Automatic correspondence between 2D and 3D images of the breast

Thomy Mertzanidou

A dissertation submitted in partial fulfillment of the requirements for the degree of **Doctor of Philosophy**

of

University College London.

Centre for Medical Image Computing University College London

2012

I, Thomy Mertzanidou, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Radiologists often need to localise corresponding findings in different images of the breast, such as Magnetic Resonance Images and X-ray mammograms. However, this is a difficult task, as one is a volume and the other a projection image. In addition, the appearance of breast tissue structure can vary significantly between them. Some breast regions are often obscured in an X-ray, due to its projective nature and the superimposition of normal glandular tissue. Automatically determining correspondences between the two modalities could assist radiologists in the detection, diagnosis and surgical planning of breast cancer.

This thesis addresses the problems associated with the automatic alignment of 3D and 2D breast images and presents a generic framework for registration that uses the structures within the breast for alignment, rather than surrogates based on the breast outline or nipple position. The proposed algorithm can adapt to incorporate different types of transformation models, in order to capture the breast deformation between modalities. The framework was validated on clinical MRI and X-ray mammography cases using both simple geometrical models, such as the affine, and also more complex ones that are based on biomechanical simulations. The results showed that the proposed framework with the affine transformation model can provide clinically useful accuracy (13.1mm when tested on 113 registration tasks). The biomechanical transformation models provided further improvement when applied on a smaller dataset. Our technique was also tested on determining corresponding findings in multiple X-ray images (i.e. temporal or CC to MLO) for a given subject using the 3D information provided by the MRI. Quantitative results showed that this approach outperforms 2D transformation models that are typically used for this task. The results indicate that this pipeline has the potential to provide a clinically useful tool for radiologists.

Acknowledgements

I would like to thank everyone that helped me complete my studies over the last few years. First, I would like to express my gratitude to my primary supervisor, Dave Hawkes, for giving me the opportunity to study in the group and for his advice, support and encouragement throughout my PhD work. His optimism and enthusiasm have always been valuable help. My deepest thanks also to John Hipwell, who has been my primary co-supervisor. He has provided invaluable advice throughout my studies and he has always been available for discussions both regarding the general direction of my work and the numerous day-to-day issues that were coming up. Moreover, his previous research and programming experience contributed considerably to the final outcome of my thesis. More specifically, I am most grateful for his help in the X-ray simulation, the ray-casting algorithm and the link between the GPU solver and the ITK toolkit. I would like to thank my Second supervisor, Simon Arridge, for the interesting and helpful discussions. I would also like to thank my MSc project supervisor, Simon Prince, who initially motivated me to pursue a PhD.

I feel very fortunate that my PhD work has been part of the EU project, HAMAM. In addition to being funded, I was given the opportunity to attend meetings and be part of fruitful discussions with the other project partners and get valuable feedback from clinicians. I would like to thank the Radboud University Nijmegen Medical Centre and the Charite Universitätsmedizin Berlin for providing the large number of clinical cases and MeVis Medical Solutions for the visualisation of one of my algorithms into the workstation prototype.

I am particularly grateful to all my co-authors for their contributions and all those in CMIC who have kindly shared their expertise with me. Jorge Cardoso for the tissue classification algorithm, Andrew Melbourne for the pectoral muscle segmentation, Lianghao Han for the breast modelling, Christine Tanner for the real compression data and Lily Zhang for her help in data processing.

Many thanks to colleagues and friends in CMIC and other departments at UCL for sharing coffees, lunches and drinks, that made my time at UCL enjoyable. Special thanks to Fred for his endless patience and support and to my close friends in Greece, Haroula, Katerina and Dimos, for sharing happy times during holidays and keeping in touch all this time.

Finally, I would like to thank my family. My parents for supporting me throughout all these years of studying, my brother and sisters for always being there for me and Lucas, Lydia and Theo for making me a very happy aunt over the last two years.

Contents

1	Intr	duction	16
	1.1	Breast Anatomy	16
	1.2	Breast imaging modalities in clinical practice	17
		1.2.1 X-ray mammogram acquisition	18
		1.2.2 DCE-MRI acquisition	20
	1.3	Problem statement	22
	1.4	Contributions	22
	1.5	Report structure	23
2	Lite	ature Review	24
	2.1	Image registration	24
	2.2	Intra-modality breast image registration	29
		2.2.1 Registration of X-ray mammograms	29
		2.2.2 Registration of DCE-MRI breast images	31
	2.3	Multimodal breast image registration	31
		2.3.1 MRI/X-ray mammography registration	31
		2.3.2 Other breast imaging modalities	35
	2.4	Multimodal registration of other deformable organs	35
	2.5	Learning soft tissue deformations	37
		2.5.1 Breast biomechanical modelling	37
		2.5.2 Statistical Deformation Models	37
	2.6	2D/3D registration of non-deformable organs	38
	2.7	Our approach	39
3	X-ra	y simulations from an MR volume	41
	3.1	Breast tissue classification	41
		3.1.1 Literature	41
		3.1.2 Our approach	44
	3.2	Perspective projection of the 3D volume	46
		3.2.1 Literature	46
		3.2.2 Our approach	47

Contents	
Contents	

	3.3	Experiments	48
	3.4	Discussion	50
4	MR	I to X-ray registration using an affine transformation	54
	4.1	Registration Framework	54
		4.1.1 Similarity measures	57
		4.1.2 Optimisation	60
	4.2	Experiments on simulated mammograms	62
		4.2.1 Recovering a known affine transformation	63
		4.2.2 Recovering a real breast deformation	66
	4.3	Experiments on real mammograms	69
		4.3.1 Validation using visual assessment on digitised mammograms	69
		4.3.2 Validation using radiologists' annotations on FFDM	69
	4.4	Discussion	76
5	MP	I to X-ray registration using an allinsoidal breast model	78
5	5.1	Breast Statistical Deformation Models	78
	5.2	An ellipsoidal breast model	80
	53	Mammographic compression simulation	81
	5.4	Building a deformation model using PCA	82
	5.5	Non-rigid registration framework	84
	5.6	Experiments	87
	5.7	Discussion	91
(MD		02
0	MIK.	Detient encoded historical simulations	92
	0.1	Patient-specific biomechanical simulations	92
		6.1.2 Metarial properties	92
		6.1.2 Compression simulation and boundary conditions	94
	62	Integration of the EEM simulations in the registration framework	95
	6.3	Experiments	101
	6.4	Discussion	101
7	Rela	nting findings between X-ray mammograms via an MR volume	105
	7.1	Analysis	105
		7.1.1 Ipsilateral registration between CC and MLO mammographic views	106
		7.1.2 Temporal mammogram registration	108
	7.2	Experiments	109
		7.2.1 Ipsilateral registration between CC and MLO mammographic views	109
		7.2.2 Temporal mammogram registration	111

	7.3	Discussion	111
8	Con	clusion	114
	8.1	Summary and conclusions	114
		8.1.1 X-ray simulations from an MR volume	114
		8.1.2 MRI to X-ray registration using an affine transformation	115
		8.1.3 MRI to X-ray registration using an ellipsoidal breast model	115
		8.1.4 MRI to X-ray registration using a patient-specific biomechanical model	116
		8.1.5 Relating findings between X-ray mammograms via an MR volume	116
	8.2	Future work	116
Lis	st of p	publications	121
Appendices 12			
A	Sum	mary of datasets	122
	A.1	MRI and X-ray film mammograms	122
	A.2	MRI breast compression volunteer data	122
	A.3	MRI and FFDM with known annotated lesions	122
	A.4	MRI and FFDM with clips	123
		A.4.1 Data with MR and X-ray compatible clips	123
		A.4.2 Data with X-ray compatible clips inserted after biopsy	123
B	Com	nputation of an X-ray attenuation volume from the MRI	124
С	Clin	ical validation on findings poorly or not visible on FFDM	126
Bi	ibliography 130		

7

List of Figures

1.1	Example MLO view mammograms of breasts with increasing breast density from left	
	(a) to right (d)	19
1.2	An example of a transverse slice coming from (a) a pre-contrast MRI and (b) the sub-	
	traction image of the pre-contrast from the post-contrast image, showing an enhancing	
	lesion (Invasive Ductal Carcinoma).	21
2.1	Illustration of the image registration problem. The goal is to calculate the transformation	
2.1	T(x), in order to align the two images. The image on the left stays stable during registra-	
	tion (fixed/target), while the image on the right is being resampled (moving/source)	25
2.2	Registration framework followed by Behrenbruch et al.	33
		00
3.1	X-ray simulation process.	42
3.2	(a) An example of a pre-contrast breast MRI and (b),(c) increasing magnification views	
	around the red cross that illustrate the partial volume effect. Due to the limited resolution	
	of the MR scanner, one voxel can belong to more than one different tissue types; in this	
	case both fibroglandular and fat.	43
3.3	Perspective projection geometry.	48
3.4	(a) Pre-contrast MRI, (b) X-ray attenuation volume using manual thresholding and	
	histogram-based classification, (c) using the EM-MRF algorithm, (d) Simulated X-ray	
	mammogram from the undeformed volume using manual thresholding, (e) using EM-	
	MRF. The two rows correspond to two patients. The red cross indicates the position of a	
	corresponding coordinate in each image	49
3.5	(a) Pre-contrast MRI of a patient and (b) the corresponding histogram of the intensities. $\ .$	50
3.6	Breast tissue classification results for the patient volume displayed in Figure 3.5, using	
	the k-means algorithm, the EM algorithm without and with the MRF regularisation	51
3.7	The X-ray attenuation volumes and the simulated mammograms produced using the EM	
	algorithm only and the EM-MRF. The pre-contrast MRI and the real mammogram are	
	given for comparison. Note that at this stage there is no registration or deformation	
	of the volume before projection. The simulated mammograms come from the original	
	undeformed MR volume	52

List of I iguics	List	of	Fig	ures
------------------	------	----	-----	------

3.8	(a) Pre-contrast MRI with strong bias field and (b) the X-ray attenuation volume. The	
	inhomogeneity correction was not successful for this case	52
4.1	Overview of the 2D/3D registration framework. The processes are illustrated in blue and	
	the data in red. The output of the registration is used for mapping between the MRI and	
	the X-ray image.	55
4.2	(a) Projection geometry used in the registration, showing the X-ray source, the breast volume V and the projection plane. Before registration the volume V is positioned such	
	that $f = 660mm$. This is extracted from the DICOM header of the real mammogram.	
	(b) Illustration of the initial position of the breast volume in V prior to registration for a	
	CC view	57
4.3	Illustration of the initial position before registration for the (a) CC and (b) MLO view of	
	one case. Left: Target mammogram in the registration space and right: overlaid with the	
	DRR outline shown in green, created at iteration 0 of the registration process	58
4.4	Illustration of the reprojection error used by Masutani et al. [Masutani et al., 1997] and	
	Hipwell et al. [Hipwell et al., 2003] for the evaluation of 2D/3D registration tasks	63
4.5	Illustration of the coordinate system used in the experiments with simulated mammo-	
	grams (CC views). The parameters of the affine transformation in Table 4.1 were chosen	
	according to this coordinate system.	65
4.6	Results of the simulated X-ray mammogram experiments for all patients (1 case per	
	patient). The images show the intensity differences between the target image and the	
	projection of the source, before and after registration for all the similarity measures tested	
	(MSD, NCC, GD and MI). For MI the results are shown using a chequerboard, as the	
	MRI projection is not a simulation of an X-ray and thus the intensities cannot be directly	
	subtracted from the target image. Each square alternates the intensities between the two	
	images: the target mammogram and the projection of the source after registration	67
4.7	Illustration of the real compression MR images acquired for one of the volunteers. The	
	volume on the left corresponds to the uncompressed breast. The rest of the volumes from	
	left to right have increased amount of compression applied from the lateral to the medial	
	direction.	68
4.8	Coronal slices of the data used for evaluation for 3 volunteers (a)-(c). From left to right	
	in each image: MRI slice before and after compression	68
4.9	Registration results on real data (5 cases, one per row). From left to right: projection	
	of the source volume before registration, after registration and real X-ray mammogram.	
	The red cross indicates the position of a corresponding coordinate in each image. \ldots .	70
4.10	Histogram of the registration errors calculated from 113 registration tasks. In red is	
	shown the median value (13.1mm) and in green the outliers.	71

4.11	Histograms of the registration errors of Figure 4.10, displayed individually for the CC and MLO views. Whilst the median error for MLO mammograms is marginally higher than for CC view mammograms (13.5mm versus 12.9mm as indicated in red) the distribution of these errors is broadly similar.	72
4.12	CC (on the left) and MLO (on the right) mammogram of patients 1 and 2. The X-ray mammogram annotation is shown in red and the projection of the MR annotation in green. (a) Patient 1 was diagnosed with an Invasive Ductal Carcinoma; the registration error is 8.12mm for the CC and 14.43mm for the MLO view. (b) Patient 2 diagnosed with Ductal Carcinoma In Situ; the registration error is 6.85mm for the CC and 2.63mm for the MLO view.	72
4.13	Registration result for CC (on the left) and MLO (on the right) mammograms of patient 3. (a) illustrates the raw mammograms and (b) the registered DRRs. A magnification view is given for both views on the bottom right corner of the raw mammograms. The evaluation in this case was done using the clip location. The clip is $2mm$ long as displayed in the magnification view; the location in the X-ray mammogram is illustrated by the high intensity region (red arrow) and the projection of the clip location in the MR is shown in green (green arrow). The registration error is $9.9mm$ for the CC and $7.7mm$ for the MLO view.	73
4.14	Two cases for which the registration error was high. CC view of patient 4 (error $28.45mm$) and MLO view of patient 5 (error $20.24mm$), both diagnosed with Invasive Ductal Carcinoma. The corresponding MR volumes and annotations are shown in Figure 4.15.	73
4.15	Corresponding MRIs and annotations of the patients illustrated in Figure 4.14. Patient 4 (left) and patient 5 (right).	74
4.16	Two MLO view mammograms that were excluded from validation, because the lesion annotation was too large.	75
4.17	The MRIs of two patients with large folding artifacts. These cases were excluded from validation.	76
4.18	Total histogram of the registration errors. In red are shown the cases that were excluded from validation due to artifacts. The histogram in green corresponds to the cases used for validation and is the same as in Figure 4.10. The total histogram of all cases is shown in blue.	77

 for by the affine registration	5.1	(a) MRI of patient p1; this is chosen as the space where all deformations are mapped for the SDM. (b) MRI of another patient p2 used for the construction of the SDM. Initially compressions are applied to all patients and then these are mapped to p1, where PCA will be used to extract the main modes of variation. (c) Deformed MRI of p1, after applying the displacement field of p2, mapped into the space of p1 using an affine registration. The anomalous shape of the breast in (c) is caused by the large difference in breast shape between p1 and p2 (particularly obvious in the coronal view) which is not compensated	
 5.3 (a) Image of the ellipsoid used for modelling and (b) its extracted mesh showing the ellipsoid between two plates, before compression	5.2	for by the affine registration. \dots To create an SDM model, the population-based approach uses the two mapping processes illustrated. M_1 is used to map the deformation fields from the individual patients in the training set to the common space of one selected patient. PCA analysis is performed in space B of this selected patient. M_2 is used to map the extracted components to the new patient space. The patient-specific methods do not use any mapping process, as spaces A, B and C are the same and they change for each individual case. The ellipsoidal method proposed in this chapter uses only M_2 , as space A and B are the space of the mean ellipsoid.	79 80
 5.4 (a) The ellipsoid image superimposed with the displacement field extracted from a compression simulation and (b) the corresponding mesh under compression	5.3	(a) Image of the ellipsoid used for modelling and (b) its extracted mesh showing the ellipsoid between two plates, before compression.	81
 5.5 Illustration of the ellipsoidal image when varying the first three principal components. For each one of the four squares the top row corresponds to the coronal view and the bottom to the axial. A superimposed pattern is used for each view for better visualisation of the deformations. The first component (b) shows the effect of varying the amount of compression, the second (c) the anisotropy ratio (notice change in extension in the axial view) and the third (d) the breast shape under compression (notice the breast shape in the coronal view). 5.6 (a) An example of an X-ray attenuation volume superimposed with the displacement field of the mean PCA component and (b) the same volume after the mean deformation. 5.7 Illustration of the percentage of the variation captured by <i>N</i> number of PCA components. Three components (red) explain 52% of the data, while ten components (green) explain 63%. 5.8 Registration results on 2 patients. For each case, the left image shows the real mammogram with the centre of mass of the MR annotation after alignment (red cross). The right image in each pair shows again the real mammogram, but as well as the projected MR annotation, the mask of the X-ray annotation is also shown in green.(a) p1 (error: 7.8mm) and (b) p2 (error: 14mm). 	5.4	(a) The ellipsoid image superimposed with the displacement field extracted from a com- pression simulation and (b) the corresponding mesh under compression.	82
 5.6 (a) An example of an X-ray attenuation volume superimposed with the displacement field of the mean PCA component and (b) the same volume after the mean deformation. 85 5.7 Illustration of the percentage of the variation captured by <i>N</i> number of PCA components. Three components (red) explain 52% of the data, while ten components (green) explain 63%	5.5	Illustration of the ellipsoidal image when varying the first three principal components. For each one of the four squares the top row corresponds to the coronal view and the bottom to the axial. A superimposed pattern is used for each view for better visualisation of the deformations. The first component (b) shows the effect of varying the amount of compression, the second (c) the anisotropy ratio (notice change in extension in the axial view) and the third (d) the breast shape under compression (notice the breast shape in the coronal view).	83
 5.7 Illustration of the percentage of the variation captured by <i>N</i> number of PCA components. Three components (red) explain 52% of the data, while ten components (green) explain 63%	5.6	(a) An example of an X-ray attenuation volume superimposed with the displacement field of the mean PCA component and (b) the same volume after the mean deformation	85
 5.8 Registration results on 2 patients. For each case, the left image shows the real mammogram with the centre of mass of the MR annotation after alignment (red cross). The right image in each pair shows again the real mammogram, but as well as the projected MR annotation, the mask of the X-ray annotation is also shown in green.(a) p1 (error: 7.8mm) and (b) p2 (error: 14mm)	5.7	Illustration of the percentage of the variation captured by N number of PCA components. Three components (red) explain 52% of the data, while ten components (green) explain 63%	88
· · · · · · · · · · · · · · · · · · ·	5.8	Registration results on 2 patients. For each case, the left image shows the real mam- mogram with the centre of mass of the MR annotation after alignment (red cross). The right image in each pair shows again the real mammogram, but as well as the projected MR annotation, the mask of the X-ray annotation is also shown in green.(a) p1 (error: 7.8mm) and (b) p2 (error: 14mm)	89

5.9	Registration results on 2 patients. For each case, the left image shows the real mammo-	
	gram with the centre of mass of the MR annotation after alignment (red cross). The right	
	image in each pair shows again the real mammogram, with the projected MR annotation	
	and the mask of the X-ray annotation, shown in green. (c) p3 (error: $9.7mm$) and (d) p4	
	(error: 5.7mm)	89
5.10	Cases for which the registration error was high. On the left are shown the X-ray mam-	
	mograms and on the right the MRIs of (a) p9 and (b) p10	90
6.1	Example of a patient's mesh showing folding of the skin that can cause problems in the	
	topology of elements when the breast undergoes a large compression.	93
6.2	(a) Example of a surface mesh extracted from one of the clinical cases. (b) The same	
	mesh, showing also a slice of the corresponding MRI. (c) Different view of image (b),	
	illustrating the approximation of the pectoral muscle as a plane	94
6.3	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	98
6.4	Illustration of a the initial position of the plates before compression to avoid intersection	
	of the breast model nodes and the plates during registration. d is the distance of each	
	plate from the centre of mass of the volume in the XZ plane and it is given by equation	
	6.1. The grey plates illustrate where the position of the plates would be set initially	
	if only the maximum distance on the Z axis was taken into account. By considering	
	also the X axis, the initial position of the plates is set to the location of the red plates.	
	Consequently, the maximum plate separation is $2 \cdot d$.	99
6.5	Registration results for two patients with identified lesions. The X-ray annotation is	
	shown in red and the projection of the MR annotation in green; their overlap is yellow.	103
6.6	Registration results for the two patients with MR and X-ray compatible clips. The clip	
	location in the X-ray mammogram is visible as the high intensity region (and a red arrow	
	for p4). The MR annotation in shown in green. For the patient p4 we also show the	
	simulated CC X-ray mammogram (d).	103
7.1	The process of determining corresponding findings between the CC and MLO view	
	mammograms of a patient, through the MRI. The patient shown is p3	107
7.2	The process of determining corresponding findings between temporal CC view mammo-	
	grams of a patient, via the MRI. The centre of the annotation in the mammogram t_1 (1)	
	is mapped onto the mammogram t_2 (5), illustrated as green. The mammogram acquired	
	at t_2 contains a clip, illustrated as red. Their overlap is yellow. The patient shown is p5	108
7.3	CC to MLO registration results for patient p8. The X-ray annotation of the CC view and	
	its projection on the MLO are shown in green. The X-ray annotation on the MLO view	
	is shown in red. The error is $0.2mm$.	110

Figures
Figures

7.4	CC to MLO registration results for patient p10. The clip on the CC view its projection
	on the MLO are shown in green. The X-ray clip on the MLO view is shown in red. The
	error is 5.1mm
7.5	2D registration errors for the temporal CC and MLO view registrations of 6 patients. Our
	method is illustrated in red and it is compared against a 2D affine transformation (green)
	and the initial 2D error, without registration (blue). The clip cases are patients p5 and p6. 111
7.6	Temporal registration results for the CC view of patient p1. The X-ray annotation at
	timepoint t_1 is shown in green, as well as its projection on the second mammogram. The
	X-ray annotation of p1 at timepoint t_2 is shown in red. Their overlap is yellow. The error
	is 0.08mm
7.7	Temporal registration results for the MLO view of patient p6. The X-ray annotation at
	timepoint t_1 is shown in green, as well as its projection on the second mammogram. The
	clip of p6 at timepoint t_2 is shown in red. Their overlap is yellow. The error is $6mm$ 112
8.1	Illustration of the registration errors achieved by the three proposed transformation mod-
	els for the MRI to X-ray mammography registration (affine, E-SDM and FEM). The
	errors correspond to the ones previously presented in Table 6.3
C.1	The mean responses of the radiologists to questions 1 and 2. Question 1: On a scale of
	one to five, how visible is the lesion in the X-ray mammogram? Question 2: On a scale
	of one to five, how confident are you that your predicted lesion location is close to, i.e.
	within approximately 20mm of, the correct location in the X-ray mammogram? \dots 128

List of Tables

4.1	Distributions from which we have randomly chosen the parameters of the affine transfor-
	mation, in order to deform the MR volume and create the target images. The coordinate
	system that we refer to is shown in Figure 4.5
4.2	Reprojection error results (in mm) of the experiments on simulated X-ray data using all
	four similarity measures. The experiments were carried out using the MRIs of 5 patients.
	For every patient, there were two known affine transformations applied (aff. 1 and aff.
	2), as explained in section 4.1
4.3	Reprojection error (in mm) for the eight volunteers before and after registration 68
5.1	Parameters that vary between the different compressions and their distributions 82
5.2	Overview of the proposed framework that uses an ellipsoidal breast model and biome-
	chanically simulated compressions
5.3	Registration error (in mm) of this method (Ellipsoid-Statistical Deformation Model, E-
	SDM) and comparison with the volume-preserving affine transformation
6.1	Overview of our patient-specific FEM-based registration framework
6.2	Comparison between the Regular Step and the Regular Step Gradient Descent optimisa-
	tion schemes
6.3	Registration error (in mm) of our FEM transformation method and comparison with
	the Ellipsoid-Statistical Deformation Model (E-SDM) and the volume-preserving affine.
	The clip cases are p4 and p5. The last three columns in the table correspond to the mean
	value, the standard deviation and the confidence intervals of the mean values for a 95%
	confidence level. For a value c in the last column the confidence interval is given by:
	$[mean - c, mean + c]. \dots \dots \dots \dots \dots \dots \dots \dots \dots $
7.1	Registration errors in mm for the CC to MLO registrations on ten cases. The clip cases
	are p9 and p10
C.1	Experiment description given to the radiologists. This was created by the HAMAM
	partners Radboud University Nijmegen Medical Centre and University College London. 127

C.2	Mean distance (in mm) of the radiologists' annotation from their mean position (vari-
	ance) for the five cases that had two or more annotations both without and with the
	linking cursor

Chapter 1

Introduction

Breast cancer is the most common type of cancer worldwide¹ and is rated fifth in the list of the most common causes of cancer death². In 2008 it was reported to be the most common type of cancer amongst women in the UK. Almost one third (31%) of all new cancer cases diagnosed in women are breast cancer³.

Although there is increasing awareness that current diagnostic tools can lead to over-diagnosis and over-treatment of benign findings that might never require clinical treatment, it is generally accepted that early detection and accurate diagnosis are crucial for the patient's prognosis. This is supported by the fact that malignant disease is more likely to be treated effectively when detected at an early stage. In that regard, a number of different imaging modalities are used in clinical practice (section 1.2). Radiologists often need to localise corresponding findings in each of the images, to fully exploit the complementary information provided by the different modalities. However, this is a difficult task, due to the highly deformable nature of the breast, the different appearance of the breast structures across modalities and their often different dimensionality (2D or 3D).

The goal of multimodal breast image registration techniques is to identify correspondences between the different modalities automatically, in order to aid radiologists in the detection, diagnosis and management of breast cancer. More specifically, the work described in this thesis is focused on determining corresponding findings between Magnetic Resonance Imaging (MRI) and X-ray mammography. We also explore the benefit of using the MRI as an intermediate modality for establishing correspondences between different X-ray mammograms.

In the following sections we first describe the breast anatomy and the use of different imaging modalities in clinical practice. Then we introduce our research problem and the challenges associated with it.

1.1 Breast Anatomy

The anatomy of the female breast is mainly based on the studies that were carried out on cadavers [Cooper, 1840]. A later study used ultrasound to further investigate the structures of the ducts and their

¹World Health Organization International Agency for Research on Cancer (2008)."World Cancer Report".

²World Health Organization (February 2006). "Fact sheet No. 297: Cancer".

³Cancer Research UK, UK Breast Cancer incidence statistics 2008.

characteristics [Ramsay et al., 2005].

A commonly accepted diagram of the breast anatomy is described in [Ramsay et al., 2005], where some aspects (for example the number of ducts) may vary between different studies and subjects. The breast consists mainly of fat and fibroglandular tissue and it is covered by skin. Breast cancer develops within the cells of the fibroglandular structures. This fibrous network consists of milk ducts, lobes, blood vessels and is supported by Cooper's ligaments.

The appearance of the fibroglandular structures varies significantly across patients. The proportion of fibroglandular tissue with respect to the total breast volume is known as breast density. It has been shown that increased breast density is strongly related to higher risk of developing cancer [Oza and Boyd, 1993]. The appearance of the two tissue types (fibroglandular and fat) varies also across time for a single subject. For example, the amount of fat can change, and more importantly the glandular tissue can also vary as it is affected by hormonal changes, such as childbirth and use of post-menopausal hormone replacement therapy. In general, breast density decreases after menopause.

1.2 Breast imaging modalities in clinical practice

X-ray mammography is widely used for screening. Women are advised to obtain a mammogram after the age of 50 in the UK on a regular basis, that is currently three years. The only cases where an alternative modality is used, are those of younger women with a family history or high genetic risk of breast cancer. Then, Dynamic Contrast Enhanced-MRI (DCE-MRI) is used instead, as it has been shown to increase sensitivity [Leach et al., 2005]. DCE-MRI is currently recommended for annual screening of high risk women, aged between 30 and 49, by the National Institute for Health and Clinical Excellence [NICE, 2006]. The advantages are that it provides a 3D image of the breast and functional information, as it indicates the areas in the breast with increased blood flow. The 3D image can be used to avoid any ambiguity caused by the projective nature of mammography.

After a suspicious region is detected in a screening mammogram, a woman is usually advised to have a new mammogram taken for diagnosis. The imaging modalities used at this stage vary according to the needs of each individual case. X-ray mammography can be repeated, with a scope this time to focus on a certain region of the breast. 2D ultrasound is widely used in clinics, as when combined with an X-ray image it can help distinguish between cysts, benign and malignant lesions. Moreover, it is easy to acquire, low-cost and it does not pose any risk to the patient, as there is no exposure to radiation. DCE-MRI is also frequently used to detect and diagnose lesions that are not visible in X-ray mammography, or further investigate mammographically detected lesions.

After a woman is diagnosed with breast cancer, new images are acquired to monitor the staging and facilitate the treatment planning and follow-up. In these cases, the modalities that are most commonly used are DCE-MRI and X-ray mammography.

Finally, it is worth mentioning some additional breast imaging modalities that are used mainly for research purposes and they are also being introduced into clinical practice. These provide certain advantages compared to previous techniques. Digital Tomosynthesis obtains a series of X-ray images from different angles and thus gives coarse 3D information of the tissue structures. Another modality is Automated Whole Breast 3D ultrasound, which provides 3D information of the breast, in two different views. This is different to using a 2D or 3D ultrasound probe, as the breast is held stable between four planes and the breast movement is minimised using a membrane. Its main advantages are the 3D view of the breast, the suitability for use on dense breasts, combined with the fact that it is easy and inexpensive in contrast with MRI. Lastly, as the breast consists of soft tissue, Optical Tomography is an imaging modality that is currently used in research. The images are produced as a result of the light that is transmitted and scattered inside the breast volume. Although the image resolution is very poor, we can extract useful information on blood volume and tumour oxygenation. Magnetic Resonance Spectroscopy (MRS) is another modality that is mainly focused on diagnosis and staging and it provides biochemical information of a lesion, rather than structural. Positron Emission Mammography (PEM) is a nuclear medicine imaging technique that provides information of the metabolism of the tumour and is typically used for monitoring only advanced or recurrent tumours.

Given this overview of the different modalities used in clinical practice we can see that X-ray mammography and DCE-MRI are used commonly to investigate the high-risk population and symptomatic patients. More specifically MRI is used as complementary modality to mammography to investigate lesions that are not clearly visible in X-ray mammograms, due to the superimposition of fibroglandular structures and also the increased breast density of some patients. Another clinical scenario where MRI is used in addition to X-ray imaging is for further investigation of lesions that are mammographically detected. In these cases MRI is used to further characterise the lesion from the functional information provided regarding the blood flow concentration and also to help treatment and surgical planning. Therefore establishing correspondences between these two modalities automatically could aid radiologists in the detection, diagnosis and management of breast cancer.

Due to the different nature of the two modalities and the fact that the breast is a highly deformable organ, determining correspondences is a difficult task for radiologists. For the clinical scenarios mentioned above, the clinical advisory board of the HAMAM project [HAMAM, 2012] has indicated that an accuracy in the order of 10mm would provide valuable information, as the mapping of the MRI enhanced area onto the X-ray mammogram would provide a small region of interest inside which the radiologist would identify the corresponding location in the X-ray image. Higher accuracy would be required if the applications included image-guided interventions. To have a better understanding of the images involved in the MRI to X-ray registration, we describe in more detail the image acquisition of an X-ray mammogram and an MRI in the next two sections.

1.2.1 X-ray mammogram acquisition

As mentioned above, X-ray mammography is the breast imaging technique that is most commonly used. During this process, the woman's breast is exposed to a low dose of X-rays. According to Beer-Lambert's law for a single X-ray wavelength, the attenuated energy I transmitted through the breast, when considering N different tissue types, is given by equation:

$$I = I_0 \cdot e^{-\sum_{i=1}^{N} (\mu_i \cdot z_i)}$$

$$(1.1)$$



Figure 1.1: Example MLO view mammograms of breasts with increasing breast density from left (a) to right (d).

where I_0 is the incident photon energy, μ_i is the X-ray attenuation coefficient of the tissue type *i* and z_i the corresponding tissue thickness. Therefore, the contrast between different tissue types in an X-ray mammogram is a result of both the difference in the thickness and the X-ray attenuation coefficients between the fat, the fibroglandular and the tumorous tissue.

Experimental results [Johns and Yaffe, 1987] conducted to determine the X-ray attenuation coefficients of the three tissue types across different energies, showed that fat has the lowest attenuation at all energies. The other two tissue types gave very similar results, with tumorous tissue having a slightly higher attenuation coefficient than the normal fibroglandular tissue, especially for lower energies. The difficulty in X-ray mammography is the ability to distinguish between the normal fibroglandular and tumorous structures. This is the reason why images from dense breasts are more difficult to interpret. Ducts, lobes and tumours appear brighter in the mammogram (higher attenuation), while adipose or fatty tissue is darker (lower attenuation). Consequently, the contrast between tumour and fibroglandular tissue is less than between tumour and fat. The difference in density between breasts can be seen clearly in Figure 1.1.

Apart from the difference between the X-ray attenuation coefficients of the soft tissue types imaged in mammography, there are additional factors that influence the contrast and the image quality of an X-ray mammogram. One of these is the X-ray beam spectrum that is used. On the one hand, beams with low photon energy are required for increased contrast, as the difference in X-ray attenuation between the two tissue types is higher at low energies. On the other hand, in this case the radiation dose to the breast tissue is higher, posing an increased risk to the patient. Therefore, the optimal X-ray beam spectrum used in mammography is a compromise between high image contrast and low radiation dose.

The image quality of an X-ray image is further affected by the X-ray quantum noise and the scattered radiation, which is partially reduced in mammography by using an anti-scatter grid. Moreover, the quality is limited by the final image spatial resolution and it is affected by the intensity variation of the X-ray beam. As rays are emitted from a point source in a conical shape, their intensity reaching the detector varies (via the inverse-square law), producing an X-ray image with varying spatial density. Furthermore, due to the geometry of the anode target in the X-ray tube that is angled, the intensity of the X-rays emitted is greater towards the cathode and lower towards the anode, as the X-rays travel through more material in that direction before being emitted. This is known as the anode heel effect.

In conventional film mammography, the final images are generated as a result of the exiting X-rays that are absorbed by the X-ray film. The use of this analogue process is currently being replaced by Full Field Digital Mammography (FFDM). In July 2011, 85% of the UK breast screening units had at least one digital mammography imaging scanner⁴ and this is expected to fully replace film mammography in the future. FFDM uses an electronic digital detector instead of film, which facilitates the transfer of images and enables their storage without degrading image quality.

