View metadata, citation and similar papers at core.ac.uk

brought to you by CORE

provided by Radboud Repository

Radboud Repository

Radboud University Nijmegen

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.
http://hdl.handle.net/2066/132712

Please be advised that this information was generated on 2017-12-05 and may be subject to change.

Medical Image Analysis 18 (2014) 674–683

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

MRI to X-ray mammography intensity-based registration with simultaneous optimisation of pose and biomechanical transformation parameters

Thomy Mertzanidou [a,*], John Hipwell [a], Stian Johnsen [a], Lianghao Han [a], Bjoern Eiben [a], Zeike Taylor [b], Sebastien Ourselin [a], Henkjan Huisman [c], Ritse Mann [d], Ulrich Bick [e], Nico Karssemeijer [c], David Hawkes [a]

[a] Centre for Medical Image Computing, University College London, Gower Street, WC1E 6BT London, UK
[b] Department of Mechanical Engineering, The University of Sheffield, Mappin Street, S1 3JD Sheffield, UK
[c] Diagnostic Image Analysis Group, Radboud University Medical Centre, P.O. Box 9101, 6500 HB Nijmegen, The Netherlands
[d] Department of Radiology, Radboud University Medical Centre, P.O. Box 9101, 6500 HB Nijmegen, The Netherlands
[e] Department of Radiology, Charite Universitätsmedizin Berlin, 10117 Berlin, Germany

ARTICLE INFO

Article history:
Received 22 April 2013
Received in revised form 10 March 2014
Accepted 14 March 2014
Available online 27 March 2014

Keywords:
Breast imaging
Multimodal registration
FEM-based transformation model
Mammography

ABSTRACT

Determining corresponding regions between an MRI and an X-ray mammogram is a clinically useful task that is challenging for radiologists due to the large deformation that the breast undergoes between the two image acquisitions. In this work we propose an intensity-based image registration framework, where the biomechanical transformation model parameters and the rigid-body transformation parameters are optimised simultaneously. Patient-specific biomechanical modelling of the breast derived from diagnostic, prone MRI has been previously used for this task. However, the high computational time associated with breast compression simulation using commercial packages, did not allow the optimisation of both pose and FEM parameters in the same framework. We use a fast explicit Finite Element (FE) solver that runs on a graphics card, enabling the FEM-based transformation model to be fully integrated into the optimisation scheme. The transformation model has seven degrees of freedom, which include parameters for both the initial rigid-body pose of the breast prior to mammographic compression, and those of the biomechanical model. The framework was tested on ten clinical cases and the results were compared against an affine transformation model, previously proposed for the same task. The mean registration error was 11.6 ± 3.8 mm for the CC and 11 ± 5.4 mm for the MLO view registrations, indicating that this could be a useful clinical tool.
© 2014 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/3.0/).

1. Introduction

X-ray mammography (MG) is routinely used both as a screening and a diagnostic tool. DCE-MRI is often used as a complementary modality to MG, for instance in the following cases (Mann et al., 2008):

- Problem solving when findings from conventional imaging (MG or ultrasound) are inconclusive – for example lesions that are only visible on a single mammographic view, or multiple MR enhancing lesions – as DCE-MRI has been shown to have high sensitivity but poor specificity (Morrow et al., 2011). In the case of multiple enhancing lesions, MRI to MG registration could help disambiguate between them, by identifying their corresponding MG position.
- Screening in patients with dense breasts who are at increased risk and more frequent screening in patients at high genetic risk of cancer, in particular those known to be more radiosensitive (Leach et al., 2005).
- Monitoring and assessing the tumour response of patients treated with neoadjuvant chemotherapy.
- Staging of women with breast cancer; in particular for women with dense breasts and patients with histologic evidence of invasive lobular carcinoma.
- Determining the primary lesion, when this is not visible in mammography, for example axillary metastases with an occult primary tumour (de Bresser et al., 2010). In this instance the sensitivity of MRI is high (90%) but specificity is relatively low

* Corresponding author. Tel.: +44 2076790177.
E-mail address: t.mertzanidou@cs.ucl.ac.uk (T. Mertzanidou).

http://dx.doi.org/10.1016/j.media.2014.03.003
1361-8415/© 2014 The Authors. Published by Elsevier B.V.
This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/3.0/).

(30%) and could potentially be improved with correlation to MG.
- Imaging after breast conservative therapy, for example to evaluate possible residual disease or further evaluate suspected recurrence.

Many studies have compared the performance of MG and MRI to determine which modality is most appropriate at each step of the patient's care pathway. In short the relative merits of these and other breast cancer imaging modalities are wide ranging and complex (Hooley et al., 2011). Despite this, there have been no comprehensive studies, to our knowledge, that have assessed the clinical benefits of combining MG and MR imaging, to affect a clinical decision with respect to breast cancer care. This can be attributed to the lack of tools currently available to enable simultaneous, multi-modal image assessment. In their prospective view of breast cancer imaging for detection and diagnosis, however, Karellas and Vedantham (2008) suggest that: "From the technological perspective, in addition to improvements with each modality, we are likely to observe an increasing trend towards multi-modality systems that combine the relative strengths of each modality". This paper addresses this issue by describing a method to spatially correlate breast MG and MRI. Finally, another potential application of MRI to MG registration is its use to facilitate X-ray stereotactic biopsy for patients that have a second-look mammography following an MRI, for lesions that were initially considered mammographically-occult. In this case, identifying correspondences between MRI and MG could be a valuable tool to enable mammographically guided, rather than MRI-guided biopsies to be performed; the latter generally being more expensive, time consuming and not widely available (Heywang-Köbrunner et al., 2009).

Identifying corresponding regions between X-ray mammograms and MR images can be problematic, due to the differences in dimensionality and image appearance, and the large breast deformation between the two modalities. Women are lying prone in the MR scanner with their breasts pendulous, while during X-ray mammography acquisitions women are standing with their breast compressed between two plates. In addition, for MG, typically two images are acquired, one Cranio-Caudal (CC) and one Medio-Lateral Oblique (MLO) view. An automated MRI to X-ray registration algorithm would be a valuable tool that could help radiologists in the diagnosis and management of breast cancer.

Previously, authors have used feature-based techniques for this task (Behrenbruch et al., 2003;Marti et al., 2004). However the selection of corresponding, distinctive features from MR and X-ray images is particularly problematic for breast images, due to the lack of anatomically distinctive, internal landmarks.

A patient-specific FE modelling approach that simulates mammographic compression was initially proposed by Ruiter et al. (2006). This implementation used the breast outline for alignment rather than the intensities in the two images. The registration was performed in two stages: In the first step a plate compression was applied and in the second a breast outline alignment was achieved by applying additional displacements on the surface nodes of the breast model. Hopp et al. (2013) recently extended this approach by introducing one additional final step, where the rotation of the breast about the anterior-posterior axis was estimated using an intensity-based optimisation. Similarly, another FEM-based approach with a contact model was proposed (Lee et al., 2013), which also employed an iterative intensity-based registration framework. However, this was limited to a 2D rigid-body registration of the simulated projection to the X-ray mammogram.

Biomechanical models have also been used previously to simulate large mammographic compressions but were not applied to the calculation of MRI to X-ray correspondences (Samani et al., 2001; Pathmanathan et al., 2004; Chung et al., 2008). In all these approaches the material parameters of the breast tissue were taken from the literature, from studies on ex-vivo tissue samples. Han et al. (2012) proposed a method for in vivo parameter estimation, using a framework that incorporates an implementation of the FEM modelling on the Graphics Processing Unit (GPU). This approach can be further extended to FEM-based registration tasks, which are computationally expensive and hence prohibitively time consuming to perform with current commercial FE solver packages.

