ESTIMATING AND RESOLVING UNCERTAINTY IN CARDIAC RESPIRATORY MOTION MODELLING

D. Peressutti[⋆]*, E.-J. Rijkhorst*[†‡]*, D. C. Barratt*[‡]*, G. P. Penney*[⋆]*, A. P. King*[⋆]

[⋆]Division of Imaging Sciences and Biomedical Engineering, King's College London, U.K.
[†]Dept. of Physics and Medical Technology, VU University Medical Center, Amsterdam, The Netherlands.
[‡]UCL Centre for Medical Image Computing, University College London, U.K.

ABSTRACT

We present a novel method for cardiac respiratory motion estimation in image-guided interventions. The technique combines a pre-procedure affine motion model and intra-procedure real-time images to estimate and correct for the respiratory motion of the heart. As well as making motion estimates, the model is able to quantify the uncertainty in these estimates. This uncertainty is resolved using a Bayesian approach based on a prior probability from the motion model and a likelihood term derived from the intraprocedure images.

The proposed method is validated using MR-derived motion fields and simulated 3D real-time echocardiography data for 4 volunteers and compared to 3 other motion estimation techniques. Our Bayesian approach shows improvements in respiratory motion estimation for each volunteer of 7.0%, 4.6%, 7.7% and 5.3% respectively, compared to the use of a motion model only.

***Index Terms*—** Respiratory motion, modelling, MRI, echocardiography, Bayesian estimation

1. INTRODUCTION

The aim of image-guided cardiac interventions is to provide the cardiologist with accurate real-time information about the location of catheters and instruments inside the heart. This is achieved by combining intra-procedure real-time imaging such as echocardiography (echo) or X-ray, with a pre-procedure roadmap of the subject's heart, typically derived from magnetic resonance (MR) or computed tomography images. Respiratory motion can cause a misalignment between the static roadmap and the underlying anatomy. Such motion is often assumed to be periodic, but is in fact more complex than this, and varies with intra-cycle variability (i.e. the difference between exhalation and inhalation, known as the hysteresis effect) and inter-cycle variability (i.e. the difference in the motion between one cycle and another)[1].

Subject-specific motion models have been proposed as a solution to the problem of respiratory motion [2]. In such techniques, a surrogate signal of the true motion is acquired at the same time as the images that will be used to form the model, and the same or a related surrogate is acquired during the intervention. The motion model estimates the 3D motion of the heart as a function of the surrogate value. However, motion models are limited by the low number and nature of the surrogates that can be used, either in the pre-procedure or intra-procedure environment (e.g. respiratory bellows, MR-navigator echoes, spirometry, chest movement). Single surrogate models cannot describe the variability of the motion, and represent the motion during an average cycle only. Furthermore, errors and uncertainty in measuring the surrogates to form and to apply the model will cause errors in the estimated motion.

Therefore there is interest in using real-time imaging data to refine the motion estimates. Due to its high temporal resolution, non-invasive nature and low cost, echo is an attractive imaging technique to provide real-time information about the motion of the heart [3]. In [4] and [5], echo images alone were used to estimate respiratory motion. The results in [5] showed that this approach was not accurate or robust enough to produce reliable motion corrections, especially in areas of the heart not covered by the images' field of view.

In order to combine the strengths of respiratory motion models and real-time imaging data, [5] proposed a technique in which motion was estimated using echo images, but the range of motions permitted was constrained by a motion model. However, even though improvements in estimation were achieved, the method could not adapt to breathing variability nor did it have any measure of confidence in the motion estimates produced.

This paper presents a novel technique for respiratory motion estimation based on a combination of a motion model and echo images. The main novelty lies in using a motion model that incorporates a measure of confidence in its estimates, allowing it to adapt to variable breathing patterns according to the information present in the imaging data. This concept of uncertainty has been previously utilised in registration of statistical shape models to limited field of view data [6]. Here, we apply a similar concept to uncertainty in respiratory motion estimates produced by motion models. To evaluate the technique proposed, we use simulated echo images and compare our technique with estimates made using model-only [2] and image-only techniques [5] as well as the combination of model-image data without a confidence measure [5].

2. METHOD AND MATERIALS

An overview of the approach and its validation are given in Figure 1. In this paper, we test our technique using dynamic 3D echo images simulated from a segmented high resolution MR image combined with realistic motion fields derived from dynamic MR images. Furthermore, in order to have a gold standard set of motion fields for validation purposes, we build the motion model from an 'artificial' MR dataset produced by warping a reference exhale dynamic MR image using the estimated non-rigid motion fields.

