

Recursive Bayesian Estimation of Respiratory Motion using a Modified Autoregressive Transition Model

Ashrani Aizzuddin Abd. Rahni^{b,a}, Emma Lewis^a and Kevin Wells^a

^aCentre for Vision, Speech and Signal Processing, University of Surrey, Guildford, GU2 7XH;

^bUniversiti Kebangsaan Malaysia, Malaysia

ABSTRACT

Compensation for respiratory motion has been identified as a crucial factor in achieving high resolution Nuclear Medicine (NM) imaging. Many motion correction approaches have been studied and they are seen to have advantages over simpler approaches such as respiratory gating. However, all motion correction approaches rely on an assumption or estimation of respiratory motion. This paper builds upon previous work in recursive Bayesian estimation of respiratory motion assuming a stereo camera observation of the motion of the external torso surface. This paper compares the performance of a modified autoregressive transition model against the previously presented linear transition model used when estimating motion within a 4D dataset generated from the XCAT phantom.

Keywords: Nuclear Medicine, Respiratory motion, Recursive Bayesian Estimation, Modified Autoregression

1. INTRODUCTION

Previous work has shown the feasibility of performing recursive Bayesian estimation of respiratory motion in the form of particle filtering by evaluation using XCAT.¹ Briefly, recursive Bayesian estimation involves the estimation of a hidden state at a discrete time index k , \mathbf{x}_k , given all observations up until that time, $\mathbf{z}_{1:k} \equiv \{\mathbf{z}_1, \dots, \mathbf{z}_k\}$. In the proposed application the state, \mathbf{x}_k , is some representation of the time varying configurations of internal organs whereas the observable, \mathbf{z}_k , represents a stereo camera capture of the anterior surface of the torso.

Recursive Bayesian estimation incorporates both a correspondence or measurement model \mathbf{b} between the hidden state, \mathbf{x}_k , and the observable, \mathbf{z}_k , and a transition model \mathbf{a} that describes the temporal evolution of the state, \mathbf{x}_k . The models can be formally defined as:

$$\mathbf{x}_k = \mathbf{a}_{k-1}(\mathbf{x}_{k-1}, \mathbf{v}_{k-1}) \leftrightarrow f(\mathbf{x}_k | \mathbf{x}_{k-1}), \quad (1)$$

$$\mathbf{z}_k = \mathbf{b}_k(\mathbf{x}_k, \mathbf{w}_k) \leftrightarrow g(\mathbf{z}_k | \mathbf{x}_k), \quad (2)$$

\mathbf{v}_{k-1} and \mathbf{w}_k are transition and measurement noise respectively and f and g are the respective pdfs of the transition and measurement models. Recursive Bayesian estimation incorporates both models by assuming a hidden Markov model view of the estimation framework (Fig. 1).

When both the transition and measurement models are linear and their respective pdf are Gaussian, the models can be expressed as:

$$\mathbf{x}_k = A\mathbf{x}_{k-1} + \mathbf{d} + \mathbf{v}_{k-1}, \quad (3)$$

$$\mathbf{z}_k = B\mathbf{x}_k + \mathbf{b}_0 + \mathbf{w}_k, \quad (4)$$

where \mathbf{d} and \mathbf{b}_0 are intercept vectors so that the respective hyperplanes represented by both models do not have to go through the origin. When both models are linear, the state estimate, $\tilde{\mathbf{x}}_k$ has a Gaussian distribution and the Kalman filter is the optimal solution to the recursive Bayesian estimation problem.² The posterior pdf of the state can thus be described exactly by its mean value, $\hat{\mathbf{x}}_k$, and covariance, P_k . Particle filtering on the other hand, represents the posterior pdf of the state as a set of N Monte Carlo samples or “particles” of the pdf, $\{\mathbf{x}_k^i\}_{i=1}^N$. The measurement model then gives a weight ω_k^i to each particle indicating their likelihood. When the posterior pdf has a parametric representation such as in the case of a Gaussian, particle filtering will not be optimal. However, when there is no analytical representation of the posterior such as for generic transition