During image acquisition and in order to increase image quality and avoid motion and scatter artifacts, the breast is extended, compressed and immobilised between two planes. In the UK, there are two images acquired, one of the Cranio-Caudal view (CC) and one of the Medio-Lateral Oblique (MLO).

1.2.2 DCE-MRI acquisition

Another frequently used modality for breast screening and cancer diagnosis is DCE-MRI. The nature of MR images differs significantly from that of X-rays. MR images are produced as a result of the nuclear magnetic resonance properties of the hydrogen atoms in the human body. Hydrogen atoms can be found in water and fat, which account for a large percentage of the human body composition.

The nucleus of the hydrogen atoms consists of a single proton. When the subject is inside the MR scanner, the average magnetic moment of all protons is aligned with the static magnetic field B_0 . During MRI, a sequence of several radio-frequency (RF) pulses is applied. These are described by the angle of the net magnetisation with respect to the main magnetic field direction, which is called flip angle. When the RF pulses are introduced, the protons absorb some of the transmitted energy, which then flips the spin direction of the hydrogen atoms. The frequency of the RF pulses at which the spin direction changes is known as the resonance or the Larmor frequency and it is given by the equation:

$$\omega_0 = \gamma \cdot B_0 \tag{1.2}$$

where γ is the gyromagnetic ratio that is specific for each type of nucleus.

After the RF pulses are turned off, this energy is released at different rates and the protons return to their equilibrium state. The difference between these relaxation times produces the contrast amongst the different tissue types. The MR signal produced by the transmission of this energy is finally transformed from the original spatial frequency space (called *k-space*) to the image space using a Fourier Transform.

There are three types of MR images that can be generated according to the relaxation processes that occur. T1-weighted images reflect the difference in relaxation times and the recovery of the equilibrium magnetisation along the longitudinal axis. This relaxation time is also known as spin-lattice. Similarly, T2-weighted images reflect the relaxation times within the transverse plane, which is caused by the varying magnetic fields of the moving protons around their neighbours. T2 is also known as spin-spin relaxation time. Finally, proton density or spin echo images are not influenced by the T1-, T2-relaxation times and the magnetic field inhomogeneities. The images produced reflect only the amount of spins in

⁴NHS Breast Cancer Screening Programme http://www.cancerscreening.nhs.uk/breastscreen/digital-mammography.html



Figure 1.2: An example of a transverse slice coming from (a) a pre-contrast MRI and (b) the subtraction image of the pre-contrast from the post-contrast image, showing an enhancing lesion (Invasive Ductal Carcinoma).

the tissue, therefore the number of hydrogen atoms found in the human body.

One factor that influences the imaging quality of an MRI is the difference between the MR signals of fat and water. This difference is known as chemical shift and it can introduce image artifacts that obscure interesting soft tissue structures. This is particularly relevant for breast imaging, as breast is largely composed of fat tissue. These artifacts can be reduced by several different algorithms known as fat suppression techniques. More details regarding the MR image acquisition can be found in the books [Dendy and Heaton, 1999] and [McRobbie et al., 2003].

It has been shown that for the breast case, the contrast produced by the T1 and T2 signals is not sufficient to differentiate between healthy and malignant tissue [Kopans, 1998]. DCE-MRI uses an intravenous contrast agent (usually gadolinium) that is injected into the patient and additional images are acquired. Contrast agents pass through the patient's vascular system to the breast and their presence there changes the relaxation time of the vascular structures. In the areas where tumours develop, there are new abnormal blood vessels generated (angiogenesis) and also leakage of vessels into the extracellular space of the tissue. Consequently the gadolinium concentration is higher there than in normal blood vessels. As a result, the images that are acquired after the injection of the agent can be subtracted from the pre-contrast MRI and the subtraction images highlight the location of the neo-vasculature associated with tumours.

Additional information that can be extracted from the DCE-MRI sequence, apart from the subtraction images, is the washout pattern of the contrast agent. If a sufficient number of post-contrast images are acquired, then by plotting the MR signal of a region over time, we can distinguish between healthy and malignant tissue. Normal breast tissue enhances gradually and to a lesser degree, while tumours enhance a lot and rapidly.

During MRI, the woman is lying prone and the breasts are pendulous under gravity inside two breast coils. An example of a slice coming from a pre-contrast MR volume is shown in Figure 1.2(a) and the corresponding slice of the subtraction volume, between the pre- and the post-contrast image, is given in Figure 1.2(b).

It is clear from the above that the images produced using MRI and X-ray mammography differ significantly in various aspects. Firstly, they are images of different dimensionality and they are produced

as a result of inherently different processes (X-ray attenuation versus magnetic resonance). Moreover, the position and compression of the breast in those two procedures are also different. This combined with the highly deformable nature of the breast tissue, further complicates the registration process. Finally, the image resolution differs too, as MRI exhibits coarser resolution. For a typical mammogram this is in the range of $[(0.05 - 0.1) \times (0.05 - 0.1)]mm^2$, while for the MRI it is $[(0.6 - 1.0) \times (0.6 - 1.0) \times (1.0 - 3.0)]mm^3$.

1.3 Problem statement

The problem that is addressed in this study is the difficult task that the radiologists encounter when attempting to determine corresponding positions in MR and X-ray images of the breast. The goal of this work is to develop a framework for the alignment of MRI to X-ray mammography images. This could aid radiologists in determining corresponding regions, as this is a challenging task and these modalities are frequently used in the management of breast cancer, due to the fact that they provide complementary information. Having built such a framework, we also investigate its use for the alignment of different X-ray mammograms from a given patient, by making use of the 3D image of breast structure provided by the MR volume.

1.4 Contributions

The contributions of this work are:

- Development of a generic framework for the alignment of MRI and X-ray breast images based on their intensities, or in other words the internal breast structures, rather than only the breast outline, the distance from the chest wall or the nipple position.
- Investigation and comparison of three different transformation models, that attempt to capture the large breast deformation, with increasing complexity: from a simple geometrical model, such as the affine, through an ellipsoidal model and to a patient-specific biomechanical model of the breast.
- Validation of the developed framework with a large number of clinical datasets, using cases with identified lesions in both modalities, annotated by experts, and MR and X-ray visible clips that have been inserted at the position of mammographically detected lesions.
- Integration of a FEM-based transformation model inside the registration framework, using simultaneous optimisation of both the pose and the biomechanical simulation parameters.
- Use of the same registration framework to determine correspondences between X-ray mammograms of the same patient via the 3D information of the MR volume. This task was previously approached as a one-to-one correspondence task, although it is one-to-many.

1.5 Report structure

Chapter 2 contains the literature review of relevant studies. We present the different intra- and intermodality registration techniques, with a particular focus on breast image registration.

Chapter 3 presents the X-ray image simulation technique using the MRI, which is an essential part of the registration process.

The following chapters describe the three different transformation models used and the results when applied to clinical datasets.

In chapter 4 we use a volume-preserving affine transformation model to approximate the mammographic breast compression. The main motivation for using this geometrical model is the low number of degrees of freedom and the ease with which it can be incorporated into clinical practice.

To achieve better accuracy and model more accurately the breast deformation, we propose learning the space of possible breast deformations by using an ellipsoidal shape of an average size and applying biomechanically simulated compressions (chapter 5). The main modes of variation are then extracted using Principal Component Analysis (PCA). The main advantage of using an ellipsoidal as opposed to a patient-specific deformation model, or an incorporation of the biomechanical model inside the iterative registration scheme, is that a single generic model is created once for all patients, eliminating the need for model creation on a patient by patient basis.

As a third transformation model, we investigate the use of a patient-specific biomechanical model of the breast, created from the MRI of the patient (chapter 6). This is enabled through a transformation model that uses a fast explicit Finite Element solver, which runs on the graphics card. The iteratively updated parameters include both parameters of the biomechanical model and the boundary conditions, and also rigid transformation parameters of the breast geometry model.

Finally, in chapter 7 we investigate the use of the same registration framework for a different application, that of establishing correspondences between 2D X-ray mammograms, using the 3D information provided by the MRI.

Chapter 8 includes the conclusions and potential extensions in future work.

Chapter 2

Literature Review

This chapter contains a review of the different methods that have been introduced in the literature and are related to the research of this study, which is focused on MRI to X-ray mammography registration and its applications. The review is structured as follows: Section 2.1 introduces the image registration problem and the classification of existing algorithms, according to their methodology. Section 2.2 refers briefly to the algorithms introduced to solve the intra-modality breast image registration problem, such as between X-ray mammograms of the same patient. Then, section 2.3 contains the different multimodal breast registration approaches; the methods that were introduced for MRI – X-ray registration are analysed in more detail, as this is the topic of our study. Since breast is a deformable organ, it is important to review how the same problem is approached for other deformable organs, apart from the breast. These are included in section 2.4. In the next part (section 2.5), we review some techniques that are linked to multimodal registration of deformable organs, such as biomechanics and statistical deformation models. Finally, as the problem we aim to solve is a 2D-3D registration problem, the last part (section 2.6) contains the different approaches that solve the same task for structures that are not deformable, such as blood vessels, bones etc.

2.1 Image registration

If we consider two different images of the same scene, object or human organ, that are captured at different times, from different views or using a different imaging modality, then these images will have different coordinate systems. Consequently, integrating the information about the object that is captured and determining corresponding locations is difficult. The term "image registration" refers to the process of aligning these two images, or in other words bringing them into the same coordinate system.

Figure 2.1 illustrates the problem of a simple image registration task. The problem that registration attempts to solve is to find a spatial transformation that maps the coordinates of one of the images to the other, so that the intensities of corresponding structures can be directly compared. Images are a discrete representation of a continuous scene. In other words, the intensities at the pixel positions form a grid of the continuous space. Therefore when we apply a transformation at a point in one of the images, the new transformed position does not necessarily coincide with a pixel centre in the other image. To solve this problem, we need to use interpolation to find the intensity at a specific non-grid position. Consequently,



Figure 2.1: Illustration of the image registration problem. The goal is to calculate the transformation T(x), in order to align the two images. The image on the left stays stable during registration (fixed/target), while the image on the right is being resampled (moving/source).

one of the images during registration needs to be resampled. This image is called the *moving* or *source* image, while the other one remains unchanged and it is called *fixed* or *target*.

To re-state now the problem in mathematical terms and looking again at Figure 2.1, an image registration algorithm is looking for the transformation **T**. For a 2D problem, the coordinates at the position (x', y') are given by:

$$(x', y') = \mathbf{T}(x, y) = (x, y) + \mathbf{u}(x, y)$$
 (2.1)

where $\mathbf{u}(x, y)$ is the displacement vector. The transformation \mathbf{T} can vary in terms of complexity. Registration algorithms can be classified into different categories according to the type of the transformation that they use to align the images. Below are given some example classes with ascending complexity/degrees of freedom:

• **Rigid Registration**: In this registration problem the two images are aligned using a rigid-body transformation, that incorporates a rotation plus a translation. This is alternatively known as Euclidean transformation and when applied on a 2D point the new coordinates are given by:

$$\begin{bmatrix} x'\\y'\end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta)\\\sin(\theta) & \cos(\theta) \end{bmatrix} \cdot \begin{bmatrix} x\\y\end{bmatrix} + \begin{bmatrix} t_x\\t_y\end{bmatrix},$$
(2.2)

where θ is the rotation angle and (t_x, t_y) the translation components. In 2D the degrees of freedom are 3 and in 3D they are 6 (three for rotations about each axis and three for translations).

• Affine Registration: An affine transformation can be seen as a rigid-body plus scaling along the different directions and shearing. The main characteristic is that it preserves parallelism; so parallel lines are transformed to parallel lines. The effect when applied on a 2D point is given by:

$$\begin{bmatrix} x'\\y' \end{bmatrix} = \begin{bmatrix} \alpha_{00} & \alpha_{01}\\\alpha_{10} & \alpha_{11} \end{bmatrix} \cdot \begin{bmatrix} x\\y \end{bmatrix} + \begin{bmatrix} t_x\\t_y \end{bmatrix}, \qquad (2.3)$$

where $\alpha_{00}...\alpha_{11}$ incorporate the rotation as in equation 2.2, the scaling for the two different axes (in α_{00} and α_{11}) and the shearing (in α_{01} and α_{10}). In 2D, the degrees of freedom are 6 (four for the matrix coefficients $\alpha_{00}...\alpha_{11}$ and and two for the translations) and in 3D they are 12. • **Deformable Registration**: This transformation allows complex deformation of objects and it can incorporate ten to millions of degrees of freedom. This type of registration is widely used in medical imaging, as human organs can be highly deformable. Solving for the transformation of a highly deformable object is not a computationally tractable task, as in the previous cases. There are several techniques in the literature that approach this registration problem.

One of the most widely used transformations are the Free-Form Deformations (FFD) based on Bsplines, that was introduced by Rueckert et al. [Rueckert et al., 1999]. The authors propose a novel method for registration, where they apply a mesh on one of the images and the points that form this grid act as control points of a cubic B-splines equation. By simply then moving those control points, the pixels of the images follow the same deformation as a B-spline curve. The equation of this deformation in 3D, given a volume $\Omega\{(x, y, z)|0 \le x < X, 0 \le y < Y, 0 \le z < Z\}$ and an $n_x \times n_y \times n_z$ mesh of control points, is given by:

$$\sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \phi_{i+l,j+m,k+n},$$
(2.4)

where $i = \lfloor x/n_x \rfloor - 1$, $j = \lfloor y/n_y \rfloor - 1$, $k = \lfloor z/n_z \rfloor - 1$, $u = x/n_x - \lfloor x/n_x \rfloor$, $v = y/n_y - \lfloor y/n_y \rfloor$, $w = z/n_z - \lfloor z/n_z \rfloor$ and B_l is the *l*-th basis function of the B-splines. The main advantage of the FFD method is that B-splines are locally controlled and thus can capture complex, local non-rigid deformations. For this reason, one has to use a global transformation (such as rigid, or affine) before applying the FFD. Other non-rigid registration methods include the use of thin-plate splines [Meyer et al., 1997], optical flow [Kumar et al., 1996], demons [Thirion, 1998], a poly-affine transformation [Arsigny et al., 2005] or a curvature-based scheme [Fischer and Modersitzki, 2004].

Another method that is widely used is the fluid registration, introduced by Christensen et al. [Christensen et al., 1996]. The object in this case is treated as a fluid, whose new position is calculated as a result of an image-derived force. More specifically, in parts of the image where the similarity measure is low, the force is high and so the the pixels in the image are displaced following a fluid deformation. In contrast with the B-splines FFD, this method follows a physics law to deform the object, treating it as fluid. The main disadvantage is that it makes use of the Navier-Stokes equation, which is computationally expensive to solve. Christensen et al. use the Successive Over Relaxation method to solve it. Other efficient ways to compute it involve using a Full multi-Grid approach ([Crum et al., 2005] and [Freeborough and Fox, 1998]), the Minimum Residual algorithm [Wollny and Kruggel, 2002], a Convolution filter ([Bro-Nielsen and Gramkow, 1996] and [D'Agostino et al., 2003]) and finally a method similar to the Fast Fourier Transform [Cahill et al., 2007].

Fluid registration belongs to a class of techniques called diffeomorphic registration approaches, that have the advantage of providing transformations that preserve the topology of the image. In other words they generate a one-to-one mapping between the images, without discontinuities in the displacement fields. Recent advances in non-rigid registration techniques propose the use of other diffeomorphic approaches, such as those based on a log-Euclidean framework ([Arsigny et al., 2006], [Ashburner, 2007], [Arsigny et al., 2009], [Vercauteren et al., 2009]). One of the most popular ones is using the demons registration approach [Thirion, 1998] inside the log-Euclidean framework [Vercauteren et al., 2009]. The authors have also extended this work to provide a computationally efficient inverse-consistent transformation [Vercauteren et al., 2008]. Finally, Klein et al. [Klein et al., 2009] recently conducted an evaluation study of non-linear deformation algorithms applied to human brain MRI registration.

In all the above cases, we assume that the registration algorithm is used to align images of the same dimension (either 2D or 3D). Another category of image registration techniques is the one that is applied on 4D images. In this case, the fourth dimension is considered to be time and these algorithms are applied in order to register image sequences that are acquired during some process that changes in time (such as heart images that follow the cardiac cycle, or lung images that follow the breathing cycle). Finally, another category is the one used to register images of different dimensions (2D/3D registration). This type of registration is the topic of this research work and it is further analysed in the next sections of the literature review as well as in the following chapters that describe our methodology in detail.

Another categorisation of image registration techniques can be performed according to which characteristic of the images is used for alignment.

- Feature-based Registration: This type of registration is based on the detection and extraction of corresponding points in the two images. These points are also known as *features* or *landmarks*. The concept is that the registration algorithm uses only extracted points from the two images in order to align them and the rest of the image information is discarded. The most important and difficult step in this kind of registration is the feature selection. This can be done by various ways: automatically, by extracting highly distinctive features in each scene, manually by selecting corresponding points, or finally (especially in medical imaging) by using *fiducial markers* attached to the patient. The selection of the most appropriate landmark set from the images is crucial for the success of the registration and it is not a trivial task. Automatic algorithms do not always perform well for all the different cases, manual extraction is laborious and subjective, while the use of fiducial markers is not always possible and is not used in general in clinical practice. On the other hand, these algorithms are generally quicker than the intensity-based alternative, as they make use of only a limited number of points for registration. A special type of feature-based registration technique involves registering only the surface of a certain object/organ, excluding any internal structures. This category is applied for example for registration of images of the prostate and also brain or tumour images, assuming that the region of interest is restricted to the surface only.
- Intensity-based Registration: The registration methods in this category use the intensity distributions in the images for alignment. They use an appropriate similarity measure to compute the similarity between the two images and the goal is to maximise it by iteratively updating the parameters of the transformation. The similarity measure is chosen according to the nature of the images involved. More details about this type of registration and its components are given in the

remainder of the thesis, as this is the approach that is used in our study. The main advantage of this approach is that by using the whole images there is no need for segmentation or extraction of landmarks, which usually are a source of error. As a result, these algorithms are generally more robust than feature-based techniques but they can be slower, since they process all the pixels in both images and not only a specific set of points. Nevertheless, recent technology advances now allow implementation of programs on Graphics Processing Units (GPUs) and thus rapidly perform computationally expensive tasks in parallel offering unprecedented speed-up. Intensity-based registration can benefit from this technology and therefore the computational cost associated with it might not be an issue in the near future.

Finally, regardless of which type of all the above registration algorithms is used, a very important part of any study is **validation**. The purpose of validation is to provide a quantitative assessment of how well the algorithm performed. In other words, we want to assign a numerical value to the registration's performance, other than the visual estimation. If the ground truth transformation between the two images is known, then the *Target Registration Error* can be simply calculated at every point x in the target image as the difference:

$$d = ||T_R(x) - T_{GT}(x)||,$$
(2.5)

where $T_R(x)$ is the transformation as a result of the registration algorithm and $T_{GT}(x)$ is the ground truth transformation. The main problem is that in most cases this transformation is not known. One method of acquiring these ground truth transformations is by generating simulated data.

Another way to validate the results is done by computing the distance between corresponding points in the target image and points in the transformed source image, after applying the transformation that was acquired as a result of the registration algorithm:

$$d = ||T_R(x_s) - x_t||, (2.6)$$

where x_s is a point in the source image and x_t a corresponding point in the target. The set of corresponding points is defined before registration and they can be obtained either by fiducial markers attached on the imaging object/organ or by an expert (radiologist). Inevitably manual annotations of images, include themselves an error even when performed by experts, as it is generally difficult to achieve high accuracy of corresponding point locations. An estimate of this error can be obtained by using annotations made by several different observers. This error is known as inter-observer variability. Even when annotations are performed by one expert, at different time-points, these also include an error, which is known as intraobserver variability. Due to the uncertainty introduced in manual annotations, these correspondences are commonly refered to as *gold standard*, rather than ground truth correspondences.

Apart from the two above methods, it is also common in medical imaging to use gold standard transformations for validation. This case is similar to the first one, but instead of using the ground truth transformation, it uses the result of a registration algorithm that is known to give high accuracy. This is named gold standard transformation and is considered to be very close to the ground truth.

All above validation metrics are computing a target registration error based on a distance metric.

For certain applications in medical imaging, such as image segmentation, it is common to use overlap metrics instead. The most popular overlap metric is the Dice Coefficient (DC). For two overlaping regions A and B, the DC is given by:

$$DC = \frac{2|A \cap B|}{|A| + |B|}.$$
(2.7)

For a perfect agreement between the two regions, the metric has a value of 1, while a value of 0 indicates no overlap between A and B. Apart from the use of these metrics for the evaluation of segmentation algorithms, they can also be used to evaluate registration techniques, for example when radiologists provide gold standard corresponding regions, rather than points between the images. In this case the two regions A and B need to have the same size, in order for equation 2.7 to give sensible results. This issue is discussed also later, when we first use radiologists' annotations for the validation of our registration algorithm (section 4.3.2). Also, as overlap metrics are equal to zero when there is no overlap, these do not provide a measure of misregistration in these cases. This is particularly important when the annotated regions are small, as their overlap after registration can often be zero.

2.2 Intra-modality breast image registration

2.2.1 Registration of X-ray mammograms

X-ray mammogram registration tasks can be divided into three main categories according to the problem that they attempt to solve: *ipsilateral* registration that aims to register CC and MLO views of the same breast, *bilateral* registration that approaches the registration between the right and left breast of the same view mammograms, from the same patient, and *temporal* or *longitudinal* registration that registers mammograms of the same view and the same breast, acquired at different time points. The latter registration task is particularly important in medical imaging, as when radiologists examine an X-ray mammogram for breast cancer they routinely compare it with previous ones and look for changes in the tissue structures. There are various methods in the literature that approach these tasks in different ways. Here, we briefly refer to the most representative ones.

In terms of feature-based techniques and bilateral registration, authors have mainly extracted control points automatically at locations such as the nipple position and the breast boundary ([Yin et al., 1991], [Yin et al., 1994], [Mendez et al., 1998]).

For temporal registration, various feature-based techniques have been proposed, where distinctive features have been extracted manually or automatically from the internal structures of the tissue ([Sallam and Bowyer, 1994], [Vujovic and Brzakovic, 1997]). For the automatic feature extraction authors have proposed a curvature measure [Marti et al., 2002], wavelets [Marias et al., 2005] or the Moravec interest operator [Kumar et al., 2001]. Intensity-based methods have also been proposed for temporal alignment; for example elastic registration [Periaswamy and Farid, 2003] and free boundary conditions for region matching [Richard and Cohen, 2003]. Also, *hybrid* methods have been introduced that combine aspects of both feature and intensity based methods ([Wirth et al., 2002], [Bakic et al., 2004]). Finally, a parametric model was proposed by Snoeren et al. [Snoeren and Karssemeijer, 2007] in order to register film to Full-Field Digital Mammograms (FFDM). An evaluation of temporal X-ray mammogram registration techniques has been done by van Engeland et al. [van Engeland et al., 2003], where the authors compared four different temporal registration methods and concluded that an intensity-based 2D affine registration outperforms simple feature-based techniques and also a thin-plate splines intensity-based approach. The same conclusion was also made in a later study [Pereira et al., 2010] that compared a 2D affine and a B-spline transformation with the annotations made by radiologists expert in mammographic interpretation. Nevertheless a more recent validation of intensity-based methods [Diez et al., 2011] showed that a B-spline transformation used in a multi-resolution scheme outperforms other transformation models, including the affine.

The problem of ipsilateral registration, between the CC and MLO view mammograms, has been mainly approached using various distance transformations. To locate the corresponding position of a finding in one of the mammographic views to the other, authors have used the distance from the nipple position and then texture information extracted from that region ([Paquerault et al., 2002], [van Engeland et al., 2006]). Zheng et al. compared three different distance metrics for the same task and found that a straight strip area that is perpendicular to the line connecting the nipple to the pectoral muscle performed best [Zheng et al., 2009]. These techniques are widely used in Computer Aided Detection (CAD) systems.

Kita et al. [Kita et al., 2001] proposed the use of the extracted curved epipolar lines to match findings between the two views, in a similar way that epipolar lines are extracted from two cameras that capture the same scene in Computer Vision applications. To account for the different breast compression between the two views, the authors proposed using a reconstruction of a simplified breast model. Although the breast compression model is only an approximation of the real breast deformation, the novelty of this technique, compared to other algorithms, is that it attempts to capture the 3D deformation of the breast, rather than use a 2D approximation that is inappropriate for this task.

Another approach that models the CC to MLO breast deformation was recently proposed [van Schie et al., 2011] for a different application, the registration of ipsilateral tomosynthesis views. The authors used a simplified semi-spherical model to simulate the breast deformation and predict the corresponding locations in the images.

In terms of validation, the results were most commonly assessed visually or by the accurate matching of a restricted region in the image that included a lesion. A novel approach was introduced by Hipwell et al. [Hipwell et al., 2007], where instead of using "one-to-one" correspondences of points in the 2D mammograms (as the evaluation in all the above techniques), the authors take into account that these are actually "one-to-many" correspondences, due to the projective nature of the images. The authors used biomechanical simulations as known 3D deformations in order to validate 2D mammography registration algorithms. In a related approach, Qiu et al. [Qiu et al., 2008] proposed the use of a FEMbased framework to map a lesion from the two views X-ray mammograms to the MRI of the patient and thus use the 3D position of the lesion in the MRI to estimate correspondences between temporal mammograms. Finally, the effect of the breast thickness variation in mammography registration was studied on phantom experiments [Richard et al., 2006]. Reviews about X-ray mammogram registration and in general breast image registration can be found by Guo et al. [Guo et al., 2006] and in the book of Suri and Rangayyan [Suri and Rangayyan, 2006].

2.2.2 Registration of DCE-MRI breast images

As we have seen in the description of the DCE-MRI acquisition (section 1.2.2), there is initially one MR image acquired before the injection of the contrast agent and then additional acquisitions take place to obtain the post-contrast images. In theory, the breast volumes in all cases remain in the same position. In practice, there is usually either a slight or a greater variation that causes artifacts in the subtraction images, due to patient motion or breathing.

There are several methods in the literature that attempt to solve this registration problem. One of the most popular is non-rigid registration using B-splines [Rueckert et al., 1999], as was discussed in more detail in section 2.1. A modification of this approach included a volume conservation constraint, which was proven to give better results in the region of the tumour, when registering DCE-MR images [Rohlfing et al., 2003]. Another popular registration technique applied to DCE-MRI is fluid registration ([Christensen et al., 1996], [Crum et al., 2005], [Cahill et al., 2007]). Other non-rigid registration techniques used for DCE-MRI registration were also discussed in section 2.1 ([Meyer et al., 1997], [Thirion, 1998], [Kumar et al., 1996]).

Hayton et al. [Hayton et al., 1997] used a pharmacokinetic model for registration. This refers to an intensity transformation approach, rather than spatial. The effect of MRI motion correction on the pharmacokinetic parameter estimation was recently studied by Melbourne et al. [Melbourne et al., 2011].

In terms of validation, breast biomechanical modelling played an important role in the DCE-MRI registration task. By simulating displacements of the breasts using the pre-contrast images, the authors have generated simulated volumes, for which the real displacements were known and thus they could be used for validation of the registration algorithms ([Schnabel et al., 2003], [Tanner et al., 2007]).

2.3 Multimodal breast image registration

The different modalities that are used for breast imaging were discussed above in section 1.2. Here, we present in more detail the algorithms that attempt to solve the MRI - X-ray mammography registration problem, followed by the ones that were introduced to register breast images from other modalities.

2.3.1 MRI/X-ray mammography registration

This is a difficult registration task, due to the nature of the images acquired (two 2D views with high contrast versus one 3D volume of low resolution). Also, the fact that the breast is a highly deformable organ increases the level of difficulty, given that it is significantly compressed during mammogram acquisition, while it is left uncompressed in a prone position during MRI. Therefore, a tool that would automatically provide correspondence between the two modalities could aid radiologists in breast cancer diagnosis and management, and provide the enabling technology for routine multi-modal Computer-Aided Diagnosis [Yuan et al., 2010].

There are three main methods introduced in the literature that approach this problem in different ways ([Marti et al., 2004], [Behrenbruch et al., 2003], [Ruiter et al., 2006]). These are explained in de-

tail below.

Marti et al. 2004, "2D-3D correspondence in mammography"

Marti et al. introduced one of the first techniques to register X-ray mammograms with MR volumes [Marti et al., 2004]. Their method includes a 2D-2D registration of the images (by using a projection of the MR volume) and also a method to find correspondences between certain Regions Of Interest (ROI) in the 2D mammogram and the 3D MRI.

Regarding the 2D-2D registration, the authors first create a 2D image of the breast from the undeformed MR volume, by parallel projection. Initially they roughly align the two images using a 2D similarity transformation (rotation, translation and scale) and mutual information (MI). Then, they achieve a more accurate registration result, by using a non-linear line operator to detect the linear structures in the two images. The goal is to match those linear structures, based on a measure of their maximal curvature.

After the 2D-2D registration task is complete and as far as the 2D-3D registration is concerned, this is achieved by using the same linear structures matching. The difference is that this time, the second 2D image is each individual MRI slice, rather than the projection image of the MR volume. If correspondences are found on the slice of the MRI (by minimising a distance metric), then the 3D position is automatically known.

The method was tested on a single patient. The evaluation of the results was based on similarity measures of the detected matched features, taking into account their position, width, orientation and the projection angle of the volume.

As explained before, this technique is based on finding feature correspondences of linear structures. Feature-based registration algorithms have the limitation that a mismatched feature could result in a very large error. Therefore these techniques are generally less robust and are less suitable for clinical use. Other limitations of this study are the size of the dataset and the evaluation method used, which did not include the use of ground truth or gold standard correspondences.

Behrenbruch et al. 2003, "Fusion of contrast-enhanced breast MR and mammographic imaging data"

Behrenbruch et al. introduced in 2003 a method to map findings in the X-ray mammograms to the DCE-MR images [Behrenbruch et al., 2003]. As discussed in the introduction, MRI provides a 3D volume of the breast, but the spatial resolution is not as high as in X-ray images. Therefore we cannot identify the precise 3D position of small structures, such as microcalcifications, which would be valuable for identifying the 3D location of Ductal Carcinoma in Situ (DCIS). This method attempts to solve this localisation problem, as it aims to find the MRI corresponding location of a point in the X-ray mammogram and in this way define the 3D location of microcalcifications that are visible only in an X-ray image. Therefore this technique is designed to map the X-ray coordinates into the MR volume.

The registration technique that is followed is feature-based and it contains two steps that are summarised in Figure 2.2. Initially, the authors use the two X-ray mammograms (CC and MLO views) and two parallel projections of the MR-volume in the same direction as the X-ray images. Then, at the first stage of the registration, they use a curvature measure to correlate the boundary points of the film mam-



Figure 2.2: Registration framework followed by Behrenbruch et al.

mogram with those of the projections from the MRI volume. This curvature measure depends on an accurate segmentation of the breast boundary, which is achieved using an intensity-based search along the image, a set of morphological operations for smoothing and finally a spline interpolation. By using this technique, the extremities that are usually present in an X-ray mammogram are deformed in order to match the shape of the boundary that is produced from the projection of the MR volume. This step provides a rough alignment of the two images.

At the second and final stage of the registration method, a wavelet-based feature detection is used, to identify correspondences between the two different modalities. The extracted features are matched using various criteria, such as scale localisation, orientation, feature area, relative motion and neighbourhood localisation. Finally, the registration is done using a combination of internal and boundary landmarks which are matched using a thin-plate spline warping technique. After registration and to specify the 3D location of a region of interest from the two 2D views, they use a 3D reconstruction method that takes into account the breast compression.

This technique was tested on 14 patients, who were diagnosed with different types of breast cancer. The registration error varied between 2mm for the best cases to around 10mm for the worst. The validation was straight-forward for the cases where the lesion position was clear in both modalities, but it was poor when in one of the images the lesion could not be seen.

The proposed method of Behrenbruch et al. has the advantage of providing a promising solution to the very important problem of 3D location of the microcalcifications, by registering X-ray mammograms and DCE-MR images. Since the microcalcifications are not visible in the MRI and thus their location in the 3D space is unknown, a fusion of an X-ray image and the MRI volume using registration can provide this information. On the other hand, it also incorporates several limitations. Firstly, a prerequisite for this feature-based registration method is that the corresponding features should be present in both the modalities. This requirement makes questionable the potential use of the algorithm as a clinical tool. This would be mostly valuable in cases where there are no clear correspondences between the modalities, such as for women with dense breasts, where findings are often obscured. Moreover the possibility of mis-matched features remains, as in all feature-based techniques.

Ruiter et al. 2006, "Model-Based registration of X-ray Mammograms and MR Images of the Female Breast"

The last of the three registration methods was proposed by Ruiter et al. ([Ruiter, 2003], [Ruiter et al., 2006]). The main innovation that the authors introduce in their work is that they take into account the deformation of the breast during image acquisition of the X-ray mammogram, by using a biomechanical model of the breast and compression simulations. In all the previous approaches, the MR volume is projected undeformed to produce a 2D image that can later be registered with the 2D mammogram. In this method, a new step is added into the process. Before the projection of the MR-volume, the authors simulate the compression of the breast during X-ray mammography, using a biomechanical Finite Element (FE) model. This way, instead of deforming the 2D projection images, to recover the 3D breast compression, the biomechanical modelling allows the 3D FE mesh to be deformed instead.

In this approach, the first step is to discretise the MR volume, in order to generate a 3D FE mesh of the breast. This is done based on a segmentation of the MRI into different tissue types. There are three different categories used: the skin, the fat and the glandular tissue. The next step is to assign material properties to them, that define the behaviour of the breast under compression. There are also other parameters that need to be set at this stage, such as the coordinate system and the boundary conditions that will be applied in order to create a deformation. Then, a deformation of the FE mesh is simulated, in a similar way to the one that occurs during X-ray mammography (compression of the breast between two plates). The resulting parallel projection of the deformed volume can then be directly compared with the mammogram. The amount of compression is updated until the chest wall to nipple distance of the simulated projection matches the one in the real mammogram. To achieve an accurate matching of the boundaries in the 2D images, the authors apply additional displacements on the surface nodes that force the 3D mesh to stretch in the medial-lateral direction.

