

An Adaptive and Predictive Respiratory Motion Model for Image-Guided Interventions

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

Interdisciplinary Medical Imaging Group, Division of Imaging Sciences, King's College London, U.K. andrew.king@kcl.ac.uk

Abstract. This paper describes a technique for constructing and applying an adaptive respiratory motion model for motion-correction of roadmaps in image-guided soft tissue interventions. The model has the capability to adapt to changes in breathing pattern, such as deep or fast breathing, which normally would result in a decrease in the accuracy of the model. It operates 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. A predictive technique is proposed which enables the amplitude of a breathing cycle to be predicted before the cycle has finished. The predicted amplitude is used to adapt the model to the current breathing pattern. The combined technique is shown to improve modelling performance by up to 31% in seven volunteer cardiac datasets. The predictive technique is also validated on three patient datasets, and shows mean errors of 0.6mm-3.7mm.

1 Introduction

Organs in the chest and abdomen, such as the heart, lungs and liver, undergo significant motion during respiration. The accuracy of guidance information during image-guided interventions on these organs is therefore reduced due to the assumption of static organ position and shape. To overcome this problem, motion models have been proposed for a range of interventions. However, our previous work has suggested that the accuracy of such models is 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.

Respiration is the result of two physiological factors: movement of the rib cage caused by the rib cage muscles and contraction of the diaphragm leading to an expansion of the lower chest [2]. The literature often distinguishes 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. For example, as the subject breathes more deeply or quickly, rib-cage breathing becomes more dominant [3].

Many respiratory motion models are based on a single respiratory parameter, or surrogate, such as the 1-D motion of the diaphragm (e.g. [1, 4]). We

postulate that the lower accuracy of such models in the presence of changes in breathing pattern 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. Previous work has addressed this limitation by introducing extra parameters [5], which implicitly attempt to capture and model the variation. However, when using motion models during image-guided interventions it is often difficult to acquire these extra parameters to apply the model. In [6] the use of extra parameters derived from the first parameter (amplitude and gradient) was proposed, but no method to estimate the amplitude in real-time was described. In this paper we assume that normal, deep and fast breathing states generally have different respiratory amplitudes and propose an extended single parameter model that can automatically adapt to changes in amplitude.

2 Method and Materials

2.1 Method

Figure 1 illustrates in general terms the processes involved in forming and applying the adaptive model. The model is formed from preprocedure imaging data together with a single respiratory surrogate. It can predict the motion of the target organ during the procedure given knowledge of the current surrogate value and the estimated amplitude of the current breathing cycle. The adaptive model consists of a number of basic sub-models, each of which describes the motion of the organ during cycles of different amplitude ranges. When the model is being formed the amplitudes are classified by post-processing the surrogate signal. A predictive technique is used to estimate the amplitude during model application.

In this paper we demonstrate the formation of an adaptive heart model from magnetic resonance imaging (MRI) data, using a pencil beam navigator applied on the diaphragm as the respiratory surrogate. However, the technique is general and can be applied to single parameter respiratory models for many applications.




Fig. 1. An overview of the formation and application of the adaptive motion model.

Basic Respiratory Model Details of the basic sub-models are given in [1]. Briefly, two cardiac gated MRI sequences are required to form the 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° , acquired voxel size $2.19 \times 2.19 \times 2.74\text{mm}^3$, reconstructed to $1.37 \times 1.37 \times 1.37\text{mm}^3$, 256×256 matrix, 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 = 11.75ms, TE = 5.84ms, flip angle = 20° , acquired voxel size $3.81 \times 4.27 \times 8.0\text{mm}^3$, reconstructed to $2.22 \times 2.22 \times 4.0\text{mm}^3$, 144×144 matrix). This sequence is similar to that used previously to form models for motion-corrected MRI image acquisition [4]. For the data used in this paper we used 120 dynamics to form each model, resulting in a typical scan time of 2 minutes. During the dynamic scan a pencil beam navigator is applied on the right hemi-diaphragm immediately before and after each dynamic acquisition. The average of these lead and trail navigators was used as the respiratory surrogate to form the models.