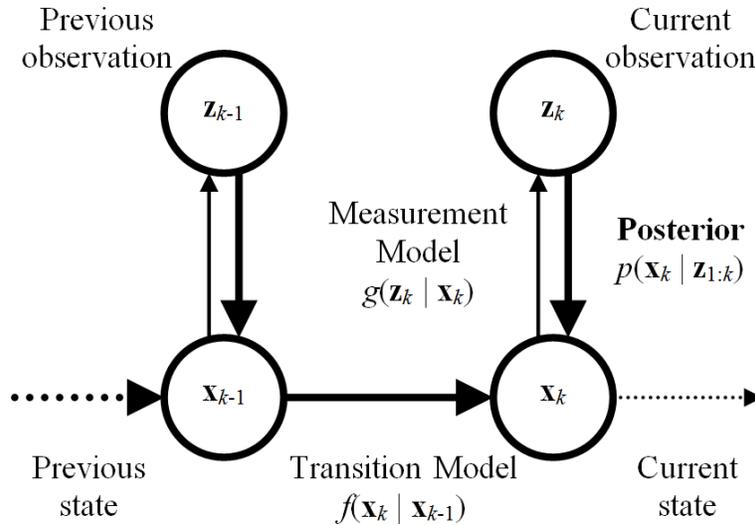


Figure 1. Diagram of the system illustrating the structure as a first order hidden Markov model. The posterior of the current state is represented by the bolded arrows and circles showing the information available at time k .

and measurement models, then particle filtering may be a better representation of the pdf than assuming that a parametric representation holds.

Recursive Bayesian estimation depends on the transition and measurement models that are used.³ When these models are found using a training dataset, they may not be generalisable to other datasets. Hence, additional adaptation methods were proposed to address this issue.¹ Due to the respiratory estimation process in nuclear medicine imaging being a case of off-line estimation (i.e. after all data is acquired), one of the methods involve spatial-temporal alignment of the training cycle and the test cycle. However, this alignment itself has uncertainties in the estimation of respiratory cycle period and amplitude. In this paper, the uncertainty in spatial and temporal scale will be accounted for in a modified transition model as described in Section 2.

2. METHODOLOGY

As in previous work,¹ the state variable consists of organ configuration, \mathbf{c}_k , from two consecutive time points. Using the following change of variables for equation 3:

$$\mathbf{x}_k = \begin{pmatrix} \mathbf{c}_k \\ \mathbf{c}_{k-1} \end{pmatrix}, A = \begin{pmatrix} A_1 & A_2 \\ I & 0 \end{pmatrix}, \mathbf{d} = \begin{pmatrix} \mathbf{d}_0 \\ 0 \end{pmatrix}, \Sigma_{\mathbf{v}} = \begin{pmatrix} \Sigma_{\mathbf{v}^*} & 0 \\ 0 & 0 \end{pmatrix}, \quad (5)$$

the linear transition model can thus be expressed as a second order autoregressive model, AR(2), with respect to organ configuration, \mathbf{c}_k :

$$\mathbf{c}_k = \sum_{p=1}^2 A_p \mathbf{c}_{k-p} + \mathbf{d}_0 + \mathbf{v}_k^*. \quad (6)$$

The modification of this AR(2) model to account for spatial and temporal scale uncertainty is thus as follows in subsections 2.1 and 2.2. As the resulting model is non-Gaussian (17), it is best applied using a particle filter (PF). However, given a sampled value of the stochastic scalars introduced in subsections 2.1 and 2.2, local linearisation can be performed, and this is described in subsection 2.3.

2.1 Accounting for Spatial Scale Uncertainty

As both transition and measurement models are linear, uncertainty in spatial scale is easily accounted for by adding a stochastic scale factor, u_k^* , in the transition model:

$$\mathbf{x}_k = u_k^* A \mathbf{x}_{k-1} + \mathbf{d} + \mathbf{v}_k, \quad (7)$$

where

$$u_k^* = \begin{cases} u_k & \text{if } u_k \geq 0 \\ 0 & \text{otherwise} \end{cases}, \quad (8)$$

and $u_k \sim \mathcal{N}(\bar{u}, \sigma_u^2)$, a Gaussian random variable with mean \bar{u} and variance σ_u^2 . This has been included previously in (Rahni et al., 2011).¹ However, here u_k is specifically defined so that it reflects the uncertainty in the amplitude ratio, r_s , where $r_s = C_n/C_\tau$, with C_n being the amplitude of a test cycle and C_τ being the amplitude of a training cycle. This uncertainty can be approximated to be Gaussian.⁴ However, the transition pdf, $f(1)$ is non Gaussian due to the stochastic scale factor, u_k^* .