We have previously investigated the performance of an intensity-based framework using simpler transformation models, such as an affine transformation (Mertzanidou et al., 2012a) and a statistical deformation model learnt from biomechanical simulations (Mertzanidou et al., 2011). In this work, we are using the same iterative optimisation framework with a new patient-specific FEM-based transformation model and we further validate its performance on a larger dataset than our preliminary study (Mertzanidou et al., 2012b).

The original contribution of our technique, compared to other approaches that used biomechanical modelling for the same application, is the use of an intensity-based registration framework with an iterative update of both the model (non-rigid, biomechanical) parameters and the rigid transformation parameters; in total seven degrees of freedom. This is achieved using a transformation module that runs on the GPU (Taylor et al., 2009), providing shorter execution times than commercial packages and enabling several hundred simulations to be performed during each registration. In summary, previously proposed methods (Hopp et al., 2013; Lee et al., 2013) perform a registration via discrete, sequential steps, whereas our approach optimises all parameters simultaneously. Therefore it can be extended or modified to include more or different parameters without altering the registration framework. Also, the parameter space is better explored, as more parameters are used and a considerably larger number of their combinations is tested during optimisation. We apply and validate our method using both CC and MLO mammograms, and we perform a direct comparison to registration using an affine transformation (Mertzanidou et al., 2012a).

2. Methodology

In a conventional intensity-based MRI to X-ray image registration algorithm – such as the one we previously proposed (Mertzanidou et al., 2012b) – the inputs to the framework are the X-ray mammogram and X-ray attenuation volume, computed from the patient's MRI. During the iterative optimisation loop, the transformation parameters $p = \{p_1, p_2, \ldots, p_N\}$ are updated, so that the similarity S between the real X-ray mammogram MG_{real} and the simulated perspective projection Ψ (i.e. ray casting) of the X-ray attenuation volume V is maximised:

$$\begin{aligned} p &= \underset{p}{\operatorname{argmax}}(S(p)) \\ &= \underset{p}{\operatorname{argmax}}(S(MG_{real}(x), \Psi(T_p(V), x)) : x \in \Omega) \end{aligned} \tag{1}$$

where x is the coordinate of each pixel within the region of interest Ω of MG_{real} and T_p is the 3D transformation.

An overview of our proposed framework is shown in Fig. 1. A key difference of this approach compared to the conventional registration pipeline described briefly above is the simultaneous optimisation of both pose parameters θ and FEM parameters ϕ during registration. In this case, Eq. (1) now becomes:

$$\begin{aligned} p &= \{\theta, \phi\} = \underset{p}{\operatorname{argmax}}(S(\theta, \phi)) \\ &= \underset{\{\theta,\phi\}}{\operatorname{argmax}}(S(MG_{real}(x), \Psi(T_{\{\theta,\phi\}}(V), x)) : x \in \Omega). \end{aligned} \tag{2}$$

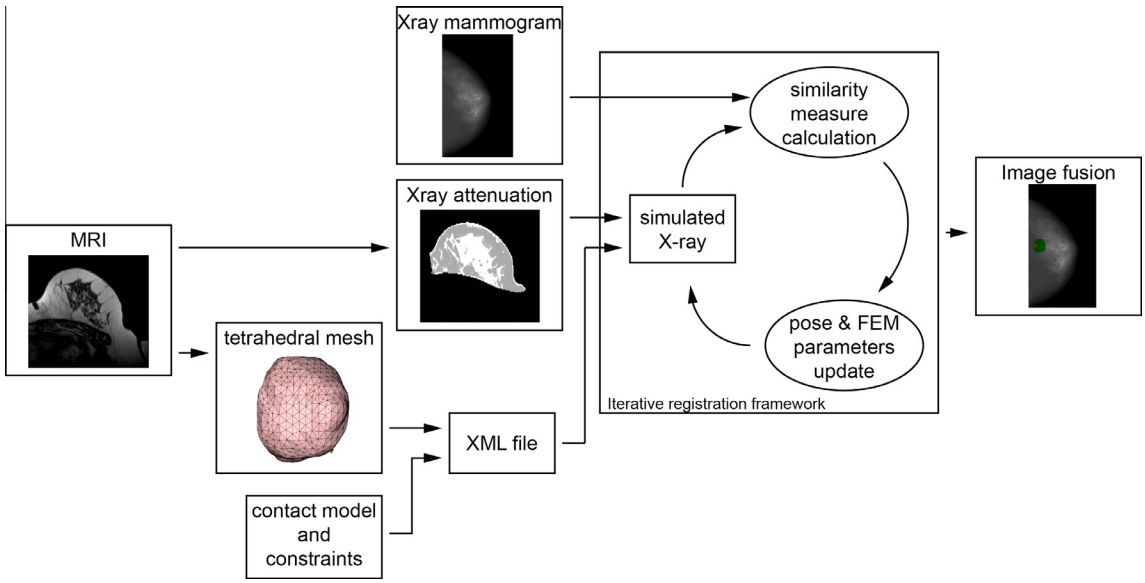

Fig. 1. Overview of the proposed registration framework.

This requires an additional input to the iterative scheme, which includes the geometry of the model, the material properties and the boundary conditions (including the dimensions and trajectory of the contact plates). In our implementation this information is stored in an XML file, as illustrated in the Fig. 1.

2.1. Patient-specific modelling

In our proposed technique, patient-specific biomechanical models are built from the pre-contrast MRI of the subject. Initially, we segment the breast volume from the background, using a simple region-growing algorithm, apply Gaussian smoothing and downsample the extracted binary mask to an isotropic volume of $[10 \times 10 \times 10]\ \mathrm{mm}^3$ resolution. This down-sampling simplifies the image for meshing purposes, reduces the computational cost of the FE solution and also produces smooth meshes, free from artefacts that can be caused, for example, from skin folding. Skin folding can occur, particularly for large breasts, due to contact with the MRI breast coil or in cases where the subject is clothed during scanning. If present the topology of the elements in this area can break down, as these undergo a large deformation. The surface mesh of the whole breast is extracted using a VTK[1] implementation of the marching cubes algorithm and the tetrahedral elements are extracted using the open-source software package TetGen.[2] A typical breast model of the ten used in this study consists of approximately 2,500 elements and 800 nodes. This choice lies approximately in the middle of the previously proposed range, that was used for the same registration task: the number of elements varied in the literature from one or two hundred (Ruiter et al., 2006; Lee et al., 2013) to tens of thousands (Hopp et al., 2013).

We use a nearly incompressible and hyperelastic neo-Hookean model (Han et al., 2012). This is transversely isotropic, to account for the reinforcement of biomechanical properties from fibre-like connective tissues in a preferred direction, as previously observed by Tanner et al. (2011). The strain-energy function, as in (Taylor et al., 2009), is given by the equation:

$$\begin{aligned}\Psi &= \Psi^{ISO} + \Psi^{TRANS} \\ &= [\frac{\mu}{2}(\bar{I}_1 - 3) + \frac{\kappa}{2}(J-1)^2] + \frac{\eta}{2}(\bar{I}_4 - 1)^2,\end{aligned} \tag{3}$$

where Ψ^{ISO} and Ψ^{TRANS} are the isotropic and anisotropic components respectively and $\bar{I}_1$ is the first principal invariant. $J = \det(F)$, where F is the deformation gradient. μ and κ are the shear and bulk moduli, $\bar{I}_4$ is the pseudo-invariant of the modified right Cauchy-Green deformation tensor, which defines the fibre orientation and η is an additional material parameter that controls its stiffness enhancement. η is incorporated into the transformation model described in the next section. In this study, we assume that the fibre enhancement is along the anterior-posterior direction, thus allowing the breast to expand more in the medial–lateral direction for a CC view compression, for example. The use of a non-linear instead of a linear model and the incorporation of an anisotropic behaviour are important for our application, as the compression is simulated using a contact model and therefore it is less constrained than when the displacements of the surface nodes are known a priori, as in (Tanner et al., 2011).

