0.52/-0.53	0.6/0.35/0.11
	Medio-lateral axis rotation (degrees)	0.04/0.06/0.11	0.0/0.28/0.45
Vol. B:	Anterio-posterior translation (mm)	0.0/*/0.11	0.21/*/1.37
	Superio-inferior translation (mm)	-0.55/*/-0.44	0.06/*/1.59
	Medio-lateral axis rotation (degrees)	0.06/*/0.09	0.37/*/0.23
Vol. C:	Anterio-posterior translation (mm)	0.12/0.12/0.12	0.04/0.26/0.89
	Superio-inferior translation (mm)	-0.36/-0.59/-0.64	1.23/0.23/0.2
	Medio-lateral axis rotation (degrees)	0.03/0.03/0.05	0.05/0.41/0.29
Vol. D:	Anterio-posterior translation (mm)	0.14/0.14/0.19	0.2/0.11/4.27
	Superio-inferior translation (mm)	-0.45/-0.42/-0.4	0.79/0.73/1.89
	Medio-lateral axis rotation (degrees)	0.13/0.11/0.09	0.29/0.88/0.43
Vol. E:	Anterio-posterior translation (mm)	-0.06/-0.05/0.03	0.55/0.37/1.82
	Superio-inferior translation (mm)	-0.49/-0.42/-0.45	0.89/2.77/1.66
	Medio-lateral axis rotation (degrees)	0.0/0.0/0.01	0.68/0.11/0.6
Vol. F:	Anterio-posterior translation (mm)	0.04/0.02/0.08	0.12/0.18/0.05
	Superio-inferior translation (mm)	-0.69/-0.4/-0.43	0.82/0.46/0.92
	Medio-lateral axis rotation (degrees)	0.08/0.07/0.03	0.33/0.21/0.16
Vol. G:	Anterio-posterior translation (mm)	*/0.05/0.26	*/0.09/1.52
	Superio-inferior translation (mm)	*/-0.25/-0.34	*/0.17/2.53
	Medio-lateral axis rotation (degrees)	*/0.06/0.03	*/0.06/2.15
Pat. A: **	Anterio-posterior translation (mm)	0.1/0.14	0.44/0.77
	Superio-inferior translation (mm)	-0.75/-0.74	0.63/1.39
	Medio-lateral axis rotation (degrees)	0.13/0.02	0.02/0.89

TABLE III

QUANTIFICATION OF THE VARIATION OF THE DOMINANT PARAMETER VALUES FOR DIFFERENT SUB-MODELS. TWO VALUES ARE SHOWN FOR EACH SUB-MODEL: THE MEAN SLOPE OF THE PARAMETER VALUE THROUGHOUT THE BREATHING CYCLE; AND THE AMOUNT OF HYSTERESIS, DEFINED AS THE MAXIMUM DIFFERENCE BETWEEN THE INSPIRATION AND EXPIRATION VALUE AT ANY POINT IN THE BREATHING CYCLE.
*: ONLY TWO SUB-MODELS WERE FORMED FOR THESE SUBJECTS DUE TO LACK OF DATA IN ONE OF THE SUB-MODELS.
**: ONLY TWO SUB-MODELS WERE FORMED FOR THIS SUBJECT BECAUSE THE DYNAMIC MRI SCAN WAS ACQUIRED FOR ONLY TWO DIFFERENT BREATHING PATTERNS.