Once the deformation of the breast is recovered in registration, this can be used both for mapping MR coordinates to the X-ray mammogram and vice versa. For the inverse process, the 3D position of a lesion that appears in the mammogram can be determined by finding the corresponding position in the simulated deformed image and then applying the inverse transformation to go back to 3D coordinates.

This framework was tested on six patients. The mean displacement of the lesion centre when this was projected from the MRI on the mammogram was 4.3mm ($\sigma = 1mm$). When this was mapped from the two mammographic views within the MRI the mean registration error was 3.9mm ($\sigma = 1.7mm$).

The results indicate that this method outperforms all the previous approaches. Moreover the methodology provides a tool for simulating realistic mammograms from MR volumes, as the breast mammographic compression is taken into account for the first time, although the projection is performed using parallel rather than perspective projection. The use of physically realistic compression simulations is a significant contribution in the literature of multimodal breast imaging registration. However in a more recent semi-automated implementation of this approach, the authors reported values of $11.8 \pm 6.5mm$ and a mean overlap of $63 \pm 40\%$ for 11 subjects (CC view only) [Hopp et al., 2012], which indicates

possibly the effect that manual interaction has on the results and also possibly the variability depending on the datasets used.

Although patient-specific biomechanical modelling can produce more realistic deformations, the manual steps involved in building a different model for each patient and the variability in the results depending on the material properties, the meshing and the simulation techniques used, make this frame-work less suitable for clinical use.

Finally, recently there has also been another study [Lee et al., 2011] that used a similar approach to Ruiter et al., but was further improved in order to drive the registration based on the image intensities rather than the breast outline and the nipple position. Furthermore, in this work the material properties, the amount of compression and the pose of the breast are optimised inside the iterative framework according to the matching of the structures in the two images. Following this approach increases the possibilities of an accurate matching between the two modalities. Nevertheless, in these preliminary results, it is not clear which are the parameters that are optimised and whether these include the initial positioning of the patient before compression. Finally, the validation was performed using only a breast phantom. The registration error was $1.35 \pm 0.06mm$.

As we can see from all the previous techniques, although there has been some significant progress in solving the MRI/X-ray registration problem, the need for an accurate, reproducible and well validated technique still remains.

2.3.2 Other breast imaging modalities

The references in the literature that attempt to register multimodal breast images are very limited. Instead of registering multimodal images that are acquired at different times, there are various techniques that propose the simultaneous acquisition of the two images. Consequently the images are already coregistered. Although these methods are not applied in routine clinical practice, they provide valuable information about how the breast tissue and its structures appear in different modalities.

Such devices for simultaneous acquisition were introduced for the combinations of digital tomosynthesis with ultrasound data [Kapur et al., 2002] and Full-Field Digital Mammography with 3D ultrasound [Kapur et al., 2004]. Co-registration of breast images can also be used for biopsy guidance, as the use of information from a second modality could improve accuracy [Piron et al., 2003]. In the same concept, a similar device was used to acquire a PEM image directly after an X-ray mammogram [Murthy et al., 2000].

For the registration of PEM to MRI, Krol et al. used a FEM-based approach [Krol et al., 2004]. The main limitation of this method is that it requires fiducial markers visible in both modalities to be attached to the breast; these are inconvenient in routine clinical practice.

2.4 Multimodal registration of other deformable organs

As we have seen so far, multimodal registration of breast images is a research field that has developed only in the last few years and therefore there are not many techniques proposed in this area. Nevertheless, multimodal registration of other human organs has been an active research field for longer and several different approaches have been proposed. These techniques are usually organ specific (such as brain, carotid artery, forearm, kidney, prostate, colon and liver) and also application specific (biopsy, radiotherapy, image guided surgery etc). In this section we present only some representative techniques that are mostly related to the registration of soft tissue organs and in particular the liver. There are plenty of registration techniques that are applied also to organs with rigid structure (such as brain, bones, neck). We will review those that refer specifically to the 2D/3D registration problem in section 2.6.

For the registration of liver images, several authors propose techniques that mainly rely on parts of the liver that have a more rigid structure, such as the vessels and the surface. For example, we refer to three different techniques that register MR to CT data. [Porter et al., 2001] and [Aylward et al., 2002] find a transformation between the different modalities by segmenting the blood vessels and optimising their correspondences. [Voirin et al., 2002] develop a feature-based method, using features on the liver surface. These approaches are not expected to perform well for the breast case, since the fibroglandular structures and vessels are not always visible in both modalities and also it is not sufficient to register the images based only on information about the surface/skin of the breast.

Boundaries are also used for registration between CT and US data of the kidney [Leroy et al., 2004]. It is worth mentioning that in this technique, the authors apply a preprocessing step to the data, in order to increase the similarity between them. In that regard, the CT images are blurred, while speckle is removed from the US data. The similarity measure used is the Correlation Ratio (CR).

In another method that registers US and MR images of the liver [Penney et al., 2004], the authors use a tracking device to capture the motion of the ultrasound probe and thus the conditions under which the images are acquired are known. Of interest for our application is the use of a pre-processing step they apply to the images before registration. They convert the intensity images to "probability images", where each pixel represents the probability of existence of corresponding vessel structures between the images at that point. In this way, the "probability images" can be used directly for registration, as both their values now represent similar quantities.

Finally, one multimodal registration method that concerns the liver and could be useful for other soft tissue organs like the breast, was proposed for registration between CT and 3D ultrasound data [Wein et al., 2008], which is related to the work previously described [Penney et al., 2004]. The main concept is that the authors use CT data to simulate ultrasound-like images and then use them for registration with the original ultrasound. Then, for registration, they use the slice of the ultrasound that has the highest entropy, to ensure that it incorporates high vascularity in order to be more similar to a CT image. Their main contribution to the literature is the use of a novel similarity measure, named *Linear Correlation of Linear Combination* (LC^2). This new measure has two desired properties; it is independent of brightness and contrast changes in the US image and also insensitive to how much the two main physical effects that produce the US image contribute to the intensity, which is important for the specific application.
2.5 Learning soft tissue deformations

As the breast is a highly deformable organ, it would be valuable for registration to acquire some prior knowledge about how the breast deforms. There are generally two main techniques that enable us to do that; the use of biomechanical simulations (section 2.5.1) and the use of Statistical Deformation Models (SDM), discussed in section 2.5.2.

2.5.1 Breast biomechanical modelling

As we have seen previously and in terms of registration, breast biomechanical modelling has been used to simulate the mammographic compression that occurs for the MRI – X-ray mammography registration ([Ruiter et al., 2006], [Hopp et al., 2012], [Lee et al., 2011]) and for validation of both DCE-MRI registration ([Schnabel et al., 2003], [Tanner et al., 2007]) and X-ray mammography registration algorithms [Hipwell et al., 2007]. Regarding multimodal breast registration, Rajagopal et al. used biomechanical modelling and proposed a new reference state of the breast [Rajagopal et al., 2008], to which all other modalities can be registered, in order to establish correspondences between them. Biomechanical modelling was also used in combination with an FFD registration technique in order to register prone to supine MR images and map pre-operative information about the location of the lesion to the intra-operative MR image ([Lee et al., 2010], [Han et al., 2011]).

Biomechanical models have also been used previously to simulate large mammographic compressions but were not tested for MRI/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. proposed a method for *in-vivo* parameter estimation, using a framework that incorporates a GPU implementation of the FEM modelling [Han et al., 2012]. This work can be further extended to FEM-based registration tasks, which were so far computationally expensive. This did not allow their use in clinical practice.

Other applications of breast biomechanics include image guided surgery or biopsy ([Azar et al., 2000], [Azar et al., 2002], [Carter et al., 2008], [Carter et al., 2009]) and breast elastography ([Sinkus et al., 2000], [Samani et al., 2007], [Washington and Miga, 2004]). An evaluation study of FEM-based large compression simulations [Tanner et al., 2011] showed an average error in the range of 4.1 - 6.7mm according to the material properties used. Finally, an analysis of the factors that influence the accuracy of the various breast biomechanical models can be found in [Tanner et al., 2006].

2.5.2 Statistical Deformation Models

Statistics have been widely used in image registration, in order to provide some prior information of the shape or the appearance of the object/organ in the images. This is done by using a population of training images, from which we can acquire some prior knowledge of the variation between them and use that later, with the goal of parametrising the deformation space and assisting the registration algorithm to converge faster and more accurately. There are two main categories of statistical deformation models in the literature; those that are based on *registration* and those that are based on *biomechanical simulations*.

The ones that are based on registration [Rueckert et al., 2003], use the same principle as the Active

Shape Models [Cootes et al., 1995] that are widely used in the Computer Vision community. The main difference is that instead of using manually selected landmarks to describe the shape of interest, they use the control points of the grid that controls the B-splines deformation. By registering the organs (in this case: brains) using an FFD registration algorithm, the authors automatically establish point correspondences. This technique can be used to acquire a mean model of the object of interest and use it later as a starting point for registration. However this method cannot be applied to the breast, due to the high variation of the breast anatomy across the population.

An alternative method is the use of a Statistical Deformation Model that is created using biomechanical simulations. This can be population-based, when the compressions are applied on a population of real breast images, or patient-specific, when the compressions are applied on one image of a specific patient. A population-based approach was used for the breast for multimodal registration [Tanner et al., 2008] and is based on a similar method introduced for prostate registration [Mohamed et al., 2002]. Patient specific biomechanical modelling of the prostate deformations was also proposed more recently [Hu et al., 2008]. The advantages and disadvantages of population and patientspecific Statistical Deformation Models are further discussed later in chapter 5, as one of our proposed transformation models uses a deformation model that is learnt from biomechanical simulations.

2.6 2D/3D registration of non-deformable organs

As we have seen so far, there is a limited number of approaches in the literature that refer to 2D/3D multimodal registration of breasts or other deformable organs. In this section, we present some representative methods that attempt to solve the 2D/3D registration problem for organs that are not deformable, such as the spine, the vessels etc. The reason for reviewing these techniques is that an initial step of alignment could be achieved using those methods for application to the breast. The review of the similarity measures and the optimisation schemes used is also relevant.

From all the registration techniques, we exclude those that are feature-based as the type of the features extracted in every case is application- and organ-specific. For breast images, extracting automatically 2D and 3D features is an open research field and there is not an established method available. Apart from the feature-based and the intensity-based techniques that are explained below, 2D/3D registration can also be done based on reconstruction and gradients ([Tomazevic et al., 2003], [Markelj et al., 2008], [Mitrovic et al., 2011]). The last two techniques are not analysed any further here, as they are not suitable for our research topic. These have been mainly used for registration of rigid structures such as bones, where the gradient information is stronger in projection images and also a partial reconstruction based on a small number of projection images, depending on the application, is possible. In mammography, the breast is compressed in different ways for the two views, CC and MLO, and also the contrast is produced from soft tissue structures and is often weak or obscured. Consequently, these methods are not discussed further in this review.

Intensity based techniques use the volume from the 3D imaging modality (MRI or CT) in order to produce a 2D projection image, called *Digital Reconstructed Radiograph* (*DRR*). The method followed to create this DRR aims to create an image that resembles as much as possible the image from the

2D image modality (usually X-ray or fluoroscopy). The registration is then performed by optimising a similarity measure between the DRR and the 2D image. The most appropriate similarity measure for this registration task varies according to the application.

Weese et al. used 2D-3D registration in order to align pre-operative CT scans with intra-operative fluoroscopy images of the spine [Weese et al., 1997]. As a similarity measure they used pattern intensity, due to its robustness in the presence of new objects in one of the images, such as the catheter during surgery. A similarity measure comparison study for the same type of images [Penney et al., 1998] shortly after showed that the most appropriate and robust similarity measures are pattern intensity and gradient difference. An intensity-based technique applied to the registration of 3D Magnetic Resonance Angiograms with 2D X-ray Digital Subtraction Angiograms [Hipwell et al., 2003] has shown that pattern intensity, gradient difference and gradient correlation are the similarity measures that are most appropriate for this task. A correlation-based measure and more specifically cross correlation has been also used for the registration of 3D CT scans of the cranium, with 2D X-ray images [Murphy, 1997]. Gradient and correlation based similarity measures were also proposed by [Lemieux et al., 1994]. Another clinical application of 2D-3D registration has been used in radiotherapy, for the alignment of portal images with CT volumes. Such approaches were followed by Gilhuijs et al., using a distance metric [Gilhuijs et al., 1996], and Khamene et al., using normalised local correlation [Khamene et al., 2006]. The most recent reviews in this area concern 2D/3D registration methods proposed for image-guided interventions [Markel] et al., 2012] and also a comparison of the optimisation techniques used for the same application [van der Bom et al., 2011].

The main advantage of intensity-based techniques is that they do not require any segmentation or extraction of points from the images and as a result are generally more automated and potentially more robust. On the other hand, the computational time is quite high and in cases of image guided surgery or radiotherapy this can be crucial. The production of DRRs is also a computationally expensive step and in order to accelerate it, authors have used various techniques, such as a restricted region of interest for projection or a different rendering algorithm [Weese et al., 1999]. Perspective projection algorithms are further discussed in chapter 3, where we describe in detail the methodology that we follow to simulate X-ray images from the MRIs.

2.7 Our approach

The subject of this study is the development of a framework for establishing correspondences between 2D and 3D images of the breast. More specifically, we propose a registration pipeline for MRI to X-ray mammography registration and we also investigate its application for finding correspondences between 2D X-ray mammograms, using the 3D information provided by the MRI.

We approach the MRI to X-ray mammography registration problem using an intensity-based method. There are two main reasons for not following a feature-based technique. Firstly, the robustness in comparison to intensity-based approaches is poor, as a misregistered set of points can result in a high registration error. Therefore such techniques cannot be easily integrated into clinical practice. Secondly, the 2D/3D feature selection from breast images still remains an open research field, mainly due to the large variation of the fibroglandular structures in their appearance across modalities, time and individuals. The selected features should be distinctive areas of normal structures in the breast, as the algorithm should not rely only on the presence of lesions.

The 3D-2D matching process is a poorly constrained, ill-posed problem and thus the optimisation is prone to terminate in local, rather than global, minima. Keeping the number of degrees of freedom of the transformation model to a minimum is therefore essential to provide robustness. Subsequently, our strategy is motivated by keeping a low number of degrees of freedom for the transformation model used and also developing a framework that can be easily incorporated into clinical practice.

We propose an MRI to X-ray mammography registration framework that can adapt to incorporate different types of transformation models. We investigate three different transformation models with increased complexity and we validate their performance on clinical cases. All our transformation models are integrated in the Insight Segmentation and Registration Toolkit [ITK, 2003], making the pipeline flexible to adapt different similarity measures and optimisation schemes. In the next chapter we describe a key component of our registration framework, which is the X-ray image simulation from the MRI.

Chapter 3

X-ray simulations from an MR volume

As the MRI is a 3D volume image, while the X-ray is a 2D projection, we facilitate the registration task by simulating an X-ray image using the MRI. The simulated 2D image can then be directly compared with the real X-ray mammogram inside the iterative registration process.

There are two main steps involved in the simulation process. Firstly, the MR intensities are transformed to X-ray attenuation. This is needed because the intensities in the two modalities represent different physical properties of the breast tissue. The contrast in an X-ray image is a result of the difference in the attenuation coefficients between the fatty and the fibroglandular tissue, as discussed in section 1.2.1. Therefore, prior to the intensity transformation, the MR voxels are classified into these two tissue categories (section 3.1). After transforming the MRI to an X-ray attenuation volume, the second part of the simulation process is the perspective projection of the 3D volume to a 2D image (section 3.2). An illustration of the simulation process is given in Figure 3.1. Section 3.3 contains the results of the X-ray simulation process at each step and a comparison of our classification method with other techniques.

3.1 Breast tissue classification

3.1.1 Literature

Tissue classification and segmentation has been an active research field in medical imaging during the past years. There are many different approaches in the literature, especially for the classification of brain tissue into white, grey matter and Cerebro-Spinal Fluid (CSF). There are two main categories of algorithms, non-parametric and parametric, that are further discussed below. It is worth mentioning that there is also a class of algorithms that uses prior shape information of the structure to be segmented; we will not refer further to these, as for the breast case the fibroglandular tissue structures vary significantly across different subjects.

Non-parametric methods classify the image voxels based on their intensity and without assuming any model of the intensity distribution of each class. Wei et al. [Wei et al., 2004] used manual thresholding to obtain a density estimation from the breast MRI and compare this with that conventionally computed from X-ray mammography. For the same application, Nie et al. [Nie et al., 2008] used a fuzzy c-means technique. This is similar to the *k-means* clustering that classifies the voxels into k clusters by assigning them to their nearest mean and iteratively updating the mean of each cluster with its new



Figure 3.1: X-ray simulation process.



Figure 3.2: (a) An example of a pre-contrast breast MRI and (b),(c) increasing magnification views around the red cross that illustrate the partial volume effect. Due to the limited resolution of the MR scanner, one voxel can belong to more than one different tissue types; in this case both fibroglandular and fat.

centroid. The advantage of incorporating fuzzy logic into the k-means algorithm [Chen and Giger, 2004] is that it allows voxels to belong to more than one tissue type (so for example a voxel can be 80% fibroglandular and 20% fat). This is a phenomenon that is observed in all MRIs (particularly of the breast), it is caused by the limited resolution of the MR scanner and is known as the *partial volume effect*. This is illustrated in Figure 3.2. Hipwell et al. [Hipwell et al., 2007] used manual thresholding to obtain the histograms of the segmented tissue types and then applied Gaussian smoothing on them to account for the partial volume effect.

Parametric methods assume that the intensity histogram of each tissue class follows a probability distribution and the image voxels are classified to different classes by fitting the underlying distributions (usually Gaussians) to the data. This is done by following an Expectation-Maximisation (EM) algorithm, explained in further detail in section 3.1.2, that iteratively updates the parameters of the probability distribution. This update can be based on the current estimate of either the Maximum Likelihood (ML) [van Leemput et al., 1999] or the Maximum A Posteriory (MAP) parameters, if a spatial prior is available. Using the spatial prior in the same framework is an advantage of these techniques, especially for the brain tissue classification, for which anatomical priors can be extracted from atlases. Other advantages are the ability to incorporate into the model a bias field correction and a regularisation according to the classification of the neighbouring voxels. Our approach includes both features and is explained in more detail in the next section 3.1.2. It has been observed that the intensity histograms do not always follow a Gaussian distribution. To take into account this variation, authors have proposed using mixtures of Gaussians [Ashburner and Friston, 2005] or power-transformed mixtures [Lee et al., 2009] instead. The parametric approaches provide a probability that each voxel belongs to one of the classes/tissue types. Therefore, the final classification takes into account the partial volume effect, as voxels can belong to more than one tissue type.

3.1.2 Our approach

As in [van Leemput et al., 1999], our method integrates an intensity model, a spatial regularisation scheme and bias field inhomogeneity correction in the same framework. The incorporation of spatial information has been shown to improve classification results in the past as it provides robustness to noise and it allows the use of anatomical information. Specifically for breast tissue classification, MRF regularisation is considered an appropriate choice due to the anatomy of the fibroglandular tissue. Since this is connected in a tree-like structure inside the breast, our hypothesis is that the voxels containing glandular tissue are more likely to appear connected to other glandular voxels rather than isolated inside the fat (and similarly for fat voxels).

The intensity model assumes three classes (for glandular, fat tissue and background) and the bias field is modelled using a third order polynomial basis function. Instead of considering Gaussian distributed intensities corrupted by a multiplicative bias field, log-transformed intensities are used to make the bias field additive. For K classes let $\theta_k = \{\mu_k, \sigma_k\}$ denote the normal probability distribution with mean μ_k and variance σ_k^2 of a voxel belonging to class k and let $z_i = e_k$ be the tissue type of voxel i, where e_k is the unit vector with the k-th component equal to 1 and the others equal to zero. For J basis functions $\phi_j(x)$, $C = \{c_1...c_J\}$ denotes the bias field parameters. The probability density for voxel i, with intensity y_i , given it belongs to class k is:

$$f(y_i|z_i = e_k, \Phi_y) = G_{\sigma_k}(y_i - \mu_k - \sum_j c_j \phi_j(x_i)),$$
(3.1)

where $\Phi_y = \{\theta_1, ..., \theta_k, C\}$ are the intensity model parameters and $G_{\sigma}()$ is a normal distribution with mean zero and standard deviation σ . The model parameters are optimised using an EM algorithm under a Maximum Likelihood formulation. Due to the large variation of glandular structures in the breast across the population, there are no anatomical priors available. If *m* is the iteration number, then the ML estimation gives:

$$\mu_k^{(m+1)} = \frac{\sum_{i=1}^n p_{ik}^{(m+1)} (y_i - \sum_{j=1}^J c_j^{(m)} \phi(x_i))}{\sum_{i=1}^n p_{ik}^{(m+1)}}$$
(3.2)

$$(\sigma_k^{(m+1)})^2 = \frac{\sum_{i=1}^n p_{ik}^{(m+1)} (y_i - \mu_k^{(m+1)} - \sum_{j=1}^J c_j^{(m)} \phi(x_i))^2}{\sum_{i=1}^n p_{ik}^{(m+1)}},$$
(3.3)

where

$$p_{ik}^{(m+1)} = \frac{f(y_i|z_i = e_k, \Phi_y^{(m)})f(z_i = e_k)}{\sum_{j=1}^K f(y_i|z_i = e_j, \Phi_y^{(m)})f(z_i = e_j)}.$$
(3.4)

The intensity model alone can only give accurate results when the different distributions are well separated. This is not the case for the glandular and fat tissue due to many voxels containing both tissue types (partial volume effect). The use of an MRF regularisation scheme improves the overall robustness of the model parameter estimation and provides spatial consistency. Voxels are thus classified based also on the current classification of the neighbouring voxels. In this case, equations 3.2 and 3.3 remain the same, while 3.4 is now given by:

$$p_{ik}^{(m+1)} = \frac{f(y_i|z_i = e_k, \Phi_y^{(m)})f(z_i = e_k|p_{N_i}^{(m)}, \Phi_z^{(m)})}{\sum_{j=1}^K f(y_i|z_i = e_j, \Phi_y^{(m)})f(z_i = e_j|p_{N_i}^{(m)}, \Phi_z^{(m)})},$$
(3.5)

where

$$f(z_i = e_k | p_{N_i}^{(m)}, \Phi_z^{(m)}) = \frac{e^{-\beta U_{mrf}(e_k | p_{N_i}^{(m)} \Phi_z^{(m)})}}{\sum_{i=1}^{K} e^{-\beta U_{mrf}(e_j | p_{N_i}^{(m)} \Phi_z^{(m)})}},$$
(3.6)

with $\Phi_z = \{G, H\}$ being the MRF model parameters and $U_{mrf}(z|\Phi_z)$ the energy function that depends on Φ_z . G and H are $K \times K$ matrices that define the transition energy between classes. Further details of the bias field parameter estimation and detailed explanations of the other equations can be found in [van Leemput et al., 1999].

The above regularisation makes the classification more robust to noise and to isolated misclassified voxels (e.g. isolated voxels classified as fat and surrounded by glandular tissue). Instead of estimating the MRF parameters from the image as in [van Leemput et al., 1999], we use a two-level MRF with its parameters derived from the anatomical properties of the breast [Cardoso et al., 2011]. In the first level, the interclass MRF energy is the same for all classes, thus the MRF only adds global spatial consistency and robustness in the parameter estimation. In the second stage, after the EM converges, the MRF energy matrices (G and H) are altered in order to include more anatomical knowledge (e.g. the cost of having glandular tissue next to the background is higher than having fat next to the background) and the classification is restarted again until convergence. This modification allows an unbiased and robust parameter estimation in the first step followed by a second step that enforces more anatomical knowledge and topological constraints.

The values of the MRF energy matrix are chosen empirically, in order to produce realistic X-ray mammogram simulations. More specifically, in our implementation we have used as matrix G the 3×3 matrix:

	Fat	Gland	Back
Fat	0	be	0
Gland	be	0	ba
Back	0	ba	0

At the first stage of the EM-MRF classification be = ba = 0.15. For the second stage, the parameter ba is altered to ba = 6, to increase the cost of having glandular tissue (Gland) next to background (Back). G is used as the energy matrix for the regularisation within the transverse plane of the volume. For MR volumes with isotropic voxels H = G. Otherwise $H = r \cdot G$, where r is the ratio of the slice thickness over the voxel size within the transverse plane.

The implementation of this algorithm was done by M-Jorge Cardoso, working on brain MRI tissue classification [Cardoso et al., 2011]. It was adapted for application to breast MRI (initialisation of the two tissue type distributions and set of the costs between them) by the author.

The only requirement of the classification algorithm is that the pectoral muscle is pre-segmented from the MRI. This is needed because automated intensity-based segmentation methods are prone to error for this task. This boundary is not well-defined in the majority of the cases, especially when the glandular tissue is very close to the chest wall, or when organs with intensities similar to fat (such as the liver) are adjacent to the rib cage. We use for this task a semi-automated pre-processing method, where the user defines landmarks on the boundary between the pectoral muscle and the breast through which a parametric B-spline surface is subsequently fitted.

After classifying the voxels into fibroglandular and fatty tissue, we then calculate the X-ray attenuation volume intensities. In our first implementation we follow the method proposed in [Hipwell et al., 2007], where these classes are weighted with different factors in order to simulate the difference in X-ray attenuation. The intensity of voxel i in the X-ray attenuation volume is given by:

$$w_G \cdot P_G^i + w_F \cdot P_F^i \tag{3.7}$$

where w_{class} (w_G and w_F) are the weights of each tissue type and P_{class}^i is the classification result for voxel *i*, for each one of the classes ($0 \le P_{class}^i \le 1$). The choice of the most appropriate weights was performed empirically. The goal was to produce simulations with similar contrast to X-ray mammograms. In our later implementation the X-ray attenuation volume is calculated using the methodology described in appendix B, removing the need for the empirical weights w_G and w_F . Our experiments in sections 4.2.1, 4.2.2 and 4.3.1 use equation 3.7. All remaining experiments follow the methodology described in appendix B.

3.2 Perspective projection of the 3D volume

There are two different methods for projecting a 3D volume to 2D, the *parallel* and the *perspective* projection. As the name indicates, the parallel projection assumes parallel rays. Therefore if we assume that the direction of projection is z, then the intensity at each 2D position of the image I_{2D} can be calculated as the integral of the intensities of 3D image I_{3D} along the z direction:

$$I_{2D}(i,j) = \sum_{k=0}^{N_Z - 1} I_{3D}(i,j,k)$$
(3.8)

where *i* and *j* are the indices along the *x* and *y* directions respectively and N_Z is the size of the 3D image along the *z* direction.

Parallel projection was used by all previous MRI/X-ray mammography registration techniques, but this provides only an approximation of the real process. As we know, the rays in X-ray mammography are emitted from a point-like X-ray source (the anode) and thus are not parallel. Perspective projection techniques assume that rays are emitted from a point source. The geometry and related work are described below.

3.2.1 Literature

Several different perspective projection algorithms were proposed in the literature, mainly applied to rigid 2D/3D registration tasks, as the ones reviewed in section 2.6. Ray-casting is a technique that is widely used. To generate the projection image, rays are cast from each pixel in the 2D image through the 3D volume to the X-ray source. The pixel value is then calculated as the integral of the ray intersections with the 3D grid. This technique is explained in more detail in the next section 3.2.2, as we have also incorporated it in our framework. Ray-casting is generally accepted as a methodology that produces high-quality projections and is frequently used as the reference to which other methods

are compared. Nevertheless, its main disadvantage is that it is computationally expensive and this can be critical especially in applications such as image-guided interventions and radiotherapy. Several acceleration techniques have been proposed with a scope to overcome this limitation. These include the use of data that are calculated from pre-computed projections before registration, as in graphics rendering ([LaRose et al., 2000], [Russakoff et al., 2003]), or are updated progressively during registration [Rohlfing et al., 2005].

An alternative to ray-casting is the shear-warp factorisation [Lacroute and Levoy, 1994] that avoids calculating the integrals along each ray by warping the 3D volume to a sheared space, where parallel projection gives the same DRR as the perspective projection in the original space. The method was applied on 2D/3D registration by Weese et al. [Weese et al., 1999]. This technique is mainly used for applications where the image quality of the projection image is not of great importance, as the resampling steps involved in warping the 3D volume can lead to loss of detail in finer structures.

Another alternative to ray-casting is the direct projection of each voxel in the 3D image on the 2D plane and the computation of the final DRR by interpolation on the projection plane. This methodology is called splatting [Westover, 1990]; it was used in medical 2D/3D registration by Birk-fellner et al. [Birkfellner et al., 2005] and was later further accelerated by a GPU implementation [Spoerk et al., 2007]. The main disadvantage of this method is the possibility of obtaining a 2D image with spatially varying intensities, as the interpolation is only performed on the 2D space.

3.2.2 Our approach

The registration of MR volumes to X-ray mammograms is a process that can be performed off-line. The registration result (ie. the transformation parameters) can then be stored in the workstation and used by the radiologist when assessing the image findings. Therefore, although the registration time is important for any clinical application, it is not as critical as in image guided procedures. At the same time, the image quality of the simulated projection image should be high, as the similarity between this projection and the real mammogram drives the registration. For these two reasons, we use in our framework a ray-casting algorithm for the perspective projection of the 3D volume.

The geometry of a ray-casting projection is shown in Figure 3.3. To simulate a mammogram, the X-ray attenuation volume V is positioned above the projection plane. The distance between the plane and the X-ray source S is called *focal length* and is typically around f = 660mm. This information is recorded in the digital DICOM images. For each grid position on the 2D plane $P_{i,j}$ we define the ray by the point $P_{i,j}$ and the direction \vec{d} :

$$\vec{d} = S - P_{i,j} = \overrightarrow{P_{i,j}S}$$
(3.9)

We then calculate the coordinates $r_1...r_n$ of the ray at regular intervals along the ray's path through the volume V. These are shown in Figure 3.3 as red crosses. As the ray passes through the 3D volume, these ray coordinates occur at non-grid positions. We therefore need to use interpolation at each location. Lemieux et al. [Lemieux et al., 1994] used tri-linear interpolation. In this case, the value at each location point is calculated from the nearest eight neighbouring grid positions. To accelerate this computationally expensive process, Penney et al. [Penney, 2000] used bilinear interpolation at each intersection of the



Figure 3.3: Perspective projection geometry.

ray with the slices of the 3D volume (so four neighbouring positions are used for interpolation). Our implementation uses the same principle, as this is more efficient computationally, without affecting significantly the image quality. Finally, the intensity at the pixel position $P_{i,j}$ is given by the integral of the intensities at the above intersections:

$$I_{2D}(i,j) = \sum_{k=1}^{n} V(r_k)$$
(3.10)

3.3 Experiments

In our initial registration experiments with data from five patients [Mertzanidou et al., 2010b] we have used manual thresholding for the breast tissue classification (as Hipwell et al. [Hipwell et al., 2007]). In [Mertzanidou et al., 2010a] we then proposed the methodology described in section 3.1.2. Figure 3.4 shows a visual comparison between the two classification methods on two patients. The data used in these experiments are described in section A.1.

In the figure we can see that the main advantage of the proposed EM-MRF method is that the final simulated mammogram contains more details of the glandular tissue. It is also fully automated. Other advantages over manual thresholding are the fact that this gives reproducible results, it takes less than one minute to process each breast volume and it requires minimal pre-processing interaction: the only requirement is that the pectoral muscle is pre-segmented from the MRI, as discussed above.

Here we investigate further the benefits of the proposed EM-MRF method, by showing results in comparison with the *k*-means algorithm and the use of the intensity model of section 3.1.2 alone, without the MRF regularisation. Figure 3.5 shows the pre-contrast MRI of a patient and the corresponding histogram. This is an example histogram which illustrates the difficulty in segmenting the two tissue types, as these do not appear clearly separated in the image histogram. In this instance, using intensity



Figure 3.4: (a) Pre-contrast MRI, (b) X-ray attenuation volume using manual thresholding and histogram-based classification, (c) using the EM-MRF algorithm, (d) Simulated X-ray mammogram from the undeformed volume using manual thresholding, (e) using EM-MRF. The two rows correspond to two patients. The red cross indicates the position of a corresponding coordinate in each image.



Figure 3.5: (a) Pre-contrast MRI of a patient and (b) the corresponding histogram of the intensities.

information only (as in the *k-means* algorithm or the fitting of Gaussian distributions) does not produce accurate results, as illustrated below.

Figure 3.6 shows the classification results of the three techniques tested. The *k-means* algorithm gives binary segmentations of the fat and the glandular tissue, and as we can see in Figure 3.6(a) the classification is clearly affected by the bias field in the MRI. The benefit of including the MRF regularisation to incorporate neighbouring voxel information is also evident from the images of Figure 3.6. We can see that the MRF provides smoother areas across the images, removing erroneous voxel classifications as glandular tissue, when these are all surrounded by fat. This assumption is particularly useful for the breast case, as glandular tissue appears mostly as a tree-like structure (and therefore connected) inside the breast. Moreover, smoothing reduces the effect of noise and results in increased contrast between the glandular and the fat tissue, which is desirable when simulating a mammogram, as the glandular structures are then more visible. Figure 3.7 shows the simulated X-ray attenuation volumes and the simulated mammograms that both methods produce. More examples of simulated mammograms using the EM-MRF method are given in the next chapters, as results of the MRI/X-ray registration process.

Finally, it is worth mentioning that although the bias field correction works reasonably well in our framework, there are certain cases, where the bias field is very strong, which results in failure of the classification in certain areas. An example is shown is Figure 3.8. In breast MRI the coils are very close to the breast surface and therefore the intensity inhomogeneity can be significant. Future work could investigate the use of alternative bias-field correction methods for those cases. A comparison of existing algorithms specifically for breast MRI was performed by Makarau et al. [Makarau et al., 2010], where the *mean-shift* algorithm was shown to outperform other techniques.