To form a basic respiratory model, each dynamic acquisition is automatically registered to the high resolution scan using an affine intensity-based algorithm. The motion of the target organ is modelled as polynomial functions of navigator value. Inspiration and expiration are modelled separately to enable the model to capture hysteresis effects. Once the model has been formed, it can produce an estimate of the motion from end-expiration to the current respiratory position given a navigator value and a breathing direction. The basic respiratory model can predict cardiac respiratory motion to within 2-4mm [1].

Adaptive Model An adaptive model, M^a , consists of a number of basic sub-models, $M_i, i = 1 \dots N$. The sub-models are formed by classifying the registration result for each dynamic according to its amplitude. We measure amplitude for each *half-cycle*, i.e. we compute the differences between the respiratory surrogate values of adjacent extreme (i.e. either end-expiration or end-inspiration) positions. Based on the half-cycle amplitude, each registration result is classified into one of N categories, each of which represents a particular range of amplitude values. Separate sub-models are formed for each category based on the classified registration results. For the experiments presented in this paper, we used three sub-models, covering 5-20%, 20-50% and 50-100% of the global maximum amplitude respectively. Half-cycles with an amplitude less than 5% of the maximum were rejected as they were probably due to noise in the surrogate signal.

To ensure that the adaptive model captured the full range of breathing motions, a breathing protocol was used each time the dynamic MRI scan was ac-

quired. For the first 40 dynamics, volunteers were given no special breathing instructions. For the next 40 dynamics, they were instructed to breathe quickly. For the final 40 dynamics, they were instructed to take deep breaths.

Predicting Amplitude At each sample point k we define $z_k = (v_k, \nabla v_k)$ to be a tuple comprising the respiratory surrogate value and its gradient. We denote the actual (unknown) next extreme position of the surrogate at sample point k by x_k and its current estimate by \hat{x}_k . This is illustrated in Figure 2(a). The technique for estimating the next extreme position is based on recursive Bayesian estimation [7]. 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

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

where Z_k is the set of all surrogate values/gradients up to and including position k (i.e. $z_1 \dots z_k$).

In (1), $p(x_k|Z_{k-1})$ is the prior probability of the surrogate value of the next extreme position. This reflects any knowledge we have about likely amplitudes before the half-cycle begins. We make the assumption that the respiratory amplitude will remain approximately constant from half-cycle to half-cycle. Therefore the value of the next extreme position is assumed to be the same as the previous one of the same type (i.e. end-expiration or end-inspiration). The form of the prior probability used was a Gaussian distribution centred on the previous extreme value,

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

where x_{prev} is the surrogate value of the previous extreme position of the same type and σ_p is an empirically determined constant (we used a value of 5). In principle more than one previous extreme value could be used in this term but for the experiments presented in this paper we used only the last extreme, x_{prev} .

$p(z_k|x_k)$ is the likelihood. This is a model of how a half-cycle of a given amplitude will produce measurements of the surrogate value and gradient. The likelihood estimate is based on previous surrogate data processed to establish the relationship between the current gradient and the distance to the next extreme position. Figure 2(b) shows a plot illustrating a sample relationship, which is modelled as a sine wave of a known amplitude and wavelength. The sine wave allows us to predict a gradient given a possible extreme surrogate value. The form of the likelihood for a possible extreme value x_k is given by

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

where ∇v_k is the actual gradient at the current point, and $\nabla \hat{v}_k = A \cdot \sin\left(\frac{(x_k - v_k)\pi}{\lambda}\right)$ is the estimated gradient given that the next extreme position will be x_k . A and

λ are the amplitude and wavelength of the sine wave. The standard deviation of the Gaussian function in (3), σ_l , indicates how much tolerance there is in the gradient estimates. This was empirically determined, and we used a value of 3. Note that $p(z_k|x_k)$ will have two peaks, corresponding to the two intersections of the current gradient with the sine wave.