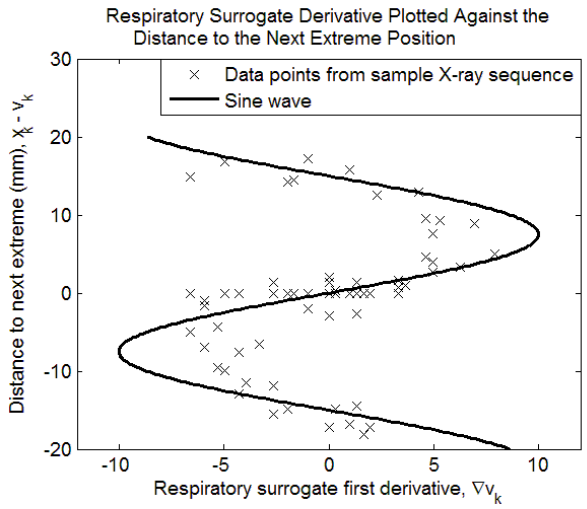

Fig. 6. Determining the relationship between the current surrogate first derivative and the distance to the next extreme position for X-ray surrogate data. The superio-inferior translation of the diaphragm was tracked in a sequence of fluoroscopic X-ray images. This tracking data was processed to determine true extreme positions, where the sign of the first derivative of the tracking data changed. For each sample point, the derivative and the distance to the next extreme was computed. The sine wave is an approximate fit to this data and allows a distance to the next extreme position to be estimated based on a surrogate's derivative value. This relationship must be determined once for each surrogate signal type.

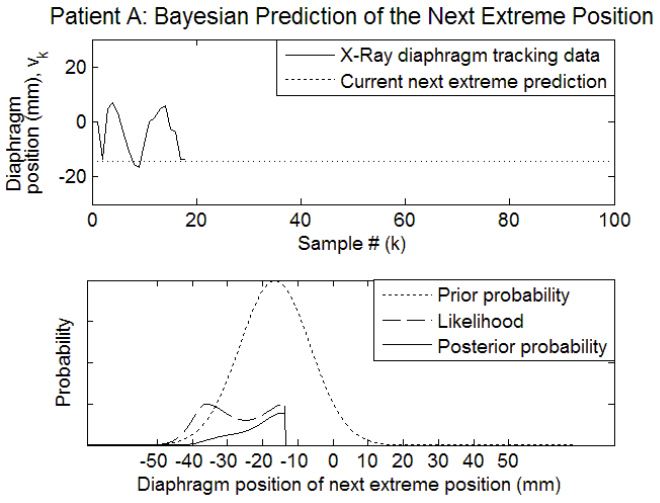

Fig. 7. Validating the amplitude prediction technique on X-ray diaphragm tracking data from patient A: Top - the diaphragm tracking data up to the current X-ray frame; bottom - using the prior and likelihood to compute the posterior probability of the next extreme position. The predicted position is at the peak of the posterior probability (7.2mm). This prediction is shown as the dotted line in the top figure.

ciple more than three sub-models could be used, but we believe that the limited benefit that this would bring would be outweighed by the increased MRI scan-time that would be necessary to acquire the data. Currently the proposed technique has a minimal effect on the clinical workflow due to the short additional scan time required.

We have used a formulation based on recursive Bayesian estimation to predict the next extreme respiratory position based on measurements of the current surrogate value and its

first time derivative. An alternative, widely applied, predictor is the Kalman filter [19], which predicts the true state of a variable given noisy observations. In fact, recursive Bayesian estimation can be seen as a generalisation of the Kalman filter. We believe that in our application there are difficulties with using the Kalman filter that make the more general theory of recursive Bayesian estimation more appropriate. The Kalman filter involves defining an observation model which maps from the true state to observed (noisy) values. The noise model for this mapping is assumed to be normally distributed. However, in our formulation, as we can see from Figure 3, a single surrogate derivative value can map to two possible extreme values, the near and far possibilities. Therefore, a one-to-one observation model could not be defined unless a non-normal noise model was used. Recursive Bayesian estimation allows us to define a more complex model for the likelihood, as illustrated in Figure 7.

An alternative approach to forming the N sub-models for different amplitude ranges would have been to form a two parameter model based on the surrogate value and the amplitude [12]. However, such an approach may have required a larger amount of data to fill the parameter space sufficiently to enable production of an accurate two parameter model. We chose instead to use a small number of single parameter sub-models and interpolate between them. The breathing protocol for the dynamic scan was designed to acquire enough data of different amplitudes to form accurate single parameter models of the motion.