2.2 Accounting for Temporal Scale Uncertainty

On the other hand, uncertainty in temporal scale can be incorporated by the fact that the transition model is equivalent to a second order autoregression. The frequency content of multidimensional AR(2) model can be seen by considering the 1D case:

$$c_k = \sum_{p=1}^2 a_p c_{k-p} + d_0 + v_k^*, \quad (9)$$

where the terms parallel those in the multidimensional case (6). If the process variable is mean-adjusted and only the deterministic portion of the process considered, it can be rearranged as:

$$c_k^* = \sum_{p=1}^2 a_p c_{k-p}^*, \quad (10)$$

where c_k^* is the mean-adjusted process variable*. Equation 10 is thus a second order recurrence relation. The characteristic polynomial^{5,6} for (10) is then:

$$\lambda^2 - a_1 \lambda - a_2 = 0, \quad (11)$$

whose complex conjugate roots are:

$$\lambda_1 = \gamma \exp(i\Omega), \lambda_2 = \gamma \exp(-i\Omega). \quad (12)$$

The roots (12) thus implies that $a_1 = 2\gamma \cos(\Omega)$ and $a_2 = -\gamma^2$. A second order recurrence relation (10) is stable when $\gamma < 1$. The general solution of the recurrence relation is thus:⁶

$$\begin{aligned} c_k^* &= \alpha_1 \lambda_1^k + \alpha_2 \lambda_2^k \\ &= \alpha_1 \gamma \exp(i\Omega k) + \alpha_2 \gamma \exp(-i\Omega k), \end{aligned} \quad (13)$$

where α_1 and α_2 are two constants. From this solution (13), it can be seen that an AR(2) process can represent a sinusoid exactly, as:

$$\sin(\Omega k + \phi) = \frac{1}{2i} (\exp(i(\Omega k + \phi)) - \exp(-i(\Omega k + \phi))). \quad (14)$$

Hence by extension the angular frequency of any 1D AR(2) process is given by the phase of the complex conjugate roots (12) of its characteristic polynomial (11).

To relate a 1D AR(2) process to a multidimensional AR(2) process, it is useful to look at the companion matrix A_c of (10) defined as:

$$A_c = \begin{pmatrix} a_1 & a_2 \\ 1 & 0 \end{pmatrix}. \quad (15)$$

Hence, the roots λ_i of the characteristic polynomial (11) are the eigenvalues of the companion matrix (15).⁵ If a vector $\mathbf{X}_k = [c_k^*, c_{k-1}^*]^T$ is defined, then using the companion matrix (15), the 1D second order recurrence

i.e. $c_k^ = c_k - \hat{c}$, where \hat{c} is the mean.

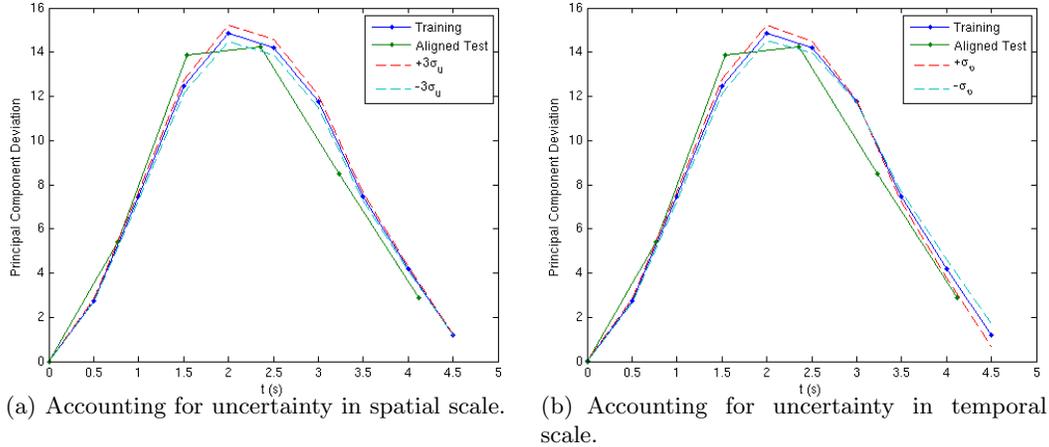


Figure 2. Figs. (a) and (b) shows the effect of uncertainty in spatial and temporal scale respectively (dashed lines). An aligned test cycle is shown for comparison.