For each subject, a high resolution 3D MR image and a sequence of low-resolution dynamic MR images were acquired. The

This work was funded by EPSRC programme grant EP/H046410/1. Erik-Jan Rijkhorst undertook this research whilst funded by EPSRC grant EP/F025750/1. Dean Barratt is funded by a Royal Academy of Engineering/EPSRC Research Fellowship.

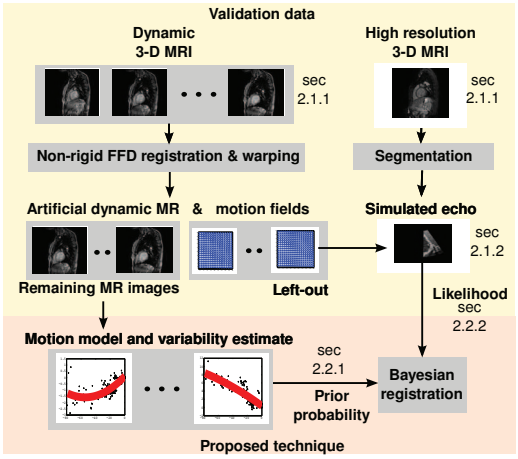


Fig. 1. Overview of proposed technique and its validation using a leave-one-out test. Real dynamic MR images are non-rigidly registered to a common reference, and warped using the resulting motion fields to create artificial, but realistic, dynamic MR images with known gold standard motion fields. Each artificial MR image is left out in turn. The non-rigid motion field of the left out dynamic is combined with a segmented high resolution MR image to produce a simulated echo image. The remaining dynamic MR images are affinely registered and used to form the motion model. The model is combined with the echo image in the proposed Bayesian approach.

dynamic images were near real-time and ECG-gated, so the dynamic sequence represents the motion of the heart due to breathing only. The superior-inferior displacement of the dome of the right hemidiaphragm was chosen as the surrogate signal, simulating a real 1-D MR navigator echo. The dynamic image having the highest surrogate value was selected as the end-exhale reference image.

Using a hierarchical free-form deformation (FFD) algorithm [7], each dynamic image was registered to the reference image forming a gold-standard set of motion fields. Combining a segmentation of the high resolution MR image and the non-rigid motion fields, a 4D set of simulated echo images was created. This forms the real-time imaging data used in testing our technique. In addition, the same non-rigid motion fields were used to warp the reference MR image to create the set of artificial MR images used to form the model.

The motion model was formed by performing affine registrations between the reference image and a masked version of the artificial dynamic MR images (the mask covered the four chambers and the major vessels of the heart). This leads to the formation of an 'average cycle' motion model that associates the values of the respiratory surrogate signal with a defined affine deformation [2].

Our technique makes use of the model within a Bayesian framework, using the information from the model as the prior knowledge of the most probable position (with associated uncertainty), and the echo images to form the likelihood of the estimated position. The final motion estimate is obtained by maximizing the posterior distribution. Details are provided in section 2.2.

2.1. Materials

2.1.1. MR acquisition

All the 3D MR images were acquired using a 1.5T Philips Achieva MR scanner. The MR sequences are as follows:

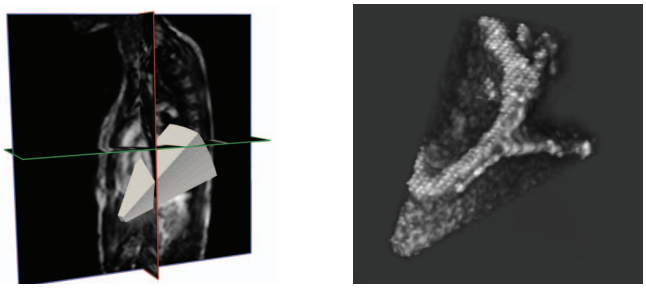


Fig. 2. Example of echo beam shape and positioning overlaid onto MR(left) and a 3D rendering of a simulated 3D echo image(right).

High resolution 3-D : 3-D balanced TFE, cardiac gated at late diastole, respiratory gated at end-expiration, typically 120 sagittal slices, $TR = 4.4ms$, $TE = 2.2ms$, flip angle= 90°, acquired voxel size 2.19x2.19x2.74 mm^3, reconstructed voxel size 1.37x1.37x1.37mm^3, acquisition window $\approx 100ms$, navigator window $5mm$, scan time approximately 5 minutes;

Dynamic 3-D : 3-D TFEPI, cardiac gated at late diastole, typically 20 slices, $TR = 10ms$, $TE = 4.9ms$, flip angle $= 20^\circ$, acquired voxel size 2.7x3.6x8.0mm^3, reconstructed voxel size 2.22x2.22x4.0 mm^3, TFE factor 26, EPI factor 13, TFE acquisition time $267.9ms$.

