#### UNIVERSITY OF CALGARY

Estimation of Three-Dimensional Breast Features from Standard Two View Mammograms

by

Charlotte Curtis

A THESIS

## SUBMITTED TO THE FACULTY OF GRADUATE STUDIES IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF MASTER OF SCIENCE

### BIOMEDICAL ENGINEERING GRADUATE PROGRAM

CALGARY, ALBERTA

December, 2010

© Charlotte Curtis 2010



The author of this thesis has granted the University of Calgary a non-exclusive license to reproduce and distribute copies of this thesis to users of the University of Calgary Archives.

Copyright remains with the author.

Theses and dissertations available in the University of Calgary Institutional Repository are solely for the purpose of private study and research. They may not be copied or reproduced, except as permitted by copyright laws, without written authority of the copyright owner. Any commercial use or re-publication is strictly prohibited.

The original Partial Copyright License attesting to these terms and signed by the author of this thesis may be found in the original print version of the thesis, held by the University of Calgary Archives.

Please contact the University of Calgary Archives for further information: E-mail: <u>uarc@ucalgary.ca</u> Telephone: (403) 220-7271 Website: <u>http://archives.ucalgary.ca</u>

### Abstract

Mammography is the current gold standard imaging modality for breast cancer diagnosis, providing detailed 2D projection images of the compressed breast. However, unprocessed mammograms are unable to provide 3D information due to significant distortion and a limited number of views.

This thesis presents a method to estimate the external shape and internal features of the breast from two 2D mammograms. The distortion resulting from compression is removed by registration with 2D projection images created from an MR reference image; this method is validated by a simulated model. The skin surface is reconstructed by fitting ellipses at twenty equally spaced coronal slices of the breast. Internal features visible on both views are reconstructed by orthogonal backprojection.

This algorithm was tested on four patient data sets, and shown to be successful for all eight breasts. Results and applications of this work to a novel microwave imaging technique (TSAR) are discussed.

### Acknowledgements

First and foremost, I would like to thank my supervisors, Dr. Elise Fear and Dr. Richard Frayne. They gave me guidance, encouragement, and inspiration, and provided motivation when my own was in short supply. I am also very grateful for the assistance of Dr. Bobbie Docktor, who cheerfully answered my endless questions about breast image acquisition and interpretation.

All of the staff and students in the TSAR, AEG, and MR groups have been invaluable as both friends and mentors. I am thankful to be a part of such groups, and it is because of them that I have grown to love the University of Calgary and my work here.

My parents, Bo and Cathy, and my brother, Sandy, have all been incredibly supportive of my persistence in remaining in school all this time. It was they who encouraged me to be inquisitive about the world, and it is they who continue to inspire me to exceed my own expectations.

I would like to thank Andrea Dukeshire for all her patience during the writing of this thesis, and for introducing me to the mountains. She has opened my eyes to what Alberta has to offer, and has been the best friend and roommate I could ask for.

A special thank-you goes to Ryan Harrison. Throughout this work he has been there for me, providing moral support in the hard times and sharing laughter in the good, and making sure that I was always well fed on gourmet home cooked meals.

Finally, I would like to acknowledge NSERC, iCORE, and the Alberta Government, who made this work possible through their generous financial support.

## Table of Contents

A	BSTRA	ACT	ii
Ac	CKNO	WLEDGEMENTS	iii
TA	ABLE	OF CONTENTS	iv
Lı	ST OF	F TABLES	vi
Lı	ST OF	FIGURES	vii
Lı	ST OF	F ACRONYMS V	iii
Lı	ST OF	F Symbols	ix
Gi	LOSSA	ARY	x
1	Inte 1.1	RODUCTION Motivation	<b>1</b> 1
	$1.2 \\ 1.3$	Objective	4 4
2	BAC	KGROUND AND LITERATURE REVIEW	7
	2.1	Digital Image Representation	7 7 13
	2.2	X-ray Mammography	15
	2.3	Magnetic Resonance Imaging	17
	2.4	Registration of Mammograms and MR Images	19
	2.5	Reconstruction of Mammograms in 3D	23
	2.6	Summary	24
3	MAN	MMOGRAM/MR IMAGE REGISTRATION	26
	3.1	Preprocessing Algorithm	26 28 29 31 32 33 34
	3.2	Registration Algorithm	35 36

		3.2.2 Mutual Information Registration
	3.3	Registration Evaluation
	3.4	Validation Using Simulated Images
	3.5	Results
	3.6	Discussion
4	MAI	MMOGRAM RECONSTRUCTION 63
	4.1	Image, Contour, and Landmark Alignment
	4.2	Skin Surface Estimation
	4.3	Internal Feature Reconstruction
	4.4	Reconstruction Evaluation
	4.5	Results
	4.6	Discussion
5	Con	NCLUSION 84
	5.1	Summary of Contributions
	5.2	Limitations $\ldots \ldots \ldots$
	5.3	Future Work and Potential Applications
		5.3.1 Statistical Deformation Model
		5.3.2 Applications to TSAR
		5.5.5 Surgical Planning
R	EFERI	ENCES 91
А	Sof	TWARE FRAMEWORK 97
В	Alg	GORITHMS 10
	B.1	Fuzzy Connected Object Extraction
	B.2	Hermite Interpolation
	B.3	The Elastic Body Spline Transformation
	B.4	B-Spline Interpolation $\ldots \ldots \ldots$

## List of Tables

$2.1 \\ 2.2$	Pixel and voxel sizes of mammographic and MR images	11 19
$3.1 \\ 3.2$	Modifications made to parameters to optimize result	$50\\50$
$4.1 \\ 4.2$	Maximum absolute errors in mm for each data set and direction Average absolute errors in mm for each data set and direction	75 75

# List of Figures

1.1	Overview of processing algorithm	6
$2.1 \\ 2.2$	Sampling a continuous one-dimensional signal	$\frac{8}{9}$
2.3	Illustration of image resolution vs. pixel resolution (400% scale) $\ldots$	11
2.4	Representation of a 3D image as a stack of slices	11
2.5	Quantization of a 1D continuous signal	13
2.6	Anatomical directions and world coordinates relative to left and right breasts	14
2.7	Anatomical planes of the human body	14
2.8	Illustration of mammographic image acquisition of the right breast	16
2.9	Patient positioning within an MR scanner	18
2.10	Registering a moving image (A1) to a fixed image (A2)	20
3.1	Overview of the preprocessing algorithm	27
3.2	Illustration of image rotation	28
3.3	Creation of an MR projection image	29
3.4	Segmentation of an image into background and breast regions	31
3.5	Ordering the set of points lying on the breast/air boundary	32
3.6	Constraint regions for landmark detection	34
3.7	Preprocessed images with contour spline in red, landmarks in green	35
3.8	Overview of the registration algorithm	36
3.9	MLO mammogram aligned to MRPI using landmarks	37
3.10	An overview of the iterative registration method	38
3.11	Mapping the moving image into the fixed image space	41
3.12	Simulated images deformed via finite element modelling	47
3.13	Simulated images prior to registration	47
3.14	Results from registration of simulated images	48
3.15	Registered simulated images following rotation of MR volume	49
3.16	Registration results for sample MLO data set (final $MI = 0.7264$ )	51
4.1	Overview of the reconstruction algorithm	64
4.2	Acquisition angles of MLO mammograms	65
4.3	Sparse wireframe of the 3D breast shape	66
4.4	Illustration of the eccentric anomaly $(t)$ of an ellipse	68
4.5	Sample skin surface reconstruction	70
4.6	Reconstruction of internal feature points	70
4.7	Skin outline comparison at two slice locations	73
4.8	Skin Reconstruction Errors vs. $x$ location	74
4.13	Comparison between feature reconstructed from mammograms and corre-	
	sponding feature seen on MR	80
51	Images of data set 100704B from different modulities	٩U
0.1	images of data set 100/041( from different modalities	30

# List of Acronyms

Notation	<b>Description</b>	<b>Page</b>
CC	cranial-caudal	1
CE	contrast enhanced	17
CHREB	Conjoint Health Research Ethics Board	46
CT	computed tomography	19
EBS	elastic body spline	34
EM	electromagnetic	2
FE	finite element	43
FOV	field of view	24
L-BFGS-B	limited-memory Broyden Fletcher Goldfarb Shanno minimization algorithm with simple bounds	40
MI	mutual information	39
MLO	medial-lateral oblique	1
MR	magnetic resonance	1
MRPI	magnetic resonance projection image	24
PDE	partial differential equation	100
PPI	pixels per inch	8
PSF	point spread function	8
RF	radio frequency	15
SPECT	single photon emission computed tomography	19
TSAR	tissue sensing adaptive radar	2

# List of Symbols

Notation ${}^{1}H$	<b>Description</b> Hydrogen-1	Page 1
$\gamma$	angle of the MLO mammogram plane relative to the upright body (where vertical = $0^{\circ}$ )	60
θ	angle of the X-ray beam vector during MLO mammogram acquisition (where vertical = $0^{\circ}$ )	26
$ \begin{array}{l} f(x,y,x) \\ f(i,j,k) \\ f_s \end{array} $	continuous image discrete image sampling frequency of a discrete signal	5 5 6
$T_s$	sampling period of a discrete signal	6
(x,y,z)	the Cartesian directions of the world coordinate system	5
(i, j, k)	the discrete integer indices of an image correspond- ing to the Cartesian directions $(x, y, z)$	5