3.4 Discussion

In this chapter we have presented the X-ray simulation process in detail and the intermediate steps that are involved. Regarding the breast tissue classification, our approach produces automatically segmentations that are visually realistic, while performing bias-field correction. Moreover, the simulated X-rays contain



Figure 3.6: Breast tissue classification results for the patient volume displayed in Figure 3.5, using the k-means algorithm, the EM algorithm without and with the MRF regularisation.



Figure 3.7: The X-ray attenuation volumes and the simulated mammograms produced using the EM algorithm only and the EM-MRF. The pre-contrast MRI and the real mammogram are given for comparison. Note that at this stage there is no registration or deformation of the volume before projection. The simulated mammograms come from the original undeformed MR volume.



Figure 3.8: (a) Pre-contrast MRI with strong bias field and (b) the X-ray attenuation volume. The inhomogeneity correction was not successful for this case.

more fibroglandular structures than manual thresholding.

Quantitative validation of this algorithm is problematic due to the lack of a ground truth classification of the breast tissue. The fine nature of the ductal network, and relatively coarse resolution of the MRI in comparison, mean that the partial volume effect is substantial. For this reason manual delineation of the ductal network to produce such a ground truth data set is impractical. The choice of the EM-MRF approach for this application therefore, was made based on its advantages over manual thresholding discussed above, and also because it has been successfully applied and validated for the classification of different tissue types visible in MR images (albeit brain tissue classification) by van Leemput et al. [van Leemput et al., 1999]. We believe the use of physical X-ray reference data in the simulation process creates images with intensity characteristics close to that of real X-ray mammograms. However given the lack of an appropriate validation strategy, this assumption must be evaluated, alongside our choice of the other registration components, with reference to the final registration error obtained by the experiments described in the following chapters. An alternative automated method targeted for this task was recently proposed by Gubern-Merida et al. [Gubern-Merida et al., 2011].

Regarding the perspective projection methodology, ray-casting produces high quality images, with a higher computational cost in comparison to other techniques. In the next two chapters, where we combine the volume transformation with the ray-casting we will review techniques that accelerate this process.

Using the methodology described in this chapter we can obtain simulated X-ray images that can be directly compared with the real mammograms. Nevertheless, there are several difficulties that still remain and are associated with this comparison. Firstly, even if we assume that we know the exact deformation of the breast that occurs during mammographic compression, the fibroglandular structures in the simulated X-ray image will not have the same appearance as in the real mammogram. The main reason for this is that the image is generated from the MR signal of the tissue, which is a different physical property than the X-ray attenuation and therefore the exact relationship between them is not known. Other aspects that differ between the two images are the spatial resolution and the factors that contribute to the degrading of the image in real X-ray acquisition, such as the quantum noise of the Xray source and the scatter. These factors were reviewed in section 1.2.1, where the X-ray mammogram acquisition was described in detail. Despite the remaining differences between the images, the X-ray simulation process facilitates the MRI to X-ray registration task, providing images that can drive the registration, when using an appropriate similarity measure between them. This is evident both in the visual and in the quantitative results of the registration experiments described in the next chapters.

Chapter 4

MRI to X-ray registration using an affine transformation

In this chapter we describe in detail our 2D/3D registration framework and the experiments carried out for the alignment of MR volumes to X-ray mammograms. In all these experiments we use an affine transformation model to approximate the breast deformation between the prone position in the MR scanner and the compression between two plates during X-ray mammography acquisition. As described in section 2.1, the affine transformation consists of a rigid-body transformation plus scaling and shearing. Our hypothesis is that the rigid-body transformation can account for the initial positioning of the breast between the X-ray source and the detector, while the scaling and shearing can approximate the deformation caused by the plates' compression. Although the real breast deformation is more complex, the goal of these first experiments is to investigate whether a global registration with a simple geometrical model can give useful clinical accuracy. The main advantage of the affine transformation is the low number of degrees of freedom (twelve) which is expected to add robustness to the ill-posed 2D/3D matching problem.

The chapter is structured as follows. Section 4.1 describes the general 2D/3D registration framework and its components. The next two sections (4.2 and 4.3) describe the experiments carried out on simulated and real mammograms respectively. Finally, advantages and limitations of this approach are discussed in section 4.4.

4.1 Registration Framework

An overview of the registration framework is given in Figure 4.1. The main registration components are further described in the next sections. The inputs to the registration pipeline are the fixed, or target image (the real X-ray mammogram) and the moving, or source image (the X-ray attenuation volume estimated from the MRI). Prior to registration, the user specifies a *metric* (similarity measure), to be used to assess the similarity between the two images. During the registration the value calculated by the metric is used by the *optimiser*, which iteratively updates the *transformation* parameters. In this chapter, we confine our investigations to the affine transformation. Then, the *interpolator* is used to compute the values of the moving image at the pixel positions of the fixed image, by projecting the 3D volume into



Figure 4.1: Overview of the 2D/3D registration framework. The processes are illustrated in blue and the data in red. The output of the registration is used for mapping between the MRI and the X-ray image.

2D. The resulting image is used again for comparison with the target, where a new iteration begins. The projection from 3D to 2D is performed by the interpolator. The process is repeated until the stopping criteria are satisfied (for example the similarity measure is a maximum).

The main difference of the 2D/3D process in comparison to a conventional intensity-based registration pipeline is that the interpolator component has been modified in order to incorporate the ray-casting from the X-ray source to the 2D detector. In a registration task between two images of the same dimensionality, every grid position of the fixed image is being transformed and the intensity at that point is given by the interpolation of the intensities at the neighbouring pixel or voxel positions in the moving image. In the 2D/3D case, the interpolation occurs as the ray transverses the moving volume grid, in the way described in section 3.2.2. Consequently, the target mammogram is compared with a 2D projection of the X-ray attenuation volume at each iteration.

There are different options for combining the transformation and the projection of the 3D volume. One would be to resample the 3D grid into a new, transformed volume and then perform the ray-casting through the new grid. The main drawback of this method is the high computational cost associated with an extra 3D interpolation followed by ray-casting. Moreover, the transformed 3D volume is not needed, since we only use the simulated projection to compare it with the real mammogram. A different method makes use of the fact that the affine transformation preserves parallelism. In other words, the ray that passes through the affine-transformed volume remains a straight line. For this reason, instead of transforming the 3D grid, we can transform instead the ray (ie. the projection plane and the X-ray source). The same approach was used previously for rigid registration tasks ([Penney et al., 1998], [Hipwell et al., 2003]).

We have seen the effect of an affine transformation on a volume in section 2.1. The 3D affine

transformation has twelve parameters; these include translations (t_x, t_y, t_z) , rotations (r_x, r_y, r_z) , scaling (s_x, s_y, s_z) and shear $(sh_{xy}, sh_{yz}, sh_{xz})$. We create an individual matrix for each one of these parameters and concatenate them into one matrix T:

$$T = T_{translation} \cdot (T_{rotation} \cdot (T_{scaling} \cdot T_{shearing})). \tag{4.1}$$

Instead of optimising the matrix coefficients of T, we optimise the affine parameters defined above. This allows us to easily include the volume preservation constraint as explained below.

The source volume of the registration framework is the MRI that is acquired in the prone position. From this point onwards we refer to the MRI breast shape as the *undeformed* breast. Our first experiments, that are described in section 4.2.1, have shown that without any constraints the breast volume can increase during registration [Mertzanidou et al., 2010b]. This happens due to the fact that the 2D/3D alignment is an ill-posed problem, and thus during optimisation the scaling in the direction of the projection, if it is not constrained, can be trapped in local minima or lead to physically unrealistic expansion or contraction of the breast. To avoid such non-physical volume variations, for example an expansion for CC views in the superior-inferior and also in the posterior-anterior direction, we include a volumepreservation constraint, by ensuring that the product of all scaling factors across the three dimensions is unity ($s_x \cdot s_y \cdot s_z = 1$). This is done by constraining the scaling on the direction of the projection (superior-inferior for a CC view) to be

$$s_z = \frac{1}{s_x \cdot s_y}.\tag{4.2}$$

This constraint removes one degree of freedom from the optimisation process, reducing the size of the search space and potentially enhancing the robustness of the registration [Mertzanidou et al., 2010a].

A good initial position of the volume before registration is important as it can lead to fast convergence and it also reduces the likelihood of the optimisation getting trapped in local minima. The projection geometry used is shown in Figure 4.2. The distance between the X-ray source and the detector can be extracted from the DICOM header of the Full-Field Digital Mammograms (FFDM). The initial translation of the volume in the direction of the projection (z axis) ensures that the volume is positioned on top of the detector. The volume is also translated in the perpendicular plane (xy plane) such that the centre of mass of the volume is projected onto the centre of mass of the real mammogram. Let $O_V = [0, 0, 0]$ be the origin of the volume V and O_P the origin of the projection plane. Then if f is the distance between the X-ray source S and the projection plane, the X-ray source is positioned such that:

$$S = [M_{Vx}, M_{Vy}, -(f - d_{Vz})],$$
(4.3)

where $M_V = [M_{Vx}, M_{Vy}, M_{Vz}]$ is the centre of mass of the volume and d_{Vz} is the size of the z dimension of the volume V in mm. The origin of the projection plane O_P is:

$$O_P = [M_{Vx} - M_{Tx}, M_{Vx} - M_{Ty}, d_{Vz}],$$
(4.4)

where $M_T = [M_{Tx}, M_{Ty}, M_{Tz}]$ is the centre of mass of the target image (real mammogram).

Regarding the initial orientation of the volume before registration, this is illustrated for the CC view in Figure 4.2(b). The initial rotation, scaling and shear parameters are set to zero. For the MLO view, the



Figure 4.2: (a) Projection geometry used in the registration, showing the X-ray source, the breast volume V and the projection plane. Before registration the volume V is positioned such that f = 660mm. This is extracted from the DICOM header of the real mammogram. (b) Illustration of the initial position of the breast volume in V prior to registration for a CC view.

volume is rotated about the y axis (r_y =45°), to account for the positioning of the X-ray source and the detector and also about the z axis to account for the in-plane rotation of the breast that usually occurs in the MLO view mammograms. We experimentally determined that an initialisation of r_z =30° produces a good initial position for the registration. An example of the DRRs generated for one case at iteration 0, for CC and MLO view registrations, is given in Figure 4.3.

In the next two sections we describe in detail the different similarity measures and optimisation techniques that we tested in our framework, as part of the work described in [Mertzanidou et al., 2010b].

4.1.1 Similarity measures

As the name of this registration component indicates, the similarity measure provides a quantitative measure of the similarity between the fixed image and the resampled moving volume, after it has been transformed (and in our case projected). There are several different similarity measures that can be used, such as *Mean* or *Sum Squared Differences*, (*Normalised*) *Cross Correlation*, *Gradient Difference*, *Gradient Correlation*, (*Normalised*) *Mutual Information* etc. The similarity measures that were used in our experiments are further explained later in this section. The choice of the most appropriate similarity measure is application-specific. Each one of them has its own characteristics that make it most suitable for a certain registration problem. For example, some perform better in cases where the images have similar intensity ranges or others when the image intensities are not related with a linear relationship (like in multimodal registration).

We have seen in the literature review (section 2.6) that authors have used various similarity measures for 2D/3D registration. Between these, it has been shown that gradient and correlation based measures perform best. A comparison of different metrics that has been done by Penney et al. [Penney et al., 1998], showed that apart from Gradient Difference, another metric that performs well





(b) MLO view

Figure 4.3: Illustration of the initial position before registration for the (a) CC and (b) MLO view of one case. Left: Target mammogram in the registration space and right: overlaid with the DRR outline shown in green, created at iteration 0 of the registration process.

is *Pattern Intensity*. Nevertheless, this was due to the application of that specific registration task, where in one of the image there was an interventional instrument present that did not appear in the second image. Pattern intensity has been proven to be robust to these changes. As in our case there are no additional objects present in one of the images, the suitability of this similarity measure is not examined. Instead, since the intensity ranges between the two images are similar for the first set of experiments (section 4.2.1), apart from the Normalised Cross Correlation and Gradient Difference metrics, we also tested the performance of Mean Squared Differences. Moreover, we used Mutual Information, which is a popular measure in multimodal registration. In this case, we do not need to simulate an X-ray at each iteration of our algorithm, as the relationship between the images that are compared does not have to be linear. Instead we use directly a perspective projection of the segmented MR volume.

In our first set of experiments, that are further described in section 4.2.1, we have tested the performance of the similarity measures that are most suitable for this 2D/3D registration task, as discussed. The description of these similarity measures is given below.

Mean Squared Differences (MSD) This similarity measure computes the pixel-wise differences of the intensities between the two images and returns the mean of these squared differences across the whole image or the region of interest that is specified. In mathematical terms, for two images A and B, this similarity measure is given by:

$$MSD(A,B) = \frac{1}{N} \sum_{i=1}^{N} (A_i - B_i)^2,$$
(4.5)

where A_i is the intensity at the *i*-th pixel in image A, B_i is the *i*-th pixel's intensity in image B and N is the number of pixels in the region of the two images that is considered for registration. For a perfect

registration result, this measure gives a zero value if the intensities in the two images are identical.

This metric is very sensitive to differences in images at alignment and subsequently is only used for registration of images that come from the same modality, acquired using identical imaging parameters, and assumes that there is no tissue change between the imaging time points. The purpose of testing the performance of this metric in our case is that for the first set of experiments, in section 4.2.1, we only use simulated X-ray images and thus this metric should be able to drive correctly the registration. For the experiments on real X-ray mammograms this similarity measure is not expected to perform as well. **Normalised Cross Correlation (NCC)** This metric assesses the similarity between two images by computing the pixel-wise cross-correlation between them and normalising the result by dividing by the square-root of the auto-correlation of the two images. The equation for images *A* and *B* is:

$$NCC(A,B) = \frac{\sum_{i=1}^{N} ((A_i - \overline{A}) \cdot (B_i - \overline{B}))}{\sqrt{\sum_{i=1}^{N} (A_i - \overline{A})^2 \cdot \sum_{i=1}^{N} (B_i - \overline{B})^2}}.$$
(4.6)

where \overline{A} and \overline{B} are the mean values of the intensities in image A and B respectively. Higher values in correlation indicate greater similarity and since the result is normalised, a perfect registration result gives NCC(A, B) = 1.

Normalised cross-correlation is a similarity measure that is suitable for images that come both from the same and also from different modalities with a linear intensity relationship, due to the normalisation done when dividing with the autocorrelations of the images. Although this measure is sensitive to the presence of additional structures in one of the two images, this case does not appear in any of our experiments and so this similarity measure is considered to be appropriate for our task.

Gradient Difference (GD) Gradient-based measures use the gradient instead of the intensity information of the two images. The gradient images on the x and the y direction are computed using the Sobel operator. More specifically, gradient difference operates on the differences of those gradient images, that are given from the equations below:

$$I_{diffH}(i,j) = \frac{\partial A}{\partial i} - \alpha \frac{\partial B}{\partial i}$$
(4.7)

$$I_{diffV}(i,j) = \frac{\partial A}{\partial j} - \alpha \frac{\partial B}{\partial j}, \qquad (4.8)$$

where A and B are the original images and α is a scaling factor. The gradient difference is then given from:

$$GD(A,B) = \sum_{i,j} \frac{A_h}{A_h + (I_{diffH}(i,j))^2} + \sum_{i,j} \frac{A_v}{A_v + (I_{diffV}(i,j))^2},$$
(4.9)

where A_v and A_h are constants and are equal to the variances of the gradient images I_{diffV} and I_{diffH} respectively. As we can see from the above equation, the better the registration of the images, the higher the value of the similarity measure.

Gradient Difference minimises the effect of the lower spatial frequencies of the images, as the filtering enables the alignment based on the edges of the structures. This is a desirable property for the registration of breast images, as we are interested in aligning the fibroglandular structures that create edges when projected from the 3D volume to 2D.

Mutual Information (MI) Mutual information estimates the similarity between two images by calculating their entropy. For a random variable A (that in our case is the image intensity of the image A), the entropy is given by:

$$H(A) = \sum_{a} p_A(a) \log p_A(a),$$
 (4.10)

where $p_A(a)$ is the probability distribution of the intensities in image A. If the intensities of two images A and B are independent, then their joint entropy H(A, B) is equal to the sum of their entropies H(A) and H(B). Otherwise, the difference between those values can give us a measure of the dependency between the two images. This measure is called mutual information and it is given by the following equation:

$$MI(A, B) = H(A) + H(B) - H(A, B)$$
(4.11)

or alternatively:

$$MI(A,B) = \sum_{a,b} p_{AB}(a,b) \log \frac{p_{AB}(a,b)}{p_A(a)p_B(b)}.$$
(4.12)

The higher the mutual information is, the greater the dependency between the two images. The main characteristic of this measure is that the type of dependency between the two variables does not have to be specified and that makes it a widely used similarity measure for multimodal registration.

This metric was introduced for image registration by Viola and Wells ([Viola and Wells, 1997], [Wells et al., 1996]) and [Collignon et al., 1995]. Similar approaches were also proposed at the same time ([Studholme et al., 1996], [Studholme et al., 1999]). The marginal and joint probability densities are estimated by drawing a random sample of elements S and using the following equation:

$$p_A(a) \approx P^*(a) = \frac{1}{N_S} \cdot \sum_{s_j} K(a - s_j)$$
 (4.13)

where K is a window function that integrates to 1 (usually the Gaussian density function) and N_S is the number of the samples used. $P^*(a)$ is known as the Parzen window density estimate. Finally, the entropy of variable A is computed using a second intensity sample R and is given by:

$$H(a) = \frac{1}{N_R} \cdot \sum_{r_j \in R} P^*(r_j) \cdot \log P^*(r_j) \approx \frac{1}{N_R} \cdot \sum_{r_j \in R} \log P^*(r_j),$$
(4.14)

where N_R is the number of samples in R.

Mutual information does not assume a linear relationship between the two images. Therefore, when using this similarity measure we do not need to simulate an X-ray at each iteration of the registration (by segmenting the MRI and computing an X-ray attenuation volume), as the two images that will be compared do not need to have intensities in the same range. Instead, we use a simple perspective projection of the original MR intensities.

4.1.2 **Optimisation**

Generally speaking, optimisation is the process of searching for the best element, between a number of alternatives. In mathematical terms, optimisation techniques aim to maximise or minimise a function by determining its maximum or minimum value respectively. For every optimisation problem, we need

to define the *function*, whose extremum we want to locate, the *variables* or *parameters* that we vary during this search and the range of the search space. Optimisation is an iterative process that maximises/minimises the defined function by varying the parameters in a specified range and finding the optimum values.

In the image registration problem, the function is the similarity measure (as discussed in section 4.1.1) and the variables are the parameters of the transformation that we use (as described in section 2.1). At each iteration of the registration the parameters are updated to give a new instance of the moving image. This is then compared with the target image and the result of the similarity measure (cost function) is passed to the optimiser, where a new iteration begins.

There is a wide range of different optimisation algorithms, such as the Gradient Descent, Conjugate Gradient, Evolutionary Optimisers and Powell Optimiser. They all have several advantages and drawbacks and the choice is application specific. When the function does not have a "noisy" behaviour (rapid changes in certain directions, for small changes of the variables) and the initial position is close to the total minimum/maximum, then the most commonly used optimiser is the Gradient Descent. The behaviour of this optimiser is explained in more detail below.

Gradient Descent optimiser The Gradient Descent optimiser (also known as steepest descent) is searching for a function's minimum (maximum) by moving towards the negative (positive) direction of its gradient. We assume that we have an one-dimensional function f(x) that is defined in a specified range of x and also its gradient can be calculated (or approximated) in the same area. Then, the gradient descent method is based on the fact that f(x) decreases faster if one moves from a certain point P to the direction of the negative gradient of the function at that point. According to this, the optimiser starts from an initial estimate of the function's minimum at point x_0 and iteratively updates this position with a new estimation $x_1, x_2, x_3...$, where x_i is given from the equation:

$$x_i = x_{i-1} - \epsilon_i \nabla f(x_{i-1}), \tag{4.15}$$

and ϵ_i is a small positive value that defines the step of the optimiser which can change in each iteration. As f(x) decreases in the direction of its gradient, then for a sequence of points x_n : $f(x_0) \ge f(x_1) \ge$ $f(x_2)$... etc, corresponding to a number of iterations, the optimisation will converge to a minimum.

In image registration, the function to be minimised/maximised is the similarity measure and the variables are the transformation parameters. The gradient descent method is widely used, especially for metrics that give smoothly variable values for small changes in the parameters of the transformation. In our experiments, we are using a variation of this optimisation technique, that follows an update scheme based on a user-defined step-size. For a set of X parameters to be optimised, each parameter x is updated according to:

$$x_{i} = x_{i-1} \pm \frac{\nabla f(x_{i-1})}{w(x)} \cdot \frac{step}{\sqrt{\sum_{k=1}^{X} (\frac{\nabla f(k_{i-1})}{w(k)})^{2}}}$$
(4.16)

where $\nabla f(x)$ is the magnitude of the gradient of the similarity measure with respect to parameter x and w(x) is a scalar weight factor that controls the relative magnitude of the step size step for each parameter. The *step* size decreases during optimisation if the direction changes, indicating that the similarity measure is close to the minimum value. The optimisation terminates when the magnitude of the gradient, for all parameters, is smaller than a pre-defined tolerance value.

The gradient descent optimiser does not perform well for metrics that have either a noisy behaviour, various local minima or a long narrow valley that leads to the global minimum. In these cases we can use other alternatives, such as an evolutionary optimiser as described below.

Evolutionary Optimiser Evolutionary optimisers approach an optimisation problem in a way that is inspired by biological evolution, as their name indicates. At each iteration, this optimiser selects randomly a new position for the variables inside the parameter space. This choice is controlled by a probability function that is centred at the current value of the variables and it either grows or shrinks, according to the following update rule:

$$\overline{x_{t+1}} = \overline{x_t} + a_t \cdot \overline{r_t} \tag{4.17}$$

$$a_{t+1} = \begin{cases} a_t \cdot c_{grow}, & \text{if } f(\overline{x_{t+1}}) < f(\overline{x_{opt}}) \\ a_t \cdot c_{shrink}, & \text{otherwise} \end{cases}$$
(4.18)

$$\overline{x_{opt}} = \begin{cases} \overline{x_{t+1}}, & \text{if } f(\overline{x_{t+1}}) < f(\overline{x_{opt}}) \\ \overline{x_{opt}}, & \text{otherwise} \end{cases}$$
(4.19)

where c_{grow} , c_{shrink} , a_0 and x_0 are given as inputs and $\overline{r_t} = N(\overline{0}, \overline{I})$ is a random vector with isotropic normal distribution of zero mean and one variance. There are several different variations of the above notation in the literature. In our experiments, we use the modified version of Styner et al. [Styner et al., 2000], where the normal density function is updated by using a matrix A that relates to the covariance matrix $\Sigma = A \cdot A^T$, instead of a simple scalar a_t .

Evolutionary optimisers are well-suited for optimisation of functions with noisy and random behaviour. In our experiments, we are using this optimiser in combination with the mutual information (MI) similarity measure, as the gradient descent optimiser did not perform well with the MI.

4.2 Experiments on simulated mammograms

As a first step of our MRI to X-ray mammography registration study and in order to ensure that the proposed framework performs reasonably in a simple registration task, we used only simulated X-rays, rather than real mammograms, as the target images for the following experiments. These simulated images were created using the method described in chapter 3. This way, the similarity between the generated DRR at each iteration of the algorithm and the target image is high if the transformation is recovered correctly. The only difference in the way that the target image is created in comparison with the DRR at each iteration, is that it contains Poisson noise, to account for the the quantum noise that occurs when photons are emitted from any X-ray source, as discussed in the previous chapter (section 3.4).

In the experiments described below, we used the reprojection error ([Masutani et al., 1997], [Hipwell et al., 2003]) for evaluation. The main concept is illustrated in Figure 4.4. All the points p_i



Figure 4.4: Illustration of the reprojection error used by Masutani et al. [Masutani et al., 1997] and Hipwell et al. [Hipwell et al., 2003] for the evaluation of 2D/3D registration tasks.

in the volume are projected on the 2D plane to give the points u_i , by using the ground truth affine matrix G. Then, those points are reprojected back into 3D space, producing a line L_i , using the affine matrix T resulting from the registration process. For each voxel coordinate in the volume P_i , the reprojection distance D_i is equal to the minimum distance of the point P_i to the line L_i . For a perfect registration the line L_i passes through p_i and thus the distance is zero. For the whole volume, the reprojection error is given by the equation:

$$D = \sqrt{\frac{\sum_{i=1}^{N} D_i^2}{N}},$$
(4.20)

where N is the number of voxels in the volume. The use of this metric is mostly suited for vascular imaging applications, where the imaging object is positioned approximately in the middle of the distance between the X-ray source and the detector. In mammography the breast is close to the detector, so the reprojection error gives almost identical results to the 2D distance of the points (lesions or landmarks) in the plane of the detector. Consequently, we later use the 2D distance as the error metric.

In the experiments described in section 4.2.1, the goal was to test different similarity measures to identify the one that performs best, while we recover a known 3D affine transformation. Then, in section 4.2.2, we use simulated mammograms generated from real breast compressions, to evaluate how accurately an affine transformation can approximate the breast compression in a 2D/3D registration task.

4.2.1 Recovering a known affine transformation

The goal of this set of experiments was to examine the performance of several similarity measures (and optimisers where necessary) in a 2D/3D registration, in order to recover a known 3D affine transformation of the breast volume. In all the experiments we have used simulated CC view X-rays from the MR volume. Initially, we applied a known affine transformation to the MRI. This deformed volume was used to to simulate an X-ray image that would be the target image of the registration. The original, undeformed MRI was used as the source image. These experiments were part of the work described in [Mertzanidou et al., 2010b].

In these experiments we have used MR images of five patients, described in section A.1. The

Affine parameter	Distribution
translation in X, along plate position	Normal, $\mu = 0, \sigma = 18mm$
translation in Y, direction of the rays	Normal, $\mu = 0, \sigma = 5mm$
translation in Z, closeness	Uniform, $min = -15mm$, $max = 15mm$
rotation in Y, in-plane	Uniform, $min = -15^{\circ}$, $max = 15^{\circ}$
rotation in Z, rolling	Uniform, $min = -11^{\circ}$, $max = 11^{\circ}$
scaling in X, expansion	Uniform, $min = 0.5$, $max = 0.9$
scaling in Y, compression	Uniform, $min = 1.1$, $max = 1.5$
shear between X and Y axis	Normal, $\mu = 0, \sigma = 0.25$

Table 4.1: Distributions from which we have randomly chosen the parameters of the affine transformation, in order to deform the MR volume and create the target images. The coordinate system that we refer to is shown in Figure 4.5.

MR volumes have voxel dimensions $[1.33 \times 1.33 \times 2.5]mm^3$ and a coronal slice orientation. The simulated mammograms have a pixel size of $[1 \times 1]mm^2$. Although real mammograms have typically finer resolution, the simulated ones are limited by the MRI. Given that the MR voxels are not smaller that 1mm at each direction, we have used $[1 \times 1]mm^2$ for the 2D pixel dimensions. This way we ensure that the accuracy is not affected, while at the same time the computational cost of the ray casting algorithm is reduced, as rays are cast from each pixel position to the X-ray source for the mammogram simulation.

The choice of the appropriate affine transformation parameters was done based on the parameters that were used by Tanner et al. [Tanner et al., 2008] for the construction of breast Statistical Deformation Models. To approximate the compression of the breast, we used scaling along the two in-plane axes. To approximate the volume conservation constraint when a breast is compressed, the choice of the scaling parameters ensured that expansion on one direction leads to compression by the same amount on the other direction. The distributions from which we randomly chose the affine parameters are shown in Table 4.1. These include rolling, in-plane rotation, translation in 3 directions, scaling and shearing. The coordinate system that these parameters refer to, is illustrated in Figure 4.5.

We have performed in total 40 experiments $(5 \times 2 \times 4)$, for which we have used 5 patients, 2 random affine transformations per patient from the distributions in Table 4.1 and the 4 similarity measures that were discussed in section 4.1.1. The choice of two affine transformations was made in order to perform a certain number of registration tasks that would indicate the similarity measures that are most appropriate for this task. As the number of available MRIs is five, we consider that 40 registration tasks are adequate for this first set of experiments. The optimiser used was evolutionary for the Mutual Information and Gradient Descent for the other metrics. Some representative results for all patients are shown in Figure 4.6. The detailed reprojection error results are given in Table 4.2 for all four similarity measures.

The results on the first set of data indicate that simulating an X-ray image performs better than a direct projection of the MR intensities and the use of Mutual Information. Both visual and numerical



Figure 4.5: Illustration of the coordinate system used in the experiments with simulated mammograms (CC views). The parameters of the affine transformation in Table 4.1 were chosen according to this coordinate system.

Table 4.2: Reprojection error results (in mm) of the experiments on simulated X-ray data using all four similarity measures. The experiments were carried out using the MRIs of 5 patients. For every patient, there were two known affine transformations applied (aff. 1 and aff. 2), as explained in section 4.1.

	patient 1		patient 2		patient 3		patient 4		patient 5	
	aff. 1	aff. 2								
before reg.	11.14	11.60	10.14	15.9	24.73	12.39	8.93	15.95	21.77	15.22
after reg.(GD)	1.54	6.55	4.03	1.98	5.91	11.09	0.61	5.76	13.73	1.60
after reg.(NCC)	3.42	1.07	3.57	2.31	6.78	8.69	0.91	5.73	62.4	6.0
after reg.(MSD)	16.54	2.78	16.95	20.3	29.53	13.52	13.73	24.13	17.78	7.34
after reg.(MI)	20.82	40.82	45.94	36.40	50.82	11.26	29.87	29.00	82.25	72.56

results show that the registration using MI has failed. Between the rest of the three similarity measures, Normalised Cross Correlation and Gradient Difference performed best. In these two cases, the mean reprojection error across all 5 patients was reduced from 14mm before registration to 4.34mm for the GD and 4.27mm for the NCC. We can see in Table 4.2 that for both metrics, the error decreased after registration for all cases, apart from one (patient 5, affine 1). Looking closely at this case in Figure 4.6 (row 5), we can see that this patient has a breast that is almost entirely composed of fatty tissue. As a result, the registration algorithm is expected to terminate in local minima, since there are few fibroglandular structures to contribute to the image similarity and so the registration is only affected by the alignment of the breast boundaries. This case was excluded from the calculation of the mean reprojection error.

The results in Figure 4.6 demonstrate that GD and NCC performed best. MI is the only metric that failed to register the images in most of the cases. MSD has not performed as well in general and the images were in most cases aligned only on the breast boundaries. This is a metric that will not be useful when using real X-ray mammograms, due to its sensitivity to scaling or offset intensity differences between the images. In terms of performance, NCC converged in less iterations than the other similarity measures. The experiments indicate similar performances for the tested similarity measures, as the ones reported for other 2D/3D registration tasks in the literature (section 2.6).

4.2.2 Recovering a real breast deformation

The goal of this second set of experiments was to evaluate how well the affine transformation performs when registering the undeformed MR volume to a simulated X-ray that was created using an MR image of a real breast compression. In these experiments we used a series of real MR compressions of the breast from 8 volunteers, described in [Tanner et al., 2011] and section A.2, in the lateral to medial direction. Figure 4.7 illustrates the images acquired for one volunteer. In our experiments we used the undeformed MRIs (for the source images) and those that correspond to the maximum amount of compression (to simulate the target X-ray images). Figure 4.8 shows examples of the MR compression images (volume sizes: $1mm \times 1mm \times 2.5mm^3$). The gold standard correspondences in this case were estimated by manually picking 3D landmarks between the undeformed and the compressed MRI. These experiments were part of our work described in [Mertzanidou et al., 2010a] and [Mertzanidou et al., 2012a] and they included the volume preservation constraint.

The mean reprojection error [Hipwell et al., 2003] for these experiments was reduced to 3.83mm (with a standard deviation of 1.59mm) after registration, from an initial 11.58mm (std 6.65mm) misalignment. All results are shown in Table 4.3. The results indicate that the affine transformation can give a clinically useful accuracy when registering 3D volumes to X-rays that have been generated using real compression images. These are promising results, although the error is expected to be higher when using real X-ray mammograms, as the similarity between the simulated DRR at each iteration and the target mammogram will change significantly.



Figure 4.6: Results of the simulated X-ray mammogram experiments for all patients (1 case per patient). The images show the intensity differences between the target image and the projection of the source, before and after registration for all the similarity measures tested (MSD, NCC, GD and MI). For MI the results are shown using a chequerboard, as the MRI projection is not a simulation of an X-ray and thus the intensities cannot be directly subtracted from the target image. Each square alternates the intensities between the two images: the target mammogram and the projection of the source after registration.



Figure 4.7: Illustration of the real compression MR images acquired for one of the volunteers. The volume on the left corresponds to the uncompressed breast. The rest of the volumes from left to right have increased amount of compression applied from the lateral to the medial direction.



Figure 4.8: Coronal slices of the data used for evaluation for 3 volunteers (a)-(c). From left to right in each image: MRI slice before and after compression.

Table 4.3: Reprojection error (in mm) for the eight volunteers before and after registration.

	v1	v2	v3	v4	v5	v6	v7	v8	mean	std
before reg.	5.87	20.1	17.69	19.85	4.01	8.6	5.97	10.56	11.58	6.65
after reg.	2.67	6.24	5.96	3.83	1.52	4.0	2.98	3.49	3.83	1.59
after reg. (std)	1.9	2.4	3.2	1.1	0.6	1.9	1.4	1.8	-	-
after reg. (max)	6.6	9.9	12.1	5.2	2.6	7.6	5.3	7.2	-	-

4.3 Experiments on real mammograms

In the following sections we present the results of the volume-preserving affine transformation model on the registration between real MR and X-ray images. Determining gold standard correspondences in real cases is not a trivial task and therefore requires the expert opinion of radiologists. The following sections describe in detail the experiments carried out on two different datasets. There is also a clinical validation experiment that was carried out as part of the HAMAM project [HAMAM, 2012] in collaboration with other partners. This is described in appendix C.