Regarding the different tissue types that can be considered, our approach assumes one homogeneous tissue type, rather than assigning different material properties, for example, to the fibroglandular tissue, fat, skin and pectoral muscle. The advantage of our simplified model is that it avoids convergence problems associated with performing dynamic FE simulations using multi-material models with the highly variable tissue material properties defined in the literature. In relevant previous work other authors have also proposed homogeneous tissue types (Chung et al., 2008; Hopp et al., 2012) and in experimental work published by Ruiter (2003), no significant effect on the results was observed when different tissue models were used.

Regarding the modelling of the pectoral muscle, we use a planar approximation for the interface between the breast and the pectoral muscle (the pectoral fascia) and allow the nodes to slide along this plane. This implementation is not expected to introduce significant errors into our modelling, as the pectoral muscle is not visible in the CC view mammograms. For the MLO views, the pectoral muscle is excluded from the area where the similarity measure is calculated, but this simplification is expected to be less accurate for the MLO view. An advantage of this approximation is the fact that it can avoid meshing problems due to topological irregularities in the segmentation of the chest wall when this intensity boundary is poorly defined in the MRI.

Finally, the plate compression is simulated using a frictionless contact model (Han et al., 2012). The main advantage of using a contact model instead of applying displacements on the surface

[1] http://www.vtk.org/.
[2] http://tetgen.berlios.de/.

nodes is that the interaction with the plates is modelled explicitly. Moreover, this method avoids artefacts on the breast surface that can occur if the displacements applied on neighbouring nodes are different. Regarding the modelling of friction between the compression plate and the breast skin, there is no experimental study that illustrates either the effect that this has on the breast deformation when mammographic compression is applied, or the friction coefficient that best describes the breast-plate interaction. Whilst friction is clearly present, in the absence of any relevant data, we assume that a frictionless model provides a reasonable first approximation to the interaction between the breast and the compression plates. We investigate the effect of this assumption in Section 3.2.

The geometry of the model is stored in an XML file, which is used as an additional input into the registration pipeline, described in detail in Section 2.3.

2.2. Transformation model

The transformation model consists of seven parameters which are iteratively updated during registration. Four of these account for the positioning of the breast before compression. More specifically these are:

- Two translations, t_x and t_y, within the plane perpendicular to the direction of the projection (*XY* plane).
- Two rotations, one for the rotation of the breast about the anterior–posterior axis *Y*, θ_y, (roll) and one about the superior–inferior axis *Z*, θ_z, (in-plane rotation).

Fig. 2 illustrates the coordinate system used for an example CC compression.

The remaining three transformation parameters control the material properties and the compression simulation of the FEM deformation. These are:

- Amount of compression, *D*, – constrained between: no compression (0%) and 90% of the maximum distance between the nodes in the direction of the projection.
- Stiffness anisotropy ratio, ρ, – constrained between $[0-512]$ (range taken from the literature (Tanner et al., 2011)).
- Poisson's ratio, ν, – constrained between $[0.45-0.499]$.

The compression is simulated using an equal and opposite displacement of both compression plates. The optimised parameter *D* is the distance between the two plates. The stiffness anisotropy ratio, ρ, is related to the parameter η in Eq. (3): $\rho = \eta/E$, where *E* is Young's modulus. In other words, ρ is the ratio of η, to underlying stiffness, *E*, and therefore the tissue is either isotropic (for $\rho = 0$) or transversely isotropic (for $\rho > 0$). In the latter case the tissue is stiffer in one specified direction than in others. Higher values of ρ correspond to a stiffer preferred direction. We assume that the breast tissue is homogeneous with Young's modulus $E = 4kPa$. As

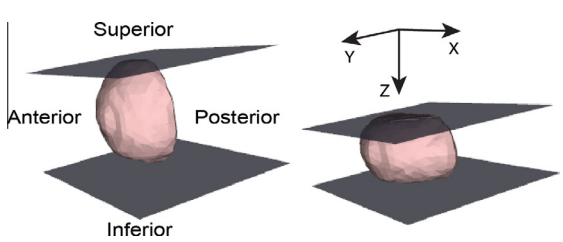

Fig. 2. Illustration of an example CC view compression of a mesh generated from a subject's MRI and the axes used for the FEM transformation geometry.

we apply displacements on the compression plates, rather than forces, and we are interested in the displacement distribution, rather than the stress distribution, the value of *E* is not critical for our application. The value has been chosen to lie within the range previously used for breast compression simulations (Han et al., 2012). Finally, the Poisson's ratio, ν, controls the volume change under compression.

The optimised parameters $p = \{t_x, t_y, \theta_y, \theta_z, D, \rho, \nu\}$ are part of a transformation module that is implemented using the Insight Toolkit,[3] without requiring the geometry model to be reloaded at each iteration of the algorithm. This implementation also provides the flexibility to use different similarity measures and optimisation techniques inside the iterative registration process described below.

2.3. Registration framework

Before registration, the MRI intensities are transformed to X-ray attenuation using the methodology previously described in (Mertzanidou et al., 2012a). In this approach, the MRI intensities are initially segmented into fat and fibroglandular tissue and are then mapped to a new X-ray attenuation volume. This new volume and the real X-ray mammogram are the inputs to the registration pipeline.

The breast volume is positioned above the detector and the distance between the X-ray source and the detector is fixed and extracted from the DICOM file of the mammogram ($f = 660$ mm). The initial translation parameters, t_x and t_y, are set such that the centre of mass of the volume is projected onto the centre of mass of the real mammogram. This provides a good initial position for registration, which is important for ensuring the optimisation scheme converges to a global minimum, and it does not require manual interaction. The amount of compression is initialised to a 50% plate displacement, the Poisson's ratio to $\nu = 0.498$ and the stiffness anisotropy ratio to $\rho = 250$, which is the midpoint of the range specified in the previous section.

The rotation parameters are initialised to $\theta_y = 0^\circ$ and $\theta_z = 0^\circ$ for the CC view mammogram registrations, while for the MLO view the roll is set to $\theta_y = 45^\circ$, to account for the different direction of the projection. This initial roll angle is extracted from the mammogram's DICOM header. The in-plane rotation is set experimentally to $\theta_z = 30^\circ$, as the breast in MLO view mammograms appears to have an in-plane rotation. This value was determined empirically, by visual inspection.

To avoid resampling the 3D volume into the transformed position and then ray-casting using this new volume grid, the transformation is performed as the ray traverses the 3D grid of the undeformed, moving volume. More specifically, during the registration process we use ray-casting from the 2D target space through the 3D grid of the moving image and integrate the intensities of each transformed intersection of the ray with the 3D grid. For point x_i, the intersection of the ray with the volume grid of the moving image, the transformation is given by the equation:

$$T(x_i) = T_{2rigid}(T_{non-rigid}(T_{1rigid}(x_i))) \tag{4}$$

where $T_{1rigid}(x_i) = T_{\{t_x,t_y\}}(R_{\theta_z}(x_i))$ and $T_{2rigid}(x_i) = R_{\theta_y}(x_i)$. The non-rigid transformation $T_{non-rigid}$ is the interpolated displacement at the current position x_i and is computed by the FE solver at the current parameter position. Note the order of the transformations in the above equation, as the transformation is defined from the "compressed" to the "uncompressed" breast space.