# Glossary

Notation	Description	Page
8-connected	the eight pixels that touch a centre pixel at all four edges and corners	29
affine transform	a linear coordinate transform followed by transla- tion that preserves colinearity of points and ratios of distances between points	101
aliasing	a type of signal corruption resulting from under- sampling where high frequency signals appear to be low frequency	6
anisotropic	the property of having different values when mea- sured in different directions	8
axilla	the armpit region of the body	31
bit	the smallest unit of information in digital computing capable of two states $(0 \text{ and } 1)$	10
complex permittivity	describes the ability of a material to store (imagi- nary part) and dissipate (real part) EM field energy	2
conductivity	describes how electric current is transmitted through a material	2
continuous	a range of quantities on a continuum (without breaks between values)	5
discrete	a range of quantities constrained to a finite se- quence (with breaks between values)	5
gold standard	definitive and benchmark diagnostic test	1
island	a small region of one pixel intensity surrounded by a large region of another intensity	28
landmark	a clearly identifiable point within an image	24
microcalcifications	small crystals of calcium that develop in breast ducts and can indicate malignancy	14
pixel pixel-wise	the smallest discrete component of a 2D image performing an operation on each individual pixel of an image	8 32

Notation	Description	Page
quantization	converting continuous values to digital values	5
registration	transforming an image into the coordinate system of another such that corresponding objects within the images are aligned	3
sampling	taking a series of values from a continuous signal to form a discrete signal	5
segmentation	a region of an image containing a single feature or (verb) the creation of such a region	27
signal spline	some quantity that changes over space or time a set of polynomials that are smoothly connected at certain data points	5 30
threshold	creating a binary image by assigning a single in- tensity value to pixels above a given level and a different value to those below or (noun) the criti- cal intensity value	28
voxel	a volume element of a 3D image analogous to a 2D pixel	8

### Chapter 1

## Introduction

#### 1.1 Motivation

Breast cancer is the most common cancer affecting women, with an estimated 22,700 new cases occurring each year in Canada alone [1]. Although overall mortality rates are declining, breast cancer continues to be the second highest cause of cancer death among women [1]. As treatments are more successful if cancer is detected in its early stages, accurate diagnosis and monitoring of disease progression are crucial to patient survival and well-being.

Various imaging techniques are used for screening, detection, and diagnosis of breast cancer, with the current gold standard or definitive test being X-ray mammography [2]. Mammography is a 2D imaging technique that involves compressing the breast and obtaining X-ray images in each of the cranial-caudal (CC) and medial-lateral oblique (MLO) directions [3]. However, mammography has been shown to have low sensitivity and specificity among pre-menopausal women and women with dense breasts [4]. Mammography is also an uncomfortable procedure, and the use of ionizing X-rays carries health risks [3].

As a result of these limitations of mammography, other imaging modalities such as magnetic resonance (MR) imaging and ultrasound are used to assist in the diagnosis of symptomatic patients [4, 5, 6]. MRI provides a 3D representation of the breast based on the response of <sup>1</sup>H protons in a changing magnetic field, whereas ultrasound creates an image based on the reflections from sound waves in a localized area [4]. Since different modalities rely on different properties of the tissue to generate an image, examining mul-

 $\mathbf{2}$ 

tiple modalities can provide diagnostic information that might be missed if only a single imaging technique were used [4]. Because of the advantages of combining information from different sources, there is considerable interest in developing new imaging methods.

One alternative approach, microwave imaging, examines the electromagnetic (EM) properties (complex permittivity and conductivity) of breast tissues. This technique was first proposed by England and Sharples in 1949, who found significant differences in EM properties between breast fat and carcinoma regions [7]. Further work supported these findings, indicating that microwave imaging has strong potential for detecting tumours in fatty regions of the breast [8]. However, recent studies have shown that only a 10% difference exists between normal glandular tissue and cancerous regions, suggesting that microwave imaging techniques require further refinement to detect tumours in non-fatty regions [9].

Two methods of obtaining an image from microwave data have been proposed: tomographic and radar. Tomographic imaging aims to reconstruct a complete map of the EM properties of the breast by recording the transmitted waves, whereas radar based techniques focus the reflected waves to detect and localize tumours [3, 10]. Using very low power microwave energy, both methods show promise as safe, low-cost, and pain-free methods of detecting tumours at early stages of development [3].

As a technology currently in its infancy, microwave breast imaging has both exciting potential and significant challenges. One such challenge is the ill-posed and non-linear reconstruction task required to form a tomographic image from the acquired electromagnetic data [11]. In radar based techniques, such as the tissue sensing adaptive radar (TSAR) system developed at the University of Calgary, antenna placement and scan pattern are important factors in image quality and must be modified on a patient-specific basis [12]. As TSAR has matured and moved into the early stages of clinical testing, the need for *a priori* structural information for each patient has been identified. A similar need has been identified to address the ill-posedness of tomographic reconstruction [13, 14]; however, the remainder of this work will focus on applications to TSAR.

Since current breast cancer care protocols specify the acquisition of mammograms in almost all cases, mammographic images have the potential to be a convenient source of *a priori* information [2]. Furthermore, the ability to directly compare the results of microwave imaging methods to the current gold standard would assist with the validation of these new diagnostic techniques.

The challenge in using mammograms as a source of *a priori* information or validation lies in the problem of obtaining 3D data from only two 2D images, both of which are independently distorted due to the compression required. In this work, MR images are used to assist with the undistortion of mammograms through the use of image registration.

The two mammographic views are oriented at approximately 45° to each other, and do not provide enough information to completely reconstruct the full 3D volume of the breast such as the image produced by MR imaging. However, several assumptions about the breast shape can be made to obtain a skin surface estimate, which would benefit TSAR imaging by allowing for a scan pattern to be determined on a patient-specific basis as well as providing information essential to creating an image from microwave signals.

In addition to the skin surface estimate, lesions that are detectable on both mammographic views can be reconstructed in 3D space. Clinical evaluation of TSAR images would benefit from prior knowledge of expected lesion locations, and the ability to determine such regions of interest from mammography would provide valuable multi-modal diagnostic power. For this application, high precision of lesion location is not required. In addition, current TSAR images are capable of resolving objects no less than 1 cm apart, which is considerably larger scale than mammography. As a result, this project aims to detect a region of interest with quadrant-level accuracy. In addition to TSAR applications, the methods described in this thesis have potential roles in several other clinical and research problems. The registered images have value in themselves, as they provide the ability to directly compare regions of interest on MR and mammogram images. The deformation resulting from registration could also provide information on the biomechanical properties of the tissues. Finally, reconstruction of mammographic features in 3D could assist in locating tumours during surgical procedures.

#### 1.2 Objective

The main objective of this project is to obtain a 3D estimation of the external shape and internal feature locations of the breast from two 2D X-ray mammograms. The work required to achieve this goal can be separated into two specific aims:

- 1. Develop a robust and validated algorithm to remove distortion from mammograms through registration with MR images.
- 2. Develop a method to reconstruct and visualize the shape of the breast and the location of relevant features in 3D with quadrant-level accuracy.

Furthermore, an emphasis on automation, ease of use, and speed will be placed on each of the algorithms developed in these specific aims.

#### 1.3 Thesis Outline

Overall, this thesis describes a method to generate a 3D estimate of the skin surface and the 3D location of visible lesions from mammograms as illustrated in the flow chart of figure 1.1. The mammographic image acquisition process requires compressing the breast, resulting in a distorted image. To estimate the original 3D shape of the breast, a reference shape is required; in this work, an MR image is used. A projection image through the MR volume and its corresponding mammogram are registered in 2D, producing as intermediate results two undistorted mammograms (CC and MLO views). These images are then aligned in 3D space according to their acquisition angles, and edge splines defining the skin/air boundary are used as inputs to an ellipse-based skin surface estimation algorithm. Finally, any features that are identifiable on both of the mammograms are located in 3D space by orthogonal backprojection.

As a result of the processing methods described in this thesis, two 3D objects are obtained: a skin surface estimate and a region of interest. The skin surface can then be used to determine a TSAR scan pattern and assist with image formation, while the 3D region of interest can be used to target a specific area and to evaluate suspicious regions with multiple modalities.

With the exception of the finite element model described in §3.4, all of the methods described in this thesis were developed by the author.

Chapter 2 introduces essential background on breast imaging modalities, the representation of images as continuous and discrete signals, and discusses previous work on breast MR/mammographic image registration and 3D reconstruction of mammograms.

Chapter 3 describes the methods used to register X-ray mammograms to MR projection images. Results are presented showing successful registration for both MLO and CC views and on simulated and acquired data.

Chapter 4 describes the methods used to reconstruct the external surface and internal features of X-ray mammograms in 3D space and presents examples of visualization results.

Chapter 5 summarizes and concludes this work, detailing contributions and future applications with a focus on TSAR imaging.



Figure 1.1: Overview of processing algorithm

## Chapter 2

## Background and Literature Review

The field of medical imaging is vast and diverse, with image formation resulting from interactions ranging from X-rays being blocked by tissues to radioactive decay. This chapter begins with an introduction to image representation as discrete and continuous signals as well as the breast imaging modalities used in this work. Following this essential background information, a review of the relevant literature on 3D mammogram reconstruction is provided.

#### 2.1 Digital Image Representation

Digital images, while creating an intuitive picture that can be displayed on a computer screen, are also represented mathematically as signals. Signals are functions of independent variables that change over a range of values; this range may be either continuous or discrete.

In this work, signals representing images are either functions of two spatial variables (2D) or of three spatial variables (3D). Processes occurring in the physical world, such as interactions between X-ray beams and the body, are modelled as continuous functions of space such as f(x, y, x), where (x, y, z) are coordinates in space [15]. However, computers are only able to store, process, and display discrete signals containing a finite number of values such as f(i, j, k), where (i, j, k) are integer indices.

#### 2.1.1 Continuous Signal Discretization

Conversion from a continuous to a discrete signal requires sampling and quantization. Sampling is a broad term covering a range of methods to take representative values from the continuous signal and discard the rest, while quantization refers to converting the value to a number that can be stored digitally [15].

#### a) Sampling

Regular sampling records the value of the continuous signal at regular intervals; for a 1D signal, the sampling period  $T_s$  or  $\Delta t$  is the length of that interval in seconds (minutes, hours, etc.), and the number of samples collected in one second is the sampling rate or frequency  $f_s = \frac{1}{T}$ . This is illustrated in figure 2.1, where the discrete signal (shown as blue dots) is a sampled version of a continuous sinusoid (black curve).