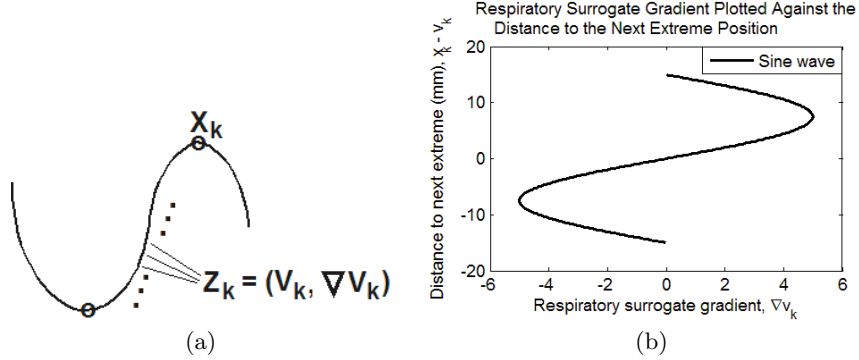


Fig. 2. (a) A sequence of sample points k within a half-cycle of respiratory surrogate values. At each point the surrogate, v_k , and its gradient, ∇v_k , are recorded, and an estimate, \hat{x}_k , is made of the true surrogate value of the next extreme position, x_k . (b) Estimating the distance to the next extreme position using the current gradient.

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 .

Model Interpolation Given the current surrogate value, the adaptive model produces an estimate for each motion parameter by interpolating between the estimates of its sub-models based on the current amplitude prediction. 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 . Now we define the adaptive estimate, \hat{T}_j , of each motion parameter j as

$$\hat{T}_j = \phi(\tilde{T}_j, \hat{x}_k), j = 1 \dots 12 \quad (5)$$

where $\tilde{T}_j = (T_j(M_1, v_k), \dots, T_j(M_N, v_k))$, 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.

2.2 Materials

The adaptive model experiments were performed in the 1.5T Philips Achieva MRI scanner at Guy’s Hospital in London. The seven volunteers all gave informed consent. The predictive technique experiments were performed on a Philips BV Pulsera single plane cardiac X-ray set. Three patients underwent cardiac catheterisation procedures and were sedated throughout. X-ray image sequences were acquired showing the motion of the patients’ diaphragm over several breathing cycles. All patients gave informed consent.

2.3 Experiments

To validate the formation of the adaptive model we used a leave-one-out test. Each adaptive model was formed from the results of registering 120 dynamic acquisitions to a high resolution volume. We formed 120 separate adaptive models by leaving out each of the registration results in turn. Each of these models was used to predict the registration result for its left-out dynamic. The predicted registrations were compared with the actual registration results. Errors were computed at ten clinically relevant anatomical landmarks.

Using this validation, we tested three different model formation techniques: a single basic model, an adaptive model with perfect knowledge of the amplitudes and an adaptive model using the amplitude prediction technique. By forming these three model types we aimed to validate two separate factors. By comparing the first and second techniques we tested the effectiveness of forming an adaptive model rather than a basic model. By giving the model perfect knowledge of the amplitudes we excluded the performance of the predictive technique. By comparing the second and third techniques we test the impact of imperfect knowledge of the amplitudes on the overall accuracy of the adaptive model.

One of our aims is to apply the adaptive model to motion-correction in X-ray guided catheterisations [8, 1]. This would involve acquiring the respiratory surrogate by tracking diaphragm motion in fluoroscopic X-ray images. Therefore, as a further validation of the amplitude prediction technique, we tracked diaphragm motion in X-ray images acquired from 3 clinical cases. The tracking was performed by defining a rectangular region of interest in the X-ray images, and finding the 1-D translation that minimised the mean sum of squared intensity differences within this region between each X-ray image and a reference image [1]. The tracking data was retrospectively processed to find the true extreme positions, which were used as a gold standard. The results of applying the amplitude prediction 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.

3 Results

The leave-one-out test results for the adaptive models formed for seven volunteers are shown in Table 1. A sample parameter plot from the adaptive model for volunteer D is shown in Figure 3.