At each iteration the model is transformed using a rigid transformation and an FE compression simulation, and it is projected into 2D using a perspective ray-casting projection. The similarity

[3] The Insight Segmentation and Registration Toolkit (ITK), http://www.itk.org.

measure used is normalised cross-correlation (NCC) and it is calculated over a region around the mammogram, that excludes the pectoral muscle in the MLO view. In our implementation, the user specifies a line to define the pectoral muscle boundary using manual interaction. In previous work (Mertzanidou et al., 2012a) we evaluated different similarity measures, and found that NCC generated the most accurate registrations. The use of NCC assumes a linear intensity relationship between our simulation (Mertzanidou et al., 2012a) and the real mammogram, therefore for digital mammograms we pre-process the raw DICOM data by computing the logarithm of the pixel intensities. The optimisation scheme is hill climbing. The value of each parameter p at iteration i is given by:

$$p^i = p^{i-1} \pm \frac{step}{w(p)} \tag{5}$$

where $w(p)$ is a scalar weight factor that controls the relative magnitude of the step size *step* for each parameter. At each iteration one parameter is updated, that which results in the largest increase of the similarity measure, at the current relative step size. The parameter p_i is only updated if the similarity increases and the *step* is decreased if the similarity does not improve using the current parameters. The optimisation terminates when the parameter *step* becomes smaller than a pre-defined value.

In our current implementation the algorithm requires approximately 2 h for each registration, on a single core, 64-bit machine, with a 2.8 GHz processor. A typical registration task usually converges within 30 iterations (approximately 420 simulations). The performance could be further optimised to include a GPU implementation of the ray-casting algorithm. Currently, for a given 3D deformation field provided by the FE solver, the combination of the ray casting and the transformation of the 3D volume requires approximately 18 s, and the FE simulation of a plate compression an additional 9 s.

3. Experiments

3.1. Validation of the proposed registration framework

For validation, we used clinical MRI and X-ray mammography data (both CC and MLO views) from ten patients. The images were acquired approximately at the same timepoint to avoid significant volume change of the fat and the fibroglandular tissue. We have used the pre-contrast T1-weighted images for processing. The MR images of six cases had a voxel size of $[0.7 \times 0.7 \times 1.3]\,\text{mm}^3$, two had $[0.7 \times 0.7 \times 2]\,\text{mm}^3$ and the remaining two had $[0.9 \times 0.9 \times 1]\,\text{mm}^3$ voxel size. The X-ray mammograms of eight patients had pixel sizes of $[0.1 \times 0.1]\,\text{mm}^2$, one was $[0.07 \times 0.07]\,\text{mm}^2$ and the last $[0.08 \times 0.08]\,\text{mm}^2$. All mammograms were resampled by a factor of 10 for registration, to more closely match the MRI resolution and reduce the computational cost of the ray-casting algorithm, used for projection of the 3D volume to 2D.

Eight of the above patients had lesions visible in both the MRIs and in the CC and MLO view mammograms. The annotations of these lesions were used as gold standard correspondences between the modalities. MRI annotations, based on the full dynamic contrast sequence, were performed using a sphere centred around the lesion, while X-ray annotations included either a disc or a free-form shape. Two patients had an MRI and X-ray compatible clip that was used as a known corresponding point between the modalities.

For each registration, the error was calculated as the 2D Euclidean distance between the centre of mass of the annotation/clip position in the X-ray mammogram and the centre of mass of the MRI annotation/clip position projected into 2D at the final registration position. We consider this metric more appropriate than an overlap measure for our application, as the size of the annotations (Table 1) can vary significantly both between different patient pathologies and between the two modalities, since they measure different physical properties of the tissue.

The detailed registration results for all cases are shown in Table 1, where our approach is compared against an affine transformation (Mertzanidou et al., 2012a). For the CC view the FEM transformation has a mean error of 11.6 ± 3.8 mm, which is comparable with the performance of the affine approach (10.6 ± 6.33 mm), while the standard deviation is lower, which indicates that the results of the FEM approach are more consistent. For the MLO view the FEM algorithm performs better with a mean error of 11 ± 5.4 mm, compared to 13.3 ± 5.9 mm for the affine transformation. Fig. 3 illustrates the error figures in two graphs, one for the CC and one for the MLO view. Table 1 and Fig. 3 show that our proposed FEM-based transformation model performs similarly for the CC and MLO view registrations, which was also observed in our previously proposed affine transformation.

In addition to the registration errors, we have included in Table 1 the lesion radius, r, for each patient measured from the original (uncompressed) prone MRI. The clip cases are excluded from this calculation. Comparing the lesion radii, r, with the registration errors, d, we can see that for six out of eight cases (75%) the MLO registration error satisfied $d < 2 \cdot r$. Similarly, we observe the same behaviour for five out of the eight (62.8%) CC view registration tasks. This indicates that there is an overlap of each lesion in the two modalities after registration in 69% of cases. Note that for this calculation we ignore any expansion of the lesion after the volume compression, which would increase the observed overlap.

Example registration results are shown for three patients with annotations in Figs. 4 and 5 and the two patients with clips in Fig. 6.

Table 1 indicates that four registration tasks out of twenty gave errors larger than 15 mm and one gave an error larger than 20 mm. The results for this patient are illustrated in Fig. 7.

Finally, we have investigated the correlation of the registration errors to different factors: the size of the MR volume, the breast density measured from the MRI, the lesion size measured from both the MRI and the X-ray mammograms and finally the distance of the lesion from both the pectoral muscle and the centre of mass of the mammogram, both measured using the CC views. Our results showed that there was no significant correlation with any of these factors, as all p values were $p > 0.05$. The lowest p value was $p = 0.06$ that corresponds to 94% confidence intervals for the correlation between the errors of the CC view registrations and

Table 1
Registration error, d, (in mm) of our FEM transformation method and comparison with an affine transformation (Mertzanidou et al., 2012a). The clip cases are patients p4 and p5. The last column, r, corresponds to the radius of the annotation (in mm) measured from the original (uncompressed) prone MRI.

	CC view		MLO view		r
	FEM	Affine	FEM	Affine	
p1	8.1	14.6	10.3	11.9	6.9
p2	7.3	13.5	12.6	7.2	12.4
p3	11.9	3.7	11.4	9.4	9.3
p4	6.9	9.9	14.5	7.7	–
p5	15.4	23.4	2.3	18.9	–
p6	12.9	4.6	15.0	20.3	3.7
p7	16.0	8.5	4.8	6.7	2.8
p8	12.1	9.5	5.9	12.4	3.7
p9	8.3	2.9	20.5	23.5	6.5
p10	17.2	15.4	13.2	15.3	10.3
Mean	11.6	10.6	11.0	13.3	6.9
Std	3.8	6.33	5.4	5.9	3.4