Figure 2.1: Sampling a continuous one-dimensional signal

It is desirable to collect the smallest number of samples possible that can still accurately be reconstructed to form the continuous signal, as lower frequency sampling rates require less complex equipment and reduce storage and processing requirements [15]. However, if a signal is not sampled at a high enough rate, it can be corrupted by aliasing, a phenomenon where the digital signal resembles a lower frequency continuous signal [15]. An example of undersampling a high-frequency sinusoid is shown in figure 2.2, with the original continuous signal shown in black, the sampled signal marked as blue dots, and the reconstructed aliased signal shown in blue.

The minimum sampling rate required to avoid aliasing is twice the frequency of the highest frequency component in the continuous signal; this threshold is known as the Nyquist rate [15]. For most natural signals, the exact frequency distribution is not known, and sampling rates higher than Nyquist are used to ensure high frequency components are accurately captured.



Figure 2.2: Undersampling of a continuous signal results in aliasing

For higher dimensional signals (2D and 3D images), the theory described above is simply extended to the appropriate number of dimensions, with variations in space replacing variations in time. In the case of a 2D image, a regular grid (the digital sensor) composed of equally spaced rectangles is used to capture the continuous signal. The values at the centres of each  $\Delta x \times \Delta y$  rectangle are assigned to a discrete value f(i, j)[15]. $\Delta x$  and  $\Delta y$  are analogous to the sampling period  $T_s$  of the 1D signal in figure 2.1, and they may be different for each dimension, resulting in anisotropic pixels. The corresponding inverse of this value, analogous to the sampling frequency  $f_s$  of a 1D signal, is pixel resolution, measured in pixels per inch (PPI).

The spatial resolution of a 2D signal is measured in line pairs per mm (lp/mm) and describes the ability of the imaging system to distinguish differences between two closely adjacent features [15]. The spatial resolution is affected by all the components of the imaging system, such as the scattering of photons through the volume of interest, and decreased spatial resolution manifests as increased "blurriness" in the final image.

The equation describing the amount of blur introduced by an imaging system is the point spread function (PSF) [15]. The PSF is the image that is obtained from imaging a single point (mathematically, the ideal Dirac  $\delta$  function). The width of the blur or "spread" of the resulting image determines the spatial frequency and the Nyquist rate required for accurate discretization [15]. Figure 2.3 illustrates the difference between image resolution and pixel resolution. Both figures 2.3c and 2.3d exhibit blurring and are unable to resolve the fine details at 5 lp/mm, but only figure 2.3c is undersampled and exhibiting aliasing.

3D images are represented as a stack of 2D image slices through the cross section of the volume of interest as illustrated in figure 2.4. The spacing between these cross sections is referred to as the slice thickness and becomes the third sampling period $\Delta z$ . The $\Delta x \times \Delta y \times \Delta z$  rectangle is known as a voxel, short for "volume element" [15]. 3D image signals can be processed either as a volume f(i, j, k), or as a series of 2D images, f(i, j) at locations  $k = [0...k_n]$ .

As the images in this work are clinical images acquired using standard breast imaging protocols, it is assumed that the pixel/voxel size is sufficient for the purposes of each modality and system resolution. The pixel and voxel sizes of the images used in this work are summarized in table 2.1.







Figure 2.4: Representation of a 3D image as a stack of slices

	Mammography	MRI
$\Delta x = \Delta y (\mathbf{mm})$	0.094 - 0.07	0.43 - 0.39
Slice thickness (mm)	N/A	1.12

Table 2.1: Pixel and voxel sizes of mammographic and MR images

#### b) Quantization

Like sampling, quantization is a necessary step in converting from continuous to discrete signals that may potentially introduce errors. While a continuous signal may have an instantaneous value consisting of any real number to an infinite level of precision, digital signals are limited to values that have a finite range and precision [16].

The fundamental storage unit of digital information is the bit, defined as being either "on" (1) or "off" (0) [16]. The data types used to store digital information are composed of several bits, and both the number of bits and the representation of the data define the range of values and precision of the data. The smallest data type is a boolean, occupying a single bit; images composed of this data type would be completely binary (black and white) [15].

Numeric data types can be divided into two major categories: integers (both signed and unsigned) and floating point numbers [16]. Integer types can only contain whole numbers, and their range is dependent on the number of bits allocated and the sign. For example, an 8-bit integer can be one of  $2^8 = 256$  values; if it is unsigned, the values range from 0 to 255; if signed, the values range from -128 to 127.

Floating point numbers, such as  $6.02214 \times 10^{23}$ , use three integers to represent fractions in scientific notation: the sign (in this case, +), the significand or fraction (602214), and the exponent (+23) [17]. This allows for a much broader range of numbers, but only a fixed precision. In this example case of a 32-bit float, the data type is capable of describing a number up to seven significant digits.

The digital images used in this work use 16-bit unsigned integers to store data; thus, a given pixel can have an integer value in the range [0, 65 535], typically scaled so that the lowest intensity signal of the image is 0 and the highest is 65 535. As only integer values are used, quantization errors could result from subtle changes in the continuous signal as illustrated in figure 2.5. However, sub-integer changes in intensity are very small relative

to the overall range and are unlikely to be physiologically relevant.



Figure 2.5: Quantization of a 1D continuous signal

#### 2.1.2 Coordinate Systems

Several different coordinate systems and terms describing anatomical directions are used throughout this thesis. Cartesian coordinate systems are used to describe mathematical methods, whereas anatomical terms are used to discuss results.

The world coordinate system (x, y, z) has units of millimetres and provides a reference for the physical size of the image. The origin (0, 0, 0) of this coordinate system is defined at a constant location on the imaging unit, whereas the origin of the resultant image is defined at the smallest (x, y, z) coordinate. This image origin marks the voxel index [0, 0, 0] or pixel index [0, 0]; as a result, increasing indices correspond to an increase in physical dimension.

Figure 2.6 illustrates the orientation of the world coordinate system and direction of positive rotation relative to both the right and left breasts. Anatomical directions are

also indicated.



Figure 2.6: Anatomical directions and world coordinates relative to left and right breasts with MLO image planes indicated by dashed lines

As the world coordinate system is defined relative to the imaging unit, 2D mammography results in coordinates in the (x, y) plane only, as it is the imaging unit that is rotated to obtain the different views.

The anatomical planes of the body are referenced with respect to a person standing in neutral posture. These planes are used to describe cross-sectional slices through the body (shown in figure 2.7).



Figure 2.7: Anatomical planes of the human body

#### 2.2 X-ray Mammography

Mammography, the current gold standard breast cancer diagnostic and screening method, is an X-ray based imaging modality. X-rays are high energy photons generated in X-ray tubes by bombarding a heavy metal anode with a stream of electrons [16]. While X-rays are emitted in all directions, the emitter is shielded to allow only a directional beam to exit. The X-ray beam is aimed through the volume of interest towards an X-ray detector and as it passes through various different tissues the beam energy is attenuated due to both absorption and scattering. The amount of attenuation is dependent on the atomic number of the material, the density and thickness of the material, and the photon energy [16].

Photons that are not significantly attenuated increase the exposure of the detector, resulting in darker regions [16]. Tissues that largely attenuate the photons appear as bright regions; in this manner, the X-ray image that is formed represents a "shadow" of the internal tissues of the volume of interest. In other words, image contrast is a result of differential attenuation by the various tissues undergoing examination [16].

If the energy of a given photon is sufficiently high, it has the potential to ionize an atom or cause an electron to break free from its orbit [16]. This can damage DNA and cause significant biological effects such as cancer. Hydrogen, the easiest atom to ionize, has the potential to be ionized at energies equal to or greater than 13.6 eV, and as diagnostic X-rays are in the 20-200 keV range, all X-ray imaging is potentially dangerous [15].

Lower energies are more readily absorbed by tissues, increasing the chance of damage due to ionization. Therefore, appropriate X-ray energies must be carefully chosen to optimize image contrast while minimizing radiation dose [16].

Most applications of X-ray imaging involve contrast between air, bone, and soft tissue,

as these have very different attenuating properties and result in high contrast images [16]. Mammography, however, seeks to discriminate between different types of soft tissue. To achieve this, low energy X-rays (20-30 keV) are used, as they are more easily attenuated by soft tissues.

As previously mentioned, lower energy X-rays result in higher radiation dose to the patient. To reduce this effect, the breast is compressed between two parallel plastic plates up to 50% of its original thickness [16]. This also improves image quality by reducing scatter of the photons and placing the detector closer to the X-ray tube, resulting in a very high resolution image capable of detecting small features such as microcalcifications [16].

In Canada, two views of each breast are acquired during mammographic imaging: one with the X-ray beam parallel to the cranial-caudal (CC) axis, and one with the beam at an oblique angle (medial-lateral oblique or MLO) running parallel to the pectoral muscle [18]. These angles and the direction of compression are shown in figure 2.8.



Figure 2.8: Illustration of mammographic image acquisition of the right breast

#### 2.3 Magnetic Resonance Imaging

MR imaging is not currently part of standard breast cancer management; however, it is occasionally used as an adjunct to mammography, particularly for younger women with dense breasts and women at high risk [19]. In addition to providing 3D images with excellent soft tissue contrast, MR imaging has the ability to show dynamic functionality of the breast through the use of contrast agent injections [19]. Since MR relies on different tissue properties than X-ray imaging techniques, breast MR and mammography are complementary modalities.

The three major components of an MR imaging system are the main magnet, the gradient coils, and the radio frequency (RF) coil, which work together to produce an image resulting from electromagnetic interactions with certain atomic nuclei (predominantly <sup>1</sup>H protons) [16].

The main magnet is a strong superconducting electromagnet (in this work, generating a field strength of 1.5 Tesla) large enough to surround the patient, causing some of the spins of atomic nuclei within the body to tend to align to the magnetic field [15]. The gradient coils provide very small magnetic field gradients in each of the three orthogonal directions, allowing for localization of the signal [16]. The RF coil is placed around the object of interest to measure the signal; in this case, the patient lies prone on the scanner with the breast pendant inside the coil, causing distortion in shape due to gravity (figure 2.9).