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	2.02 +/- 0.91	2.0 +/- 0.99	1.99 +/- 1.01	1.5%
Max.	1.92 +/- 0.93	1.78 +/- 0.91	1.78 +/- 0.91	7.3%
Vol. B: Mean	1.81 +/- 1.06	1.58 +/- 0.96	1.61 +/- 0.95	11.0%
Max.	2.03 +/- 1.18	1.62 +/- 1.0	1.71 +/- 1.01	15.8%
Vol. C: Mean	1.18 +/- 0.6	1.02 +/- 0.54	0.98 +/- 0.52	16.9%
Max.	1.48 +/- 0.66	1.2 +/- 0.63	1.14 +/- 0.58	23.0%
Vol. D: Mean	1.39 +/- 0.84	1.21 +/- 0.77	1.17 +/- 0.72	15.8
Max.	1.44 +/- 0.9	1.13 +/- 0.49	1.09 +/- 0.48	24.3%
Vol. E: Mean	1.73 +/- 0.92	1.69 +/- 0.93	1.67 +/- 0.94	3.5
Max.	1.98 +/- 1.04	1.86 +/- 0.98	1.87 +/- 0.99	5.5%
Vol. F: Mean	1.7 +/- 0.78	1.46 +/- 0.77	1.45 +/- 0.73	14.7 %
Max.	1.69 +/- 0.71	1.13 +/- 0.51	1.16 +/- 0.57	31.4%
Vol. G: Mean	2.68 +/- 1.55	2.13 +/- 1.3	2.24 +/- 1.41	16.4
Max.	1.48 +/- 0.6	1.2 +/- 0.61	1.13 +/- 0.59	23.6%

Table 1. 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 dataset 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 amplitude prediction technique experiments, the mean errors for the prior-only and Bayesian techniques were 3.3mm and 2.0mm respectively for patient A, 9.0mm and 3.7mm for patient B and 1.4mm and 0.6mm for patient C. Figure 4(a) illustrates the use of the prediction technique on patient A. The gradient of the tracking data at the current frame is low, so the likelihood has two peaks representing possible next extreme positions (the near and far possibilities). The prior is a Gaussian distribution centred on the expected next extreme position (i.e. that it will be the same as the previous one). The posterior probability is the product of the prior and likelihood. The current peak posterior probability indicates a predicted next extreme position of -6.2mm, as indicated in the top graph. Figure 4(b) shows the errors throughout the tracking sequence.

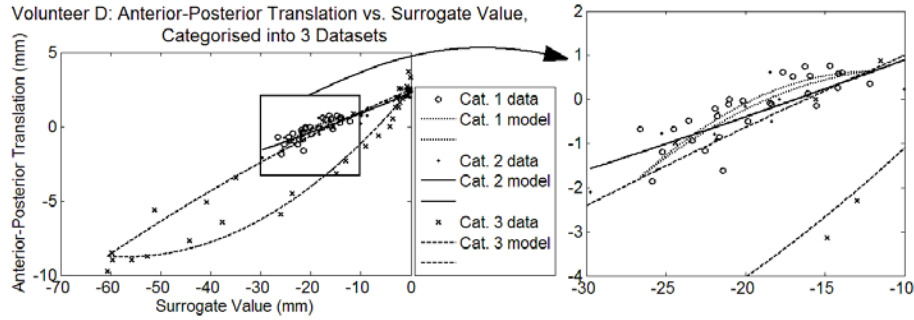


Fig. 3. Plot of the antero-posterior translation from an adaptive model formed for vol. D. The registration results used to form the models are indicated by the markers 'o', '.' and 'x', representing small, medium and large amplitude cycles respectively. The curves represent the motion model fit to this data for each sub-model. There are two curves for each sub-model: one for inspiration and one for expiration [1]. Note the strong hysteresis effect for cycles of large amplitude. 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 model for the wrong amplitude was applied for this volunteer.

4 Discussion

We have described a technique for constructing single parameter motion models that has the capacity to adapt to different breathing patterns. In the results presented in Table 1, comparing the first and second techniques 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 similar for the two approaches. In all seven cases the predictive-adaptive technique gave improvements over the single model, with a maximum improvement of 31.4%. The maximum improvements were typically found during 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.