For this study, 4 volunteer datasets were processed. For each dataset, the dynamic sequences consisted of 300 images: 100 acquired during normal breathing, 100 during fast breathing and 100 during deep breathing. The short acquisition time put restrictions on the coverage of the dynamic scans, therefore the field of view covered most, but not all, of the four chambers of the heart.

2.1.2. Dynamic B-Mode echo simulation

In this work we used simulated 3D echo images to test our algorithm. This enabled us to have knowledge of the gold standard motion fields for validation purposes, which we would not have had using real echo data. The images were simulated using the realistic fast ultrasound propagation model proposed in [8]. The method requires maps of the 3D acoustic impedance and absorption coefficients corresponding to different tissues. These were obtained by manually segmenting the high resolution MR image into blood pool, myocardium, lungs and liver. The values of acoustic impedance and absorption coefficients were assigned to the segmented regions according to [8]. To simulate a realistic speckle noise, two different textures, one for the myocardium and one for the other tissues, were used. These textures were obtained using a texture quilting technique [9] applied to real 3D echo patches of the myocardium and the blood pool, acquired from a volunteer.

After transforming the high-resolution MR volume to the coordinate system of the dynamic reference MR image using a rigid-body transformation, the segmentation was warped to each dynamic inhale position using the FFD motion fields from the non-rigid registrations. The acoustic impedance and absorption maps were then assigned and the dynamic echo images simulated. A trans-apical acoustic window and a realistic echo beam shape were adopted with a field of view (FOV) that is typical of images acquired in real-time from a modern echo machine (azimuthal angle = 30°, inclination angle = 30°, radial depth = $130mm$). The images covered most of the

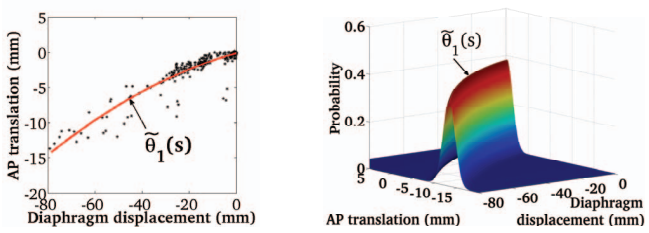


Fig. 3. Example of model and associated uncertainty for the 1^{st} affine parameter (anterior-posterior translation): model fit (left) and corresponding Gaussian distribution (right). The distribution represents how the data points (i.e. motion estimates for each dynamic MR image) vary from the fitted curve (the solid red line).

left ventricle in the exhale state. An example of a simulated image (voxel size 2.0x2.0x2.0mm^3) is shown in Figure 2.

2.2. Bayesian motion estimation

We adopt a Bayesian approach to combine prior knowledge in the form of a MR-derived motion model with real-time echo images.

The posterior probability describes the state of certainty about a given phenomenon, which in our case is the affine transformation θ that describes the respiratory motion of the heart. Bayes' law states that the posterior probability is proportional to the product of the prior probability and the likelihood,

$$p(\theta|I,s) = \frac{p(I|\theta,s)\cdot p(\theta|s)}{p(I|s)}. \quad (1)$$

The prior probability $p(\theta|s)$ describes the degree of certainty of the motion before the echo images are considered and is determined by the motion model and the surrogate value s. The likelihood $p(I|\theta,s)$ encompasses the information carried by the new data, namely, by the dynamic echo imaging data I.

As a simplification, $p(I|s)$ is assumed to be uniform. The desired motion estimate $\hat{\theta}$ is obtained in a *maximum a posteriori* manner,

$$\hat{\theta} = \arg\max_{\theta,s}\{p(\theta|I,s)\}. \quad (2)$$

The surrogate value was optimised first and a hill climbing optimisation algorithm was then used to estimate $\hat{\theta}$. We now describe the prior and likelihood terms in more detail.

2.2.1. Prior Probability

The prior probability is derived from the MR-derived motion model. This is formed from affine registrations between the reference dynamic MR image and the artificial dynamic MR images. This process results in 12 affine parameters for each dynamic image. 2^{nd} order polynomial curves are fitted in a least squares sense to each affine parameter as a function of the respiratory surrogate signal [2]. Therefore, given a value s of the surrogate acquired in the intra-procedure setting, the model outputs a twelve parameter affine transformation, which we denote $\tilde{\theta}(s)$.