4.3.1 Validation using visual assessment on digitised mammograms

Our algorithm was initially tested on MR images and digitised film mammograms (CC views) of a dataset of five subjects from a high-risk population (section A.1). As gold standard correspondences were not available for this dataset, the results were evaluated visually. Figure 4.9 shows the registration results. Although we cannot obtain a quantitative validation, we can see that the similarity between the real mammogram and the projection of the source volume after registration has greatly improved, with the breast volume expanding to the medial and lateral direction. These experiments were part of the work described in [Mertzanidou et al., 2010a].

4.3.2 Validation using radiologists' annotations on FFDM

In these experiments we used clinical datasets of MRIs and FFDMs (section A.3), both CC and MLO views, that were acquired approximately at the same time point¹. The voxel resolution of the MRIs varied, as the images were acquired from three different scanners. The majority had a resolution of either $[0.9 \times 0.9 \times 1.0]mm^3$ or $[0.6 \times 0.6 \times 1.3]mm^3$; whilst one had a resolution of $[0.7 \times 0.7 \times 2]mm^3$. The original resolution of the X-ray mammograms was $[0.1 \times 0.1]mm^2$ for all of the cases apart from one that was $[0.085 \times 0.085]mm^2$; they were subsampled by a factor of 10 for registration to match the MRI resolution and reduce the computational cost associated with the ray-casting.

The patients had a range of different pathologies. One of the patients had an MR and X-ray compatible clip inserted after breast biopsy (section A.4.1). This was used as ground truth correspondence. The rest of the patients had clearly visible findings in both modalities. These were annotated and the annotations were used as gold standard correspondences for validation. For the annotated data, the MR findings were marked using one or multiple spheres, while the X-ray images using either a disk, or more frequently a free-form shape defining the outline of the finding.

As it is generally harder to annotate the 3D images accurately, the spheres did not always represent the finding's actual volume, but were rather centred around it. As a result we consider as the most appropriate error metric the distance between the centres of the annotated regions. In all the results shown below, the registration error is the 2D Euclidean distance between the centres of mass of the X-ray annotation and the projection of the MR annotation, after being deformed with an affine transformation.

For validation we have performed in total 113 registration tasks, including both CC (n = 55) and MLO (n = 58) view mammograms. These came from 49 patients, some of which had multiple studies

¹In most of the cases these were acquired the same day, overall within a month.



Figure 4.9: Registration results on real data (5 cases, one per row). From left to right: projection of the source volume before registration, after registration and real X-ray mammogram. The red cross indicates the position of a corresponding coordinate in each image.



Figure 4.10: Histogram of the registration errors calculated from 113 registration tasks. In red is shown the median value (13.1mm) and in green the outliers.

of MRI and X-ray mammogram pairs acquired at different times. The total histogram of the registration errors is given in Figure 4.10. The median registration error including all cases is 13.1mm. As it can also be seen in the histogram, for a limited number of cases (n = 3), the registration was not successful for reasons that are further discussed below. We consider these cases as outliers. The percentage of the registrations that had an error larger than 30mm is 2%. Using the median instead of the mean value is more informative of the registration error's central tendency, as it represents better the value around which the majority of the registration errors are clustered. The outliers affect the computation of the mean error value, while the median is less sensitive to their presence.

Figure 4.11 shows the distributions of the registration errors for the CC and MLO view mammograms separately. The histograms and the median values illustrate that overall the algorithm is slightly less accurate for the MLO view, although the difference between the two median values is negligible (12.9mm for the CC and 13.5mm for the MLO view).

Figures 4.12 and 4.13 illustrate the results for 3 patients (6 registrations), for which the registration was considered to perform well. In the first two cases it is clear that the appearance of the lesion shape in the two modalities varies significantly. We can also see that even in cases where the error in mm seems rather high (Figure 4.12(a): 14.43mm and Figure 4.13: 9.96mm), there is still an overlap between the two annotations and the registration result can give a good indication about the location of the lesion on the X-ray.

The results of two cases for which the registration did not perform well are shown in Figure 4.14. Figure 4.15 shows the two corresponding MR volumes to illustrate the difficulty in registering these lesions.

As we can see in the figures illustrating the results, the annotation of the projected MR lesion is generally larger than the corresponding annotation in the X-ray mammogram. Overall, the mean radius



Figure 4.11: Histograms of the registration errors of Figure 4.10, displayed individually for the CC and MLO views. Whilst the median error for MLO mammograms is marginally higher than for CC view mammograms (13.5mm versus 12.9mm as indicated in red) the distribution of these errors is broadly similar.



(a) Patient 1

(b) Patient 2

Figure 4.12: CC (on the left) and MLO (on the right) mammogram of patients 1 and 2. The X-ray mammogram annotation is shown in red and the projection of the MR annotation in green. (a) Patient 1 was diagnosed with an Invasive Ductal Carcinoma; the registration error is 8.12mm for the CC and 14.43mm for the MLO view. (b) Patient 2 diagnosed with Ductal Carcinoma In Situ; the registration error is 6.85mm for the CC and 2.63mm for the MLO view.


Figure 4.13: Registration result for CC (on the left) and MLO (on the right) mammograms of patient 3. (a) illustrates the raw mammograms and (b) the registered DRRs. A magnification view is given for both views on the bottom right corner of the raw mammograms. The evaluation in this case was done using the clip location. The clip is 2mm long as displayed in the magnification view; the location in the X-ray mammogram is illustrated by the high intensity region (red arrow) and the projection of the clip location in the MR is shown in green (green arrow). The registration error is 9.9mm for the CC and 7.7mm for the MLO view.



Figure 4.14: Two cases for which the registration error was high. CC view of patient 4 (error 28.45mm) and MLO view of patient 5 (error 20.24mm), both diagnosed with Invasive Ductal Carcinoma. The corresponding MR volumes and annotations are shown in Figure 4.15.



Figure 4.15: Corresponding MRIs and annotations of the patients illustrated in Figure 4.14. Patient 4 (left) and patient 5 (right).

of the MR lesions is 11.7mm, while that of the X-ray annotations is 9.9mm. This small difference in size can be caused by the difficulty in annotating the 3D volumes as discussed above, and the different contrast mechanisms in MRI and X-ray. When the MR lesions are projected into 2D after transformation, the mean radius of these areas is 15.9mm. This increase in size is expected from the global affine transformation model, since the lesions are not modelled separately as rigid objects. Therefore, although their volume is preserved, an expansion of the breast in the medial-lateral direction (for a CC view) would result in an expansion of the mean radius of the lesion when this is projected into 2D.

Overall, the main characteristic of the cases where the registration did not perform well are patients with large fatty breasts and cases with large, irregular-shaped lesions. In our experiments, the results did not show any correlation between the breast density or size and the registration error. Nevertheless, the analysis of the cases that had errors larger than 25mm showed that seven out of nine cases were large breasts. Their mean volume size was 46% larger than the mean breast size of all cases ($1.55 \cdot 10^6 mm^3$ against $1.06 \cdot 10^6 mm^3$). Moreover, five of these cases were also fatty breasts, with an average of 76% by volume of fatty tissue; the range for all the breast volumes in the dataset is 58%-78% (mean: 71%).

For the case of very fatty breasts, there is not enough information (ie. glandular tissue) to drive the registration and subsequently the optimisation is more likely to terminate in local minima. The registration of these cases is of less clinical importance, as for the radiologists establishing correspondences between the images is more challenging in cases of dense rather than fatty breasts. We also expect the affine transformation assumption to be less accurate for the large-size breasts, as in these cases the breast undergoes large anisotropic deformations that cannot be approximated accurately by our affine transformation model. Finally, we consider that the cases with large irregular-shaped lesions are not suitable for validation, since the contrast mechanisms are different between the two modalities and hence the lesion appearance differs significantly between them; as a result, their centres of mass would not necessarily correspond.

Although not suitable for validation, if accurate alignment of large lesions is possible, then we can extract useful information regarding tumour heterogeneity that could be useful for therapy planning and



Figure 4.16: Two MLO view mammograms that were excluded from validation, because the lesion annotation was too large.

response to treatment.

One of the main issues that arises from the validation of this alignment task on clinical cases, is how the gold standard correspondences between the two images are determined. Although in clinics the registration would only be useful to the radiologists in cases where lesions are not easily identified in both modalities, clearly visible lesion cases are used in this study to provide quantitative results. Nevertheless, due to the different nature of the images involved, the appearance of a lesion in a projection X-ray mammogram differs significantly from the enhanced 3D area in the MRI. Moreover it is not always straightforward to identify the lesion boundaries. Consequently, the 2D Euclidean distance could give results that are not representative of the actual correspondence. We believe that the cases that would be best suited for validation are the ones with small lesions or clip data, as the one displayed in Figure 4.13.

The total number of clinical cases that we originally had access to was 76 patients. From these, we excluded from validation the ones that had artifacts. For example, we excluded those that had annotation problems (8 cases), either because of lesions that were mis-matched between the MRI and the X-ray mammogram, or because of very large lesion annotations like the ones shown in Figure 4.16, as the Euclidean distance is not meaningful for these patients. There were also 4 cases excluded for which our breast segmentation carried out before registration failed to produce reasonable results around the chest wall, due to non-clearly defined boundaries between the pectoral muscle and the breast tissue.

The rest of the patients that were excluded from validation are patients with breast folding artifacts in the MRI. Two examples of these cases are shown in Figure 4.17. We found that for a large number of cases the breast was significantly deformed in the MRI. One reason why folding was so frequent in our dataset is that, according to the protocol of the clinic, women were advised to wear t-shirts during MRI acquisition to limit subtraction artifacts coming from motion of the breast before and after the injection of the contrast agent. Furthermore, a light breast compression was also sometimes applied, to reduce motion. As a result, this can cause folding artifacts in the MRI, particularly for large breasts. Our algorithm cannot compensate for folding of the breast and we believe these cases will also be problematic for previously published methods which attempt to deform the breast volume from the prone position to the compression between two plates. Subsequently these cases were excluded from validation. An

4.4. Discussion



Figure 4.17: The MRIs of two patients with large folding artifacts. These cases were excluded from validation.

alternative option for reducing breast motion artifacts without introducing folding would be volume registration between the pre- and the post-contrast image [Rueckert et al., 1999].

As our algorithm does not require significant manual interaction, we ran the registrations for all the above cases. The histogram of the registration errors is shown in Figure 4.18. We also display in the same figure the histogram of the cases used for validation and the total histogram for comparison. We can see from the plots that the errors for the cases that contained artifacts are spread equally throughout the total error range and account for all the errors that are larger than 45mm. These experiments were part of the work described in [Mertzanidou et al., 2012a].

4.4 Discussion

In this chapter we have described in detail our registration framework and the experiments carried out using the volume-preserving affine transformation model to approximate the breast deformation when compressed by two plates. This is the first method proposed that uses the structures within the breast for alignment, rather than surrogates based on breast outline or nipple position, and that has been tested on a meaningfully large number of datasets. Its main advantages are the ability to be easily integrated in clinics and also to provide reproducible results with minimal pre-processing interaction. The only interactive step is selecting landmarks on the pectoral muscle boundary. With an alternative automated method incorporated for pectoral muscle delineation, such as that recently proposed by [Gubern-Merida et al., 2011], the pipeline would be fully automated. For each registration, the algorithm requires around 20 minutes, on a single core, 64 - bit machine, with a 2.8GHz processor. We believe that this can be further improved with a GPU-based ray-casting algorithm, as this part is the most computationally expensive component.

The experiments on simulated data indicate that a simple affine transformation model can approximate a real breast compression in a 2D/3D registration task with a mean error of 3.83mm, when tested on eight cases. These results are promising for demonstrating the suitability of the affine transformation for this task, however these cases contain limited rotation, shear and more complex non-rigid deformations



Figure 4.18: Total histogram of the registration errors. In red are shown the cases that were excluded from validation due to artifacts. The histogram in green corresponds to the cases used for validation and is the same as in Figure 4.10. The total histogram of all cases is shown in blue.

which we expect from many real mammographic compressions.

The results of 113 registration tasks on clinical cases show that our algorithm can be applied in clinical practice giving useful accuracy (median 13.1mm). The results show a comparable accuracy to patient-specific biomechanical modelling ([Hopp et al., 2012], mean 11.8mm). Nevertheless we cannot directly compare the two statistics as the two techniques were tested on different data sets. In future work it would be beneficial to compare our method with patient-specific biomechanical modelling on the same dataset for a more rigorous comparison, as the accuracy can vary with breast shape, size and the pathology or other features used for validation. In chapter 6 we present experimental results on a smaller dataset, where the affine transformation is compared against a personalised FEM-based transformation model.

The experiments in this chapter give us an indication of the problems associated with the 2D/3D registration process and the limitations of affine transformations. These are caused by the fact that registration is driven by the similarity measure that is calculated in 2D and thus any displacements in the direction of the projection cannot be recovered (there are multiple 3D deformations that will generate identical projections/DRRs). This difference in dimensionality when using a 2D cost function to optimise a 3D transformation, is also the reason why the registration is prone to terminate in local minima. The low number of degrees of freedom make this algorithm less likely to be trapped in local minima and this is a principle that we take into account in our next approach, when using a different transformation model.

The next two chapters explain in detail how we incorporate prior knowledge about the breast deformation inside the transformation model. The new models are expected to achieve better registration accuracy, by constraining the transformation to the set of plausible deformations.

Chapter 5

MRI to X-ray registration using an ellipsoidal breast model

The main difficulty associated with the MRI to X-ray mammography registration task is the large deformation of the breast between the two different image acquisitions. Women are lying prone in the MR scanner, while they are standing, with their breast compressed between two plates, to obtain an X-ray mammogram. This complex deformation was approximated so far using an affine transformation. In this chapter we present a new approach that uses biomechanical modelling to simulate breast compressions and then incorporate them into a new transformation model that can be used for registration.

We propose the use of an ellipsoidal breast model to approximate the average breast shape and we use biomechanically simulated compressions to learn the deformations that occur during mammogram acquisition. By varying the parameters of the simulations, we simulate a wide range of compressions and then use Principal Component Analysis (PCA) to extract the main modes of variation. During registration, we optimise the coefficients of these modes and the parameters that define the initial position of the breast before compression. The proposed transformation parametrises the deformation into a lower dimensional space and captures the complex mammographic plate compression in only eleven degrees of freedom. This method has the advantage of using a transformation model with a low number of parameters, that is learnt using physically realistic breast deformations, instead of a geometrical model.

Firstly, we discuss the types of Statistical Deformation Models proposed in the literature (section 5.1) and then we describe the use of the proposed mean ellipsoidal breast model (section 5.2). The mammographic compression simulations are discussed in section 5.3 and the extraction of the main modes of variation in section 5.4. The modified, non-rigid registration framework is presented in section 5.5 and the experiments used for validation in section 5.6. Finally, section 5.7 contains the discussion.

5.1 Breast Statistical Deformation Models

The goal of our work in this chapter is to create a large number of plausible deformations that occur during mammographic compression and then use these to extract the main modes of variation. In our approach the deformations are generated from biomechanically simulated compressions. To generate these simulations we need to use a breast model (for the geometry) and a series of model parameters that



Figure 5.1: (a) MRI of patient p1; this is chosen as the space where all deformations are mapped for the SDM. (b) MRI of another patient p2 used for the construction of the SDM. Initially compressions are applied to all patients and then these are mapped to p1, where PCA will be used to extract the main modes of variation. (c) Deformed MRI of p1, after applying the displacement field of p2, mapped into the space of p1 using an affine registration. The anomalous shape of the breast in (c) is caused by the large difference in breast shape between p1 and p2 (particularly obvious in the coronal view) which is not compensated for by the affine registration.

define the boundary conditions and the mechanical properties of tissues.

Regarding the breast model, there are different approaches that can be used. Tanner et al. have followed a population-based approach [Tanner et al., 2009], where the breast compressions are applied to the MR volumes of 20 different patients. The deformation fields are then mapped into a common space of a patient's breast that is selected as reference. The assumption is that this breast has an average size and shape. The mapping into a common space is essential, as we aim to extract the main modes of variation of the deformation resulting from the simulated compression, rather than the breast shape variation across subjects. In the original Active Shape [Cootes et al., 1995] and Active Appearance Models [Cootes et al., 2001] this mapping is done using a set of points that represent corresponding anatomical locations, which are used to warp the images. For breast MRIs that come from different patients, there are no anatomical features, other than the nipple position, that can be used as corresponding locations. Tanner et al. have proposed an affine intensity-based registration of the breast MR binary masks for this mapping. Our experiments showed some issues associated with this method. These arise from the fact that breast shape variation across patients is large. The artifacts occur when the displacement fields corresponding to one patient are mapped onto another. An example image illustrating this effect is shown in Figure 5.1.

As an alternative to the population-based approach, the breast model can be extracted from the MR volume of a specific patient. This means that a new breast model needs to be created for each one of the MRI/X-ray registration tasks. In this case, all the breast compressions are applied in the same space of this particular patient. This is clearly illustrated in Figure 5.2, which summarises the differences between population- and patient-based approaches. The patient-specific approach has the advantage of using the exact geometry of the subject of interest for the simulations, while avoiding the mapping of the deformation fields both before and after performing PCA. Nevertheless, the need of repeating this



Figure 5.2: To create an SDM model, the population-based approach uses the two mapping processes illustrated. M_1 is used to map the deformation fields from the individual patients in the training set to the common space of one selected patient. PCA analysis is performed in space B of this selected patient. M_2 is used to map the extracted components to the new patient space. The patient-specific methods do not use any mapping process, as spaces A, B and C are the same and they change for each individual case. The ellipsoidal method proposed in this chapter uses only M_2 , as space A and B are the space of the mean ellipsoid.

process in a patient-by-patient basis is the main drawback of this technique.

5.2 An ellipsoidal breast model

In the approach followed in this chapter the breast geometry model is created from an ellipsoidal shape that approximates the average breast size and shape. More specifically, we are using half of an ellipsoid, whose dimensions are calculated from a population of 20 MR breast images. This approach has the advantage of using the same space for all the simulated compressions. Therefore, artifacts such as those created by the population-based model and illustrated in Figure 5.1 are avoided. Moreover, the modelling and PCA analysis are only performed once for all patients, as opposed to being repeated for each subject in the patient-specific approach. However both in population-based and in our approach, the extracted modes of variation still need to be mapped onto each patient's space for the MRI to X-ray registration (Figure 5.2). Although this can cause problems for patients with breasts asymmetries, these artifacts will only occur for this subset of patients, while in the population-based approach these are applied to all patients, because they are incorporated in the modes of variation, as they occur before the PCA analysis.

Below we describe in detail the process of creating a geometry model for the ellipsoidal case. The geometry of an ellipsoidal shape is given by:

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1$$
(5.1)

where a and b are the equatorial radii and c is the polar radius, shown in Figure 5.3(a). The radii are extracted from a population of 20 breast MR images (appendix A.1), using the the mean values. The computed values are a = b = 60mm and c = 160mm. The image is sampled at a resolution of $[1 \times 1 \times 1]mm^3$, which is similar to a typical resolution of an MR scanner that is used for mammography.

The next step in the modelling process is meshing. To extract the surface mesh, we are using the VTK¹ implementation of the marching cubes algorithm [Lorensen and Cline, 1987]. The tetrahedral

¹http://www.vtk.org/



Figure 5.3: (a) Image of the ellipsoid used for modelling and (b) its extracted mesh showing the ellipsoid between two plates, before compression.

elements are extracted using the open-source software package TetGen². The final model consists of 3, 535 nodes and 12, 056 elements. Figure 5.3(b) shows the extracted mesh.

5.3 Mammographic compression simulation

For the biomechanical compression simulations, we use a homogeneous transversely isotropic hyperelastic material [Han et al., 2012]. As we use one geometrical model that is applied to all patients, we only consider one material and do not assign different material properties to the fibroglandular and the fatty tissue. We also include anisotropy to account for the reinforcement of biomechanical properties from fiber-like connective tissues in preferred directions [Tanner et al., 2011]. We therefore allow the breast to expand more in the Medial-Lateral (ML) direction, than in the Anterior-Posterior (AP), for a CC view compression. This approach can be further justified by the fact that the breast is already extended more in the AP direction due to gravity, as the MRI is acquired in the prone position.

To create a range of different compression configurations, we vary the amount of compression, the ratio of tissue enhancement coefficient (anisotropy) and the Poisson's ratio, that controls the amount of breast volume change. For each compression the parameters are chosen randomly from the distributions shown in Table 5.1 that are taken from the literature ([Tanner et al., 2009], [Tanner et al., 2011]). The amount of compression refers to the displacement of the plates during compression. Therefore, this is 0% before compression and it would be 100% for a maximum compression, if there was no breast between them and the position of the two plates coincided.

At this stage, we do not vary the parameters that define the original position of the breast before compression. The pose that is defined from possible rotations and translations of the volume will be optimised inside the iterative registration process.

For the boundary conditions, we constrain the nodes that are close to the chest wall along the Anterior-Posterior (AP) direction, but allow unconstrained movement of the nodes in the other two di-

²http://tetgen.berlios.de/

Parameter	Distribution				
Amount of compression	Normal ($\mu = 55\%, std = 4\%$)				
Poisson's ratio	Uniform [0.45, 0.498]				
Ratio of tissue enhancement coefficient	Uniform [1, 512]				

Table 5.1: Parameters that vary between the different compressions and their distributions.



Figure 5.4: (a) The ellipsoid image superimposed with the displacement field extracted from a compression simulation and (b) the corresponding mesh under compression.

rections. As we use half an ellipsoid it is straight-forward to extract the nodes that are close to the chest wall, as they lie on a plane. To simulate the compressions, we apply displacements on two contact plates and use a GPU implementation³ of an explicit Finite Element (FE) solver [Taylor et al., 2009]. The contact model avoids artifacts on the breast surface that can be caused when applying individual displacements on the surface nodes [Ruiter, 2003]. Further explanations of the different techniques that model plate compression are given in the next chapter, section 6.1.3, where breast biomechanical modelling is discussed in more detail. Figure 5.4 shows an example of a compression applied on the ellipsoid to simulate a mammographic CC view compression.

5.4 Building a deformation model using PCA

After simulating a range of compressions on the ellipsoid model (we have used n = 100), the next step is learning the ellipsoid breast deformations using PCA. Each deformation vector \mathbf{D}_i consists of the concatenation of the displacements d in the X, Y and Z directions on a regular grid across the volume. We are using a regular spacing of 2mm across all three directions. Subsampling the deformation fields has a significant benefit in the computational cost of the eigenanalysis that follows. For a total number of m points in the grid:

$$\mathbf{D}_i = (\mathbf{d}_1, \dots, \mathbf{d}_{3m}) \tag{5.2}$$

³http://niftysim.sourceforge.net/





(c) mean $\pm 3\sigma_2$, second mode

(d) mean $\pm 3\sigma_3$, third mode

Figure 5.5: Illustration of the ellipsoidal image when varying the first three principal components. For each one of the four squares the top row corresponds to the coronal view and the bottom to the axial. A superimposed pattern is used for each view for better visualisation of the deformations. The first component (b) shows the effect of varying the amount of compression, the second (c) the anisotropy ratio (notice change in extension in the axial view) and the third (d) the breast shape under compression (notice the breast shape in the coronal view).

In a similar way to [Cootes et al., 2001], PCA is used to extract the main modes of variation and approximate any deformation field **D** by:

$$\mathbf{D} = \widehat{\mathbf{D}} + \mathbf{P} \cdot \mathbf{b} \tag{5.3}$$

where $\hat{\mathbf{D}}$ is the mean deformation field, \mathbf{P} is the matrix of the first t principal components (eigenvectors) and **b** is a vector of weights for each one of the t eigenvectors used. The eigenvectors \mathbf{e}_i and their corresponding eigenvalues λ_i are extracted from the covariance matrix:

$$\mathbf{S} = \frac{1}{n} \sum_{i=1}^{n} (\mathbf{D}_{i} - \widehat{\mathbf{D}}) (\mathbf{D}_{i} - \widehat{\mathbf{D}})^{T}$$
(5.4)

Figure 5.5 shows the the effect of varying the first 3 principal components of the learnt deformations.

The resulting displacement fields that represent the mean deformation and the deviations from it can then be mapped onto the MRI of a new patient. The mapping is done using a registration that includes scaling and translation between the binary masks of the ellipsoid model and the patient's MRI. As we

Once for	1. Extract the mean dimensions of an ellipsoid, from a
all patients	population of breast MRIs
before	2. Simulate compressions on the ellipsoidal model
registration	3. Extract the mean deformation and the main modes of variation
	using PCA
Once	1. Map the PCA components to the new patient
for each	2. Segment the pre-contrast MR to fibroglandular tissue & fat
patient	3. Simulate a volume corresponding to X-ray attenuation
Iterative	X-ray mammogram: target, X-ray attenuation volume: moving
registration	1. Cast rays through the transformed <i>moving</i> image
process	2. Calculate the similarity measure in 2D between
	the real and the simulated mammogram
	3. Update the 3D transformation parameters, according to the derivative of
	the similarity
	4. Go to step (1) until convergence

Table 5.2: Overview of the proposed framework that uses an ellipsoidal breast model and biomechanically simulated compressions.

saw in Figure 5.2, an advantage of using the ellipsoidal model is that this mapping process is only done once, after performing the PCA and extracting the main components. When using a population model, all breast shapes need to be mapped to the same space and given their large variability and the lack of anatomical correspondences, the registrations of the deformation fields can cause significant artifacts. As a result, these artifacts then cannot be excluded and they appear in the main PCA components.

5.5 Non-rigid registration framework

In this section we use the deformations that were learnt as part of our transformation model. The general registration framework remains the same as illustrated before in Figure 4.1. The only differences, as opposed to the processes described in chapter 4, are the transformation used and the way the transformation is combined with the ray-casting. These are further analysed below. Table 5.2 gives an overview of all the steps involved in the registration framework. It is given here in order to combine the information of all the methods described so far and summarise the complete process. Apart from the main registration process that is updated iteratively, the framework consists of two more parts that are performed in advance, off-line. The first is the PCA analysis, explained in section 5.4 that needs to be done only once. The second part, also performed before registration, needs to be repeated for each patient and is explained further below.

Before registration, we simulate the X-ray attenuation volume from the pre-contrast MRI, with the method discussed in chapter 3. Another step that is performed off-line is the mapping of the PCA



Figure 5.6: (a) An example of an X-ray attenuation volume superimposed with the displacement field of the mean PCA component and (b) the same volume after the mean deformation.

components to the new patient space, as the analysis was carried out in the ellipsoidal model space. The mapping is done using a registration between the binary masks of the ellipsoid and the patient's MRI. The updated parameters are translations and anisotropic scaling factors along the three axes.

Scaling the deformation fields assumes a linear relationship between the deformations of the ellipsoidal shape and the ones of the real patient. However, it is known that the breast undergoes non-linear deformations when it is being compressed. As a result, scaling only gives an approximation for the mapping between the ellipsoid and the patients' space. Nevertheless, as the dimensions of the ellipsoid are extracted from real MR images, our hypothesis is that the dimensions represent the average population and that the error introduced by this simplification is not significant for the majority of cases, apart from very small and particularly very large breasts.

Figure 5.6 shows an example of a patient's X-ray attenuation volume after being deformed with the mean component of the PCA analysis. Note that the deformation field is defined even outside the breast volume. All the principal components were extrapolated using B-splines, so that the displacement fields vary smoothly outside the breast shape. We followed this technique because the mapping from the ellipsoidal model to the patient's breast did not always ensure that all the breast volume is inside the boundaries of the ellipsoid. With this approach we avoid losing information around the breast surface from any breast voxels that are not included inside the ellipsoidal shape.

We propose a transformation model that consists of two parts, rigid and non-rigid. The rigidbody transformation determines the initial position of the breast on the detector, before compression. It includes three translations in all directions and two rotations; one for the rotation about the Anterior-Posterior axis of the breast (rolling) and one for the rotation about the Superior-Inferior axis (in-plane). The total number of degrees of freedom introduced by the rigid-body transformation is five:

$$dof_{rigid} = dof_{rotations} + dof_{translations} = 2 + 3 = 5$$
(5.5)

For the non-rigid part of the transformation, we use the PCA components that are extracted from

the ellipsoid deformations in the previous section. As we have already seen, each one of the deformation components \mathbf{D}_j is mapped onto the new patient space. In other words, the component \mathbf{D}_j is a three-dimensional image, of the same size as the X-ray attenuation volume, that contains for each voxel position x_i a vector with the learnt displacements along the three axes X, Y and Z. The parameters of the non-rigid transformation that are optimised inside the registration framework are the scalar weights b_j of each one of the m components used in the transformation. As we have seen before and given that $|b_j| \leq \pm 3\sigma_j$, a new deformation instance can be expressed as a linear combination of the deformation components. Consequently, the non-rigid transformation of a point x_i is given by the position x_i plus the deformation defined by the mean and the weighted deformation components:

$$T_{non-rigid}(x_i) = x_i + D^i + b_1 \cdot D_1^i + b_2 \cdot D_2^i + \dots + b_m \cdot D_m^i$$
(5.6)

As we can see in the above equation, the non-rigid transformation model contains m degrees of freedom:

$$dof_{non-rigid} = \{b_1, b_2, ..., b_m\} = m$$
(5.7)

The choice of the optimal number of deformation components m is usually performed according to the percentage of the variation captured by that number of components. For example, as the variance of each component is given by the corresponding eigenvalue, then using m components means that the model has the flexibility to capture the following percentage of the data:

$$\sum_{i=1}^{m} \lambda_i \\
\sum_{i=1}^{t} \lambda_i$$
(5.8)

where t is the total number of components.

Our proposed transformation model (with both the rigid and the non-rigid components) has a total number of degrees of freedom that is:

$$dof_{total} = dof_{non-rigid} + dof_{rigid} + dof_{mapping} = m + 5 + 3 = m + 8$$
(5.9)

where $dof_{mapping}$ refers to the three degrees of freedom of the scaling factors that are applied to the displacement fields, when these are mapped from the ellipsoid to the new patient space.

To combine efficiently the two types of transformation with the ray-casting algorithm we use the following technique. As we have discussed before for the affine transformation model, it is not computationally efficient to resample the 3D volume into the transformed position and then run the ray-casting algorithm on the new volume grid. Instead, the transformation is performed as the ray transverses 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. In other words, we first compute the intersections of the ray with the 3D grid of the moving volume, then transform them using the combination of rigid and non-rigid transformations and finally integrate the intensities that correspond to the transformed intersections. This way we avoid an extra 3D interpolation of the transformed 3D volume that would add computational cost and introduce additional sampling errors.

	p1	p2	р3	p4	p5	рб	p7	p8	p9	p10	mean
E-SDM CC	7.8	14.0	9.7	5.7	9.3	13.7	6.2	12.3	24.4	24.3	9.8
Affine CC	1.9	14.2	12.6	12.0	28.3	15.1	28.5	14.6	39.1	28.3	15.9

Table 5.3: Registration error (in mm) of this method (Ellipsoid-Statistical Deformation Model, E-SDM) and comparison with the volume-preserving affine transformation.

To combine the rigid and non-rigid transformation parts, it is important to respect the order that the different transformations occur, in order to avoid artifacts. As the transformation is defined from the *compressed* to the *uncompressed* breast space, the transformation of each point is described by equation 5.10. For a point x_i the transformation is given by:

$$T(x_i) = T_{2rigid}(T_{non-rigid}(T_{1rigid}(x_i)))$$
(5.10)

where $T_{non-rigid}$ is given by equation 5.6 and:

$$T_{1rigid}(x_i) = T_{translation}(R_{in-plane}(x_i))$$
(5.11)

$$T_{2rigid}(x_i) = R_{rolling}(x_i) \tag{5.12}$$

The initial position of the X-ray attenuation volume before registration ensures that the centres of mass of the X-ray mammogram and the simulated DRR coincide, as we saw in the previous chapter for the affine transformation model. Furthermore, the initialisation of the rotations for a CC and MLO view registration is the same. At iteration zero of the registration the mean component is applied to the volume, while the weights of the other components are all initialised to zero.

5.6 Experiments

For validation we used MRIs and CC view X-ray mammograms of 10 patients that were acquired approximately at the same time point (a subset of those described in section A.3). The patient group was different to that used to extract the mean ellipsoid dimensions (section A.1). As before, the registration error is the 2D Euclidean distance between the centres of the X-ray annotation and the projection of the MR annotation, after being deformed, first with our method and then following an affine transformation (chapter 4) for comparison. These experiments were part of the work described in [Mertzanidou et al., 2011].

In these experiments we used 3 PCA components, so the total number of degrees of freedom was 11. This choice was done experimentally and as a trade-off between the percentage of the deformations explained by the components and the computational time. Our experiments showed that the use of more components has a negligible effect in the final registration position and the quantitative results did not improve. More specifically, three components explain 52% of the data, while for example the use of ten components, that explains 63% of the data, provides very similar breast deformations compared to the three components but it significantly increases the computational cost. These values are illustrated in Figure 5.7.



Figure 5.7: Illustration of the percentage of the variation captured by N number of PCA components. Three components (red) explain 52% of the data, while ten components (green) explain 63%.

A summary of the results is given in Table 5.3. Overall, and excluding the two last cases where the registration failed, our method outperformed the affine transformation by several mm in all cases, apart from one (p1). Figures 5.8 and 5.9 illustrate the results on 4 cases. We can see that the projection of the MR finding was close to the X-ray annotation and even in cases where the distance was rather high (eg. p2 = 14mm), the projection was still inside the annotated X-ray area, shown in green. In all eight cases the projected MR annotation mask overlapped with that of the X-ray and the mean registration error was 9.8mm, as opposed to 15.9mm for the affine.

The two cases that gave very high errors are shown in Figure 5.10 together with the MRI annotations, to illustrate the difficulty in registering these findings. As we can see, p9 has a finding that is very close to the pectoral muscle. Our ellipsoidal model is expected to perform less well for lesions that are so close to the chest wall.

The main characteristic of the second failure case (p10), which is illustrated in Figure 5.10(b) is that this is a large breast. Our ellipsoidal model is expected to be less accurate for very large breasts, since the deformations they undergo are highly complex and non-linear. The approximation that we use when scaling the deformation fields in order to map them on the new patient is more likely to cause larger errors for larger breasts.