relation (10) can be represented as a 2D first order recurrence relation, $\mathbf{x}_k = A_c \mathbf{x}_{k-1}$. By further extension, any n dimensional first order recurrence relation is associated with a characteristic polynomial (11) of order n .⁵

By extension of the case of a 1D AR(2) model, the frequency content of a multidimensional AR(2) model is determined by the eigenvalues of the transition matrix A , which occur in complex conjugate pairs, $\lambda_j = \gamma \exp(\pm i\Omega_j)$. The phases of these eigenvalues are thus the angular frequency content of the AR(2) model. Temporal scale uncertainty can thus be incorporated by multiplying these angular frequencies by a second stochastic scalar:

$$\Omega_{j,k} = v_k^* \Omega_j. \quad (16)$$

and using $\Omega_{j,k}$ instead of Ω_j . v_k^* is defined similarly to u_k^* in (16), except in this case it denotes the uncertainty in the ratio r_T , of the period of a test cycle, T_n , to the period of a training cycle, T_τ , where $r_T = T_n/T_\tau$. The modified AR(2) model incorporating both types of uncertainties (temporal and spatial scale, represented by v_k^* and u_k^* respectively) is thus:

$$\mathbf{x}_k = u_k^* A(v_k^*) \mathbf{x}_{k-1} + \mathbf{d} + \mathbf{v}_k, \quad (17)$$

where the matrix A is now a non-linear function of v_k^* .

In Figs. 2(a) and (b), the linear transition model is simulated as a recurrence relation where the confidence intervals (dashed lines) indicate the range of variability in the trajectory when uncertainty in spatial scale (7) and temporal scale (16) respectively. In Fig. 2(a) the maximum amplitude divergence is seen around $t = 2$ s while in Fig. 2(b) there is an additional divergence at around $t = 4.5$ s. The trajectory shown is that of the first principal component when the model is formed in PCA space. For temporal uncertainty, the half range of the sampling period, with the observable at 10Hz, is equated to a $3\sigma_T$ confidence interval of a Gaussian[†]. The aligned first principal component of a test cycle is also shown.

2.3 Marginalisation using Local Linearisation

Due to the stochastic scalars introduced in subsections 2.1 and 2.2, the resulting modified AR(2) model is non Gaussian and is hence more accurately realised with a recursive Bayesian estimator that does not assume normality such as a PF. However, if this transition model is conditioned upon a certain value of u_k^* and v_k^* , then it becomes a normal AR(2) model and is thus Gaussian. Therefore, the Bayesian recursion (Fig. 1) can be linearised locally.²

[†]This is a heuristic choice of a small value of σ_T so that the effect of temporal scale uncertainty is comparable to the effect of spatial scale uncertainty (Fig. 2).