In order to describe the uncertainty of the model's motion estimates, a Gaussian distribution is fitted to the variation of each affine parameter away from the model estimate. An example is shown in Figure 3 for the 1^{st} affine parameter (i.e. anterior-posterior translation). The peak of the Gaussian corresponds to the fitted value $\tilde{\theta}_i(s)$ while the variance is a function of the residuals of the fitting method.

The prior probability is given by the product of the twelve Gaussian distributions,

$$p(\theta|s) = \prod_i \frac{1}{\sqrt{2\pi\sigma_{p_i}^2}} e^{-\frac{\left(\theta_i - \tilde{\theta}_i(s)\right)^2}{2\sigma_{p_i}^2}}, \quad i = 1,\ldots,12. \quad (3)$$

where σ_{p_i} is the standard deviation of the residuals of the model fit for the i^{th} parameter.

2.2.2. Likelihood

To express the information provided by the echo data in probabilistic terms, we used again a Gaussian distribution to describe the similarity between a dynamic echo image I and the affine transformation of an end-exhale echo image, $T(I_{ref},\theta)$,

$$p(I|\theta,s) = \frac{1}{\sqrt{2\pi\sigma_l^2}} e^{-\frac{\left(NCC(T(I_{ref},\theta),I)-\mu\right)^2}{2\sigma_l^2}}, \quad (4)$$

where $NCC(T(I_{ref},\theta),I)$ is the value of the normalized cross correlation between $T(I_{ref},\theta)$ and the echo inhale image I, $\mu = 1$ and σ_l^2 is the likelihood variance. This was determined empirically and the value used for all datasets was 0.0005.

2.3. Validation

For validation of the method proposed, three other motion estimation approaches were implemented and compared with our Bayesian technique :

model-only, given the surrogate value s, the estimate is $\tilde{\theta}(s)$;

image-only, as in [5], the estimate is given by registering the echo images only, without the constraints of a motion model. However, unlike the real situation we also provided $\tilde{\theta}(s)$ as a starting estimate for the optimization algorithm. This was done because the small field of view and the even smaller overlapping region of $T(I_{ref},\theta)$ and I made the optimisation unstable. A hill climbing optimisation algorithm was used;

simple model and image, using a method similar to that described in [5], the estimate is obtained by maximizing $NCC(T(I_{ref},\theta),I)$ but constraining θ to lie on the model, that is

$$\hat{\theta} = \tilde{\theta}(\hat{s}), \quad \hat{s} = \arg\max_s\{NCC(T(I_{ref},\tilde{\theta}(s)),I)\},$$

without any measure of uncertainty.

The models used in all the methods compared were formed using a leave-one-out method, i.e. for each dynamic image tested, the model used was formed from all dynamic images apart from the tested image.

To assess the accuracy of the different approaches, a set of K landmarks ($\approx$ 35,000) was positioned on a regular Cartesian grid within a manually delimited elliptical mask covering the four chambers and the major vessels of the heart in the end-exhale reference image, $\vec{x}_k, 1 \le k \le K$. Using the ground-truth motion fields G, obtained by means of the FFD non-rigid registration, the set of landmarks was warped to each of the H dynamic MR acquisition inhale phases, $H = 300 - 1$ for each dataset. Denoting the affine transformation of each of the techniques compared as T, the error in estimation was computed as the median of $\|G_h(\vec{x}_k) - T_h(\vec{x}_k))\|, \forall k, \forall h$. The median was adopted because of the skewness of the error distribution (see Figure 4).

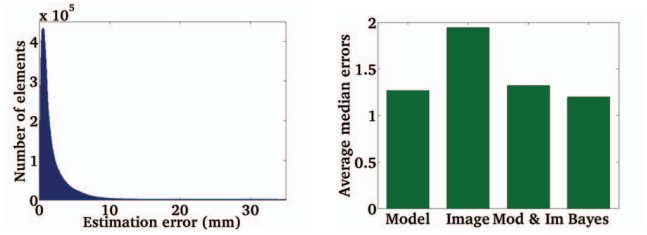


Fig. 4. Example of error distribution (left) and average values of median errors over the 4 volunteers (right).

	Model (mm)	*Image* (mm)	*Model & image* (mm)	*Bayes* (mm)
Vol. A	1.219	1.978	1.374	1.134
Vol. B	1.366	2.179	1.344	1.303
Vol. C	1.078	1.805	1.129	0.995
Vol. D	1.402	1.819	1.465	1.328
Average	1.266	1.945	1.328	1.190

Table 1. Results of the validation. For each volunteer, the median value of the errors is computed. Our approach shows the lowest error in estimation over the methods compared for each volunteer.