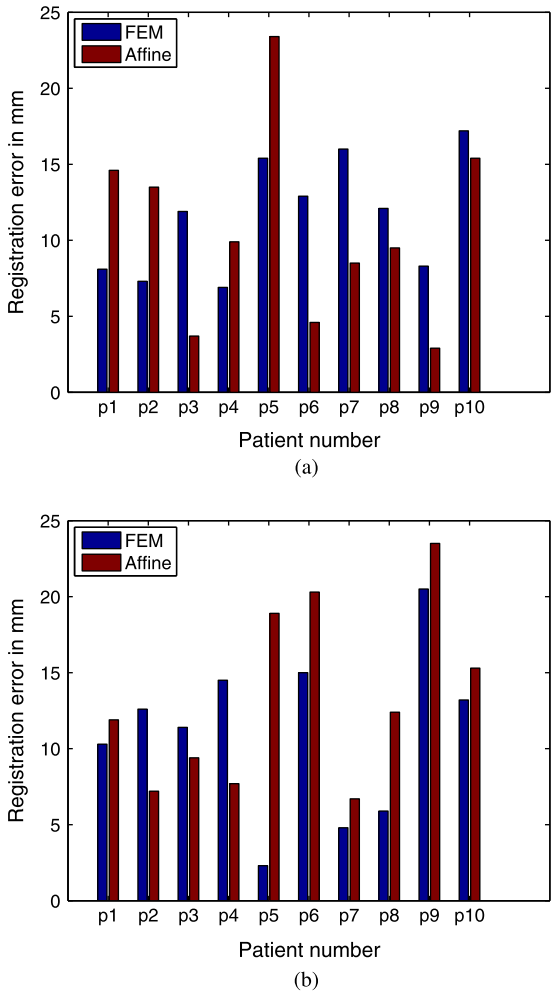

Fig. 3. The registration errors (in mm) by patient for the CC (a) and the MLO view (b).

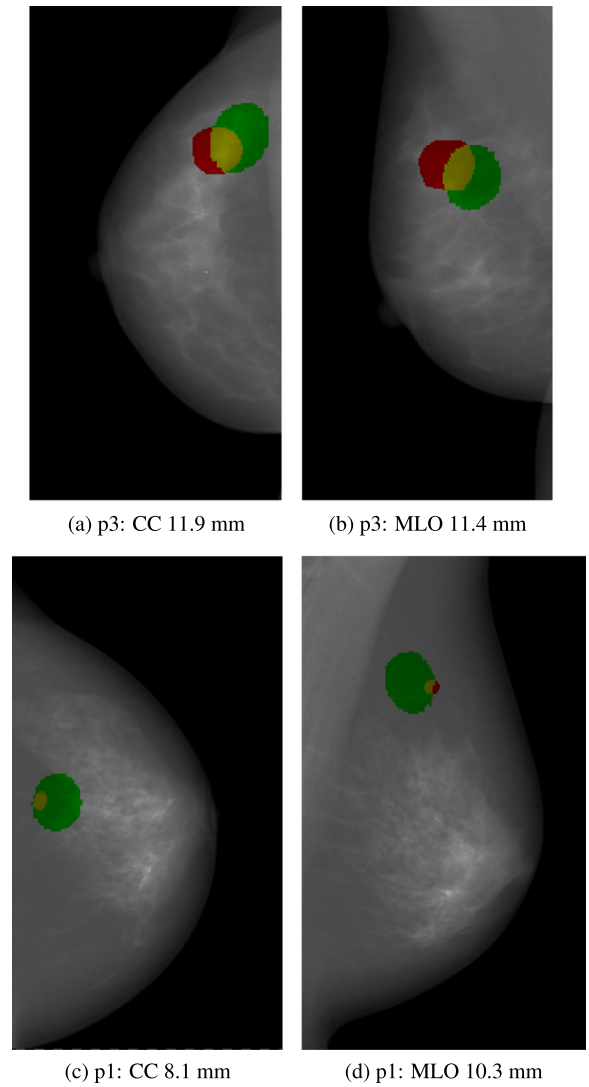

Fig. 4. Registration results for two patients, p1 and p3. The X-ray annotation is shown in red and the projection of the MR annotation in green; their overlap is yellow. Inevitably each modality can give different estimates of lesion size, but all cases show overlap. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the distance of the lesion from the pectoral muscle. This graph is shown in Fig. 8. This behaviour can be explained by the fact that lesions which are closer to the pectoral muscle have more constrained movement rather than those that are further away, and therefore the registration is more likely to give lower errors.

3.2. Breast modelling sensitivity study

As described above, the breast biomechanical modelling used in this study includes several simplifications that facilitate its use within the registration framework. The aim of the following experiments is to quantify the effect that these approximations have on the final compressed breast shapes. More specifically we conducted a sensitivity study to assess the influence of skin and friction on the simulations, as our model assumes a homogeneous tissue type and frictionless contact model with the mammographic plates.

In the next two sections we have used the breast models introduced in the previous section to perform a series of compression simulations. For these we have used the optimised registration parameters determined above. We then compared the differences in the nodal displacements between our model and that (i) with the addition of skin (Section 3.2.1) and (ii) with the addition of friction (Section 3.2.2).

3.2.1. Modelling skin

To model the skin, we added a layer of membrane elements to the model's external surface with a thickness of 2 mm and an elastic modulus that is two orders of magnitude stiffer than the breast tissue (Gefen and Dilmoney, 2007). The skin model is neo-hookean and incompressible and based on the work of (Bonet et al., 2000).

For each of the 10 patients, we have performed two compression simulations, one with our original model and one with the new model that includes the skin. The compression parameters used are those obtained for the corresponding CC view registrations without skin (or friction). Fig. 9 shows the differences in the nodal displacements when considering the two models. In addition to the 3D nodal displacement differences illustrated in blue, we have also calculated the differences on the axial plane, as these contribute most to the 2D errors, for a CC view compression. The results are summarised in Table 2.

3.2.2. Modelling friction

To quantify the effect of friction on the simulations, we have performed two sets of compression simulations, similar to the skin experiments, one with our original frictionless contact model and one with the new model that includes a friction coefficient

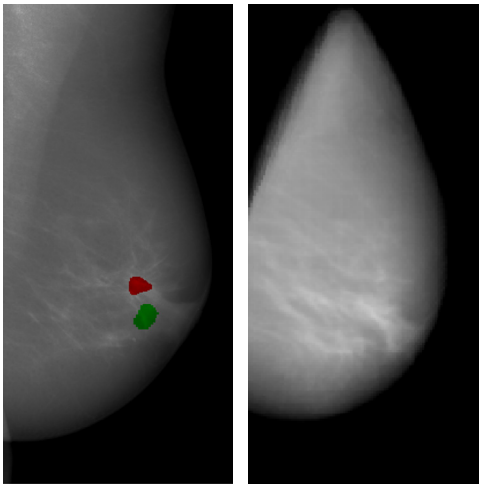

(a) p6: MLO 15 mm (b) p6: MLO sim. mammo

Fig. 5. Registration results for the MLO view of patient p6. As previously, the X-ray annotation is shown in red and the projection of the MR annotation in green. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$\mu = 0.3$ between the breast and the two mammographic plates (Hsu et al., 2011).

Fig. 10 shows the differences in the nodal displacements when using the models with and without friction. As previously, the compression parameters used were those obtained for the corresponding CC view registrations without friction (or skin), and we also show the nodal displacements on the axial plane. The results are summarised in Table 2.

4. Discussion

In Fig. 4 it is clear that the two modalities can give different estimates of the lesion size. In general, the projections of the MRI annotations appear larger than those of the X-ray images. This difference could be partially explained by the fact that the two modalities measure different physical properties of the tissue and by virtue of the fact that manual annotations are generally harder to perform accurately for 3D structures. Moreover, the size of malignant lesions on X-ray mammography is generally underestimated, while MRI measurements have been repeatedly shown to be larger, but more accurate. In our dataset, if we exclude the clip cases, the mean radius of the eight MRI annotations was 6.9 ± 3.4 mm and very similar to that of the X-ray mammograms, 6.9 ± 2.7 mm. When the lesions are deformed during registration, their radius can be reduced in the direction of the projection, due to compression, and consequently increased in the perpendicular plane. This is expected since we are using one homogeneous material for the FEM simulations and therefore the lesions are not modelled separately as stiff or rigid structures. In our experiments, we found that the mean radius of the MRI lesions when projected onto the mammograms following registration was increased to 11.3 ± 6.1 mm. Future work includes the use of different material parameters for the fat, the fibroglandular tissue, tumorous tissue, the skin (potentially) and the pectoral muscle.