During image acquisition, the RF coil or a larger whole-body RF coil emits a series of carefully timed electromagnetic pulses, which cause the protons to tip out of alignment with the main magnetic field. As these protons realign themselves, they emit a responding RF signal [15]. It is this induced signal that is measured and used to form the final 3D image.



Figure 2.9: Patient positioning within an MR scanner

Since the electromagnetic radiation used in MR imaging is in the radiofrequency (MHz) range, the energies associated are much lower than the 13.6 eV required to cause ionization [16]. As such, MR imaging is considered safer than X-ray based modalities. However, RF energy can be deposited into tissue in the form of heat, and rapidly changing magnetic fields can induce currents [16]. To avoid harmful effects from these factors, careful selection of the RF pulse sequence, such as limiting the rate of change of the gradients, is required.

In addition to potential direct damage to tissues, MR imaging can cause metal artifacts within the body to heat up or shift, thus damaging tissues indirectly. Due to this effect, patients with significant metallic devices such as pacemakers, surgical clips, and some prostheses are unable to undergo MR examination [16]. Furthermore, patients unable to withstand loud noises or enclosed spaces may not be able to tolerate MR image acquisition.

MR imaging has a complex set of parameters that must be optimized to produce an appropriate image. In particular, increasing image quality affects the time required to obtain an image, calculated for a single x location according to the simplified equation [16]:

$$t = M_y \times M_z \times TR \times N_{ex} \tag{2.1}$$

where  $M_y$  and  $M_z$  are the number of pixels in the y and z directions, TR is the repetition time between pulse sequences, and  $N_{ex}$  is the number of measurements made per pixel. Increases in any of these parameters can lead to improved image quality, but introduces the risk of patient discomfort and motion artifacts [16].

Clinical breast MR imaging is performed with the use of a contrast agent, referred to as contrast enhanced (CE) imaging. MR contrast agents contain gadolinium, a paramagnetic material that affects the magnetic response of nearby atomic nuclei and causes them to appear brighter on the resulting image [16]. In CE imaging of the breast, the contrast agent is injected into the bloodstream and a series of images are acquired to observe the rate at which the contrast-containing blood enters and exits various regions of the breast [19]. The speed at which contrast is taken up and washed out of a region can indicate malignancies [19].

A comparison between MR imaging and X-ray mammography is presented in table 2.2 below.

Parameter	MR	Mammography
Cause of contrast	Proton density and relax-	Atomic number, density,
	ation time	and thickness
Information provided	Structural and functional	Structural, including mi-
	(with CE)	crocalcifications
Spatial resolution	Low	High
Acquisition time	Minutes	Seconds
3D information	Complete	Limited to 2 views
Anatomical distortion	Gravity	Compression

Table 2.2: Comparison of breast MR imaging and X-ray mammography

#### 2.4 Registration of Mammograms and MR Images

Examining table 2.2, it is evident that mammograms and MR images provide different and complementary information. Previous researchers have noticed this fact, and several methods for combining the information from MR and mammographic images have been proposed. The most common method for determining correspondence between images is registration [20].

Image registration can be formally defined as the process of mapping points from one image (the "moving" image) to corresponding points on another image (the "fixed" image) [21]. This process is illustrated in figure 2.10, with the mapping transform represented by  $\mathbf{T}$ . Registration in its simplest form is rigid, where the moving image is simply translated and rotated until it corresponds with the fixed image. In contrast, non-rigid or deformable registration allows the moving image to deform nonlinearly until correspondence is achieved [21]. In other words, the transformation mapping points from the moving to the fixed image cannot be represented by a simple matrix. For



Figure 2.10: Registering a moving image (A1) to a fixed image (A2)

both deformable and rigid applications, there are three types of registration algorithms: landmark-based, region-based, and intensity-based [21]. Landmark-based methods involve manually or automatically selecting corresponding points on the two images; region-based methods require segmentation and identification of corresponding regions of interest; and intensity-based methods operate directly on the intensity information of the images.

Rigid registration methods are both the simplest and most rigorously validated [22]. These techniques are are well established in research and are becoming increasingly common in clinical settings, particularly in neuroimaging [20, 22]. While the simplest of these methods are the fully automated, validated and highly accurate intra-modality registration algorithms, the most common clinical applications involve registration of low-resolution functional images such as those obtained from single photon emission computed tomography (SPECT) to anatomically detailed roadmap images obtained from computed tomography (CT) or MR [22]. These images allow for brain function to be localized to anatomical features, and one such algorithm is gaining clinical acceptance for epilepsy diagnosis [23].

Non-rigid registration is facing larger resistance to clinical adoption due to issues such as difficulties in validating results as well as increased computational complexity [20]. However, validated techniques that are accurate enough for clinical applications are becoming more common while technological advances are removing the computational limitations, and it is anticipated that image registration will play a greater role in clinical image analysis in the future [22].

Breast images present a unique challenge to image registration due to the large and anisotropic deformations resulting from mammographic compression and gravitational forces during MR acquisition [6, 24]. Furthermore, the two modalities differ in dimension, resolution, dynamic range, intensity/tissue relationship, and noise level. As a result, breast image registration methods are in their infancy compared to neuroimaging and other rigid registration application areas.

So far, only three approaches to problem of registering mammograms to MR images have been published. In 2003, Behrenbruch *et al.* used a landmark-based registration algorithm to align mammograms to "simulated mammograms" created from MR images [25]. Behrenbruch's technique began with three predefined external landmarks followed by internal landmarks obtained through a wavelet-based feature detection method. Martí *et al.* took a similar approach in 2004, with the addition of an intensity-based registration technique to determine the appropriate angle for MLO projection of the MR data [26]. Both groups effectively removed the mammographic distortion through their non-rigid registration techniques.

Most recently, Ruiter *et al.* took the approach of creating a patient-specific finite element model based on an MR image, deforming it computationally to simulate mammographic compression, then creating a simulated mammogram through this compressed volume [27]. Following model creation, Ruiter's group achieved correspondence by iteratively modifying the applied deformation, creating a projection image, then using an intensity-based registration method until satisfactory alignment was achieved.

Results of these studies are difficult to compare, as all three chose different reporting methods. Martí's work focuses on estimating the projection angle of the MLO view and determining the MR slice corresponding to features seen in mammograms, but does not present quantitative results of lesion localization [26]. Behrenbruch *et al.* compare lesion locations on the registered images, reporting errors averaging 20% of the deformation displacement distance [25]. Finally, Ruiter's group measured the Euclidean distance between corresponding lesions as seen in both the 2D registered view and the 3D MR view, resulting in an average error of 4 mm [27].

The method proposed by Ruiter *et al.* produced highly accurate results, and as such is objectively the best of the three. However, the authors provide the caveat that the data set used involved mammograms acquired under compression of only 21% strain, which is on the low end of mammographic procedures [27]. Furthermore, development of a patient-specific finite element model is not a trivial task, and as such poses a barrier to clinical use of this technique.

The feature-based registration techniques described above have similar limitations. As breasts are non-uniform structures, the presence of specific internal landmarks is not guaranteed, which could lead to non-corresponding points being used as alignment points. Furthermore, pathological regions were often used as strong landmarks for alignment [25, 26]. However, it is often desirable to compare regions that do not obviously correspond such as microcalcifications seen on mammograms or regions of increased contrast agent uptake on MR images.

All three of the aforementioned studies were performed using film mammograms that were digitized. While the resulting images were of similar diagnostic quality as current digital images, information about the MLO acquisition angle and the amount of mammographic compression were lost. As a result, estimation of these parameters posed further challenges [25, 26, 27].

#### 2.5 Reconstruction of Mammograms in 3D

Estimation of 3D features from mammograms has received limited attention due to the same challenges affecting registration. Nonetheless, several groups have investigated reconstruction of microcalcifications in 3D.

The first attempt at estimating 3D features from standard two view mammograms was in 1998, when Müller *et al.* developed a method to reconstruct the 3D shape of microcalcification clusters [28]. Müller's group found corresponding microcalcification clusters on each view and backprojected along the X-ray beam trajectories, then found the 3D location as the intersection between the two [28]. Affine deformation due to compression was assumed, and the cluster was presented in 3D space without information on location within the breast.

In 2001, Yam *et al.* and Kita *et al.* expanded on this concept by estimating the curved trajectory of the X-ray beam due to mammographic compression as well as estimating the shape of the breast surface from the mammograms [29, 30]. This allowed for the localization of 3D structures with deformation taken into account, as well as provided

a reference model of the breast shape. However, this work was limited by two major assumptions: 1) the undistorted breast profile is the same as an eroded version of the compressed breast profile; and 2) the MLO view is approximately the same as the ML view [30].

More accurate results can be obtained by precise calibration of the mammographic imaging system, as demonstrated by Daul *et al.* [31]. However, as this requires a modification to standard imaging procedures, it cannot be compared with previous techniques.

#### 2.6 Summary

This chapter introduced important concepts such as the representation of images as continuous and discrete signals, as well as some of the critical differences between mammography and MR imaging. Limitations of previous breast MR/mammogram registration techniques were identified as: the need for discrete feature detection, lack of knowledge of MLO acquisition angle, and complex finite element model creation. Mammographic feature reconstruction also suffered from the lack of MLO acquisition angle information, as all previous studies found in literature assumed the two views to be at 90° to each other.

In the work described in this thesis, a combination of methods similar to those found in the literature was used, with modern automated intensity-based registration replacing previous feature-based techniques. In chapter 3, the mammograms are registered to projection images created from MR image volumes to remove the distortion caused by mammographic compression. Following this undistortion step, chapter 4 describes how features visible in the mammograms were backprojected along orthogonal X-ray trajectories to determine their locations in 3D space. Chapter 4 also introduces a novel skin surface estimation technique and quantifies the accuracy of this estimate through
comparison with a ground truth.