For the MLO view registrations, our ellipsoidal SDM approach failed to produce reasonable results for this dataset. However, in our later work with the patient-specific biomechanical modelling, described in the next chapter, we have investigated alternative optimisation schemes. Registration results for the MLO view are presented in the next chapter for all three transformation models (affine, E-SDM, patient-specific FEM) using the new optimiser for the E-SDM approach.



Figure 5.8: Registration results on 2 patients. For each case, the left image shows the real mammogram with the centre of mass of the MR annotation after alignment (red cross). The right image in each pair shows again the real mammogram, but as well as the projected MR annotation, the mask of the X-ray annotation is also shown in green.(a) p1 (error: 7.8mm) and (b) p2 (error: 14mm).



Figure 5.9: Registration results on 2 patients. For each case, the left image shows the real mammogram with the centre of mass of the MR annotation after alignment (red cross). The right image in each pair shows again the real mammogram, with the projected MR annotation and the mask of the X-ray annotation, shown in green. (c) p3 (error: 9.7mm) and (d) p4 (error: 5.7mm).





(b) p10 (24.3mm)

Figure 5.10: Cases for which the registration error was high. On the left are shown the X-ray mammograms and on the right the MRIs of (a) p9 and (b) p10.

5.7 Discussion

In this chapter we saw in detail the steps involved in the learning process of the mammographic compressions. The extracted mean deformation and the modes of variation are then used in our transformation model for registration. The proposed ellipsoidal breast model has the advantage of avoiding artifacts that can be created in the population-based models. Also, it does not require repeating the modelling process for each patient, as in the patient-specific methods. This makes our method more suitable for use in clinical practice, as there are usually manual steps involved in the modelling process, explained in more detail in the next chapter. Moreover, in comparison to patient-specific biomechanical modelling that does not perform any statistical analysis [Ruiter et al., 2006], this method provides the benefit that results are less dependent on the modelling and the FE solver that are used.

Our motivation for using this transformation model was its ability to avoid the problems associated with patient-specific modelling, while providing clinically useful registration accuracy. The experiments on a clinical dataset show an improved accuracy (9.8mm) in comparison to the affine transformation model that can be clinically useful. However, the initial results indicate a wide range of registration errors.

Finally, a potential limitation of this approach concerns how well an ellipsoidal shape and a single tissue type can approximate the shape and deformation of a real breast. To answer this question the next chapter includes testing against a biomechanical patient-specific simulation model that updates the modelling parameters inside the iterative registration process.

Chapter 6

MRI to X-ray registration using a patient-specific biomechanical model

In chapters 4 and 5 we saw two different transformation models that can be used for an intensity-based MRI to X-ray mammography registration. Both models have the advantage of being applicable to an extensive population of breast shapes, without requiring manual interaction. Nevertheless, both incorporate simplifications of the actual deformation that the breast undergoes during mammographic compression.

In this chapter we are considering a patient-specific transformation model that is based on biomechanical simulations and we compare its performance to the two other transformations. The new transformation model is built individually for each patient, using the MR volume. The novelty compared to other patient-specific methods proposed for the same task ([Ruiter et al., 2006], [Hopp et al., 2012], [Lee et al., 2011]) is the integration of the simulation inside the iterative intensity-based registration framework, as opposed to registration based on surrogates such as the breast outline and the nipple position. This is enabled by the use of an integrated transformation module that runs on the graphics processing unit [Taylor et al., 2009], providing shorter execution times than commercial packages.

Section 6.1 describes the details of the patient-specific biomechanical modelling and its challenges. Section 6.2 summarises the integration of the biomechanical simulations into the registration framework. Finally, sections 6.3 and 6.4 contain the experimental results on clinical cases and the discussion respectively.

6.1 Patient-specific biomechanical simulations

Biomechanical modelling is a powerful tool that can provide physically realistic breast deformations, as opposed to geometrical models. Nevertheless, there are several challenges that are associated with the patient-specific biomechanical modelling of a highly deformable organ with a complex anatomical structure, such as the breast. Firstly we address these problems and we briefly review how these were tackled in the literature. We then present our approach and contribution.

6.1.1 Meshing

As we saw in chapter 5, the first step in biomechanical modelling consists in extracting the surface and and the volumetric mesh from the MRI. In our previous SDM approach the MRI is a binary ellipsoidal



Figure 6.1: Example of a patient's mesh showing folding of the skin that can cause problems in the topology of elements when the breast undergoes a large compression.

mask. This simplification which uses an average breast shape has the advantage of providing a smooth mesh. However, real breast shapes often have irregularities. For example, the surface can have areas where the skin is folded, especially in parts where the breast is attached to the chest wall, such as that illustrated in Figure 6.1. This can be problematic if we attempt to apply large breast compressions, as the topology of the elements could break down. Subsequently, this would necessitate re-meshing our model after applying smaller amounts of compression in order to correct the node and element topology. Remeshing the breast model adds computational cost to the simulation process and could require manual interaction.

Another factor that can cause problems in meshing is the breast folding that is a result of the breast touching the bottom of the coil, or the patients wearing t-shirts during scanning, as we saw in Figure 4.17. Moreover, this type of folding violates the assumption that the breast is pendulous under gravity, and since these deformations cannot be recovered with tools currently available, all cases that include breast folding are excluded from this study.

The choice of an optimal number of nodes and elements is also an open research problem when generating surface and volumetric meshes. This choice is a trade-off between accuracy and robustness or simplicity of the model. Very fine meshes have the advantage of representing the real breast shape more accurately, but mean that the FE solver requires longer computational time. Also, large deformations, such as the mammographic compression, can cause topology problems when applied to very fine meshes, as discussed above.

In the literature authors have proposed a wide range of values for meshing, from very fine to very coarse meshes, but there is no experimental study that demonstrates the optimal values. When modelling breast plate compressions, the number of elements used varied from an order of one or two hundred ([Ruiter et al., 2006], [Chung et al., 2008]) to tens of thousands ([Samani et al., 2001], [Tanner et al., 2011], [Han et al., 2012]).

In our approach, before extracting the surface mesh, we initially segment the volume from the background using a simple region-growing algorithm, and then apply Gaussian smoothing and downsample the extracted binary mask to an isotropic volume of 10mm, as in [Tanner et al., 2011]. This way, we can



Figure 6.2: (a) Example of a surface mesh extracted from one of the clinical cases. (b) The same mesh, showing also a slice of the corresponding MRI. (c) Different view of image (b), illustrating the approximation of the pectoral muscle as a plane.

avoid small irregularities of the breast shape that might cause problems to the FE solver, such as those in Figure 6.1. The surface mesh is extracted using the VTK¹ implementation of the marching cubes algorithm [Lorensen and Cline, 1987] and the tetrahedral elements are extracted using the open-source software package TetGen², as in the ellipsoidal case. An example of an extracted mesh is shown in Figure 6.2. This consists of 2, 529 elements and 777 nodes. The numbers lie in the mid-range of breast mesh sizes proposed in the literature, as described above. The modelling of the pectoral muscle illustrated in the same figure is discussed in section 6.1.3.

6.1.2 Material properties

After meshing, the next step in modelling consists in assigning either one or multiple material properties to the different breast tissue types. These describe the mechanical properties of the breast and hence model the relationship between the applied stress and resulting strain. One option is to assume that the breast consists of one single tissue type, as a simplification. In this case, the elements of the breast model have the same behaviour whether they belong to fibroglandular or to fatty tissue, as in the ellipsoidal case. Alternatively, we can segment the breast tissue from the MRI into these two classes and assign different material properties to each of them. Moreover, there are other tissue types that can be considered. For example, the skin can be modelled either as a membrane or as a thin layer over the surface of the breast, tumours can be assigned specific material properties if their presence and locations are known. Finally, the pectoral muscle can also be modelled as another class. Although it is not visible in a CC view X-ray mammogram, it affects the breast tissue behaviour under compression, especially for the part that is closer to the chest wall.

Assigning different material properties to the extracted elements is not a trivial task, due to the partial volume effect and also the dependence on the size of the elements used, as these are usually

¹http://www.vtk.org/

²http://tetgen.berlios.de/

larger than the voxel size. Thus one element can correspond to more than one tissue type. One further complication is the fact that the difference in material properties between the different tissue types is not known. Although there are several material parameters that are proposed in the literature, their variation and range is high and there is no experimental study that shows the optimal values. In our experiments, sampling randomly from the distributions of the proposed values can cause convergence problems for the FE solver used and can introduce the need for re-meshing. Both issues involve manual interaction to overcome them and therefore do not allow the modelling process to be automated.

When modelling large breast compressions authors proposed mainly non-linear material models, as linear elastic models are considered less accurate for large deformations. More specifically, a hyperelastic, (nearly) incompressible material was proposed by [Samani et al., 2001], [Ruiter et al., 2006], [Chung et al., 2008], [Han et al., 2012], [Lee et al., 2011] and [Hopp et al., 2012]. [Tanner et al., 2011] used a linear elastic model, although in the experiments described, the displacements of all the surface nodes were known and constrained. Therefore for this particular application the effect of the material model is expected to have less influence on the results.

Regarding the material classes used, [Samani et al., 2001], [Ruiter et al., 2006], [Han et al., 2012] and [Tanner et al., 2011] used three different classes: *fibroglandular*, *fat* and *skin*. [Chung et al., 2008] and [Hopp et al., 2012] proposed one homogeneous tissue type instead. [Ruiter et al., 2006] and [Tanner et al., 2011] reported an anisotropic behaviour of the breast when applying mammographic compressions, with a reduced elongation in the Anterior-Posterior direction and an increased expansion in the Medial-Lateral direction, as discussed previously in section 5.3. This was incorporated in the modelling by [Tanner et al., 2011] and [Han et al., 2012].

In our implementation, we are using a nearly incompressible and hyperelastic neo-Hookean material model [Han et al., 2012] which is transversely isotropic, to incorporate the anisotropic behaviour that was previously observed and account for the reinforcement of biomechanical properties from fiber-like connective tissues in a preferred direction. The use of a non-linear instead of a linear model and the incorporation of an anisotropic behaviour are considered to be crucial for our application, as the compression is simulated using a contact model (discussed below) and not constrained displacements of the surface nodes as in [Tanner et al., 2011]. The material parameters of this model are optimised during the registration framework as we will see in section 6.2.

Regarding the different tissue types, our approach assumes a homogeneous tissue type as a simplification. The advantage of this implementation is that it avoids convergence problems of the FE solver, as shown by experimental results, due to the large variation in the material properties used in the literature. In relevant previous work other authors have also proposed homogeneous tissue types ([Chung et al., 2008], [Hopp et al., 2012]). Experimental work [Ruiter, 2003] showed no significant effect on the results when different tissue models are used instead.

6.1.3 Compression simulation and boundary conditions

The choice of boundary conditions and the technique used to simulate the mammographic compression also have an effect on the simulation results. As there is no gold standard available, a range of approaches have been adopted which produce different solutions. As in the previous points, the choice of boundary conditions and compression simulation may introduce the need of intermediate re-meshing steps.

Plate compression

The plate compression that occurs in mammography has been modelled in the literature using mainly two different techniques. The first involves applying displacements to the surface nodes of the breast mesh in the direction perpendicular to the two plates ([Ruiter et al., 2006], [Hipwell et al., 2007]). The compression plates and their interaction with the breast are not explicitly modelled in this case, but it is assumed that the displacements applied on the surface nodes have the same effect. In was reported [Ruiter, 2003] that this technique leads to artifacts on the breast surface (breast swelling) around the nodes which are adjacent to the ones that the displacements are applied to. An alternative method is to explicitly model the interaction between the contact plates and the breast tissue using either a friction [Lee et al., 2011] or a frictionless model ([Chung et al., 2008], [Han et al., 2012]).

In our approach, we are using the frictionless contact model proposed by Han et al. [Han et al., 2012]. The main advantage of using a contact model instead of applying displacements on the surface nodes is that the interaction is modelled explicitly. Moreover, this method avoids artifacts on the breast surface that can occur if the displacements applied on neighbouring nodes are different, as discussed previously. 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 or the friction coefficient that describes best the breast-plate interaction. Our assumption is that a frictionless model provides a good approximation to the actual breast compression.

Pectoral muscle

The behaviour of the breast tissue under compression is influenced by the presence of the pectoral muscle. This was modelled in the literature either by constraining the nodes close to the pectoral muscle to be fixed ([Samani et al., 2001], [Chung et al., 2008], [Hopp et al., 2012]), or by allowing them to slide along the chest wall ([Tanner et al., 2011], [Han et al., 2012], [Lee et al., 2011]).

In our implementation, we approximate the pectoral muscle by a plane and allow the nodes to slide along that plane. This approximation is not expected to introduce significant errors to 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 our 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 that might occur, as the chest wall surface often has topology irregularities. These are caused from the fact that the segmentation algorithms do not produce a smooth profile on that region, as the intensity boundaries are not clearly defined. An example mesh showing the approximation of the pectoral muscle as a plane is shown in Figure 6.2(c).

Once	1. Segment the pre-contrast MR to fibroglandular tissue & fat
for each	2. Simulate a volume corresponding to X-ray attenuation
patient	3. Approximate the chest wall using a plane
	4. Extract the surface and volumetric mesh
	5. Create an <i>xml</i> file with the geometry, the material properties and
	the constraints of the breast model
Iterative	Inputs: X-ray mammogram (target), X-ray attenuation volume (moving),
registration	breast model <i>xml</i> file
process	1. Cast rays through the deformed <i>moving</i> image
	2. Calculate the similarity measure in 2D between
	the real and the simulated mammogram
	3. Update the 3D transformation parameters, according to similarity
	4. Go to step (1) until convergence

Table 6.1: Overview of our patient-specific FEM-based registration framework.

6.2 Integration of the FEM simulations in the registration framework

In this section we use the patient-specific breast FE model inside the registration process. This is achieved via a transformation module that uses the FE solver to simulate the mammographic compression. The registration framework is the same as illustrated previously in Figure 4.1, with the modification described in the previous chapter (section 5.5) regarding the technique used to combine the non-rigid transformation with the ray-casting. The difference of the approach described here is the transformation model, that consists of biomechanical simulations of the mammographic compression.

An overview of the complete pipeline is given in Table 6.1. As we can see, this approach has more pre-processing requirements for each patient before registration than the other transformation techniques, as the breast FE model needs to be created for each subject using the methodology described in the previous section. After the modelling is completed, the information is stored in an *xml* file. This file is used as an additional input into the registration pipeline, together with the X-ray mammogram and the X-ray attenuation volume.

Figure 6.3 illustrates the axes to which the transformation parameters correspond and an example CC view compression for one of the patients. The parameters that are optimised inside the registration process are:

- Translation in X
- Translation in Y
- Rotation about the Y axis (rolling)
- Rotation about the Z axis (in-plane)



Figure 6.3: 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.

- Amount of compression constrained between: no compression (0%) and 90% of the maximum distance between the nodes in the Z direction (for a CC view)
- Ratio of tissue enhancement coefficient (Anisotropy) constrained between [0-512] (range taken from the literature [Tanner et al., 2009], [Tanner et al., 2011])
- Poisson's ratio constrained between 0.45 and 0.499

The first four parameters account for the initial positioning of the breast before compression and the last three define the amount of breast compression and the behaviour of the breast tissue. In our SDM approach, in the previous chapter, we attempted to extract the main modes of variation caused by the last three parameters. In this method, we optimise these directly for each patient.

As previously, the distance between the X-ray source and the detector is extracted from the DICOM header of the X-ray mammogram and is f = 660mm for all the mammograms using in this study. Before registration, the moving volume is positioned on top of the detector and is translated in the corresponding XY plane such that the centre of mass is projected onto the centre of mass of the real mammogram. This provides a good initial position for the registration and does not require manual interaction. The amount of compression is initialised to a 50% plate displacement, Poisson's ratio to $\mu = 0.498$ and the ratio of tissue enhancement coefficient to r = 250.

For an MLO view, the initial position changes, such that the breast has an initial rotation of 45° about the Y axis (roll) and an in-plane rotation of 20-30°. The initial value of the roll angle is extracted from the mammogram's DICOM header and the initial in-plane rotation is determined experimentally to be either 20° or 30°, depending on the mammogram of each individual case.

The order of the transformations is important for this case as was described before for the SDM transformation. For a point x_i , which is the intersection of the ray with the volume grid of the moving image, equation 5.10 remains the same:

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

where $T_{1rigid}(x_i) = T_{translation}(R_{in-plane}(x_i))$ and $T_{2rigid}(x_i) = R_{rolling}(x_i)$, but now the non-rigid transformation $T_{non-rigid}$ is not given by the PCA components. It is the interpolated node displacement at the current position x_i as computed by the FE solver at the current parameter position.



Figure 6.4: Illustration of a the initial position of the plates before compression to avoid intersection of the breast model nodes and the plates during registration. d is the distance of each plate from the centre of mass of the volume in the XZ plane and it is given by equation 6.1. The grey plates illustrate where the position of the plates would be set initially if only the maximum distance on the Z axis was taken into account. By considering also the X axis, the initial position of the plates is set to the location of the red plates. Consequently, the maximum plate separation is $2 \cdot d$.

To avoid an intersection of the breast nodes and the compression plates when the breast is rotated about the Y axis, we re-position the plates before registration such that they are equidistant from the centre of mass. This distance is equal to the maximum distance between all the nodes and the centre of the mass on the XZ plane. Figure 6.4 illustrates this process. Let $C(c_x, c_y, c_z)$ be the centre of mass of the breast volume and $N_i(n_{ix}, n_{iy}, n_{iz})$ the *i*-th node of the FE model. As we can see in the figure, the initial plate separation is $2 \cdot d$ and the distance d is given by

$$d = max(d_i^{XZ}(N_i, C)) \tag{6.1}$$

where

$$d_i^{XZ}(N_i, C) = \sqrt{(n_{ix} - c_x)^2 + (n_{iz} - c_z)^2}$$
(6.2)

The optimised parameter is the separation between the two compression plates, and its maximum value is $2 \cdot d$. This way, the deformed volume is always centred around the centre of mass. If the initial position of the plates was changing at each iteration, then the corresponding deformed volumes would not be centred around the same point. This would further complicate the optimisation process as the magnification factor of the projected mammogram would vary.

Regarding optimisation, in our previous implementations of the affine and the ellipsoidal transformation models we were using a gradient descent optimiser with a regular step size, as described in section 4.1.2. Gradient-based optimisers are very commonly used in medical image registration, providing good results when the optimised function has a smooth behaviour. When applying the same optimisation method using the FE model, the results showed that the optimisation failed to reach the global minimum, as the final alignment did not provide visually good results. Moreover, often the similarity at the last iteration was lower than that obtained for previous iterations. This behaviour can occur when the

Table 6.2:	Comparison	between	the Regular	Step	and the	Regular	Step	Gradient	Descent	optimisation
schemes.										

Regular Step	Regular Step Gradient Descent					
At each iteration i	At each iteration i					
For each parameter p	For each parameter p					
$p^i = p^{i-1} \pm \frac{step}{w(p)}$	$p^{i} = p^{i-1} + d \cdot \frac{G(p^{i-1})}{w(p)} \cdot \frac{step}{\sqrt{\sum_{k=1}^{P} (\frac{G(k^{i-1})}{w(k)})^{2}}}$					
update p^i if similarity increases						
decrease step if the similarity doesn't improve	decrease step if the direction $d = \pm 1$ changes					
$stop \ if \ step < tolerance$	$stop \ if \ gradient < tolerance$					
*G(p) : the gradient magnitude of the similarity m	easure with respect to parameter p					
*w(p): the weight of parameter p that controls the magnitude difference between the transformation						
parameters						

optimised function has a noisy profile with respect to the parameters. Moving towards the direction of the gradient does not guarantee better similarity. Our experiments showed that, especially for the FEM transformation model, where the initial position of the volume is already close to the optimal solution, the similarity measure has a noisy behaviour around it and this optimisation scheme did not provide good registration results.

To overcome this problem, we propose the use of a regular step optimiser that does not use gradients. The optimisation scheme is simple. At each iteration, the optimiser sequentially varies all the parameters by a specified factor (*step*) and keeps the value that results in a better similarity. A comparison of this optimisation scheme against the one that was used so far is given in Table 6.2. The results showed a significant improvement when the new optimiser was used.

A further small improvement was achieved when updating at each iteration of the registration only one parameter, instead of all sequentially. According to this scheme, the parameter that is updated at each iteration is the one that results in the largest increase of the similarity measure, at the current relative step size. This is also known as a simple hill-climbing [Russell and Norvig, 2003] optimisation scheme. Using this technique is a logical approach for the optimisation of our FEM-based transformation parameters, as this has the effect that the translations and rotations, which naturally contribute most to the similarity measure, are updated first. Then, the rest of the parameters are updated according to their contribution to the similarity measure, rather than their relative position in the transformation module, which would occur if they were updated sequentially.

Table 6.3: Registration error (in mm) of our FEM transformation method and comparison with the Ellipsoid-Statistical Deformation Model (E-SDM) and the volume-preserving affine. The clip cases are p4 and p5. The last three columns in the table correspond to the mean value, the standard deviation and the confidence intervals of the mean values for a 95% confidence level. For a value *c* in the last column the confidence interval is given by: [mean - c, mean + c].

	p1	p2	p3	p4	p5	mean	std	conf.
FEM CC	8.0	6.8	7.2	4.8	11.4	7.6	2.4	2.1
E-SDM CC	12.3	16.3	3.6	4.4	28.8	13.0	10.2	9.0
Affine CC	14.6	13.5	3.7	9.9	23.4	13.0	7.1	6.3
FEM MLO	12.2	11.5	10.3	11.1	6.2	10.2	2.3	2.0
E-SDM MLO	11.0	24.7	16.5	7.5	21.5	16.2	7.1	6.2
Affine MLO	11.9	7.2	9.4	7.7	18.9	11.0	4.7	4.1

6.3 Experiments

For validation we used MRIs and X-ray mammograms of five patients that were acquired approximately at the same time point. These data are described in section A.3 and they include also two clip cases (section A.4.1). In total we used three patients with identified lesions in both modalities and two with MR and X-ray compatible clips. As before, the registration error is the 2D Euclidean distance between the centres of the X-ray annotation/clip position and the projection of the MR annotation/clip position, after being deformed, first with our FEM-based method and then following an affine (chapter 4) and an SDM transformation (chapter 5) for comparison. These experiments were part of the work described in [Mertzanidou et al., 2012b].

The data selection process ensured that this dataset does not include any MRI cases with unusual breast shapes and irregularities, such as folding. This facilitates the process of building a breast model and avoids artifacts and convergence problems of the FE solver when a large breast compression is applied. The solver used is the same as in the previous chapter, a GPU implementation³ of an explicit Finite Element (FE) solver [Taylor et al., 2009].

The registration errors of all registration tasks are given in Table 6.3. Overall, we can see from the results that the patient-specific FEM-based transformation performed best. Moreover, the standard deviation of the registration errors had the lowest value. The latter means that this approach gave the most consistent results amongst the three models. These results are not surprising, as a model that is built from a patient's MRI is less likely to give large errors than a geometrical or a statistical deformation model, that are not constrained to provide physically realistic deformations. In addition to the mean values and the standard deviation, Table 6.3 contains the confidence intervals of the mean values for a 95% confidence level and assuming normally distributed errors. As the confidence intervals for all

³http://niftysim.sourceforge.net/

three transformations overlap, one can conclude that the differences in mean values are not statistically significant. However given the small number of cases, future validation tests on a larger dataset will examine whether this is a valid conclusion.

While for the CC view registrations the FEM had clearly the lowest error $(7.6 \pm 2.4mm)$, for the MLO view the registrations errors were slightly larger $(10.2 \pm 2.3mm)$ and also comparable to the affine transformation. One possible explanation is the fact that the pectoral muscle has a larger effect on the registration results of the MLO view and thus the simplification of the pectoral muscle as planar could have contributed to the method's accuracy for this view. Further validation on a larger number of cases could verify this explanation.

It is also clear from the results that for this dataset the SDM approach performed slightly worse than the affine for the MLO view. For the CC view the results are comparable, which can occur for a specific dataset, although in the previous results the SDM performed better. In these experiments we used the new optimisation technique described in section 6.2, as the one used in the previous chapter did not provide good alignment. Nevertheless, the results indicate that this approach is less reliable for this view. We can conclude that the approximation of the breast as an ellipsoidal shape has a more significant and negative effect on the MLO view registration.

Figures 6.5 and 6.6 show some example registration results of the FEM-based transformation. As was discussed in section 4.3.2 inevitably each modality can give different estimates of a lesion's size, but all the cases with annotated lesions show overlap. The difference in size as an effect of the different contrast mechanisms and the difficulty in manually annotating the lesions, cannot be avoided. Another factor that results in larger sizes of the MR annotations is the use of a homogeneous tissue type that can result in a compression of the lesion size in the direction of the compression and an expansion in the perpendicular direction. This is expected, as the lesion is not modelled explicitly in our approach and it was also previously observed in other FEM-based simulations [Ruiter, 2003]. The use of a different tissue type for tumours could overcome this problem, but would also complicate further the modelling and the FE simulation processes. Our error metric is not influenced by this lesion expansion effect, as it takes into account only the centres of mass of the different annotations.

6.4 Discussion

In this chapter we have presented a novel framework for MRI to X-ray mammography registration using an iteratively updated FEM breast compression simulation. The results on five clinical datasets show improved accuracy compared to the affine and the SDM transformation models, indicating that this could be a useful tool and potentially help in better breast cancer detection and diagnosis. Better accuracy was expected from this transformation model, as it provides subject-specific and more physically realistic compressions of the breast than the other two models.

Compared to other patient-specific FEM-based methods used for this task, quantitative results on clinical cases showed a mean error of 4.3mm on 6 cases [Ruiter et al., 2006], and in a more recent semi-automated implementation of the same approach $11.8 \pm 6.5mm$ on CC view mammograms of 11 patients [Hopp et al., 2012]. However, a meaningful comparison is not possible unless these algorithms



Figure 6.5: Registration results for two patients with identified lesions. The X-ray annotation is shown in red and the projection of the MR annotation in green; their overlap is yellow.



(a) p5 CC, 11.4mm (b) p5 MLO, 6.2mm

(c) p4 CC, 4.8mm

(d) p4 CC, sim

Figure 6.6: Registration results for the two patients with MR and X-ray compatible clips. The clip location in the X-ray mammogram is visible as the high intensity region (and a red arrow for p4). The MR annotation in shown in green. For the patient p4 we also show the simulated CC X-ray mammogram (d).

are tested on the same datasets. The main novelty of our method is the integration of a biomechanical transformation model into an intensity-based registration, where we simultaneously optimise both the pose, via four degrees of freedom, and the model parameters, via a further three degrees of freedom. This framework maximises the amount of information used by the optimisation, compared to registration methods based only on the breast outline [Ruiter et al., 2006], increasing the likelihood of the correct transformation being obtained.

The breast modelling process proposed in this chapter includes certain simplifications, such as the use of coarser resolution images to extract the meshes, the use of one homogeneous tissue type and the approximation 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 pectoral muscle segmentation. As discussed previously this can be automated in the future, and it could be integrated into the clinics, providing a fully-automated patient-specific framework for MRI to X-ray mammography alignment.

Nevertheless, large breast compressions, such as those occurring in mammography, can cause convergence problems for the FE solver. As a result, this could introduce the need for manual interaction, which would not occur with an affine or the SDM transformation model. Further validation tests should be performed to examine the applicability of this registration framework in clinical practice.

In future work it would be interesting to investigate the performance of this framework using a more complex model of the breast. This would include for example more tissue types. Classes that can be considered are the fibroglandular and the fat tissue, the skin, pectoral muscle and tumour (if present). Regarding the boundary conditions, we could also optimise the modelling of the pectoral muscle as in this implementation it is simply approximated by a plane. Another factor that could be tested is the removal of the effect of gravity before compression as proposed by [Rajagopal et al., 2008]. Finally, further work includes validation using a larger number of datasets.

Chapter 7

Relating findings between X-ray mammograms via an MR volume

In the previous chapters we presented our registration framework for mapping MR coordinates to the X-ray mammograms and we investigated the performance of three different transformation models for this task. Here we are using this framework for a different application: the mapping between X-ray mammograms, via the MR volume.

The clinical importance of this application is high. X-ray mammogram registration techniques can be used to assist radiologists in automatically determining correspondences when examining temporal mammographic images of the same view, or when they relate findings from the CC to the MLO view and vice-versa. Determining correspondences between X-ray mammograms is of interest for cases with an available MRI. For the remaining cases, our method paves the way for a model-based approach that could be performed without MRI in future. Moreover, these techniques can be integrated in CAD algorithms ([Paquerault et al., 2002], [van Engeland et al., 2006], [Zheng et al., 2009]) to reduce the search space when mapping texture features that are extracted from different mammographic views and also relate the extracted information to improve CAD performance.

In the next sections we describe the modifications made in order to adapt the registration framework for this purpose and also demonstrate its use for ipsilateral and temporal registration of X-ray mammograms.

7.1 Analysis

As we have seen in the literature review (section 2.2.1), determining corresponding regions in X-ray mammograms has been an active research field. Nevertheless, most previous techniques approach this problem as a 2D-to-2D correspondence task (so assuming one-to-one correspondences), either using a mapping based on texture and intensity measures, or using a distance transform from features such as the pectoral muscle and the nipple position. Since X-ray mammograms are projection images, any 2D transformation model is inappropriate for this task, as the 3D physical transformation of the tissue cannot be modelled by a 2D technique. This problem was approached as a one-to-many correspondence approach for the validation of X-ray mammography registration algorithms [Hipwell et al., 2007] and for relat-

ing findings between ipsilateral X-ray mammography views [Kita et al., 2001] and breast tomosynthesis views [van Schie et al., 2011].

Assuming that we have an MRI and two mammograms (either CC and MLO, or temporal) of the same patient, then we can perform two registration tasks, to relate the MRI to the X-ray mammograms independently. After having acquired the two transformations, we can use them to map coordinates from one X-ray image to the other.

7.1.1 Ipsilateral registration between CC and MLO mammographic views

An example case that illustrates the general pipeline is given in Figure 7.1. Step (1) shows the CC view of a patient's X-ray mammogram. The area illustrated in green corresponds to the detected lesion (Invasive Ductal Carcinoma). Step (2) shows the deformed MRI volume at the registration position, obtained via a FEM registration between the MRI and the CC view mammogram. The ray illustrated in red corresponds to the voxels that are projected to the centre of the green annotated area in the previous step. Step (3) shows the original undeformed MRI of the patient and the red voxels correspond to the ray illustrated in the previous figure. For this step, we are using the transformation result of the CC FEM registration. At this stage we know which voxels of the original MRI correspond to the lesion identified in the CC view. Step (4) shows the deformed MRI volume at the registration position, obtained via a FEM registration between the MRI and the MLO mammogram. For this step, we are using the transformation incorporates a rolling of the breast of approximately 45°. Finally, step (5) shows the MLO view of the patient's X-ray mammogram overlaid with the lesion identified in the MLO view in green and the projection of the lesion position from the CC view (corridor shown in the previous step) in red. As expected, the correspondence between the CC and MLO view of the patient is not one-to-one, but one-to-many.

The registration framework described in the previous chapters does not need to be altered for establishing correspondence between mammograms. However there is an additional process that needs to be followed. So far, we have seen that the output of the registration is the set of optimised transformation parameters. These parameters are then used to map regions of interest between the MRI and the two X-ray images. We have described previously how we map a position in the MRI onto the 2D X-ray mammogram. For the inverse step that is used to map the selected X-ray coordinates to the MR, we simply cast a ray from a specified 2D position on the mammogram, in this case the centre of mass of the X-ray annotation, and then extract the four neighbours of the intersections of this ray with the 3D MR grid, shown in red in step (2). We then use the FEM transformation from the registration result to map the coordinates from the compressed to the uncompressed MRI. These transformed intersections are shown in step (3) in red.

For display purposes, the corridors in Figure 7.1 have been dilated, using a binary structuring element, and subsequently smoothed using Gaussian smoothing. This process has the effect of producing a single connected region when the surface mesh is extracted. In practice, this corridor can have discontinuities, particularly during the mapping process from stage (2), the CC compressed volume, to stage (3), the undeformed MRI. In all the other results displayed in the experiments (section 7.2) the corridors



Figure 7.1: The process of determining corresponding findings between the CC and MLO view mammograms of a patient, through the MRI. The patient shown is p3.



Figure 7.2: The process of determining corresponding findings between temporal CC view mammograms of a patient, via the MRI. The centre of the annotation in the mammogram t_1 (1) is mapped onto the mammogram t_2 (5), illustrated as green. The mammogram acquired at t_2 contains a clip, illustrated as red. Their overlap is yellow. The patient shown is p5.

have not been dilated or smoothed.

7.1.2 Temporal mammogram registration

Following the same concept, we can use this framework for the alignment of temporal mammographic images of the same view. For the registration between CC view mammograms acquired at timepoints t_1 and t_2 , steps (4) and (5) in Figure 7.1 will be different. Assuming that the X-ray mammogram in step (1) was acquired at t_1 , then mapping a region in the X-ray mammogram to the undeformed MRI shown in step (3) requires the use of the CC FEM transformation between the MRI and the X-ray mammogram at t_1 . Subsequently, we can then use the CC FEM transformation between the MRI and the X-ray mammogram at t_2 to obtain a new deformed volume in step (4) and then project this on the t_2 mammogram in the final step (5). The same process can be applied to the MLO view mammograms.

The temporal mammogram registration process is valid if we assume that there is no significant change of the breast tissue during this time. This can be ensured if the two timepoints are close to each other. In practice, when there is a large time difference between the two image acquisitions the breast tissue appearance can change. For example this can be caused by changes in the amount of fatty tissue, the effect of tumour growth, or the removal of breast tissue during surgery. For these cases, and if there is an MRI acquired at both timepoints t_1 and t_2 , we can add one more step in the registration pipeline that
Table 7.1: Registration errors in *mm* for the CC to MLO registrations on ten cases. The clip cases are p9 and p10.

	p1	p2	р3	p4	p5	рб	p7	p8	p9	p10	mean	std
FEM CC-to-MLO	15.6	11.5	0.1	15.6	2.3	8.9	9.4	0.2	13.0	5.1	8.1	5.9

corresponds to the temporal alignment of the MR images. Such temporal MRI registration techniques have been recently proposed to monitor change of the breast tissue over time and tumour response to treatment ([Chittineni et al., 2007], [Li et al., 2009], [Boehler et al., 2010]).