In addition to the use of one homogeneous material, the breast modelling process used in our work includes certain approximations, such as the use of coarser resolution images to extract the meshes and the modelling of the pectoral muscle as a plane. Nevertheless, these methods contribute to a more automated approach of the breast modelling, which when combined with our registration framework, only requires one interactive step, which is the

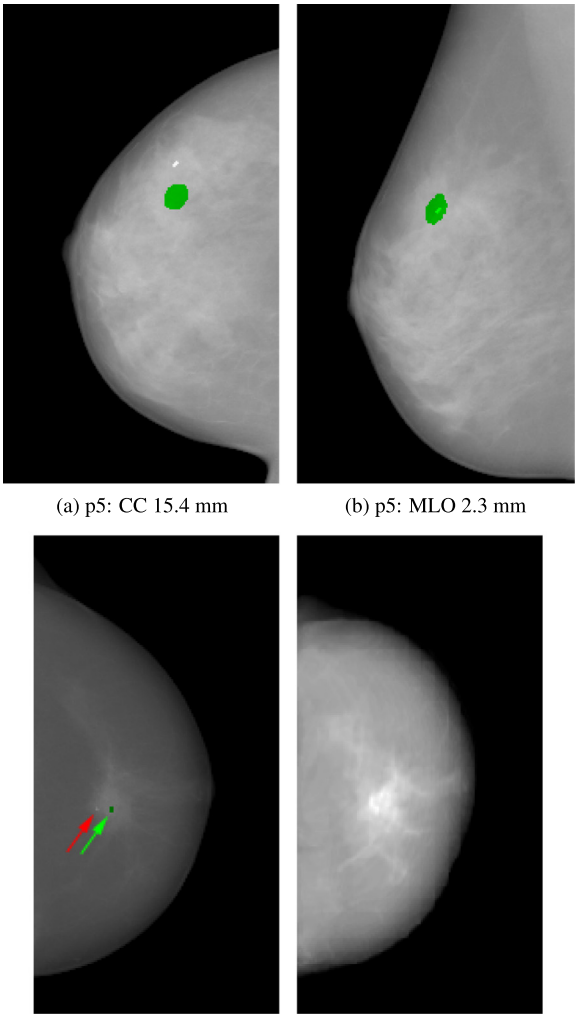

(a) p5: CC 15.4 mm (b) p5: MLO 2.3 mm

(c) p4: CC 6.9 mm (d) p4: CC sim.mammo

Fig. 6. Registration results for the two patients with MR and X-ray compatible clips, p4 and p5. The clip location on the X-ray mammogram is visible as the high intensity region (and a red arrow for p4). The MR annotation is shown in green. For the patient p4 we also show the simulated CC X-ray mammogram (d). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

pectoral muscle segmentation. This can be automated in future work, as automated methods exist (Wu et al. (2012),Gubern-Merida et al. (2011)) and our algorithm could be potentially integrated into clinics, providing a fully-automated patient-specific framework for MRI to X-ray mammography alignment. In our experiments we found that modelling the pectoral muscle using a semi-automated segmentation of the boundary between the pectoral muscle and the breast tissue, instead of a plane, does not improve registration accuracy. To investigate the effect of using this different breast geometry, we repeated all registration tasks using the new models and found that the registration accuracy was similar for the CC view (the error was 12.8 ± 8 mm, compared to that obtained with the original plane approximation of 11.6 ± 3.8 mm). However, for the MLO view using the segmented pectoral muscle shape produced an error of 16.1 ± 6.2 mm, which is considerably higher than when using the plane approximation (11.0 ± 5.4 mm).

The use of the breast modelling simplifications has an additional advantage. Although the dynamic FE solver that we use (Taylor et al., 2009) has the benefit of providing rapid computation times, large mammographic compressions, as those occurring in mammography, can cause convergence problems for dynamic solvers. The use of one homogeneous material and coarser

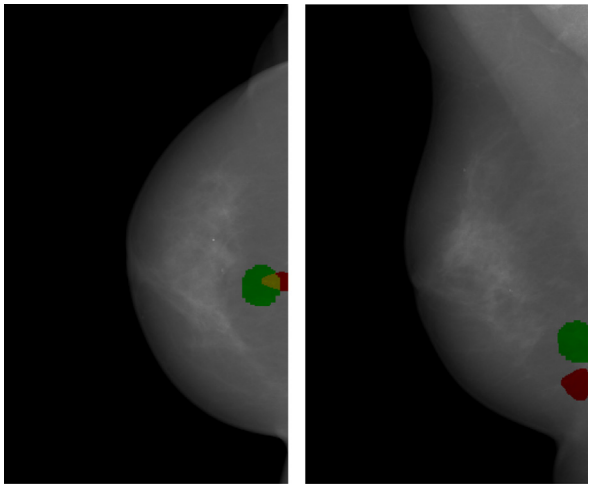

(a) p9: CC 8.3 mm (b) p9: MLO 20.5 mm

Fig. 7. Registration results for patient p9. The MLO view registration for this patient gives the largest registration error using our approach.

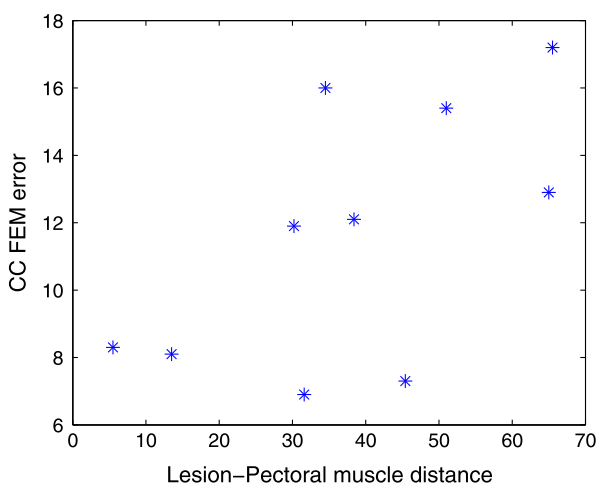

Fig. 8. Plot of the CC registration errors with respect to the distance of the lesion from the pectoral muscle, measured from the CC view mammograms. The values in both axes are expressed in mm.

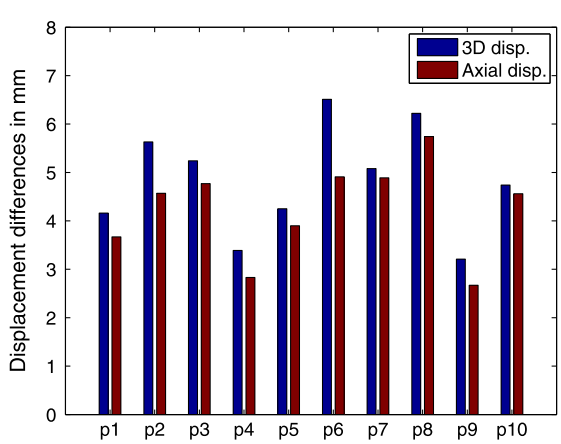

Fig. 9. Mean absolute differences in the nodal displacements (3D in blue and axial in red) using each breast model with and without skin. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2
Mean differences (in mm) in the nodal displacements (3D and axial) between using our original breast model and one considering: only skin, only friction, and both skin and friction.