By combining modern registration methods and 3D estimation techniques, several limitations of the methods found in the literature are addressed. Intensity-based registration removes the need for discrete feature detection in both modalities, while undistortion of the mammograms allows for orthogonal backprojection to be valid. This work also presents a validation method for the registration technique, addressing one of the key clinical concerns with image registration.

# Chapter 3

# Mammogram/MR Image Registration

This chapter describes the methods used to remove mammographic distortion by registering mammograms to projection images created from breast MR images. First, a series of preprocessing steps are performed on both images to detect landmarks, which are used for preliminary alignment. The mammogram is then mapped into the space of the magnetic resonance projection image (MRPI) using an intensity-based registration method. The purpose of this registration procedure is to generate an estimate of a mammogram that would have been produced in the absence of compression, allowing for the reconstruction methods of chapter 4 to be applied.

The algorithms described below were implemented using a combination of the Insight Toolkit (ITK), an open source C++ image processing library, and the Visualization Toolkit (VTK) [32, 33]. A graphical user interface was created to enable interactive modification of algorithm parameters, as determining appropriate values required applying the algorithm and inspecting the result. A more complete description of the software created for this work is available in appendix A.

# 3.1 Preprocessing Algorithm

The preprocessing algorithm, shown in figure 3.1, was used to prepare the images for registration. The only difference between the treatment of the two image types is in the first step, where a single MR volume was used to generate two 2D MRPIs at different angles. These two images were then processed as if they were individual mammograms.



Figure 3.1: Overview of the preprocessing algorithm

#### 3.1.1 MR Volume Rotation and Projection

In order to register 2D mammograms to 3D MR image volumes, 2D images approximating the fields of view (FOVs) of the mammograms were first created from the MR volume. In digital mammography, the angle of the X-ray beam vector relative to vertical is known; thus, an excellent estimate of the exact acquisition angle of the imaging plane can be assumed.

In the case of the medial-lateral oblique (MLO) FOV, the MR image volume was rotated around the x axis (defined in §2.1.2) by the negative value of the X-ray beam vector  $\theta$  contained in the metadata of the corresponding mammogram. The z and y coordinates were multiplied with a 2D rotation matrix, leaving the x coordinates unchanged.

$$\begin{bmatrix} z'\\ y' \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta)\\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} z\\ y \end{bmatrix}$$
(3.1)

While the voxel coordinates of the rotated volume were easily obtained using equation 3.1, the intensity values required interpolation and resampling onto a regular grid [34]. This was accomplished by defining a new rectangular grid large enough to contain the rotated volume and computing the intensity values of each new voxel using bilinear interpolation. This rotation procedure is illustrated in figure 3.2.



Figure 3.2: Illustration of image rotation

After the MR image volume was rotated to match the field of view of the mammogram, a projection image was created. The average voxel intensity value along the z axis (MLO) or y axis (cranial-caudal or CC) was computed and assigned to a 2D pixel. This procedure is illustrated in figure 3.3.





(a) Mean intensity projection through volume

(b) Resulting 2D image

Figure 3.3: Creation of an MR projection image

Following MRPI creation for both MLO and CC views, the rest of the preprocessing algorithm was performed separately on each of the MRPIs as well as their corresponding mammograms.

MLO MRPIs have coordinates in the xy' plane as described above, while CC MRPIs have coordinates in the zy plane. However, in mammography as well as in the 2D image data type used for processing operations, only x and y coordinates are available and spatial orientation information is lost; thus, for the remainder of this chapter, image processing operations are described for images in the xy plane.

## 3.1.2 Background Segmentation and Smoothing

To detect the anatomical landmarks and the breast/air boundary line, the image was separated into background (air) and breast regions. This was accomplished by segmentation of the background using a fuzzy connectedness algorithm [35]. Medical images, by nature of their acquisition processes, contain inherent inaccuracies [35]. As such, it can often be a challenge to segment an image into binary objects, as a given object is unlikely to have clearly-defined edges and uniform intensity values. One approach to this problem uses the notion of fuzzy "hanging togetherness" to describe the connectivity of pixels within objects.

The fuzzy connectivity algorithm, described in detail in appendix B.1, is initiated with a seed pixel (chosen in this work to be the top-left pixel). Each of the neighbours of this pixel are examined and a connectivity value is computed, and if any of these pixels are determined to be connected, their neighbours are in turn examined. This process continues until all the pixels in the image have been visited [35].

This algorithm produced a new image composed of the fuzzy connectedness values for each pixel relative to the seed pixel, ranging from 0.0 to 1.0. The image was then divided into background and breast regions by thresholding at a manually adjustable fuzzy membership value, with an intensity value of 0 assigned to the background region and 1 assigned to the breast region. The threshold value typically ranged from 0.3 to 0.6.

Following fuzzy segmentation, a median filter was used to remove scattered islands in the binary segmentation and ensure a smooth breast edge. Median filtering is a nonlinear processing technique that implements the following algorithm:

- 1. Duplicate the image to obtain an input and an output image.
- 2. Compute the median of an  $n \times n$  region centred on the first pixel of the input image.
- 3. Assign the median value to the corresponding pixel in the output image.
- 4. Shift the  $n \times n$  region to the next pixel in the input image.
- 5. Repeat steps 3-4 until the entire input image has been traversed.

For pixels at the extrema where the  $n \times n$  region extended beyond the bounds of the image, an intensity of 0, or background, was assumed. Different values of n were appropriate for different image types and pixel sizes; 5 was typically used for mammograms and 3 for MRIs.

An example of a mammogram with its corresponding smoothed segmentation is shown in figure 3.4



(a) Original mammogram







#### 3.1.3 Edge Detection

Using the segmented image as shown in figure 3.4b, the largest edge in the image was detected and assumed to be the line defining the boundary between the background and the breast. This line was created by iterating through all the pixels of the foreground (breast) region and examining its 8-connected neighbours. If there were at least one of each background and foreground neighbours, the centre pixel was assumed to be part of the boundary.

As there was no prior knowledge of the boundary shape within the segmentation image, pixel scan order (figure 3.5) was used to iterate through the pixels. However, to be able to compute curvatures (described in step 3.1.5) the points must be ordered such that a smooth and continuous path is formed by moving from one to the next. Again, this was accomplished by iterating through each of the pixels and examining its neighbours. The first pixel that had only one neighbour in the boundary set was defined as the start. From there, the centre of the  $3 \times 3$  neighbourhood was shifted to the next boundary pixel, and this process was repeated until the other end was reached.



Figure 3.5: Ordering the set of points lying on the breast/air boundary

#### 3.1.4 Spline Creation

In order to display the boundary as a continuous curve, the ordered points were used to create 1D cardinal splines for the x and y coordinates. To ensure smoothness and minimize the impact of outliers, the points were averaged in sets of N, where N is a user-determined integer typically ranging from 3 to 15 depending on image resolution. The new reduced data sets were then used as control points in the cardinal splines.

A spline can be defined as "a set of polynomials of degree k that are smoothly connected at certain data points" [36]. To enforce the smoothly connected requirement, at each data point (where two polynomials connect) the derivatives up to the  $(k-1)^{st}$  order must be equal.

For a third order cardinal spline, first-order continuity is enforced but second-order continuity is not; as such, it is not a true spline according to the definition above [36]. However, it is constructed in a similar manner, and provides a useful degree of local To construct a cardinal spline fit to n points, the set of points is divided into n-3 overlapping groups of four consecutive points each:

$$[P_1, P_2, P_3, P_4], [P_2, P_3, P_4, P_4], \cdots [P_{n-3}, P_{n-2}, P_{n-1}, P_n]$$

$$(3.2)$$

For each group, Hermite interpolation is used to create a curve segment (appendix B.2). For the cardinal spline, tangents  $P_1^t$  and  $P_2^t$  of equation B.7 are defined as  $s(P_3 - P_1)$  and  $s(P_4 - P_2)$ , respectively, where s is a real number. The choice of s, ranging from 0 to 1, controls the "tension" of the spline; If s = 0, a straight line (infinite tension) is obtained [36]. However, this is counter intuitive to the concept of tension, so a tension parameter T is defined as T = 1 - 2s, allowing a tension value of 0 to correspond to the minimal tension case [36].

For this work, a tension parameter of T = 0 was chosen. This results in a curve that is defined as having zero tension, also known as the Catmull-Rom spline.

#### 3.1.5 Landmark Point Set Creation

As shown by Behrenbruch *et al.* three anatomical landmarks can be used to assist with alignment of breast images: the curvatures where the breast meets the chest wall, referred to as axilla and rib<sup>1</sup>; and the nipple. To automatically detect these landmarks, the maximum curvature points along the spline were computed [25].

The curvature C(u) where u is a parameter ranging from 0 to 1.0 along the path of the spline, was computed as the dot product of the curvature and the normalized gradient as follows:

$$C(u) = \frac{\frac{d^2u}{dx^2} \cdot \frac{du}{dy} - \frac{du}{dx} \cdot \frac{d^2u}{dy^2}}{(\frac{du}{dx})^2 + (\frac{du}{dy})^2}$$
(3.3)

 $<sup>^{1}</sup>$ For the CC case, these landmarks do not correspond to the axilla and rib regions; however, the same labels are used for convenience

At a manually adjustable number of equally spaced points (ranging from 50 to 150 depending on spline length), the curvature was computed according to equation 3.3. The nipple was defined as the point of maximum positive curvature, and was constrained to region A as shown on figure 3.6. The axilla and rib landmarks were defined as the locations of maximum negative curvature in regions B and C respectively.



Figure 3.6: Constraint regions for landmark detection

In certain cases, such as data sets where the nipple was obscured or a portion of the chest wall was cut off, manual identification of the landmarks was required. For these situations, an interactive graphical interface was used to assist the user in selecting a location u along the spline.

#### 3.1.6 Cropping and Background Zeroing

The final stage in the preprocessing algorithm was to crop the image around the breast region and "zero" the background. Both of these procedures remove extraneous information and prevent background noise or the chest wall from influencing the registration.