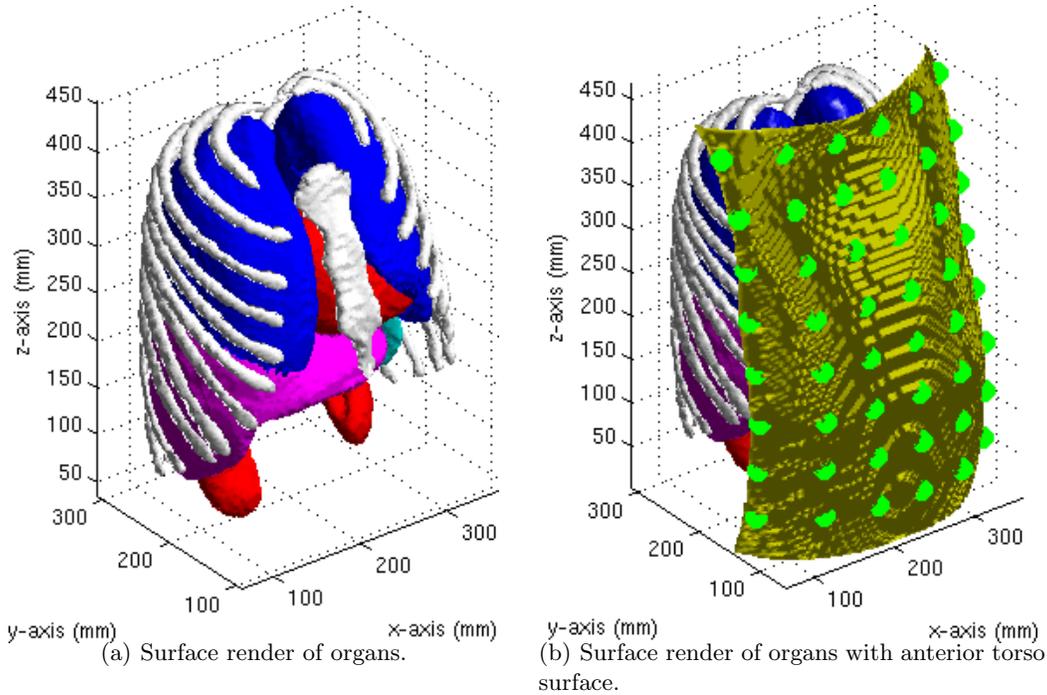


Figure 3. Fig. (a) shows the rendering of the chosen organs from XCAT, namely the ribcage (white), lungs (blue), heart (red), liver (magenta), spleen (cyan), and kidneys (red). Fig. (b) shows the rendering of the chosen organs with the anterior torso surface. Points chosen for the observable are shown as green circles (\bullet).

If a vector is defined to represent the additional random variables: $\mathbf{u}_k = [u_k^*, v_k^*]^T$, then each particle or Monte Carlo realisation of the state, \mathbf{x}_k^i , can thus be sampled from the following importance density:

$$p(\mathbf{x}_k^i | \mathbf{x}_{k-1}^i, \mathbf{z}_k, \mathbf{u}_k^i) = \mathcal{N}(\hat{\mathbf{x}}_k^i, \hat{P}_k^i), \quad (18)$$

where the mean $\hat{\mathbf{x}}_k^i$ and \hat{P}_k^i are found using Kalman filtering for each sampled value of the additional random variable, \mathbf{u}_k^i . The weight of each particle, ω_k^i , is then:

$$\omega_k^i \propto \omega_{k-1}^i \frac{p(\mathbf{z}_k | \mathbf{x}_k^i) p(\mathbf{x}_k^i | \mathbf{x}_{k-1}^i, \mathbf{u}_k^i)}{q(\mathbf{x}_k^i | \mathbf{x}_{k-1}^i, \mathbf{z}_k, \mathbf{u}_k^i)} \quad (19)$$

3. EVALUATION AND RESULTS

Following previous evaluation,¹ the state, \mathbf{x}_k , comprises of organ configurations, \mathbf{c}_k , consisting of affine deformation parameters for each organ. On the other hand, the observable, \mathbf{z}_k , consists of the 3D coordinates of 48 virtual markers on the torso surface. This representation parallel that used in (Rahni et al., 2013).⁷ A rendering of the chosen organs and the torso surface is shown in Fig. 3.

Table 1 summarises the respiratory cycles generated from XCAT that are used for evaluation. The same dataset have been used in a previous evaluation of recursive Bayesian estimation of respiratory motion.¹

As in the previous evaluation,¹ the training cycle is the same as ‘‘Cycle 2’’ in Table 1, except that a different realisation of noise is used for the observable, \mathbf{z}_k . The accuracy of the estimates $\tilde{\mathbf{x}}_k$ are evaluated as mean Euclidean distance error of organ voxels due to the estimated configuration $\tilde{\mathbf{c}}_k$ from their true positions given by the XCAT vector output facility. The results for a Kalman filter (KF), a sampling importance resampling (SIR)¹ PF using the modified AR(2) model and a locally linearised PF (LLPF) are shown in Fig. 4. In the KF, a normal AR(2) transition model is used as the purpose of the modified AR(2) model used in the PFs is to incorporate the uncertainties in spatial-temporal alignment (section 2). The spatial-temporal alignment itself is

Table 1. Respiratory Parameters for Training and Test Datasets

Cycle	Motion Amplitude (cm)		Cycle Period (s)
	Diaphragm (SI)	Chest (AP)	
1	1	0.6	5
2	2	1.2	5
3	3	1.8	5
4	4	2.6	5
5	1	0.6	3
6	2	1.2	3
7	3	1.8	3