	3D disp.	Axial disp.
Only skin	4.84	4.25
Only friction	2.3	2.06
Skin & friction	4.89	4.36

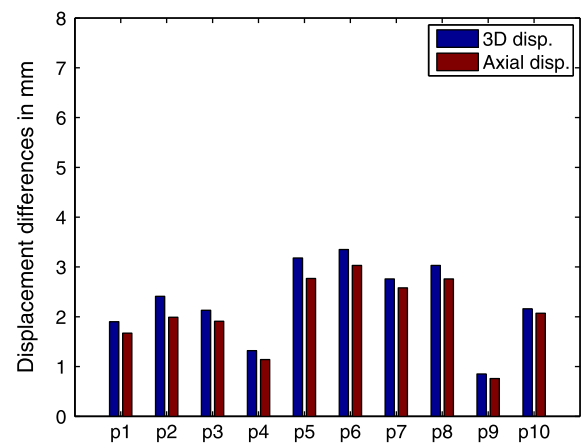

Fig. 10. Mean absolute differences in the nodal displacements (3D in blue and axial in red) using each contact model with and without friction. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

resolution meshes enables us to minimise the number of diverging simulations.

To quantify the effect of some biomechanical modelling aspects of our proposed method, we have conducted a sensitivity study in Section 3.2. Table 2 summarised the results of our experiments when using (i) the original model, (ii) the model including only skin, (iii) only friction and (iv) both skin and friction. We can see that incorporating friction into our model between the breast and the contact plates produces mean differences in the nodal displacements that are in the order of 2 mm, while the corresponding mean value for the skin is 4.25 mm for the axial plane. The results of the sensitivity analysis demonstrate that using a frictionless contact model is not expected to introduce significant differences in the breast compression simulations. Therefore ignoring friction provides a reasonable assumption for our application. Modelling the skin, instead of using one homogeneous tissue model, has a larger effect on the final nodal displacements of the breast models with a mean difference of approximately 4 mm. We found that the differences in the final breast compressed shapes are most noticeable in the periphery of the breast, as illustrated for patient 1 in Fig. 11. Future work could investigate further the effect of the skin on the final registration accuracy.

Regarding the optimisation of the biomechanical model parameters, we have optimised the position of the compression plates, rather than fixing their value according to the breast thickness value provided in the DICOM header, as these measurements can often be inaccurate. Diffey et al. (2008) found a mean difference of 10.8 mm between the maximum and minimum thickness measured in the anterior-posterior direction. In our experiments we found that the mean breast thickness after registration was 61 mm, while the mean value provided by the DICOM headers of the mammograms was 55 mm, which is within the error margins determined by (Diffey et al., 2008). There is a significant correlation between the optimised and the DICOM thickness values, with $p < 0.01$. For the remaining biomechanical model parameters, the

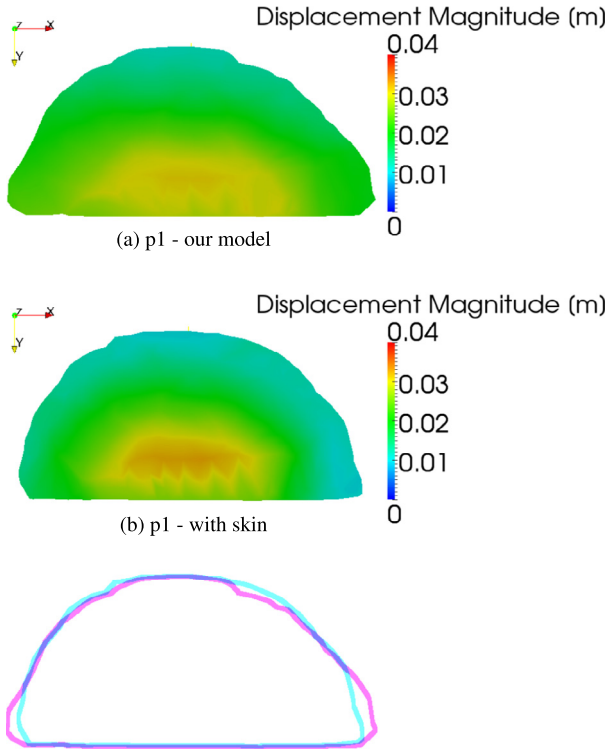

(a) p1 - our model

(b) p1 - with skin

(c) comparison of the contours

Fig. 11. Comparison of the final compressed breast shapes without (a) and with skin (b). The cross-section images correspond to the area of the breast just below the compression plate and they are colour-coded according to the magnitude of the displacements, shown here in metres. A comparison of the contours (c) shows the model without skin in magenta and the model with skin in cyan. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

optimisation results of the Poisson's ratio gave a mean value of $\nu = 0.4796 \pm 0.009$. The minimum value was 0.46 and the maximum 0.4956. Finally, the mean optimised value of the stiffness anisotropy ratio was $\rho = 250 \pm 14$, the minimum was 225 and maximum was 280.

Regarding registration accuracy and compared to other patient-specific FEM-based methods used for this task, quantitative results on clinical cases showed initially a mean error of 4.3 mm on 6 cases (Ruiter et al., 2006). A more recent semi-automated implementation of the same approach had an error of 11.8 ± 6.5 mm, on CC view mammograms only, of 11 patients (Hopp et al., 2012). Finally, in the latest automated implementation (Hopp et al., 2013), where the results were improved by optimising the roll angle using intensity-based registration, the mean error was 13.2 mm on CC view mammograms of 78 patients. However this methodology was not tested on MLO view mammograms, for which the registration is often more challenging, due to greater uncertainties regarding the positioning of the breast before compression and also the effect of the pectoral muscle on the compression simulation. Our previous affine transformation model (Mertzanidou et al., 2012a) gave a median error of 13.1 mm when tested on both CC and MLO view registrations (113 in total). However, it is difficult to compare the various techniques based on these results, unless all the algorithms are tested on the same data sets.

5. Conclusion

We have presented a framework for an intensity-based MRI to X-ray mammography registration using a novel iteratively updated FEM breast compression simulation. The results on twenty registration tasks from ten patients indicate that this could be a useful tool and potential aid to breast cancer detection and diagnosis.

Our transformation model provides a more physically realistic deformation of the breast than previous models proposed for this application but it is captured in only seven degrees of freedom. The pose of the breast, including two rotations, is defined by four degrees of freedom and the biomechanical model is specified via a further three degrees of freedom. Our method is the first to simultaneously optimise such a complete description of mammographic breast deformation and this, coupled with the low degrees of freedom, provides a tightly constrained optimisation strategy. This is illustrated in the validation experiments which provide a direct comparison between this method and use of a volume-preserving affine transformation model which has 11 degrees of freedom. The fact that both approaches produce similar results indicates that our new method is better constrained. In addition, incorporating this transformation model into an intensity-based registration framework, maximises the amount of information used by the optimisation, increasing the likelihood of the correct transformation being obtained. During the iterative optimisation loop, we observed that the breast outline is aligned initially and correspondences between structures internal to the breast are subsequently refined as the registration progresses.

Finally, future work includes further validation on a larger data set and investigation of the effect that a more accurate modelling of the breast has on the registration accuracy. This will include assigning different material properties to the fibroglandular, the adipose tissue, the tumour and the skin, as well as precise modelling of the boundary between the pectoral muscle and the breast. Some initial results are provided in this work, where a sensitivity study was performed to quantify the effect of some modelling approximations – regarding skin and friction – on the resulting breast compressions. This exploits our registration framework by enabling the influence of individual biomechanical (or other) parameters to be investigated, which in turn will lead to advances in our understanding of breast biomechanics.