Zeroing the background, or setting all of the background pixel values to zero, was accomplished using the binary segmentation image created in step 3.1.2. This image was multiplied pixel-wise with the original mammogram or MRPI, resulting in a background region that was uniformly zero and a breast region that was unchanged.

Cropping of the image was performed by removing the portions of the edge spline that extended beyond the rib and axilla landmarks, then cropping the image around the smallest rectangular bounding box containing this spline.

Example images resulting from the preprocessing pipeline are shown in figure 3.7.



(a) Mammogram



(b) MR Projection image

Figure 3.7: Preprocessed images with contour spline in red, landmarks in green

# 3.2 Registration Algorithm

The registration algorithm, as shown in the flowchart in figure 3.8, takes the outputs from the preprocessing algorithm and warps the mammogram to match the shape of the corresponding MRPI. While this appears to be simpler than the preprocessing algorithm (figure 3.1), the steps involved are more complicated and have parameters that are more difficult to estimate.



Figure 3.8: Overview of the registration algorithm

#### 3.2.1 Preliminary Landmark Alignment

Registration is performed in the world coordinate system (see §2.1.2). As the two acquisition methods are very different, the mammogram and the MRPI are unlikely to be close to each other in physical space. Landmark-based registration using the three features marked in figure 3.7 was used to align the two images and to provide a better starting image for intensity-based registration.

The transformation was implemented using an elastic body spline (EBS) coordinate transform, a landmark-based deformable registration technique [37]. The EBS transform uses the displacement between corresponding landmarks as control points and the rest of the pixel locations are interpolated using a spline modelling the physical properties of an elastic body [37]. In other words, the image is assumed to be composed of a certain material, and is allowed to deform based on how that material would behave if the deformation were fixed at the control points.

While the method described by Davis *et al.* applies to the deformation of a 3D volume, it can also be used to deform a 2D image by eliminating the third dimension in a set of independent equations [37]. The result is that the image is still modelled as an elastic body, but only two dimensions are considered when computing the spline parameters.



Figure 3.9: MLO mammogram aligned to MRPI (contour shown in red) using landmarks

Using the EBS method, the displacement of a pixel location  $\vec{x} = [x, y]$  can be calculated as follows:

$$\vec{d}(\vec{x}) = \sum_{i=0}^{N} \mathbf{G}(\vec{x} - \vec{p_i})\vec{c_i} + \mathbf{A}\vec{x} + \vec{b}$$
(3.4)

where  $\vec{p_i}$  are the coordinates of the N landmarks,  $c_i$  are spline coefficients, and  $\mathbf{A}\vec{x} + \vec{b}$  is an affine transform accounting for the bulk displacement, rotation, and scaling of the image. A more complete description of this equation and the EBS transformation is included in appendix B.3.

For this work, the three landmarks detected in the preprocessing algorithm were used to calculate the parameters of the EBS transform. The mammogram was first transformed to the space of the MRPI using the affine component of equation 3.4, then deformed pixel-wise to give a rough deformation. An example of the result of this preliminary alignment is shown in figure 3.9.

#### 3.2.2 Mutual Information Registration

The EBS transform described in the previous section computes image displacement values based on the known displacements of landmarks. For intensity-based image registration methods, no such known displacements are available. Therefore, an iterative process beginning with an initial estimate was used.

Iterative registration algorithms are treated as optimization problems [21]. Three components, described in the following sections, are used in the registration process: a) an interpolator and transform, b) a metric, and c) an optimizer. The high-level algorithm for registration is shown as a flow chart in figure 3.10.



Figure 3.10: An overview of the iterative registration method (adapted from [34])

#### a) The Interpolator and Transform

The interpolator and transform are used during each iteration to transform the coordinates of the moving image into the space of the fixed image. This transform can be as simple as a rigid rotation and translation, or it can be a complex deformation field. The interpolator is used to determine the intensity values of the transformed image, as there is unlikely to be a one-to-one correspondence between the discrete image grids [21].

For this work, complex non-linear deformations were obtained by using a B-spline basis to represent the images and deformations, providing a smoother gradation of intensity values than simpler interpolation methods [38, 39]. Recalling the definitions for image representation presented in §2.1, an image represents an underlying continuous image f(x, y). To interpolate on a B-spline basis, any value of  $f(\mathbf{x})$  that does not correspond to a discrete pixel location  $\mathbf{x}_i$  can be calculated according to:

$$f(\mathbf{x}) = \sum_{i} c_i \beta^{(3)}(\mathbf{x} - \mathbf{x}_i)$$
(3.5)

where  $\mathbf{x}$  and  $\mathbf{x}_i$  are, respectively, the continuous and discrete *n*-tuples representing the location in the image (e.g.  $\mathbf{x} = [x, y]^T$  for a 2D image);  $c_i$  are the coefficients calculated according to the upsampled method described in appendix B.4; and  $\beta^{(3)}$  is the separable sampled third-order B-spline convolution kernel for each dimension [39]. For the x coordinate, the B-spline kernel is given as:

$$\beta^{(3)}(x) = \begin{cases} \frac{1}{6}(4 - 6x^2 + 3|x|^3), & 0 \le |x| < 1\\ \frac{1}{6}(2 - |x|)^3, & 1 \le |x| < 2\\ 0 & 2 \le |x| \end{cases}$$
(3.6)

The image transform is also modelled as a cubic B-spline using a sparse regular grid of control points  $\lambda_j$ , with j knots on a user-defined grid spacing n-tuple  $\rho$  [39]. In this work, the value of  $\rho$  was generally [3, 3] to [5, 5]. The spacing of the grid is computed as:

$$\Delta \rho = \left[ \Delta \rho_x, \Delta \rho_y \right]^T = \left[ \frac{q_x - 1}{\rho_x - 1}, \frac{q_y - 1}{\rho_y - 1} \right]^T$$
(3.7)

where  $q_x$  and  $q_y$  are the dimensions of the moving image.

Each of the control points  $\lambda_j$  has an associated deformation coefficient  $\delta_j$  describing the deformation in each of the component directions. The deformation at any image point **x** can be interpolated via:

$$D(\mathbf{x}|\delta) = \sum_{j} \delta_{j} \beta^{(3)} \left(\frac{\mathbf{x} - \lambda_{j}}{\Delta \rho}\right)$$
(3.8)

where again,  $\beta^{(3)}$  is the separable cubic B-spline convolution kernel [39]. Finally, the transformation of the moving image is achieved by combining the deformation term of

equation 3.8 with a bulk rotation R and translation T:

$$\mathbf{g}(\mathbf{x}|\mu) = R(\mathbf{x} - \mathbf{x}_C) - (T - \mathbf{x}_C) + D(\mathbf{x}|\delta)$$
(3.9)

where **x** is any pixel location  $[x, y]^T$  in the fixed image and  $\mathbf{x}_C$  is the centre of the moving image. The full set of transformation parameters is given as  $\mu = \{\alpha, \beta, t_x, t_y; \delta_j\}$ , where  $\alpha$  and  $\beta$  give the Euler angles of the rotation matrix R;  $t_x$  and  $t_y$  define the translation vector T; and  $\delta_j$  are the deformation coefficients of equation 3.8.

The transform undergoing optimization maps coordinates from the *fixed* image space into the *moving* image space [21]. While counter intuitive to the task of aligning the moving image to the fixed image, this can be explained as follows (see figure 3.11):

- 1. A pixel location  $(i_f, j_f)$  on the discrete grid of the fixed image is located in physical space  $(x_f, y_f)$ .
- 2. The coordinate  $(x_f, y_f)$  is transformed using equation 3.9, yielding coordinates  $(x_m, y_m)$  in the moving image space.
- 3. The coordinate  $(x_m, y_m)$  is located on the discrete grid of the moving image.
- 4. The intensity value  $f(x_m, y_m)$  is interpolated from the adjacent pixels of the moving image and assigned to the original pixel location  $(i_f, j_f)$  in the fixed image grid.

This process is repeated for each pixel location of the fixed image grid, resulting in a new image in the same physical space as the fixed image but with intensity values corresponding to the moving image.

b) The Metric

The metric is the component of the registration system that implements a cost function to compare the moving image to the fixed image [21]. In its simplest form, this might consist of computing the pixel-wise difference between the two images. Such a metric is



Figure 3.11: Mapping the moving image into the fixed image space (adapted from [21]) not possible for registration of multiple modalities, as the intensity values are not likely to correspond [39]. Therefore, a measurement of mutual information (MI) was chosen to compare the fitness of the transformed mammogram to the MRPI.

MI provides a statistical measure of how much information one image tells about another [40]. If two images do indeed represent the same object, then it is reasonable to expect that the MI will be higher when the two images are aligned. The main advantage to this metric is that the form of the dependency between the two images does not need to be specified, allowing it to work with images of differing tissue/intensity correspondence [21].

In this work, the implementation developed by Mattes et al. was used, as it is more

computationally efficient than previous methods [39, 40]. Mattes' method, based on work by Thévenaz and Unser, uses the negative of MI as a measure of image discrepancy so that a minimizing optimizer may be used [39, 41].

The cost function, expressed as a function of the independent transformation parameters  $\mu$  of equation 3.9, is given as [39]:

$$S(\mu) = -\sum_{\kappa}^{N} \sum_{\iota}^{N} p(\kappa, \iota | \mu) \log_2 \frac{p(\kappa, \iota | \mu)}{p_M(\iota | \mu) p_F(\kappa)}$$
(3.10)

which has units of bits. p,  $p_M$ , and  $p_F$  are the joint, moving image marginal, and fixed image marginal probability distributions, respectively, and  $\iota$  and  $\kappa$  are the integer indices of the N histogram bins for each image. The probability distributions are estimated from the two images by using a B-spline based Parzen window, and are explained in more detail by both Mattes *et al.* and Thévanaz and Unser [39, 41]. In the implementation developed by Mattes *et al.* (used in this work), a random subset of pixels (ranging from 50%-100% of the total number) was used to estimate the probability distributions, increasing the computational efficiency of the registration.