7.2 Experiments

For validation we used clinical X-ray mammograms of patients for whom an MRI was also acquired. The mammograms included annotated lesions of various pathologies. We also included cases that had an X-ray compatible clip inserted after biopsy. The clip was inserted at the position of the lesion and consequently the annotations and clip positions were used as gold standard correspondences across the mammograms. The inter-observer error of the manual annotations calculated from 4 radiologists on 4 of the mammograms used in the experiments was 2.8mm.

For the evaluation, we consider the centres of mass of the annotations and/or clips. Our error metric is the minimum of the 2D Euclidean distances between the centre of mass of the annotation/clip in the second mammogram and the projection of the uncompressed and recompressed ray locus corresponding to the centre of the annotation in the first mammogram. This is the same error metric that was also proposed by [Kita et al., 2001].

7.2.1 Ipsilateral registration between CC and MLO mammographic views

In these experiments, we used data from ten patients, for which we had both MR and X-ray images acquired approximately at the same time point. Two patients had inserted clips that were used for validation (section A.4.1) and the rest had annotated lesions (section A.3).

The registration errors are given in Table 7.1. The mean registration error for this task is $8.1 \pm 5.9mm$. Figures 7.3 and 7.4 illustrate two example cases. It is clear that the projected area on the MLO view is not a straight line. Also its length can vary, depending on the position of the lesion inside the breast: for example the length of the line in Figure 7.4 is relatively short, as the lesion is positioned close the the edge of the mammogram, while the line illustrated in Figure 7.3 is longer, as the lesion is positioned closer to the centre of the breast. The advantage of our method compared to other techniques that use a distance metric from the nipple and the chest wall, is that it can represent each correspondence case individually and therefore incorporate these variations, as a patient-specific model is used for each case.



Figure 7.3: CC to MLO registration results for patient p8. The X-ray annotation of the CC view and its projection on the MLO are shown in green. The X-ray annotation on the MLO view is shown in red. The error is 0.2mm.



Figure 7.4: CC to MLO registration results for patient p10. The clip on the CC view its projection on the MLO are shown in green. The X-ray clip on the MLO view is shown in red. The error is 5.1mm.



Figure 7.5: 2D registration errors for the temporal CC and MLO view registrations of 6 patients. Our method is illustrated in red and it is compared against a 2D affine transformation (green) and the initial 2D error, without registration (blue). The clip cases are patients p5 and p6.

7.2.2 Temporal mammogram registration

In these experiments, we demonstrate the use of the same framework for the matching between temporal mammograms. We used data from six patients, four with annotated lesions at two time points (section A.3) and two with annotated diagnostic mammograms that were mapped to X-rays obtained after biopsy, with a clip inserted at the lesion's position (section A.4.2). Our method was compared against a 2D affine registration, that has been proven previously to outperform other methods ([van Engeland et al., 2003], [Pereira et al., 2010]). The experiments that use the 2D affine, intensitybased registration using normalised mutual information were performed by John Hipwell, as part of the work published in [Pereira et al., 2010]. The algorithm was previously implemented by Julia Schnabel and Daniel Rueckert based on an original algorithm developed by Colin Studholme. We also provide the errors when no registration is performed.

All results are given in Figure 7.5. The mean registration error of our approach is $3.5 \pm 3mm$, the 2D affine registration error is $6.5 \pm 4.6mm$ and no registration gives an error of $10.5 \pm 5.4mm$. As we can see, our method outperforms the 2D affine registration method. It also provides good registration results for all cases, with a maximum error of 8.5mm, while the maximum error of the 2D affine transformation technique is 14.4mm. Two example cases, one good registration result and one with a larger error, are illustrated in Figures 7.6 and 7.7. The figures illustrate the one-to-many mapping between temporal mammograms.

7.3 Discussion

In this chapter, we have demonstrated the use of our FEM-based MRI to X-ray mammography registration framework for the mapping of findings between X-ray mammograms of the same patient that have been acquired from a different view or at a different timepoint. It is the first time that a patient-specific 3D model of the breast with a physically realistic biomechanical simulation has been used for this task.



Figure 7.6: Temporal registration results for the CC view of patient p1. The X-ray annotation at timepoint t_1 is shown in green, as well as its projection on the second mammogram. The X-ray annotation of p1 at timepoint t_2 is shown in red. Their overlap is yellow. The error is 0.08mm.



Figure 7.7: Temporal registration results for the MLO view of patient p6. The X-ray annotation at timepoint t_1 is shown in green, as well as its projection on the second mammogram. The clip of p6 at timepoint t_2 is shown in red. Their overlap is yellow. The error is 6mm.

The results showed that our method outperforms the 2D affine transformation that is typically used for temporal alignment. In our experiments we assumed that the breast tissue of the patient has not changed significantly over time. In future work and if an MRI is available at both time points, an additional step could be inserted into the pipeline, that would perform a non-rigid registration between the two MRIs, to account for any changes in breast volume and structure.

Our proposed framework was also tested for the matching between CC and MLO mammograms. The errors for these registrations were larger than the those of the temporal registration. This was expected, as the difference in breast deformation between same view mammograms will inevitably be much smaller than the difference between CC and MLO views. Compared to other methods proposed in the literature for the same task, the method of Kita et al. [Kita et al., 2001] is the one that is most closely related to our approach, as it did not make use of CAD extracted features. Instead the authors used a simplified breast model reconstructed from the two mammographic views. The mean registration error in their study was 6.8mm, which is comparable to our results. Nevertheless, the two mean errors cannot be directly compared, as they were not applied to the same data sets. Future work includes validation on a larger data set and further investigation of the effect that the biomechanical modelling has on the registration accuracy, such as the use of more tissue types and more accurate modelling of the pectoral muscle.

Finally, the main drawback of this approach is that it requires an MRI of the patient to be also acquired, in order to build the 3D model. Nevertheless, this is the first method that proposes a patient-specific model with a physically realistic breast deformation for this application. Previous methods which employ 2D transformations are fundamentally inappropriate for this task. This study can lead to a model-based approach for this task, that could be performed without MRI in future.

Chapter 8

Conclusion

This thesis has presented a general framework for intensity-based registration between MR images and X-ray mammograms. For this purpose, we have proposed an EM-MRF breast tissue classification technique that is required for the simulation of an X-ray attenuation volume from the MRI (chapter 3). We then investigated the use of three different transformation models, with increased complexity, that attempt to capture the complex breast deformation between the prone position in the MR scanner and the compression between two plates during the X-ray mammogram acquisition (chapters 4, 5 and 6). Finally, we have also investigated the use of the same framework for determining correspondences between X-ray mammograms through the use of the 3D deformation model provided by the MRI of the patient (chapter 7).

In section 8.1 we summarise and discuss the findings of each method proposed in all the above thesis chapters. Finally, in section 8.2 we discuss possible limitations, we propose methods to overcome these in future work and we present different potential application areas of this framework.

8.1 Summary and conclusions

To summarise the advantages of our proposed registration framework, this is the first approach that demonstrates the alignment between clinical MRI and X-ray images of the breast based on the breast's internal structures. Moreover, this framework is generic and can incorporate any transformation model, optimisation strategy or similarity metric by virtue of its derivation from the Insight Toolkit registration methodology [ITK, 2003]. All three proposed transformation models, from geometrical, through to ellipsoidal-SDM and finally to FEM-based patient-specific, give reproducible results and parametrise the space of the complex breast deformations using a low number of degrees of freedom. Our approach is targeted for clinical use, as it could be potentially easily incorporated into clinical practice. Below we summarise the conclusions for each one of the methodologies described in this thesis.

8.1.1 X-ray simulations from an MR volume

In chapter 3 we have seen in detail our methodology for simulating X-ray mammograms using the MR volume. Although the research topic of this work is not the accurate X-ray image simulation from the MRI, this is an important component of the registration framework, as the similarity between real and simulated mammograms is used to drive the optimisation process.

The main contribution is the use of an EM-MRF breast tissue classification algorithm that was previously proposed for the classification of brain tissue MRI voxels. As it has been discussed, the validation of this technique is problematic due to the lack of a ground truth dataset with known, non-binary classifications of the voxels in the breast. Consequently, the choice of this approach was made due to its advantages over manual thresholding and specifically its ability to generate non-binary classifications, give reproducible results and produce simulated X-ray mammograms that include finer details of the fibroglandular structures than those produced using manual thresholding.

Overall, the use of the X-ray image simulation technique has provided us with the flexibility to compare different transformation models, with increasing complexity, for their suitability in MRI to X-ray mammography registration, as in all cases the simulation technique remained the same. The described framework includes certain simplifications, as for example it does not account for scattering, nevertheless it is computationally efficient and generates X-ray images that provide higher registration accuracy than other proposed methods for this task. In the next section we discuss how this methodology could be further improved and optimised to potentially increase registration accuracy.

8.1.2 MRI to X-ray registration using an affine transformation

Our proposed framework for MRI to X-ray image registration has been introduced in chapter 4, where we have also presented our experiments using a volume-preserving affine transformation model. This was the first work that investigated an intensity-based 2D/3D registration framework for this task. Previous methods have mainly used a small number of extracted features or the breast outline, while in this work we propose an alignment based on the matching of the structures within the whole breast. This method uses the complete information about the images and therefore provides a more robust alternative to feature-based techniques.

Furthermore this was the first methodology that was validated in a large number of clinical datasets. The experimental results on 113 registrations tasks, coming from 49 patients, showed that this framework with a relatively simple transformation model could potentially be clinically useful for breast cancer detection and diagnosis, providing a median registration error of 13.1mm.

8.1.3 MRI to X-ray registration using an ellipsoidal breast model

Chapter 5 describes our work that proposed a new transformation model with increased complexity, in comparison to the affine, incorporated inside the same registration framework. We have proposed an ellipsoidal SDM transformation model that uses breast biomechanical simulations to learn the space of plausible deformations that the breast undergoes under compression, by extracting the main modes of variation. As part of the transformation model, apart from the breast pose parameters, we also optimise the relative weights of the extracted modes of variation during the iterative optimisation scheme of the registration.

This approach has the advantage of using a more physically realistic transformation model than the affine, as it is based on biomechanical simulations of a plate compression. Furthermore, the approximation of the breast shape as ellipsoidal, has the advantage of providing a model that is more easily adaptable clinically, as it does not require patient-specific models to be built. The results on 10 clinical

datasets demonstrate an improvement in the accuracy (9.8mm) compared to the affine transformation model (15.9mm).

8.1.4 MRI to X-ray registration using a patient-specific biomechanical model

The final transformation model that we have proposed is a patient-specific biomechanical model of the breast (chapter 6), whose parameters are iteratively updated inside the registration framework. The main advantage of this approach is the use of a specific model for each patient, extracted from the MRI and also the use of the updated rigid transformation parameters and the biomechanical simulation parameters in the same framework, using an FE solver that runs on the GPU. This is the first approach that used a 2D/3D intensity-based framework for this task, with iteratively optimised parameters of both the pose and biomechanical modelling of mammographic compression.

The experiments on five clinical datasets show that this method outperforms the other two transformation models, although for the MLO view the results show slightly lower accuracy, which is comparable with the affine transformation. In addition to the improved mean accuracy of the patient-specific approach, the standard deviation about this mean is also lower than the other methods. This suggests (although the number of data sets is small) that this transformation model provides more consistent results. To summarise the results of all three transformation models, these are displayed in Figure 8.1.

8.1.5 Relating findings between X-ray mammograms via an MR volume

Finally, in chapter 7 we have demonstrated the use of the same registration framework for the mapping between temporal and different view mammograms of the same patient. This task was mainly approached so far as a one-to-one correspondence mapping and 2D transformation models were employed. Our technique takes into account that the correspondences are one-to-many and thus the alignment is performed based on the 3D information provided by the MRI.

Our experiments showed that this framework could also be used for this application, providing good accuracy. The registration error for the temporal mammogram registration was reduced to $3.5 \pm 3mm$, while the commonly employed 2D affine registration provided $6.5 \pm 4.6mm$ on the same dataset. Furthermore the projection of a point in one mammogram to a curve in the other, provides an illustration of the accuracy that can be achieved when employing the use of various distance metrics for matching, for example distances from the nipple or the pectoral muscle. These techniques are frequently followed in clinical practice.

8.2 Future work

In this section we discuss how this work can be further improved and extended in future work. We also propose different applications, in which the same framework could be potentially integrated.

Further validation of the registration framework

Our experiments that compared all three proposed transformation models for registration, indicate that the patient-specific method performs best. Further validation tests performed on larger datasets will ensure that these results are not specific to the dataset used. It would also be beneficial to use a common









Figure 8.1: Illustration of the registration errors achieved by the three proposed transformation models for the MRI to X-ray mammography registration (affine, E-SDM and FEM). The errors correspond to the ones previously presented in Table 6.3.

117

validation dataset in order to compare our method with other FEM-based techniques previously proposed for this task ([Ruiter et al., 2006], [Hopp et al., 2012], [Lee et al., 2011]).

Validation of the X-ray simulation technique

As discussed in chapter 3, the validation of the breast tissue classification algorithm is problematic, due to lack of ground truth data sets. The EM-MRF technique that we propose for this task and the X-ray simulation methodology that is described in appendix B, provide simulated X-ray mammograms that can drive the intensity-based registration and lead to a convergence of the optimisation algorithm. Nevertheless, both techniques incorporate the use of user-defined parameters and simplifications, such as the cost of having the class 'Fat' next to 'Glandular' tissue and the use of an effective attenuation volume which ignores beam hardening effects. In future work one could further investigate the effect that these parameters and simplifications have on the iterative registration in terms of its convergence performance and accuracy.

Further optimisation of the implementation

Future work could include further optimisation of the current pipeline implementation. For example, the use of a fully-automated pectoral muscle segmentation currently exists [Gubern-Merida et al., 2011] and therefore, if incorporated, it can provide a framework that does not require any manual interaction, enabling an easier integration into clinical practice.

Furthermore, the implementation of the ray-casting algorithm could be improved. As we saw in chapter 5, for the non-rigid registration framework at the moment we are computing the intersections of the rays with the 3D grid at each iteration of the registration. One option to accelerate this process is to save these intersections in memory, as they remain the same throughout the registration process. Moreover, a GPU implementation could be used to speed up this process, as the ray-casting is an algorithm that can be easily parallelised.

The optimisation scheme used in the registration could also be further investigated, as other optimisation strategies could lead to faster and potentially more robust convergence.

Improved breast biomechanical modelling

For the biomechanical modelling of the breast, we have made several assumptions that have contributed to an automated implementation of the pipeline and help to ensure that the FE solver converges robustly. These include the use of a homogeneous tissue type, the subsampling of the MRI before meshing, and the approximation of the pectoral muscle as a plane. In future work, one could investigate further whether the use of more tissue types (for the fibroglandular, fat, tumour and skin), the accurate modelling of the pectoral muscle, and the removal of the effect of gravity from the MRI before compression, can improve the results obtained so far.

Use of the post-contrast MR images

In the work described in this thesis, we have used only the pre-contrast DCE-MRI of the breast for registration. The use of the post-contrast image combined with the pre-contrast can give detailed information about the location and the shape of the tumour, as we have seen in section 1.2.2. In future work one could investigate further the benefit of incorporating information from the post-contrast images to assign a different attenuation coefficient to the tumorous tissue than the normal fibroglandular structures and thus enhance the contrast between them and produce more realistic X-ray simulations from the MRI.

Application in breast cancer detection and diagnosis

As it has been discussed before, the MRI to X-ray mammography registration can be used as part of a clinical tool in order to assist radiologists in integrating the information shown in the two modalities and potentially improve breast cancer detection and diagnosis. Assuming that the radiologists are using a multimodal workstation, our algorithm could be part of the viewing tool, where the radiologists would be able to navigate their cursor through the MRI and view the corresponding location on the CC and MLO view mammograms either as a point position or as a disk with a radius that equals the mean registration error extracted from a validation test. This tool would be particularly useful for the detection of mammographically occult lesions and further evaluation of mammographically detected lesions. For this application, the algorithm can run once off-line, when both images become available. The result can be stored as a transformation model, allowing the multimodal viewing to be performed in real-time.

Application in CAD algorithms

Our MRI to X-ray registration framework could be incorporated into multimodal CAD algorithms [Yuan et al., 2010], to automatically map features extracted from the MRI to both X-ray mammographic views. This automatic correspondence between features would help both in the training process of feature classifiers and also in the testing of a suspicious new finding, as the classifier would take into consideration the information extracted from the regions of interest in both modalities. Moreover, the mapping between the CC and MLO view of the mammograms, that was described in chapter 7, could also be useful in X-ray CAD algorithms, although this currently requires the patient to also have an MR image acquired.

Application in X-ray guided biopsy

Another future potential application of our framework is its use as a tool to enable X-ray guided biopsy for lesions that are not clear or visible in ultrasound and X-ray. Currently, these patients undergo MRIguided biopsy, which is expensive, time-consuming and inconvenient for clinical practice. If an MRI to X-ray registration framework were available with a sufficient level of accuracy, then a lesion visible in the MRI could be identified and mapped onto the X-ray mammogram. This would enable the radiologist to perform an X-ray guided biopsy, which is a procedure that would be faster, less expensive and more widely available than using MRI. Nevertheless this application would require better accuracy than multimodal detection and diagnosis and therefore our implementation would need further improvement and validation.

X-ray mammogram registration

We proposed a novel technique that models, in a physically plausible way, the 3D deformation that occurs between two mammographic acquisitions (temporal and different view) and uses this to establish correspondence. One limitation of our method is that it currently requires the MRI and therefore its clinical value is limited to these cases. However, this technique paves the way for a model-based approach that could be performed without MRI in future. For example, one could use a generic 3D model of each patient's breast, whose shape and dimensions would be estimated from the CC and MLO view mammograms. This model could then be used to relate the information between two X-ray projection images. Such a model would then be applicable to all patient cases and would model a physically real-istic deformation, as opposed to previous methods which employ 2D transformations, and are therefore inappropriate for this task.

List of publications

2012

- [Mertzanidou et al., 2012a] Mertzanidou, T., Hipwell, J., Cardoso, M., Zhang, X., Tanner, C., Ourselin, S., Bick, U., Huisman, H., Karssemeijer, N., and Hawkes, D. (2012). MRI to X-ray mammography registration using a volume-preserving affine transformation. *Medical Image Analysis*, 16:966-975.
- [Mertzanidou et al., 2012b] Mertzanidou, T., Hipwell, J., Han, L., Taylor, Z., Huisman, H., Bick, U., Karssemeijer, N., and Hawkes, D. (2012). Intensity-based MRI to X-ray mammography registration with an integrated fast biomechanical transformation. In *International Workshop on Breast Imaging*, pages 48-55.

2011

 [Mertzanidou et al., 2011] Mertzanidou, T., Hipwell, J., Han, L., Huisman, H., Karssemeijer, N., and 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*, pages 161-168.

2010

- [Mertzanidou et al., 2010a] Mertzanidou, T., Hipwell, J., Cardoso, M., Tanner, C., Ourselin, S., and Hawkes, D. (2010). X-ray mammography -MRI registration using a volume-preserving affine transformation and an EM-MRF for breast tissue classification. In *International Workshop on Digital Mammography*, pages 23-30.
- [Mertzanidou et al., 2010b] Mertzanidou, T., Hipwell, J., Tanner, C., and Hawkes, D. (2010). An intensity-based approach to X-ray mammography MRI registration. In *Proceedings of SPIE Medical Imaging: Image Processing*, pages 7623-106.

Appendix A

Summary of datasets

This section summarises the datasets used in all the experiments described in the thesis.

A.1 MRI and X-ray film mammograms

These are DCE-MR images and digitised film mammograms from the UK MR Breast Screening Study MARIBS [Leach et al., 2005] and the Cambridge University Hospitals NHS Foundation Trust. The MR volumes used are the pre-contrast T1-weighted images and have voxel dimensions $[1.33 \times 1.33 \times 2.5]mm^3$. The original digitised mammograms have a pixel size of $[0.1 \times 0.1]mm^2$ and were resampled to $[1 \times 1]mm^2$ for registration, as discussed in the report.

A.2 MRI breast compression volunteer data

These data were acquired in the University College London Hospital, as part of the validation study described in [Tanner et al., 2011]. They are pre-contrast T1-weighted MR images with voxel size of $[1.0 \times 1.0 \times 2.5]mm^3$. The volunteers had initially pre-contrast images acquired without any compression and then repeated scans with an increasing amount of compression from the lateral to medial direction.

A.3 MRI and FFDM with known annotated lesions

These cases were acquired in Radboud University Nijmegen Medical Centre and include DCE-MRI and FFDM from women with benign or malignant lesions that are visible in both modalities. We used the pre-contrast MRI T1-weighted images. The resolution of the MRIs varies across patients and is either $[0.9 \times 0.9 \times 1.0]mm^3$ or $[0.6 \times 0.6 \times 1.3]mm^3$. The original resolution of the X-ray mammograms is $[0.1 \times 0.1]mm^2$; they were subsampled by a factor of 10 for registration.

The identified lesions were annotated by expert imaging scientists using one or multiple spheres for the MRI. A tresholding approach was used on the contrast enhancement within the sphere/spheres to find the enhancing lesion part following Otsu's method. A disk or a free-form shape for the annotation of the X-ray mammograms.

A.4 MRI and FFDM with clips

A.4.1 Data with MR and X-ray compatible clips

These are patient cases that had both an MRI and also X-ray mammograms acquired with an MR and X-ray compatible clip at the lesion's position, in the Charite Universitätsmedizin Berlin. The voxel resolution of the MRI (T1-weighted) is $[0.7 \times 0.7 \times 2]mm^3$ and the FFDM $[0.085 \times 0.085]mm^2$. The X-ray mammograms was subsampled by a factor of 10 for registration.

A.4.2 Data with X-ray compatible clips inserted after biopsy

These cases include patients with lesions that are not easily identified on the X-ray diagnostic mammograms. These were annotated both on those images and also on the diagnostic T1-weighted MRIs. These patients had also an MR-guided biopsy where an clip was inserted at the lesion's position. Subsequently, another FFDM was acquired after biopsy, where the clip was visible. Correspondence between these temporal mammograms can be established using the clip positions and the annotations on the diagnostic images.

These cases were acquired in Radboud University Nijmegen Medical Centre. The voxel resolution of the MRI is $[0.7 \times 0.7 \times 1.3]mm^3$ and the FFDM $[0.1 \times 0.1]mm^2$. The X-ray mammograms was subsampled by a factor of 10 for registration.

Appendix B

Computation of an X-ray attenuation volume from the MRI

The following methodology is work of John Hipwell and it was published as part of our study in [Mertzanidou et al., 2012a].

After classifying each voxel in the MR volume into fibroglandular and fatty tissue, we then calculate an "effective" monoenergetic X-ray attenuation volume, H. H captures the relative non-linear attenuations of a poly-energetic X-ray spectrum, by fat and fibroglandular tissue, in a single volume. In this way we can repeatedly simulate DRRs during the iterative registration, using a simple raycasting and summation of H, without having to recompute the attenuated spectrum for each ray cast. [Robinson and Scrimger, 1991] demonstrate that that the use of a theoretical, effective attenuation and a mono-energetic beam provides good agreement with laboratory measurements. Each voxel H(i) in this volume is given by:

$$H(i) = \frac{\int_{\epsilon=0}^{\epsilon_{max}} N_0(\epsilon)\epsilon. \left\{ P_F(i)\mu_F(\epsilon) + P_G(i)\mu_G(\epsilon) \right\} d\epsilon}{\int_{\epsilon=0}^{\epsilon_{max}} N_0(\epsilon)\epsilon d\epsilon}$$
(B.1)

where $N_0(\epsilon)$ is the X-ray spectrum with respect to photon energy ϵ , $P_{F|G}(i)$ is the probability of tissue classes Fat (F) or Glandular (G) for voxel *i*, given by the EM-MRF classification and $\mu_{F|G}(\epsilon)$ is the linear attenuation of tissue classes F or G at photon energy ϵ . Details of the relevant X-ray parameters, namely the anode type and anode angle, are obtained from the mammogram's DICOM header and the manufacturer's X-ray set specifications, respectively. From these the X-ray spectrum, $N_0(\epsilon)$, can be estimated using published data [Cranley et al., 1997]. Similarly the linear attenuation coefficients, $\mu_{F|G}(\epsilon)$, can be obtained from publicly available data published by the National Institute of Standards and Technology (NIST) [Hubbell and Seltzer, 2004]. The purpose of this volume is to create an image which can be repeatedly and efficiently projected to simulate a Full-Field Digital Mammogram. Beer's law [Beer, 1852] describes the absorption of X-ray photons by distance h of a given material with linear, mono-energetic, attenuation coefficient μ . For a monochromatic X-ray beam this can be expressed as follows:

$$I = I_0 e^{-h\mu} \tag{B.2}$$

where I_0 is the incident photon energy and I is the attenuated energy. When viewing an image of X-ray

attenuation, such as a mammogram, it is common to "log invert" the raw data, I, so that intensities in the mammogram, I_M , reflect the total attenuation of X-rays reaching the detector:

$$I_M = -\ln\left(I_0 e^{-h\mu}\right) \tag{B.3}$$

$$= -\ln(I_0) - \ln(e^{-h\mu})$$
 (B.4)

$$= -\ln(I_0) + h\mu \tag{B.5}$$

Substituting the effective attenuation volume H in equation B.5 and ignoring the constant $\ln(I_0)$ term gives:

$$I_M = \int_{h=0}^{h_{max}} H(h)dh \tag{B.6}$$

which is simply a ray-casting of H.

Appendix C

Clinical validation on findings poorly or not visible on FFDM

The purpose of this set of experiments was to validate the affine transformation model pipeline in a clinical setting, as a requirement of the HAMAM project [HAMAM, 2012]. The goal was to investigate the benefit of using automatically determined correspondences in cases where the lesion was poorly or not visible in the X-ray mammogram, as these are the cases where a registration can be useful to radiologists. Therefore, the data used was from patients that had an MR-guided biopsy as the lesion was not visible in the X-ray mammogram. This was joint work with the HAMAM partners Radboud University Nijmegen Medical Centre, MeVis Medical Solutions and University College London.

The result of the registration was presented to radiologists via a "linked cursor" which enables the user to navigate through the MRI and view the projection of the MR position on the X-ray mammogram. The radiologists were asked to indicate the lesion position on the CC and MLO view of the mammograms with or without using the linked cursor. They were also asked to indicate how visible was the lesion and how confident they were for their prediction. The experiment description as given to the radiologists is given in Table C.1.

Data from 12 patients were used¹. The six radiologists participating in the experiments were randomly allocated to view half of the cases with and half without the linked cursor between the modalities.

Due to technical problems related to the data exchange and the integration to the workstation, there were several cases that were not accessed by some of the radiologists during this initial evaluation. As a result, no statistical analysis could be performed using the data. Nevertheless, the results are summarised below.

Firstly, we have calculated the projection of the MR annotation on the X-ray mammograms, to which we refer here as the *registration result*. We have calculated the mean distance between the mean radiologists' annotation and the *registration result* both without and with the use of the linking cursor. The distance was decreased using the cursor from 35.05mm to 31.37mm.

Also, to get an indication of how well the radiologists agreed on the X-ray lesion location, we have calculated the mean distance of their individual annotations to the mean annotation, so in other words

¹Data from: Radboud University Nijmegen Medical Centre

Table C.1: Experiment description given to the radiologists. This was created by the HAMAM partners Radboud University Nijmegen Medical Centre and University College London.

In this experiment, you are being asked to provide your expert opinion of the location of one or more lesions in a pair of CC and MLO X-ray mammograms, corresponding to predefined lesions visible in MRI. We will use this information to evaluate the accuracy and efficacy of our spatial correspondence software.

In half the cases you will be shown, the location of the X-ray lesion, as predicted by the current version of our registration software, will also be displayed. We ask you to consider this information before marking your own opinion of the lesion location in each of the CC and MLO views.

In the other half of the cases, you will be asked to mark your prediction of the lesion location without the benefit(!) of the registration software's prediction.

For the first few cases the lesion in the X-ray mammogram should be clearly visible, but in the later cases the lesion may not be visible at all. Either way, you are asked to make your best guess of the lesion location in the X-ray mammograms. The location you choose should correspond, to the best of your ability, to the location displayed in the MRI.

In addition to the location, you are asked to answer the following questions for each case that you assess:

1. On a scale of one to five, how visible is the lesion in the X-ray mammogram?

1 = very obvious, 5 = not visible at all.

- 2. On a scale of one to five, how confident are you that your predicted lesion location is close to, i.e. within approximately 20mm of, the correct location in the X-ray mammogram?
 - 1 = very confident, 5 = very uncertain.

Mammogram view	patient 1	patient 2	patient 3	patient 4	patient 5	
CC - without	75.76	99.71	115.91	107.70	91.95	
CC - with	66.05	59.76	22.34	28.59	82.85	
MLO - without	107.69	152.81	102.34	111.02	175.30	
MLO - with	42.86	14.80	33.34	16.16	84.56	

Table C.2: Mean distance (in mm) of the radiologists' annotation from their mean position (variance) for the five cases that had two or more annotations both without and with the linking cursor.



Figure C.1: The mean responses of the radiologists to questions 1 and 2. *Question 1*: On a scale of one to five, how visible is the lesion in the X-ray mammogram? *Question 2*: On a scale of one to five, how confident are you that your predicted lesion location is close to, i.e. within approximately 20mm of, the correct location in the X-ray mammogram?

the variance. To do this, we needed patients that were accessed by two or more radiologists both without and with the use of the linking cursor. There were five of these cases and the results are summarised in Table C.2. We can see that in all cases the mean distance was decreased with the use of the *registration result*.

Finally, the radiologists' impression for the difficulty of the cases viewed and their confidence of their annotation is summarised in Figure C.1.

Due to the small number of cases assessed, we cannot perform any statistical analysis on the results. Nevertheless, a first indication that we can obtain is that the linking cursor has reduced the spread of the radiologists' annotations and that these were closer to the registration result when the cursor was used. The radiologists were therefore influenced by the registration result.

For this dataset, as the lesion is not visible in the X-ray mammogram, there is no gold standard correspondence available. Therefore, even if a study with a larger dataset determines that the radiologists are influenced, the results cannot be conclusive about whether the registration has influenced them towards the correct lesion location. This can only be done for a dataset that has clearly defined correspondences, as the one described in the previous section. It is also worth noticing the wide range of values in Table C.2, and the large error in *mm* both using the linked cursor and without (31.37*mm* and

35.05mm respectively). The numbers indicate that establishing correspondences in this set of data is a particularly difficult task for radiologists. However, we can acknowledge the fact that the reported 2D locations indicated by the radiologists might incorporate errors related to the communication of the data across the project partners.

Bibliography

- [Arsigny et al., 2009] Arsigny, V., Commowick, O., Ayache, N., and Pennec, X. (2009). A fast and log-Euclidean polyaffine framework for locally linear registration. *Journal of Mathematical Imaging and Vision*, 33:222–238.
- [Arsigny et al., 2006] Arsigny, V., Commowick, O., Pennec, X., and Ayache, N. (2006). A Log-Euclidean framework for statistics on diffeomorphisms. In *Medical Image Computing and Computer-Assisted Intervention*, pages 924–931.
- [Arsigny et al., 2005] Arsigny, V., Pennec, X., and Ayache, N. (2005). Polyrigid and polyaffine transformations: a novel geometrical tool to deal with nonrigid deformations-application to the registration of histological slices. *Medical Image Analysis*, 9:507–523.
- [Ashburner, 2007] Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *Neuroim-age*, 38:95–113.
- [Ashburner and Friston, 2005] Ashburner, J. and Friston, K. (2005). Unified segmentation. *Neuroimage*, 26:839–851.
- [Aylward et al., 2002] Aylward, S., Jomier, J., Guyon, J., and Weeks, S. (2002). Intraoperative 3D ultrasound augmentation. In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, pages 421–424.
- [Azar et al., 2000] Azar, F., Metaxas, D., and Schnall, M. (2000). A finite model of the breast of predicting mechanical deformations during biopsy procedure. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 38–45.
- [Azar et al., 2002] Azar, F., Metaxas, D., and Schnall, M. (2002). Methods for modeling and predicting mechanical deformations of the breast under external pertubation. *Medical Image Analysis*, 6:1–27.
- [Bakic et al., 2004] Bakic, P., Richard, F., and Maidment, A. (2004). Effect of breast compression on registration of successive mammograms. In *International Workshop on Digital Mammography*.
- [Beer, 1852] Beer, A. (1852). Bestimmung der Absorption des rothen Lichts in farbigen Flüssigkeiten. *Annalen der Physik*, 162:78–88.