Our hypothesis in developing this model was that changes in respiratory amplitude would lead to different types of motion as a function of the single respiratory parameter. The results presented in this paper have supported this hypothesis: if the motion function was the same during different breathing patterns then the adaptive model would not perform any better than a single basic model. An interesting area for future work is to investigate whether other predictors could be used to adapt a motion model, instead of or as well as the amplitude. Other possibilities that we plan to investigate include the wavelength, frequency, and the offset of the surrogate signal (i.e. the value of the surrogate at the beginning of a breathing cycle). However, it should be noted that including more than one such predictor into the adaptive model would increase its complexity. For example, to include both amplitude and offset would require nine sub-models to be formed instead of three, to cover all possible combinations of values for the predictors. Consequently, this would require significantly more data and computation time to form these sub-models. In this paper we have shown how using a single predictor can produce impressive results. Evaluating the effects of further predictors and devising effective strategies to incorporate them into the model remains a significant research challenge. Our proposed technique has also offered a potential solution to the problem raised by the changing motion function. The adaptive model has the advantage of simplicity of application, in that we still only need to acquire a single parameter during the procedure. The improved accuracy of the technique could have potential benefit in a range of clinical applications. For example, in cardiac catheterisations image guidance could be employed in a number of procedures that are currently not feasible due to their higher accuracy requirement. In lung radiotherapy treatment margins could be reduced due to improved respiratory motion correction. The degree of improvement that will be achieved in other applications may vary from organ to organ, and will depend on the degree to which the organ's motion is affected by changes in breathing pattern. Intuitively it seems possible that similar improvements to those found in the heart (i.e. up to 40%) could be achieved, but this hypothesis requires further verification and research.

Acknowledgements

The authors would like to thank the members of the Division of Imaging Sciences for useful discussions and feedback.

References

[1] King, A.P., Boubertakh, R., Ng, K.L., Ma, Y.L., Chinchapatnam, P., Gao, G., Schaeffter, T., Hawkes, D.J., Razavi, R., Rhode, K.S.: A technique for respiratory motion correction in image guided cardiac catheterisation procedures. In: Proceedings SPIE Medical Imaging. Volume 6918. (2008)

[2] West, J.B.: Respiratory Physiology: The Essentials. 7th edn. Lippincott, Williams and Wilkins (2004)

[3] De Troyer, A., Estenne, M.: Coordination between rib cage muscles and diaphragm during quiet breathing in humans. Journal of Applied Physiology **57**(3) (1984) 899–906

[4] Konno, K., Mead, J.: Measurement of the separate volume changes of rib cage and abdomen during breathing. Journal of Applied Physiology **22**(3) (1967) 407–422

[5] Sharp, J.T., Goldberg, N.B., Druz, W.S., Danon, J.: Relative contributions of rib cage and abdomen to breathing in normal subjects. Journal of Applied Physiology **39**(4) (1975) 608–618

[6] Manke, D., Rosch, P., Nehrke, K., Bornert, P., Dossel, O.: 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) (September 2002) 1132–1141

[7] Shechter, G., Shechter, B., Resar, J.R., Beyar, R.: Prospective motion correction of X-ray images for coronary interventions. IEEE Transactions on Medical Imaging **24**(4) (April 2005) 441–450

[8] King, A.P., Boubertakh, R., Rhode, K.S., Ma, Y.L., Chinchapatnam, P., Gao, G., Tangcharoen, T., Ginks, M., Cooklin, M., Gill, J.S., Hawkes, D.J., Razavi, R.S., Schaeffter, T.: A subject-specific technique for respiratory motion correction in image-guided cardiac catheterisation procedures. Medical Image Analysis **13**(3) (2009) 419–431

[9] Santhanam, A.P., Imielinska, C., Davenport, P., Kupelian, P., Rolland, J.P.: Modeling real-time 3-D lung deformations for medical visualization. IEEE Transactions on Information Technology in Biomedicine **12**(2) (2008) 257–270

[10] Al-Mayah, A., Modely, J., Brock, K.K.: Contact surface and material nonlinearity modeling of human lungs. Physics in Medicine and Biology **53** (2008) 305–317