3. RESULTS

The results of the validation are shown in Table 1 and Figure 4. For all volunteers, the *image only* method had the highest error, in agreement with [5], while the *simple model and image* method showed higher prediction errors compared to the *model only* estimates for all volunteers but one. Our approach outperformed all the other techniques for each volunteer, with improvements over the *model only* estimation of 7.0%, 4.6%, 7.7% and 5.3% for volunteers A-D respectively. Mann-Whitney U tests between the *model only* and the *Bayesian* estimation errors showed a statistically significant difference for all volunteers ($p < 0.05$). Vol. B showed the lowest rate of variability among the different breathing patterns, which may explain the similar performance of three of the methods.

4. DISCUSSION AND CONCLUSIONS

We have presented a novel method for cardiac respiratory motion estimation in image-guided interventions. A MR-derived motion model and information from real-time echo images are combined in a Bayesian approach to estimate and resolve the uncertainty in motion estimation. To the authors' knowledge, this is the first work that has used such a measure of confidence to combine the information provided by the motion model and the echo imaging data.

Artificial MR images with gold standard non-rigid motion fields were used to build the motion model and estimate the uncertainty, while simulated echo images were obtained by combining the segmentation of a high resolution MR image and the gold standard motion fields. Even though the simulated images differ from real echo images, the validation approach implemented allowed us to compare the motion estimates with known and real non-rigid motion fields. We believe the use of real motion fields is crucial when validating this type of work as the methods proposed must be able to perform accurately when presented with real variations in breathing state.

The technique proposed was compared to *image only* [5], *model only* [2] and *simple model and image* [5] approaches and showed the lowest median error for each of the 4 volunteers considered. Compared to the *model only* estimates, improvements in estimation ranging from 4.6% to 7.7% were achieved.

Future work will involve testing our technique using real 3D echo images. Furthermore, experiments will be performed using different FOVs and probe locations to investigate the effects these have on the performance of the technique. Our ultimate aim is to use the motion model to guide the echo acquisition during the intervention to maximise the information gained about the variability of the breathing motion. The use of a respiratory bellows as the surrogate signal will be investigated because it can be easily acquired in the pre- and intra-procedure settings.

Although these are preliminary results, our technique has the potential to improve estimates of cardiac respiratory motion in image-guided interventions. For many cardiac procedures, even small estimation improvements are worthwhile. A further application would be to provide the cardiologist with a confidence indicator in the motion estimation, quantifying the contribution provided by the prior knowledge of the model and the real-time information from 3D echo images. Finally, this work focused on cardiac respiratory motion, but could be applied to other organs in the chest affected by respiratory motion.

5. REFERENCES

[1] J.M. Blackall, S. Ahmad, M.E. Miquel, J.R. McClelland, D.B. Landau, and D.J. Hawkes, "MRI-based measurements of respiratory motion variability and assessment of imaging strategies for radiotherapy planning," *Physics in Medicine and Biology*, vol. 53, pp. 4147–4169, 2006.

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

[3] V. Grau, C. Szmigielski, H. Becher, and J.A. Noble, "Combining apical and parasternal views to improve motion estimation in real-time 3D echocardiographic sequences," in *5th IEEE ISBI 2008*, 2008, pp. 516 –519.

[4] W. Wein, J.-Z. Cheng, and A. Khamene, "Ultrasound based respiratory motion compensation in the abdomen," in *MICCAI 2008 Workshop on Im Guid and Comp Ass for Soft-Tissue Interv*.

[5] A.P. King, C. Jansen, K.S. Rhode, D. Caulfield, R.S. Razavi, and G.P. Penney, "Respiratory motion correction for image-guided cardiac interventions using 3-D echocardiography," *Medical Image Analysis*, vol. 14, pp. 21–29, 2010.

[6] R. Blanc, E. Syrkina, and G. Székely, "Estimating the confidence of statistical model based shape prediction.," *Proc of IPMI*, vol. 21, pp. 602–613, 2009.

[7] C. Buerger, T. Schaeffter, and A.P. King, "Hierarchical adaptive local affine registration for fast and robust respiratory motion estimation," *Medical Image Analysis*, vol. 15, pp. 551–564, 2011.

[8] E.-J. Rijkhorst, D. Heanes, F. Odille, D. Hawkes, and D. Barratt, "Simulating dynamic ultrasound using MR-derived motion models to assess respiratory synchronisation for image-guided liver interventions," in *Proc of IPCAI*, 2010, pp. 113–123.

[9] A.A. Efros and W.T. Freeman, "Image quilting for texture synthesis and transfer," in *SIGGRAPH, Association for Computing Machinery.*, ACM, Ed., 2001, vol. 1496, pp. 341–346.