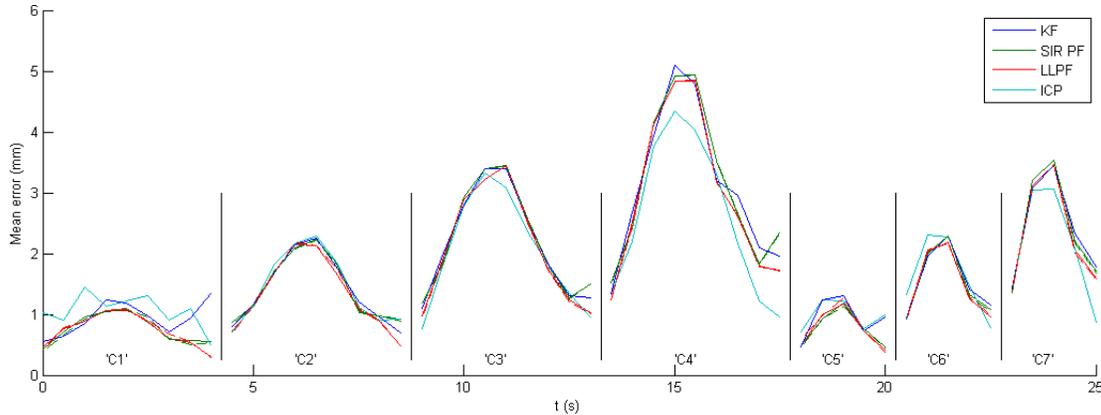


Figure 4. Mean error over all voxels in the chosen organs for each frame $k = t_k$ in all test cycles (“C1” to “C7”). The reference frame of each test cycle is excluded as there is no estimation in those frames. For the PFs, there is a second set of results using the MAP estimate of the state indicated by dashed lines.

used in all three recursive Bayesian estimators. The mean errors for affine assumption of organ deformation via ICP (iterative closest points), which underlies the estimation framework, is also shown for comparison. Table 2 on the other hand shows the mean values of the errors over all frames.

The results show that in this evaluation using XCAT, all three recursive Bayesian estimators show very similar results, partly due to XCAT being a simplified representation of respiratory motion. However, over all frames, compared to the KF, the use of a modified AR(2) in the SIR PF has reduced the mean error by 2%. The LLPF, which approximates an optimal sampling density (18), reduces the mean error even more, by 6% compared to the KF.

4. CONCLUSION

This paper showed how spatial and temporal scale uncertainty can be included in the transition model of a particle filter for motion estimation. An alternative approach is to include the uncertainties as deterministic parameters in a Kalman filter as in (Smith et al., 2012).⁸ Further investigation is needed as well as possibly incorporating non-linear models and evaluating on more clinically applicable data.

Acknowledgement

Ashrani Aizzuddin Abd. Rahni was funded by the Malaysian Ministry of Higher Education and University Kebangsaan Malaysia. Acknowledgement also to John Jones for the implementation of ICP.

Table 2. Comparison of Mean Errors for the Bayesian estimators

Error Type	Method			
	KF	SIR PF	LLPF	ICP
Mean Total Error (mm)	1.828	1.778	1.708	1.717

REFERENCES

- [1] A. A. Abd. Rahni, E. Lewis, M. J. Guy, B. Goswami, and K. Wells, "A Particle Filter Approach to Respiratory Motion Estimation in Nuclear Medicine Imaging," *IEEE Trans Nuc Sci* **58**, pp. 2276–2285, Oct 2011.
- [2] B. Ristic, S. Arulampalam, and N. Gordon, eds., *Beyond the Kalman Filter: Particle Filters for Tracking Applications*, Artech House, 2004.
- [3] Z. Chen, "Bayesian Filtering: From Kalman Filters to Particle Filters, and Beyond," tech. rep., McMaster University, 2003.
- [4] B. Blumenfeld, *Operations Research Calculations Handbook*, CRC Press, 2nd ed., 2009.
- [5] P. Cull, M. Flahive, and R. Robson, *Difference Equations: From Rabbits to Chaos*, Springer, 2005.
- [6] K. Rosen, ed., *Discrete Mathematics and Its Applications*, McGraw-Hill, 6th ed., 2007.
- [7] A. A. Abd. Rahni, E. Lewis, and K. Wells, "Characterisation of respiratory motion extracted from 4d mri," *SPIE Medical Imaging: Image Processing* **8669**, 2013.
- [8] R. Smith, A. A. Abd. Rahni, R. Jones, and K. Wells, "Recursive Bayesian Estimation for Respiratory Motion Correction in Nuclear Medicine Imaging," in *IEEE NSS/MIC Conf Rec*, 2012.