[11] Manke, D., Nehrke, K., Bornert, P.: Novel prospective respiratory motion correction approach for free-breathing coronary MR angiography using a patient-adapted affine motion model. Magnetic Resonance in Medicine **50** (2003) 122–131

[12] McClelland, J., Blackall, J., Tarte, S., Hughes, S., Hawkes, D.: Non-rigid registration based respiratory motion models of the lung using two parameters. Medical Physics **34**(6) (June 2007) 2516–2516

[13] King, A.P., Rhode, K.S., Razavi, R., Schaeffter, T.: An adaptive and predictive respiratory motion model for image-guided interventions. In: Proceedings MICCAI Workshop on Image Guidance and Computer Assistance for Soft-Tissue Interventions. (2008)

[14] Keegan, J., Gatehouse, P., Yang, G.Z., Firmin, D.: 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** (2002) 476–481

[15] MacQueen, J.B.: Some methods for classification and analysis of multivariate observations. In: Proceedings of the Fifth Symposium on Math, Statistics and Probability, University of California Press (1967) 281–297

[16] Nehrke, K., Bornert, P., Manke, D., Bock, J.C.: Free-breathing cardiac MR imaging: Study of implications of respiratory motion - initial results. Radiology **220** (2001) 810–815
[17] Jazwinski, A.H.: Stochastic Processes and Filtering Theory. Academic Press (1970)
[18] Rhode, K.S., Sermesant, M., Brogan, D., Hegde, S., Hipwell, J., Lambiase, P., Rosenthal, E., Bucknall, C., Qureshi, S.A., Gill, J.S., Rezavi, R., Hill, D.L.G.: A system for real-time XMR guided cardiovascular intervention. IEEE Transactions on Medical Imaging **24**(11) (November 2005) 1428–1440
[19] Kalman, R.E.: A new approach to linear filtering and prediction problems. Trans. AMSE J. Basic Eng. **82** (1960) 35–45