#### c) The Optimizer

An optimizer is a technique used to find the best solution to a problem that does not necessarily have a unique correct answer. Generally, optimizers determine minima (or maxima) of a cost function describing the problem at hand. As an example, a simple univariate optimization problem may require only computing the derivative and setting it equal to zero, thus finding the critical point of the curve [42].

For the cost function described by equation 3.10, a limited-memory Broyden Fletcher Goldfarb Shanno minimization algorithm with simple bounds (L-BFGS-B) was used as the optimizer. This is a modification of the popular quasi-Newton BFGS method [42].

Newton's method of optimization is an iterative algorithm that achieves fast convergence in complex problems by assuming that the cost function takes the form of a quadratic in the local region around the point of interest, then taking a step towards the minimum (or maximum) of that quadratic. Mathematically, the value of the set of parameters  $\mu$  undergoing optimization is calculated at iteration m + 1 according to [42]:

$$\mu_{m+1} = \mu_m - [\nabla^2 S(\mu_m)]^{-1} \nabla S(\mu_m)$$
(3.11)

where  $\nabla^2 S(\mu)$  is the Hessian matrix of the MI function and  $\nabla S(\mu)$  is the gradient.

The BFGS method aims to resolve two problems that can arise while implementing Newton's method: the computation expense of evaluating or inverting  $\nabla^2 S(\mu)$ , and the direction of convergence; equation 3.11 may move either "uphill" or "downhill" at locations far from the minimum [42].

Both of these problems are solved by introducing a matrix H to approximate  $\nabla^2 S(\mu)$ . While the algorithm begins with an initial estimate  $H_0$ , each subsequent iteration is updated with an efficient addition operation, with the direction of the update vector constrained to ensure convergence [42].

As H must be retained through each iteration, memory limitations can become a problem. The L-BFGS method addresses this problem by further approximating Hwith a low-rank matrix, while the introduction of bound constraints on the independent variables results in the final L-BFGS-B method used in this work [39].

### 3.3 Registration Evaluation

For registration between two images of the same modality, evaluation of "goodness of fit" can be visualized directly by computing the pixel-wise difference [34]. However, intermodality registration cannot be assessed in this manner, as pixel intensities are unlikely to correspond. In addition, breast imaging poses additional registration assessment challenges due to the highly variable nature of the tissues and a lack of internal landmarks [25]. Therefore, another method is required to determine registration accuracy. The values produced by the MI metric provide a scalar measure of fitness. These values can be plotted against iteration number to provide an understanding of how image correlation is achieved; however, as a single scalar value it can be a difficult measure to understand. As an alternative visual method of fitness assessment, the joint entropy can be plotted as a 2D image.

The MI function  $S(\mu)$  of equation 3.10 can also be expressed in terms of image entropy or randomness as follows [43]:

$$S(\mu) = E(f) + E(m|\mu) - E(f,m|\mu)$$
(3.12)

where E(f) and E(m) are the entropies of the fixed and moving images, respectively, and E(f,m) is the joint entropy of the two images; thus, maximization of MI is the same as minimization of joint entropy. The entropy function is defined in terms of the probability distribution as [43]:

$$E(x) = \sum_{i}^{X} p_X(x) log_2(p_X(x))$$
(3.13)

where x are the pixels in image X. The joint entropy is calculated similarly, substituting the joint probability distribution for the marginal distribution  $p_X$ .

For the purposes of visualization, the joint entropy is plotted using the image histograms in lieu of the estimated probability distributions. The joint entropy plot is thus an  $N \times M$  image, where N and M are the number of bins in the fixed and moving image histograms respectively. The intensity value f at each pixel location (i, j) in the entropy image is computed as [34]:

$$f_{ij} = -p_{ij} \log_2(p_{ij}) \tag{3.14}$$

where  $p_{ij}$  is calculated from the frequency count of the bins of the joint histogram between the two images:

$$p_{ij} = \frac{q_{ij}}{\sum_{i=0}^{N-1} \sum_{j=0}^{M-1} q_{ij}}$$
(3.15)

The value  $q_{ij}$  is the frequency of the bin ij in the joint histogram, calculated as the number of pixels where the fixed image has intensities falling into bin i and the moving image has intensities in bin j at corresponding locations [34]. The entropy plot therefore produces a visualization of pixel correspondence between images, regardless of modality, and the scalar MI value is the sum of all pixels in this plot. An entropy plot showing perfect correlation (i.e. a plot of the joint entropy between two copies of the same image) appears as straight y = x line, where the x and y axes correspond to the histogram bins of the fixed and moving images, respectively.

In addition to the MI value, the accuracy of registering the external shape of the breast can be calculated from binary segmentations of the undistorted mammograms and MRPI breast regions. The measure used in this work is Dice's coefficient, calculated as [44]:

$$D = \frac{2|X \cap Y|}{|X| + |Y|}$$
(3.16)

where |X| and |Y| are the sizes in physical units of the fixed and moving breast segmentations, and  $|X \cap Y|$  is the size of the overlapping region. Dice's coefficient ranges from 0 (no overlap) to 1 (completely aligned), and represents the overlapping fraction of the total area.

### 3.4 Validation Using Simulated Images

Validation of image registration is a difficult task as the ground truth is generally not available [6]. Typical methods of quantifying registration accuracy involve placing fiducials in the object of interest, performing phantom studies, or using clearly defined landmarks within the object [6].

In this work, fiducial placement was determined to be impractical as mammograms and MR images could not be acquired on the same day. Similarly, no clearly defined landmarks could be identified within the breast. Therefore, a simulated phantom was created and deformed using a corresponding coupled Eulerian-Lagrangian finite element (FE) model created by Martin Kuhlmann [45].

The FE model was defined as a hemispherical Lagrangian skin surface filled by an Eulerian fluid representative of fatty tissue as well as a sphere of denser Eulerian material to act as glandular tissue [45]. The flat surface of the hemisphere was defined as the chest wall and was not allowed to deform, while a parallel-plate displacement boundary condition was used to mimic mammographic compression in the CC view. A constant force simulating gravity was used for the MR distortion case [45].

A simulated image matching the physical dimensions of the FE model (illustrated in figure 3.12a) was created with a voxel intensity value of 200 for fatty tissue. Glandular tissue was assigned a voxel intensity of 400, and two smaller spheres with intensity values of 800 were created to represent lesions. While these lesions were not present on the FE model, it is assumed that they do not affect significantly affect breast deformation.

The point clouds defining the positions of the skin elements of the FE model before and after mechanical deformation were used to deform the simulated image. The point clouds defining the original geometry were treated as the moving landmarks of the EBS transform described in §3.2.1, while the deformed geometry point clouds were used as fixed landmarks. This procedure was repeated for both CC mammographic compression and MR gravitational distortion. Cross sections of the original and deformed images are shown in figure 3.12.

Average intensity projections through each of the deformed images were taken and the resulting 2D images were registered using the algorithm described in section 3.2. The simulated MRPI, original simulated mammogram, and joint entropy plot prior to registration are shown in figure 3.13. Similarly, the simulated MRPI, undistorted simulated mammogram, and joint entropy plot after registration are shown in figure 3.14.



Figure 3.12: Simulated images deformed via finite element modelling



(a) MRPI



(b) Original Mammogram



(c) Joint Entropy Before Registration

Figure 3.13: Simulated images prior to registration



Figure 3.14: Results from registration of simulated images

The final MI value for the registered simulated images was 2.2377, while the MI for the simulated MRPI registered to itself (auto-correlation) was 3.366; therefore, the registration result is highly accurate numerically. This can also be seen by comparing the joint entropy plots of figures 3.13c and 3.14c, where the plot prior to registration shows significant disorder and the plot following registration gains cohesiveness as well as approaches the ideal y = x line. Finally, visual inspection of the registered image shows that despite the high level of deformation applied to the two models, close correspondence was achieved.

To test the sensitivity of the registration algorithm to variations in image rotation, the simulated image was rotated around the x axis from 0 to 15° in 5° intervals. The resulting registered images, using the parameters determined to be optimal for the zero rotation case, are shown in figure 3.15 for the 5- and 10-degree rotations, while the 15degree rotation case failed to converge. These results indicate that registration accuracy is not significantly affected by rotation errors of up to 10 degrees.



Figure 3.15: Registered simulated images following rotation of MR volume

# 3.5 Results

The registration process described in the previous section was tested on a total of eight pairs of images from patients enrolled in studies 22121 and 18463 approved by the Conjoint Health Research Ethics Board (CHREB). Patients whose breasts contacted large regions of the MR coil were excluded, as the resulting images were significantly deformed and could not be used as reference images.

For each pair of images, parameters were optimized iteratively through visual assessment of results. Modifications were made according to table 3.1.

Figure 3.16 shows the registered mammogram and pre- and post-registration entropy plots of the sample MLO data set, with the MRPI and original mammogram shown in figure 3.7. The joint entropy plot following registration shows improved cohesiveness relative to the joint entropy plot prior to registration; however, a perfect x = y curve is not expected with clinical data, as the intensity distributions of the two image types do not have a one-to-one correspondence.

The final MI value, Dice's coefficients, and time taken to perform registration for each

Problem	Modification		
Optimizer failure	Relax accuracy and tolerances, de-		
	crease number of histogram bins		
Image not aligning in dark regions	Increase number of histogram bins		
Image contour too variable/smooth	Decrease/increase number of B-		
	spline knots		
Small features not aligning	Increase number of spatial samples		
	used by metric		
Image not deforming enough	Increase number of corrections		
Optimizer exits without converging	Increase number of iterations		

Table 3.1: Modifications made to parameters to optimize result

data set are presented in table 3.2.