- [Behrenbruch et al., 2003] Behrenbruch, C., Marias, K., Armitage, P., Moore, N., English, R., Clarke, J., and Brady, M. (2003). Fusion of contrast-enhanced breast MR and mammographic imaging data. *Medical Image Analysis*, 7:311–340.
- [Birkfellner et al., 2005] Birkfellner, W., Seemann, R., Figl, M., Hummel, J., Ede1, C., Homolka, P., Yang, X., Niederer, P., and Bergmann, H. (2005). Wobbled splatting - A fast perspective volume rendering method for simulation of X-ray images from CT. *Physics in Medicine and Biology*, 50:73– 84.
- [Boehler et al., 2010] Boehler, T., Schilling, K., Bick, U., and Hahn, H. (2010). Deformable image registration of follow-up breast magnetic resonance images. In *4th International Workshop on Biomedical Image Registration*, pages 13–24.
- [Bro-Nielsen and Gramkow, 1996] Bro-Nielsen, M. and Gramkow, C. (1996). Fast fluid registration of medical images. In *Proceedings of Visualisation in Biomedical Computing*, pages 267–276.
- [Cahill et al., 2007] Cahill, N., Noble, J., and Hawkes, D. (2007). Fourier Methods for Nonparametric Image Registration. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8.
- [Cardoso et al., 2011] Cardoso, M., Clarkson, M., Ridgway, G., Modat, M., Fox, N., Ourselin, S., and The Alzheimer's Disease Neuroimaging Initiative (2011). LoAd: A locally adaptive cortical segmentation algorithm. *Neuroimage*, 56:1386–1397.
- [Carter et al., 2008] Carter, T., Tanner, C., Barratt, D., Beechey-Newman, N., and Hawkes, D. (2008). Mr navigated breast surgery: Method and initial clinical experience. In *Proceeding of Medical Image Computing and Computer-Assisted Interventions*, pages 356–363.
- [Carter et al., 2009] Carter, T., Tanner, C., and Hawkes, D. (2009). Determining material properties of the breast for image-guided surgery. In *Proceedings of SPIE Medical Imaging*, volume 7261.
- [Chen and Giger, 2004] Chen, W. and Giger, M. (2004). A fuzzy c-means (FCM) based algorithm for intensity inhomogeneity correction and segmentation of MR images. In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, pages 1307–1310.
- [Chittineni et al., 2007] Chittineni, R., Su, M., and Nalcioglu, O. (2007). Breast Delineation using Active Contours to Facilitate Coregistration of Serial MRI Studies for Therapy Response Evaluation. In *IEEE International Conference on Image Processing*, pages 261–264.
- [Christensen et al., 1996] Christensen, G., Rabitt, R., and Miller, M. (1996). Deformable templates using large deformation kinematics. *IEEE Transactions on Medical Imaging*, 5:1435–1447.
- [Chung et al., 2008] Chung, J., Rajagopal, V., Nielsen, P., and Nash, M. (2008). Modelling mammographic compression of the breast. In *Proceeding of Medical Image Computing and Computer-Assisted Interventions*, pages 758–765.

- [Collignon et al., 1995] Collignon, A., Maes, F., Delaere, D., Vandermeulen, D., Suetens, P., and Marchal, G. (1995). Automated multi-modality image registration based on information theory. In *Information Processing in Medical Imaging*, pages 263–274.
- [Cooper, 1840] Cooper, A. (1840). *On the anatomy of the breast*. Longman, Orme, Green, Browne and Longmans.
- [Cootes et al., 2001] Cootes, T., Edwards, G., and Taylor, C. (2001). Active Appearance Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:681–685.
- [Cootes et al., 1995] Cootes, T., Taylor, C., Cooper, D., and Graham, J. (1995). Active shape models their training and application. *Computer Vision and Image Understanding*, 61:38–59.
- [Cranley et al., 1997] Cranley, K., Gilmore, B., Fogarty, G., and Desponds, L. (1997). Catalogue of Diagnostic X-ray Spectra and Other Data. Technical report, The Institute of Physics and Engineering in Medicine.
- [Crum et al., 2005] Crum, W., Tanner, C., and Hawkes, D. (2005). Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging. *Physics in Medicine and Biology*, 50:5153– 5174.
- [D'Agostino et al., 2003] D'Agostino, E., Maes, F., Vandermeulen, D., and Suetens, P. (2003). A viscous fluid model for multimodal non-rigid image registration using mutual information. *Medical Image Analysis*, 7:565–575.
- [Dendy and Heaton, 1999] Dendy, P. and Heaton, B. (1999). *Physics for Diafnosic Radiology*. Institute of Physics Publishing, The institute of Physics, London.
- [Diez et al., 2011] Diez, Y., Oliver, A., Llado, X., Freixenet, J., Marti, J., Vilanova, J., and Marti, R. (2011). Revisiting intensity-based image registration applied to mammography. *IEEE Transactions* onInformation Technology in Biomedicine, 15:716–725.
- [Fischer and Modersitzki, 2004] Fischer, B. and Modersitzki, J. (2004). A unified approach to fast image registration and a new curvature based registration technique. *Linear Algebra and its applications*, 380:107–124.
- [Freeborough and Fox, 1998] Freeborough, P. and Fox, N. (1998). Modeling brain deformations in Alzheimer disease by fluid registration of serial 3D MR images. *Journal of Computer Assisted Tomography*, 22:838–843.
- [Gilhuijs et al., 1996] Gilhuijs, K., van de Ven, P., and van Herk, M. (1996). Automatic threedimensional inspection of patient setup in radiation therapy using portal images, simulator images, and computed tomography data. *Medical Physics*, 23:389–399.

- [Gubern-Merida et al., 2011] Gubern-Merida, A., Kallenberg, M., Marti, R., and Karssemeijer, N. (2011). Fully automatic fibroglandular tissue segmentation in breast MRI: atlas-based approach. In *MICCAI Workshop on Breast Image Analysis*, pages 73–80.
- [Guo et al., 2006] Guo, Y., Sivaramakrishna, R., Lu, C., Suri, J., and Laxminarayan, S. (2006). Breast image registration techniques: a survey. *Medical and Biological Engineering and Computing*, 44:15–26.
- [HAMAM, 2012] HAMAM (2007-2012). European 7th Framework Program, Highly Accurate Breast Cancer Diagnosis through Integration of Biological Knowledge, Novel Imaging Modalities, and Modelling, ICT-2007.5.3. http://www.hamam-project.org.
- [Han et al., 2011] Han, L., Hipwell, J., Mertzanidou, T., Carter, T., Modat, M., Ourselin, S., and Hawkes, D. (2011). A hybrid FEM-based method for aligning prone and supine images for image guided breast surgery. In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, pages 1239–1242.
- [Han et al., 2012] Han, L., Hipwell, J., Tanner, C., Taylor, Z., Mertzanidou, T., Ourselin, S., and Hawkes, D. (2012). Development of patient-specifc biomechanical models for predicting large breast deformation. *Physics in Medicine and Biology*, 57:455–472.
- [Hayton et al., 1997] Hayton, P., Brady, M., Tarassenko, L., and Moore, N. (1997). Analysis of dynamic MR breast images using a model of contrast enhancement. *Medical Imge Analysis*, 1:207–224.
- [Hipwell et al., 2003] Hipwell, J., Penney, G., McLaughlin, R., Rhode, K., Summers, P., Cox, T., Byrne, J., Noble, A., and Hawkes, D. (2003). Intensity-based 2-D–3D registration of celebral angiograms. *IEEE Transactions on Medical Imaging*, 22:1417–1426.
- [Hipwell et al., 2007] Hipwell, J., Tanner, C., Crum, W., Schnabel, J., and Hawkes, D. (2007). A new validation method for X-ray mammogram registration algorithms using a projection model of breast X-ray compression. *IEEE Transactions on Medical Imaging*, 26:1190–1200.
- [Hopp et al., 2012] Hopp, T., Baltzer, P., Dietzel, M., Kaiser, W., and Ruiter, N. (2012). 2D/3D image fusion of X-ray mammograms with breast MRI: visualizing dynamic contrast enhancement in mammograms. *International Journal of Computer Assisted Radiology and Surgery*, 7:339–348.
- [Hu et al., 2008] Hu, Y., Morgan, D., Ahmed, H., Pendse, D., Sahu, M., Allen, C., Emberton, M., Hawkes, D., and Barratt, D. (2008). A statistical motion model based on biomechanical simulations for data fusion during image-guided prostate interventions. In *Proceedings of Medical Imaging Computing and Computer Assisted Intervention*, pages 737–744.
- [Hubbell and Seltzer, 2004] Hubbell, J. and Seltzer, S. (2004). Tables of X-Ray Mass Attenuation Coefficients and Mass Energy-Absorption Coefficients from 1 kev to 20MeV for elements Z = 1 to 92 and 48 Additional Substances of Dosimetric Interest. Technical report, Ionizing Radiation Division, Physics Laboratory, NIST, 100 Bureau Drive, M/S 8460 Gaithersburg, MD 20899-8460.

- [Johns and Yaffe, 1987] Johns, P. and Yaffe, M. (1987). X-ray characterisation of normal and neoplastic breast tissues. *Physics in Medicine and Biology*, 32:675–695.
- [Kapur et al., 2004] Kapur, A., Carson, P., Eberhard, J., Goodsitt, M., Thomenius, K., Lokhandwalla, M., Buckley, D., Roubidoux, M., Helvie, M., Booi, R., LeCarpentier, G., Erkamp, R., Chan, H., Fowlkes, J., Thomas, J., and Landberg, C. (2004). Combination of digital mammography with semi-automated 3D breast ultrasound. *Technology in Cancer Research & Treatment*, 3:325–334.
- [Kapur et al., 2002] Kapur, A., Krckerb, J., Astleya, O., Buckleya, D., Eberharda, J., Alyassina, A., Clausa, B., Thomeniusa, K., Myersa, H., Rumseya, M., Johnsona, R., and Karra, S. (2002). Fusion of digital mammography with breast ultrasound: a phantom study. In *Proceedings of SPIE Medical Imaging*, pages 526–537.
- [Khamene et al., 2006] Khamene, A., Blochb, P., Wein, W., Svatos, M., and Sauer, F. (2006). Automatic registration of portal images and volumetric ct for patient positioning in radiation therapy. *Medical Image Analysis*, 10:96–112.
- [Kita et al., 2001] Kita, Y., Highnam, R., and Brady, M. (2001). Correspondence between different view breast X-rays using curved epipolar lines. *Computer Vision and Image Understanding*, 83:38–56.
- [Klein et al., 2009] Klein, A., Andersson, J., Ardekanic, B., Ashburner, J., Avants, B., Chiang, M., Christensen, G., Collins, D., Gee, J., Hellier, P., Song, J., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P., Vercauteren, T., Woods, R., Mann, J., and Parsey, R. (2009). Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *Neuroimage*, 46:786–802.
- [Kopans, 1998] Kopans, D. (1998). Breast Imaging. Lippincott Williams and Wilkins.
- [Krol et al., 2004] Krol, A., Unlu, M., Baum, K., Mandel, J., Lee, W., Coman, I., Lipson, E., and Feiglin, D. (2004). MRI/PET nonrigid breast-image registration using skin fiducial markers. In Workshop on the Nuclear Radiology of Breast Cancer, Rome, Italy.
- [Kumar et al., 1996] Kumar, R., Asmuth, J., Hanna, K., Bergen, J., Hulka, C., Kopans, D., Weisskoff, R., and Moore, R. (1996). Application of 3D registration for detecting lesions in magnetic resonance breast scans. In *Proceedings of SPIE Medical Imaging*, pages 646–656.
- [Kumar et al., 2001] Kumar, S., Sallam, M., and Goldgof, D. (2001). Matching point features under small nonrigid motion. *Pattern Recognition*, 34:2353–2365.
- [Lacroute and Levoy, 1994] Lacroute, P. and Levoy, M. (1994). Fast volume rendering using a shearwarp factorization of the viewing transformation. In *21st annual conference on Computer graphics and interactive techniques*, pages 451–458.

- [LaRose et al., 2000] LaRose, D., Bayouth, J., and Kanade, T. (2000). Transgraph: Interactive intensitybased 2-D/3-D registration of X-ray and CT data. In *Proceedings of SPIE Medical Imaging: Image Processing*, pages 385–396.
- [Leach et al., 2005] 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., and 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). *The Lancet*, 365:1769 – 1778.
- [Lee et al., 2011] Lee, A., Rajagopal, V., Reynolds, H., Doyle, A., Nielsen, P., and Nash, M. (2011). Breast X-ray and MR image fusion using Finite Element Modeling. In *MICCAI Workshop on Breast Image Analysis*, pages 129–136.
- [Lee et al., 2010] Lee, A., Schnabel, J., Rajagopal, V., Nielsen, P., and Nash, M. (2010). Breast image registration by combining finite elements and free-form deformations. In *International Workshop on Digital Mammography*, pages 736–743.
- [Lee et al., 2009] Lee, J., Su, H., Cheng, P., Liou, M., Aston, J., Tsai, A., and Chen, C. (2009). MR Image Segmentation Using a Power Transformation Approach. *IEEE Transactions on Medical Imaging*, 28:894–905.
- [Lemieux et al., 1994] Lemieux, L., Jagoe, R., Fish, D., Kitchen, N., and Thomas, D. (1994). A patientto-computed-tomography image registration method based on digitally reconstructed radiographs. *Medical Physics*, 21:1749–1760.
- [Leroy et al., 2004] Leroy, A., Mozer, P., Payan, Y., and Troccaz, J. (2004). Rigid registration of freehand 3D ultrasound and CT-Scan kidney images. In *Proceedings of Medical Imaging Computing and Computer Assisted Intervention*, pages 837–844.
- [Li et al., 2009] Li, X., Dawant, B., Welch, E., Chakravarthy, A., Freehardt, D., Mayer, I., Kelley, M., Meszoely, I., Gore, J., and Yankeelov, T. (2009). A nonrigid registration algorithm for longitudinal breast MR images and the analysis of breast tumor response. *Magnetic Resonance Imaging*, 27:1258– 1270.
- [Lorensen and Cline, 1987] Lorensen, W. and Cline, H. (1987). Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 163–169.
- [Makarau et al., 2010] Makarau, A., Huisman, H., Mus, R., Zijp, M., and Karssemeijer, N. (2010). Breast mri intensity non-uniformity correction using mean-shift. In *Proceedings of SPIE Medical Imaging: Computer-Aided Diagnosis*, page 76242D.

- [Marias et al., 2005] Marias, K., Behrenbruch, C., Parbhoo, S., Seifalian, A., and Brady, M. (2005). A registration framework for the comparison of mammogram sequences. *IEEE Transations on Medical Imaging*, 24:782–790.
- [Markelj et al., 2012] Markelj, P., Tomazevic, D., Likara, B., and Pernusa, F. (2012). A review of 3D/2D registration methods for image-guided interventions. *Medical Image Analysis*, 16:642–661.
- [Markelj et al., 2008] Markelj, P., Tomazevic, D., Pernus, F., and Likar, B. (2008). Robust Gradient-Based 3-D/2-D Registration of CT and MR to X-Ray Images. *IEEE Transactions on Medical Imaging*, 27:1704–1714.
- [Marti et al., 2002] Marti, R., Zwiggelaar, R., and Rubin, C. (2002). Automatic point correspondence and registration based on linear structures. *International Journal on Pattern Recognition and Artificial Intelligence*, 16:331–340.
- [Marti et al., 2004] Marti, R., Zwiggelaar, R., Rubin, C., and Denton, E. (2004). 2D-3D correspondence in mammography. *Cybernetics and Systems*, 35:85–105.
- [Masutani et al., 1997] Masutani, Y., Dohi, T., Yamane, F., Iseki, H., and Takakura, K. (1997). Interactive virtualized display system for intravascular neurosurgery. In *CVRMed-MRCAS*, pages 427–435.
- [McRobbie et al., 2003] McRobbie, D., Moore, E., Graves, M., and Prince, M. (2003). *MRI from picture to proton*. Cambridge University Press.
- [Melbourne et al., 2011] Melbourne, A., Hipwell, J., Modat, M., Mertzanidou, T., Huisman, H., Ourselin, S., and Hawkes, D. (2011). The effect of motion correction on pharmacokinetic parameter estimation in dynamic-contrast-enhanced MRI. *Physics in Medicine and Biology*, 56:7693–7708.
- [Mendez et al., 1998] Mendez, A., Tahoces, P., Lado, M., Souto, M., and Vidal, J. (1998). Computeraided diagnosis: automatic detection of malignant masses in digitized mammograms. *Medical Physics*, 25:957–964.
- [Mertzanidou et al., 2010a] Mertzanidou, T., Hipwell, J., Cardoso, M., Tanner, C., Ourselin, S., and Hawkes, D. (2010a). X-ray mammography - MRI registration using a volume-preserving affine transformation and an EM-MRF for breast tissue classification. In *International Workshop on Digital Mammography*, pages 23–30.
- [Mertzanidou et al., 2012a] Mertzanidou, T., Hipwell, J., Cardoso, M., Zhang, X., Tanner, C., Ourselin, S., Bick, U., Huisman, H., Karssemeijer, N., and Hawkes, D. (2012a). MRI to X-ray mammography registration using a volume-preserving affine transformation. *Medical Image Analysis*, 16:966–975.
- [Mertzanidou et al., 2011] Mertzanidou, T., Hipwell, J., Han, L., Huisman, H., Karssemeijer, N., and 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*, pages 161–168.

- [Mertzanidou et al., 2012b] Mertzanidou, T., Hipwell, J., Han, L., Taylor, Z., Huisman, H., Bick, U., Karssemeijer, N., and Hawkes, D. (2012b). Intensity-based MRI to X-ray mammography registration with an integrated fast biomechanical transformation. In *International Workshop on Breast Imaging*, pages 48–55.
- [Mertzanidou et al., 2010b] Mertzanidou, T., Hipwell, J., Tanner, C., and Hawkes, D. (2010b). An intensity-based approach to X-ray mammography MRI registration. In *Proceedings of SPIE Medical Imaging: Image Processing*, pages 7623–106.
- [Meyer et al., 1997] Meyer, C., Boes, J., Kim, B., Bland, P., Zasadny, K., Kison, P., Koral, K., Frey, K., and Wahl, R. (1997). Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformations. *Medical Image Analysis*, 1:195–207.
- [Mitrovic et al., 2011] Mitrovic, U., Markelj, P., Likar, B., Milosevic, Z., and Pernus, F. (2011). Gradient-based 3D-2D registration of cerebral angiograms. In *Proceedings of SPIE Medical Imaging: Image Processing*, pages 79621P–79621P–8.
- [Mohamed et al., 2002] Mohamed, A., Davatzikos, C., and Taylor, R. (2002). A combined statistical and biomechanical model for estimation of intra-operative prostate deformation. In *Proceedings of Medical Imaging Computing and Computer Assisted Intervention*, pages 452–460.
- [Murphy, 1997] Murphy, M. (1997). An automatic six-degree-of-freedom image registration algorithm for image-guided frameless stereotaxic radiosurgery. *Medical Physics*, 24:857–866.
- [Murthy et al., 2000] Murthy, K., Aznar, M., Bergman, A., Thompson, C., Robar, J., Lisbona, R., Loutfi, A., and Gagnon, J. (2000). Positron emission mammographic instrument: Initial results. *Radiology*, 215:280–285.
- [NICE, 2006] NICE (October 2006). National Institute for Health and Clinical Excellence. Familial breast cancer: Quick reference guide.
- [Nie et al., 2008] Nie, K., Chen, J., Chan, S., Chau, M., Yu, H., Bahri, S., Tseng, T., Nalcioglu, O., and Su, M. (2008). Development of a quantitative method for analysis of breast density based on three-dimensional breast MRI. *Medical Physics*, 35:5253–5262.
- [Oza and Boyd, 1993] Oza, A. and Boyd, N. (1993). Mammographic parenchymal patterns: a marker of breast cancer risk. *Epidemiologic Reviews*, 15:196–208.
- [Paquerault et al., 2002] Paquerault, S., Petrick, N., Chan, H., Sahiner, B., and Helvie, M. (2002). Improvement of computerized mass detection on mammograms: Fusion of two-view information. *Medical Physics*, 29:238–247.
- [Pathmanathan et al., 2004] Pathmanathan, P., Gavaghan, D., Whiteley, J., Brady, M., Nash, M., Nielsen, P., and 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*, pages 217–224.

- [Penney, 2000] Penney, G. (2000). Registration of Tomographic Images to X-ray Projections for Use in Image Guided Interventions. PhD thesis, King's College London.
- [Penney et al., 2004] Penney, G., Blackall, J., Hamady, M., Sabharwal, T., Adam, A., and Hawkes, D. (2004). Registration of freehand 3D ultrasound and magnetic resonance liver images. *Medical Image Analysis*, 8:81–91.
- [Penney et al., 1998] Penney, G., Weese, J., Little, J., Desmedt, P., Hill, D., and Hawkes, D. (1998). A Comparison of Similarity Measures for Use in 2-D–3-D Medical Image Registration. *IEEE Transactions on Medical Imaging*, 17:586–595.
- [Pereira et al., 2010] Pereira, S., Hipwell, J., Mccormack, V., Tanner, C., Moss, S., Wilkinson, L., Khoo, L., Pagliari, C., Skippage, P., Kliger, C., Hawkes, D., and Silva, I. (2010). Automated registration of diagnostic to prediagnostic X-ray mammograms: Evaluation and comparison to radiologists' accuracy. *Medical Physics*, 37:4530–4539.
- [Periaswamy and Farid, 2003] Periaswamy, S. and Farid, H. (2003). Elastic registration in the presence of intensity variations. *IEEE Transactions on Medical Imaging*, 22:865–874.
- [Piron et al., 2003] Piron, C., Causer, P., Jong, R., Shumak, R., and Plewes, D. (2003). A hybrid breast biopsy system combining ultrasound and mri. *IEEE Transactions on Medical Imaging*, 22:1100 – 1110.
- [Porter et al., 2001] Porter, B., Rubens, D., Strang, J., Smith, J., Totterman, S., and Parker, K. (2001). Three-dimensional registration and fusion of ultrasound and MRI using major vessels as fiducial markers. *IEEE Transactions on Medical Imaging*, 20:354–359.
- [Qiu et al., 2008] Qiu, Y., Sun, X., Manohar, V., and Goldgof, D. (2008). Towards Registration of Temporal Mammograms by Finite Element Simulation of MR Breast Volumes. In *Proceedings of SPIE Medical Imaging*, pages 69182F–69182F–7.
- [Rajagopal et al., 2008] Rajagopal, V., Nash, M., Highnam, R., and Nielsen, P. (2008). The breast biomechanics reference state for multi-modal image analysis. In *International Worshop on Digital Mammography*, pages 385–392.
- [Ramsay et al., 2005] Ramsay, D., Kent, J., Hartmann, R., and Hartman, P. (2005). Anatomy of the lactating human breast redefined with ultrasound imaging. *Journal of Anatomy*, 206:525–534.
- [Richard et al., 2006] Richard, F., Bakic, P., and Maidment, A. (2006). Mammogram Registration: A Phantom-Based Evaluation of Compressed Breast Thickness Variation Effects. *IEEE Transactions* on Medical Imaging, 25:188–197.

- [Richard and Cohen, 2003] Richard, F. and Cohen, L. (2003). A New Image Registration technique with free boundary constraints: application to mammography. *Computer Vision and Image Understanding*, 89:166–196.
- [Robinson and Scrimger, 1991] Robinson, D. and Scrimger, J. (1991). Monoenergetic approximation of a polyenergetic beam A theoretical approach. *British Journal of Radiology*, 64(761):452–454.
- [Rohlfing et al., 2003] Rohlfing, T., Maurer, C., Bluemke, D., and Jacobs, M. (2003). Volumepreserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Transactions on Medical Imaging*, 22:730–741.
- [Rohlfing et al., 2005] Rohlfing, T., Russakoff, D., Denzler, J., Mori, K., and Maurer, C. (2005). Progressive attenuation fields: Fast 2d-3d image registration without precomputation. *Medical Physics*, 32:2870–2880.
- [Rueckert et al., 2003] Rueckert, D., Frangi, A., and Schnabel, J. (2003). Automatic Construction of 3-D Statistical Deformation Models of the Brain Using Nonrigid Registration. *IEEE Transactions on Medical Imaging*, 22:1014–1025.
- [Rueckert et al., 1999] Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., and Hawkes, D. (1999). Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18:712–721.
- [Ruiter, 2003] Ruiter, N. (2003). *Registration of X-ray mammograms and MR-volumes of the Female Breast based on Simulated Mammographic deformation*. PhD thesis, University of Mannheim.
- [Ruiter et al., 2006] Ruiter, N., Stotzka, R., Muller, T., Gemmeke, H., Reichenbach, J., and Kaiser, W. (2006). Model-Based registration of X-ray Mammograms and MR images of the female breast. *IEEE Transactions on Nuclear Science*, 53:204–211.
- [Russakoff et al., 2003] Russakoff, D., Rohlfing, T., and Maurer, C. (2003). Fast intensity-based 2D-3D fluoroscopy-to-CT registration of clinical data using light fields. In *IEEE International Conference* on Computer Vision, pages 416–422.
- [Russell and Norvig, 2003] Russell, S. and Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. Prentice Hall.
- [Sallam and Bowyer, 1994] Sallam, M. and Bowyer, K. (1994). Registering time sequences of mammograms using a two-dimensional image unwarping technique. In *International Workshop on Digital Mammography*, pages 121–130.
- [Samani et al., 2001] Samani, A., Bishop, J., Yaffe, M., and Plewes, D. (2001). Biomechanical 3-D Finite Element Modeling of the Human Breast Using MRI Data. *IEEE Transactions on Medical Imaging*, 20:271–279.

- [Samani et al., 2007] Samani, A., Zubovits, J., and Plewes, D. (2007). Elastic moduli of normal and pathological human breast tissues: an inversion-technique-based investigation of 169 samples. *Physics in Medicine and Biology*, 52:1565–1576.
- [Schnabel et al., 2003] Schnabel, J., Tanner, C., Castellano-Smith, A., Degenhard, A., Leach, M., Hose, D., Hill, D., and Hawkes, D. (2003). Validation of nonrigid image registration using finite-element methods: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 22:238–247.
- [Sinkus et al., 2000] Sinkus, R., Lorenzen, J., Schrader, D., Lorenzen, M., Dargatz, M., and Holz, D. (2000). High-resolution tensor mr elastography for breast tumour detection. *Physics in Medicine and Biology*, 45:1649–1664.
- [Snoeren and Karssemeijer, 2007] Snoeren, P. and Karssemeijer, N. (2007). Gray-scale and geometric registration of full-field digital and film-screen mammograms. *Medical Image Analysis*, 11:146–156.
- [Spoerk et al., 2007] Spoerk, J., Bergmann, H., Wanschitz, F., Dong, S., and Birkfellner, W. (2007). Fast drr splat rendering using common consumer graphics hardware. *Medical Physics*, 34:4302–4308.
- [Studholme et al., 1996] Studholme, C., Hill, D., and Hawkes, D. (1996). Automated 3-D registration of MR and CT images of the head. *Medical Image Analysis*, 1:163–175.
- [Studholme et al., 1999] Studholme, C., Hill, D., and Hawkes, D. (1999). An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32:71–86.
- [Styner et al., 2000] Styner, M., Brechbuhler, C., Szekely, G., and Gerig, G. (2000). Parametric Estimate of Intensity Inhomogeneities Applied to MRI. *IEEE Transactions on Medical Imaging*, 19:153–165.
- [Suri and Rangayyan, 2006] Suri, J. and Rangayyan, R. (2006). *Recent Advances in Breast Imaging, Mammography, and Computer-aided Diagnosis of Breast Cancer.* SPIE Society of Photo-Optical Instrumentation Engineers.
- [Tanner et al., 2008] Tanner, C., Hipwell, J., and Hawkes, D. (2008). Statistical Deformation Models of Breast Compressions from Biomechanical Simulations. In *International Workshop on Digital Mammography*, pages 426–432.
- [Tanner et al., 2009] Tanner, C., Hipwell, J., and Hawkes, D. (2009). Using Statistical Deformation Models for the Registration of Multimodal Breast Images. In *Proceedings of SPIE Medical Imaging*.
- [Tanner et al., 2006] Tanner, C., Schnabel, J., Hill, D., and Hawkes, D. (2006). Factors influencing the accuracy of biomechanical breast models. *Medical Physics*, 33:1758–1769.
- [Tanner et al., 2007] Tanner, C., Schnabel, J., Hill, D., and Hawkes, D. (2007). Quantitative evaluation of free-form deformation registration for dynamic contrast-enhanced MR mammography. *Medical Physics*, 34:1221–1233.

- [Tanner et al., 2011] Tanner, C., White, M., Guarino, S., Hall-Craggs, M., Douek, M., and Hawkes, D. (2011). Large breast compressions – Observations and evaluation of simulations. *Medical Physics*, 38:682–690.
- [Taylor et al., 2009] Taylor, Z., Comas, O., Cheng, M., Passenger, J., Hawkes, D., Atkinson, D., and Ourselin, S. (2009). On modelling of anisotropic viscoelasticity for soft tissue simulation: Numerical solution and GPU execution. *Medical Image Analysis*, 13:234–244.
- [Thirion, 1998] Thirion, J. (1998). Image matching as a diffusion process: An analogy with Maxwell's demons. *Medical Image Analysis*, 2:243–260.
- [Tomazevic et al., 2003] Tomazevic, D., Likar, B., T.Slivnik, and Pernus, F. (2003). 3-D/2-D registration of CT and MR to X-ray images. *IEEE Transactions on Medical Imaging*, 22:1407–1416.
- [van der Bom et al., 2011] van der Bom, I., Klein, S., Staring, M., Homan, R., Bartels, L., and Pluim, J. (2011). Evaluation of optimization methods for intensity-based 2D-3D registration in X-ray guided interventions. In *Proceedings of SPIE Medical Imaging: Image Processing*, pages 796223–1 – 796223– 15.
- [van Engeland et al., 2003] van Engeland, S., Snoeren, P., Hendriks, J., and Karssemeijer, N. (2003). A comparison of methods for mammogram registration. *Transactions on Medical Imaging*, 22:1436– 1444.
- [van Engeland et al., 2006] van Engeland, S., Timp, S., and Karssemeijer, N. (2006). Finding corresponding regions of interest in mediolateral oblique and craniocaudal mammographic views. *Medical Physics*, 33:3203–3212.
- [van Leemput et al., 1999] van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P. (1999). Automated Model-Based Tissue Classification of MR Images of the Brain. *IEEE Transactions on Medical Imaging*, 18:897–908.
- [van Schie et al., 2011] van Schie, G., Tanner, C., Snoeren, P., Samulski, M., Leifland, K., Wallis, M., and Karssemeijer, N. (2011). Correlating locations in ipsilateral breast tomosynthesis views using an analytical hemispherical compression model. *Physics in Medicine and Biology*, 56:4715–4730.
- [Vercauteren et al., 2008] Vercauteren, T., Pennec, X., Perchant, A., and Ayache, N. (2008). Symmetric log-domain diffeomorphic registration: A demons-based approach. In *Medical Image Computing and Computer-Assisted Intervention*, pages 754–761.
- [Vercauteren et al., 2009] Vercauteren, T., Pennec, X., Perchant, A., and Ayache, N. (2009). Diffeomorphic demons: Efficient non-parametric image registration. *Neuroimage: Special Issue on Mathematics in Brain Imaging*, 45:S61–S72.
- [Viola and Wells, 1997] Viola, P. and Wells, W. (1997). Alignment by maximization of Mutual Information. *International Journal of Computer Vision*, 24:137–154.

- [Voirin et al., 2002] Voirin, D., Payan, Y., Amavizca, M., Letoublon, C., and Troccaz, J. (2002). Computer-aided hepatic tumour ablation: requirements and preliminary results. *Comptes Rendus Biologies*, 325:309–319.
- [Vujovic and Brzakovic, 1997] Vujovic, N. and Brzakovic, D. (1997). Establishing the correspondence between control points in pairs of mammographic images. *IEEE Transactions on Medical Imaging*, 6:1388–1399.
- [Washington and Miga, 2004] Washington, C. and Miga, M. (2004). Modality independent elastography (mie): A new approach to elasticity imaging. *IEEE Transactions on Medical Imaging*, 23:1117–1128.
- [Weese et al., 1999] Weese, J., Göcke, R., Penney, G., Desmedt, P., Buzug, T., and Schumann, H. (1999). Fast voxel-based 2D/3D registration algorithm using a volume rendering method based on the shear-warp factorization. In *Proceedings of SPIE Medical Imaging: Image Processing*, pages 802–810.
- [Weese et al., 1997] Weese, J., Penney, G., Desmedt, P., Buzug, T., Hill, D., and Hawkes, D. (1997). Voxel-based 2-D/3-D registration of fluoroscopy images and CT scans for image-guided surgery. *IEEE Transactions on Medical Imaging*, 1:284–293.
- [Wei et al., 2004] Wei, J., Chan, H., Helvie, M., Roubidoux, M., Sahiner, B., Hadjiiski, L., Zhou, C., Chenevert, S. P. T., and Goodsitt, M. (2004). Correlation between mammographic density and volumetric fibroglandular tissue estimated on breast MR images. *Medical Physics*, 31:933–942.
- [Wein et al., 2008] Wein, W., Brunke, S., Khamene, A., Callstrom, M., and Navab, N. (2008). Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis*, 12:577–585.
- [Wells et al., 1996] Wells, W., Viola, P., Atsumi, H., Nakajima, S., and Kikinis, R. (1996). Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1:35–51.
- [Westover, 1990] Westover, L. (1990). Footprint evaluation for volume rendering. In 21st annual conference on Computer graphics and interactive techniques, pages 367–376.
- [Wirth et al., 2002] Wirth, M., Narthan, J., and Gray, D. (2002). A model for nonrigid mammogram registration using mutual information. In *International Workshop on Digital Mammography*, pages 243–245.
- [Wollny and Kruggel, 2002] Wollny, G. and Kruggel, F. (2002). Computional cost of nonrigid registration algorithms based on fluid dynamics. *IEEE Transactions on Medical Imaging*, 21:946–952.
- [Yin et al., 1991] Yin, F., Giger, M., Doi, K., Metz, C., Vyborny, C., and Schmidt, R. (1991). Computerized detection of masses in digital mammograms: Analysis of bilateral subtraction images. *Medical Physics*, 18:955–963.

- [Yin et al., 1994] Yin, F., Giger, M., Doi, K., Metz, C., Vyborny, C., and Schmidt, R. (1994). Computerized detection of masses in digital mammograms: Automated alignment of breast images and its effect on bilateral-subtraction technique. *Medical Physics*, 21:445–452.
- [Yuan et al., 2010] Yuan, Y., Giger, M., Li, H., Bhooshan, N., and Sennett, C. (2010). Multimodality Computer-Aided Breast Cancer Diagnosis with FFDM and DCE-MRI. Academic Radiology, 17:1158–1167.
- [Zheng et al., 2009] Zheng, B., Tan, J., Ganott, M., Chough, D., and Gur, D. (2009). Matching breast masses depicted on different views: A comparison of three methods. *Academic Radiology*, 16:1338– 1347.