Acknowledgements

This work was funded by the European 7th Framework Program grants HAMAM (FP7-ICT-2007.5.3, 224538) and VPH-PRISM (FP7-ICT-2011-9, 601040), and the EPSRC grants (EP/E031579/1) and MIMIC (EP/K020439/1). The authors would like to thank the Radboud University Nijmegen Medical Centre for the MRI and X-ray mammography data with annotations and the Charite Universitätsmedizin Berlin for the clip cases used in this study.

References

Behrenbruch, C., Marias, K., Armitage, P., Moore, N., English, R., Clarke, J., Brady, M., 2003. Fusion of contrast-enhanced breast MR and mammographic imaging data. Med. Image Anal. 7, 311–340.

Bonet, J., Wood, R., Mahaney, J., Heywood, P., 2000. Finite element analysis of air supported membrane structures. Comput. Methods Appl. Mech. Eng. 190, 579–595.

de Bresser, J., de Vos, B., van der Ent, F., Hulsewe, K., 2010. Breast MRI in clinically and mammographically occult breast cancer presenting with an axillary metastasis: a systematic review. Ejso 36, 114–119.

Chung, J., Rajagopal, V., Nielsen, P., Nash, M., 2008. Modelling mammographic compression of the breast. In: Proceeding of Medical Image Computing and Computer-Assisted Interventions, pp. 758–765.

Diffey, J., Hufton, A., Beeston, C., Smith, J., Marchant, T., Astley, S., 2008. Quantifying breast thickness for density measurement. In: International Workshop on Digital Mammography, pp. 651–658.

Gefen, A., Dilmoney, B., 2007. Mechanics of the normal woman's breast. Technol. Health Care 15, 259–271.

Gubern-Merida, A., Kallenberg, M., Marti, R., Karssemeijer, N., 2011. Fully automatic fibroglandular tissue segmentation in breast MRI: atlas-based approach. In: MICCAI Workshop on Breast Image Analysis, pp. 73–80.

Han, L., Hipwell, J., Tanner, C., Taylor, Z., Mertzanidou, T., Ourselin, S., Hawkes, D., 2012. Development of patient-specifc biomechanical models for predicting large breast deformation. Phys. Med. Biol. 57, 455–472.
Heywang-Köbrunner, S.H., Sinnatamby, R., Lebeau, A., Lebrecht, A., Britton, P.D., Schreer, I., 2009. Interdisciplinary consensus on the uses and technique of mr-guided vacuum-assisted breast biopsy (vab): results of a european consensus meeting. Eur. J. Radiol. 72, 289–294.
Hooley, R.J., Andrejeva, L., Scoutt, L.M., 2011. Breast cancer screening and problem solving using mammography, ultrasound, and magnetic resonance imaging. Ultrasound Quart. 27, 23–47.
Hopp, T., Baltzer, P., Dietzel, M., Kaiser, W., Ruiter, N., 2012. 2D/3D image fusion of X-ray mammograms with breast MRI: visualizing dynamic contrast enhancement in mammograms. Int. J. Comput. Assist. Radiol. Surg. 7, 339–348.
Hopp, T., Dietzel, M., Baltzer, P., Kreisel, P., Kaiser, W., Gemmeke, H., Ruiter, N., 2013. Automatic multimodal 2d/3d breast image registration using biomechanical fem models and intensity-based optimization. Med. Image Anal. 17, 209–218.
Hsu, C.M., Palmeri, M.L., Segars, W.P., Veress, A.I., Dobbins, J.T., 2011. An analysis of the mechanical parameters used for finite element compression of a high-resolution 3d breast phantom. Med. Phys. 38, 5756.
Karellas, A., Vedantham, S., 2008. Breast cancer imaging: a perspective for the next decade. Med. Phys. 35, 4878–4897.
Leach, M., Boggis, C., Dixon, A., Easton, D., Eeles, R., Evans, D., Gilbert, F., Griebsch, I., Hoff, R., Kessar, P., Lakhani, S., Moss, S., Nerurkar, A., Padhani, A., Pointon, L., Thompson, D., Warren, R., 2005. Screening with magnetic resonance imaging and mammography of a UK population at high familial risk of breast cancer: a prospective multicentre cohort study (MARIBS). Lancet 365, 1769–1778.
Lee, A., Rajagopal, V., Doyle, A., Nielsen, P.M.F., Nash, M.P., 2013. Breast lesion co-localisation between X-ray and mr images using finite element modelling. Med. Image Anal. 17, 1256–1264.
Mann, R., Kuhl, C., Kinkel, K., Boetes, C., 2008. Breast MRI: guidelines from the European society of breast imaging. Eur. Radiol. 18, 1307–1318.
Marti, R., Zwiggelaar, R., Rubin, C., Denton, E., 2004. 2D–3D correspondence in mammography. Cybernet. Syst. 35, 85–105.
Mertzanidou, T., Hipwell, J., Cardoso, M., Zhang, X., Tanner, C., Ourselin, S., Bick, U., Huisman, H., Karssemeijer, N., Hawkes, D., 2012a. MRI to X-ray mammography registration using a volume-preserving affine transformation. Med. Image Anal. 16, 966–975.
Mertzanidou, T., Hipwell, J., Han, L., Huisman, H., Karssemeijer, N., Hawkes, D., 2011. MRI to X-ray mammography registration using an ellipsoidal breast model and biomechanically simulated compressions. In: MICCAI Workshop on Breast Image Analysis, pp. 161–168.
Mertzanidou, T., Hipwell, J., Han, L., Taylor, Z., Huisman, H., Bick, U., Karssemeijer, N., Hawkes, D., 2012b. Intensity-based MRI to X-ray mammography registration with an integrated fast biomechanical transformation. In: International Workshop on Breast Imaging, pp. 48–55.
Morrow, M., Waters, J., Morris, E., 2011. Breast cancer 1 MRI for breast cancer screening, diagnosis, and treatment. Lancet 378, 1804–1811.
Pathmanathan, P., Gavaghan, D., Whiteley, J., Brady, M., Nash, M., Nielsen, P., Rajagopal, V., 2004. Location by simulating large deformations of the breast using a 3D Finite Element Model and nonlinear elasticity. In: Proceedings of Medical Imaging Computing and Computer Assisted Intervention, pp. 217–224.
Ruiter, N., 2003. Registration of X-ray mammograms and MR-volumes of the Female Breast based on Simulated Mammographic deformation. Ph.D. thesis, University of Mannheim.
Ruiter, N., Stotzka, R., Muller, T., Gemmeke, H., Reichenbach, J., Kaiser, W., 2006. Model-based registration of X-ray Mammograms and MR images of the female breast. IEEE Trans. Nucl. Sci. 53, 204–211.
Samani, A., Bishop, J., Yaffe, M., Plewes, D., 2001. Biomechanical 3-D finite element modeling of the human breast using MRI data. IEEE Trans. Med. Imag. 20, 271–279.
Tanner, C., White, M., Guarino, S., Hall-Craggs, M., Douek, M., Hawkes, D., 2011. Large breast compressions – observations and evaluation of simulations. Med. Phys. 38, 682–690.
Taylor, Z., Comas, O., Cheng, M., Passenger, J., Hawkes, D., Atkinson, D., Ourselin, S., 2009. On modelling of anisotropic viscoelasticity for soft tissue simulation: numerical solution and GPU execution. Med. Image Anal. 13, 234–244.
Wu, S., Weinstein, S., Kontos, D., 2012. Atlas-based probabilistic fibroglandular tissue segmentation in breast MRI. In: Medical Image Computing and Computer-Assisted Intervention, pp. 437–445.