The reason for the improvements can be seen from Figure 3: changing the amplitude of respiration causes a change in the relationship between organ motion and the single respiratory parameter. In Figure 3, there is a strong hysteresis effect when the subject breathes deeply, which is absent in smaller amplitude cycles. Using a normal single parameter motion model would lead to errors of up to 4mm in this parameter alone. The adaptive motion model would have much improved performance, as it would detect the large amplitude cycle and apply the appropriate sub-model. The magnitude of the improvement achieved seems to be highly subject dependent. For example, for volunteers A and E the predictive and adaptive technique did not show much improvement over the basic model, but for the other five volunteers there was a marked improvement.

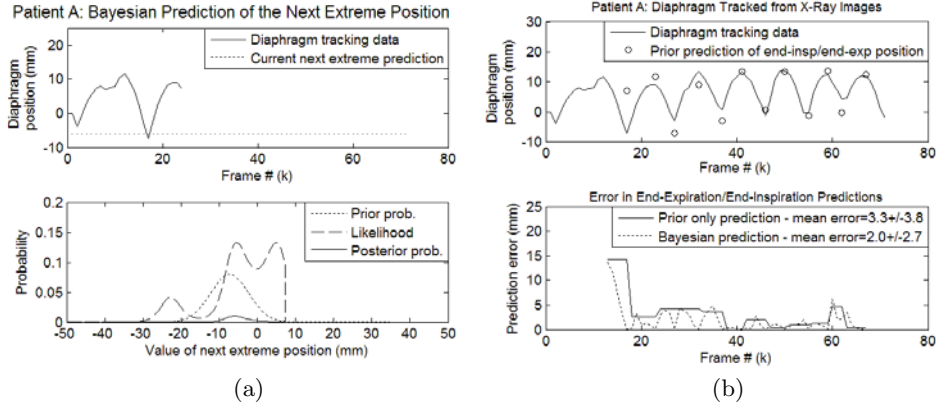


Fig. 4. Validating the amplitude prediction technique on X-ray diaphragm tracking data from patient A: (a) top - the diaphragm tracking data up to X-ray frame 24; bottom - using the prior and likelihood to compute the posterior probability of the next extreme position at X-ray frame 24. The predicted position is at the peak of the posterior probability (-6.2mm). This prediction is shown as the dotted line in the top figure; (b) top - the complete diaphragm tracking sequence overlaid with the prior predictions of the next extreme position (i.e. that they are the same as the previous extreme of the same type); bottom - prediction errors using the prior-only and the Bayesian technique throughout the tracking data sequence, compared to the actual extreme positions.

Validation 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 techniques' ability to predict the correct registration given a navigator value and an estimated amplitude. Therefore we cannot interpret the figures in Table 1 as true target registration errors. However they are useful for comparing the performances of the different modelling techniques.

The selection of the amplitude parameters used to form the sub-models (i.e. 5-20%, 20-50% and 50-100% of the global maximum amplitude) is currently based on our observation of typical respiratory amplitudes during fast, normal and deep breathing. We have not yet investigated the impact of modifying these parameters but plan to do so in future work.

The amplitude prediction results indicate that the Bayesian technique offers significant improvements over the prior-only approach. In particular, Figure 4(b) shows that in noisy parts of the signal, the prior-only approach performs poorly, but the Bayesian technique quickly recovers from the poor initial predictions (e.g. at around frame 15).

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 pat-

terns then the adaptive model would not perform any better than the basic model. The 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 will 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.

Currently we have performed clinical validation of the basic single parameter model on three cases [1]. In future work we plan to evaluate the improvement resulting from the use of the predictive and adaptive model over the basic model in further clinical cases.

Acknowledgments

This work was funded by EPSRC grant EP/D061471/1, DTI technology programme grant 17352 and Philips Medical Systems, Best, the Netherlands.

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: Proc. SPIE Medical Imaging. (2008)
2. West, J.B.: Respiratory Physiology: The Essentials. 7th edn. Lippincott, Williams and Wilkins (2004)
3. 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
4. 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
5. 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
6. 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
7. Jazwinski, A.H.: *Stochastic Processes and Filtering Theory*. Academic Press (1970)
8. 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