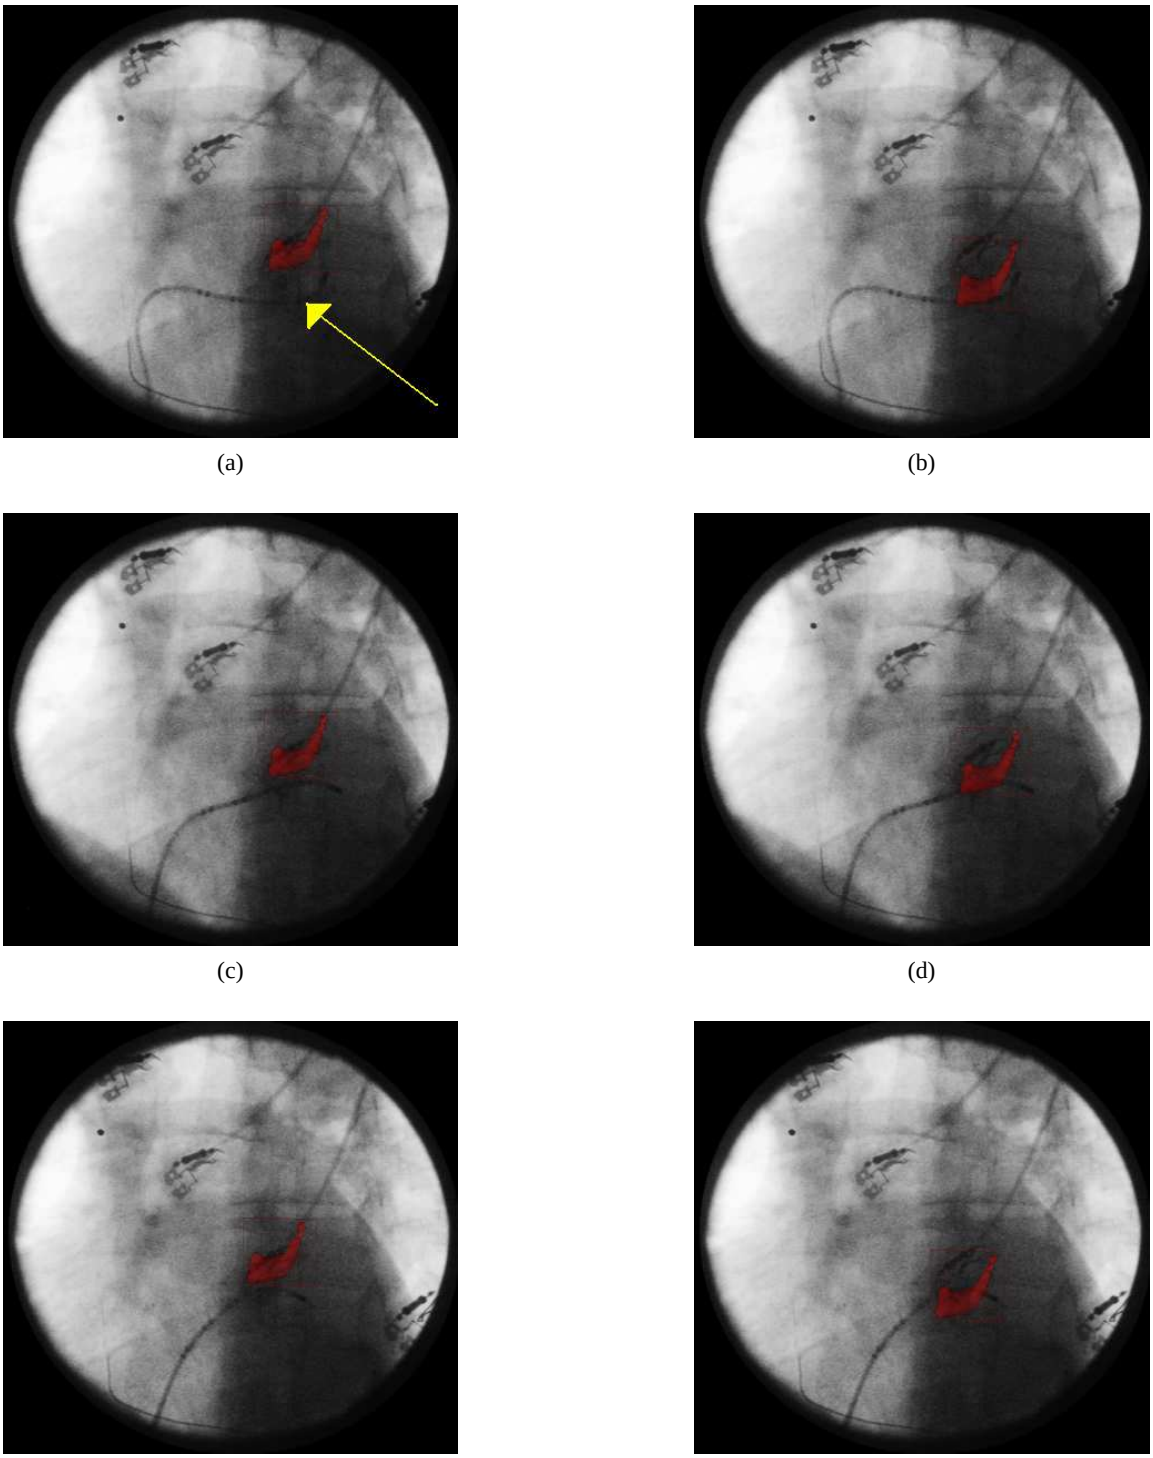

(a) (b) (c) (d) (e) (f)

Fig. 8. Motion correcting a roadmap using the predictive and adaptive model: (a), (c), (f): three sample frames at full inspiration from the X-ray sequence for patient A, showing the segmented coronary sinus overlay in red. The coronary sinus catheter is visible in the X-ray images and is indicated by the yellow arrow in (a); (b), (d), (f): the same three frames showing the coronary sinus overlay motion corrected using the predictive and adaptive model. The estimated misalignment between the coronary sinus catheter and the roadmap overlay was reduced from 7.7mm RMS to 2.8mm RMS after motion correction.