Subject	Breast	View	Final MI Value	Dice's Coeff.	Time (s)
091208	L	CC	0.5802	0.92	195
091208	L	MLO	0.5988	0.92	147
091208	R	CC	0.7086	0.97	147
091208	R	MLO	0.7264	0.92	131
091210	L	CC	0.6401	0.95	126
091210	L	MLO	0.7384	0.92	213
091210	R	CC	0.7023	0.89	106
091210	R	MLO	0.7781	0.91	71
100704	L	CC	0.5218	0.94	191
100704	L	MLO	0.7514	0.98	121
100704	R	CC	0.6594	0.92	191
100704	R	MLO	0.7453	0.95	121
Average			0.6792	0.93	147

Table 3.2: Registration Results

The MRPIs and corresponding registered mammograms from all eight data sets, along with brief descriptions, are shown in §3.5.1- §3.5.8. Visual inspection of the results indicates that a very good match of external shape has been achieved for all data sets. This is corroborated by the high values of Dice's coefficient in table 3.2, averaging 0.93.



(a) Registered mammogram with MRPI contour

(b) Joint entropy before (top) and after (bottom) registration



#### 3.5.1 Subject 091208, Left Breast

The MR image of this data set exhibits flattening of the breast at the nipple end, caused by contact with the bottom of the MR RF coil. This resulted in difficulties in locating the nipple on the MRPI, particularly the MLO view (3.17c), potentially reducing the accuracy of the landmark registration algorithm.



(a) Reference MRPI (CC view)



(c) Reference MRPI (MLO view)



(b) Undistorted Mammogram (CC view)



(d) Undistorted Mammogram (MLO view)

Figure 3.17: Result of registering data set 091208-L

# 3.5.2 Subject 091208, Right Breast

This data set was chosen as the exemplar for all the algorithm descriptions in this thesis, as the nipple is clearly visible and centred in all images. However, the registration algorithm was still unable to predict the concave contour at the bottom of the MLO view (3.18c).



(a) Reference MRPI (CC view)



(c) Reference MRPI (MLO view)



(b) Undistorted Mammogram (CC view)



(d) Undistorted Mammogram (MLO view)

Figure 3.18: Result of registering data set 091208-R

#### 3.5.3 Subject 091210, Left Breast

This data set provides a good example of the challenge of non-uniform breast shapes. In the MR image, the side of the breast is in contact with the MR RF coil, seen in the CC MRPI (3.19a), whereas the nipple on the MLO mammogram (3.19d) is downturned. As the registration algorithm has no knowledge of anatomical motion, this resulted in strange deformations in the nipple region.



(a) Reference MRPI (CC view)



(b) Undistorted Mammogram (CC view)



(c) Reference MRPI (MLO view)



(d) Undistorted Mammogram (MLO view)

Figure 3.19: Result of registering data set 091210-L

# 3.5.4 Subject 091210, Right Breast

As this is the other breast from the same subject as presented in §3.5.3, a similar breast shape is seen. However, the nipple on the MLO mammogram was not downturned, and the resulting undistorted mammogram has a smoother contour (figure 3.20d).



(a) Reference MRPI (CC view)



(b) Undistorted Mammogram (CC view)



(c) Reference MRPI (MLO view)



(d) Undistorted Mammogram (MLO view)

Figure 3.20: Result of registering data set 091210-R

# 3.5.5 Subject 100201, Left Breast

The mammograms for this data set were acquired using a newer digital mammography system, resulting in obviously different images. This demonstrates that the registration algorithm is robust to variations in image acquisition.



(a) Reference MRPI (CC view)



(b) Undistorted Mammogram (CC view)



(c) Reference MRPI (MLO view)



(d) Undistorted Mammogram (MLO view)

Figure 3.21: Result of registering data set 100201-L

This is the other breast of the previous section (§3.5.5), and thus the mammograms were acquired using the same imaging system. With the clearly visible linear structures of the mammograms, images 3.22b and 3.22d are good demonstrations of the ability of the algorithm to deform the mammograms without exceeding anatomical limits.



(a) Reference MRPI (CC view)



(b) Undistorted Mammogram (CC view)



(c) Reference MRPI (MLO view)



(d) Undistorted Mammogram (MLO view)

Figure 3.22: Result of registering data set 100201-R

#### 3.5.7 Subject 100704, Left Breast

This data set is another good example of a "typical" breast shape, with relatively consistent curvature and symmetry. However, the nipple location has shifted laterally due to gravity during MR imaging, as is seen on the CC view (3.23a). This resulted in a slight distortion of the nipple shape of the corresponding registered mammogram (3.23b).



(a) Reference MRPI (CC view)



(b) Undistorted Mammogram (CC view)



(c) Reference MRPI (MLO view)



(d) Undistorted Mammogram (MLO view)

Figure 3.23: Result of registering data set 100704-L

As the only data set containing a lesion visible on both mammograms and the MR image volume, this data set was used as the test case for internal feature reconstruction, described in chapter 4. While the lesion is clearly visible in both mammographic views, it cannot be seen in the MRPIs.



(a) Reference MRPI (CC view)



(b) Undistorted Mammogram (CC view)



(c) Reference MRPI (MLO view)



(d) Undistorted Mammogram (MLO view)

Figure 3.24: Result of registering data set 100704-R

As described in §3.3, registration accuracy can often be difficult to assess. However, creation of a simulated image containing such landmarks provides a partial validation of the registration algorithm, and the results from §3.4 clearly show that when there are corresponding features visible, registration is highly accurate. In addition, the algorithm is robust to slight variations in rotation, which may result from variations in patient positioning.

The final metric values presented in table 3.2 are considerably lower than the MI values obtained during simulated image registration, as shown in figure 3.14. This is expected, as the images themselves are much more complex and variable; for example, registration of the test case MLO mammogram of figure 3.16 to itself resulted in an MI value of 2.420. The MI values of table 3.2, as well as the appearance of the joint entropy plots, are similar to previous work performing a different multi-modality registration task [43]. This is a promising result, as it indicates that reasonable registration has been achieved.

The MI values are difficult to interpret, largely due to the variability of the images, as the auto-correlation MI values (representing an effective maximum) range from 2.0 to 3.5. Furthermore, the number of histogram bins chosen for registration had an effect on the final MI values, as this changed the pixels used to compute the joint entropy.

In 2004, Martí *et al.* used an MI registration algorithm to compare mammograms to MR projection images at various angles, obtaining a maximum MI value of 0.233 [26]. While small differences in MI values are not necessarily indicative of differences in accuracy, the average final MI value of 0.6792 obtained in this work is over twice the value obtained by Martí *et al*; suggesting that the technique presented in this thesis is more accurate than the registration technique used by Martí *et al* [26].
In contrast to the registration technique developed by Behrenbruch *et al.* in 2003, the MI method used in this work does not require the identification of corresponding internal landmarks [25]. This reduces the reliance on discrete features and removes any assumptions of feature correspondence, resulting in a more robust registration algorithm that is suitable for a wide variety of breast types.

In the literature review presented in §2.4, the method developed by Ruiter *et al.* was found to produce the most accurate registration results. However, Ruiter's technique was limited to mammographic strains of 21% or less, while the technique presented in this thesis has no such limitation [27]. Furthermore, the MI technique does not require construction of a patient-specific finite element model, allowing for a more automated and computationally efficient registration procedure.

In general, the undistorted mammograms tended to fit within the boundaries of their corresponding MRPIs, indicating a slight underestimation of the size of the breast. This is a reasonable result, as the longitudinal extension of the breast under the influence of gravity in an MR scanner exceeds the lateral expansion resulting from mammographic compression. In other words, the original mammograms tend to have smaller surface areas than the MRPIs.

For all of the processing operations, several iterations were required to determine the optimal parameters, particularly during the mutual information registration step. A slight modification of one parameter (e.g. changing the number of spatial samples from 50% to 60%) has the potential to have a significant effect on the final result, with the optimizer occasionally getting trapped in local minima. However, with repeated executions performed by an experienced operator, satisfactory registration was achieved for all data sets.

In summary, this chapter presented a preprocessing algorithm, a registration algorithm, and a validation method. The contributions of this work include:

- The incorporation of precise mammographic acquisition parameters in the generation of the MRPI.
- The ability to manually override landmark detection in difficult cases.
- The use of a robust and automated MI technique to refine registration.
- A registration validation technique incorporating a simulated image and an FE model.
- The development of a graphical framework for ease of parameter variation and instantaneous feedback.

# Chapter 4

## Mammogram Reconstruction

This chapter describes the methods used to estimate the 3D shape of the breast surface and the 3D locations of internal regions of interest using only CC and MLO mammograms. The 3D breast surface will benefit TSAR imaging by allowing for the antenna scan pattern to be determined prior to patient arrival and by providing essential structure information for image formation from the EM signals. Similarly, prior knowledge of an internal region of interest will allow for both TSAR scanning and analysis to be targeted to a subsection of the breast volume. For this purpose and due to the centimetre-scale resolving power of TSAR, quadrant-level accuracy of internal lesion localization is sufficient, where the quadrants of the breast are upper outer, lower outer, upper inner, and lower inner.

To achieve these aims, the undistorted images resulting from the methods described in chapter 3, together with the corresponding edge contours and landmarks, are first aligned spatially to form a sparse wireframe breast model. The 3D surface is then estimated by fitting ellipses at evenly spaced intervals along the length of the breast from nipple to chest wall. Internal regions of interest are identified manually on the two images and backprojected to determine their location in 3D. This procedure is illustrated below in figure 4.1.

### 4.1 Image, Contour, and Landmark Alignment

The first step in establishing correspondence between the 2D mammograms and the 3D breast is to accurately position the two mammograms relative to each other. In



Figure 4.1: Overview of the reconstruction algorithm (dotted lines indicate steps performed only on data sets with visible features)

previous work, accuracy ranged from assuming a perfect right angle between images [28] to performing calibration during image acquisition [31]. In this work, modifications to standard imaging protocols were undesired; therefore, several assumptions were made in order to align the images:

- The subject is standing perfectly vertical during mammogram acquisition.
- The midpoint of the chest wall landmarks on both images correspond to the same location.
- The angle of obliquity contained in the medial-lateral oblique (MLO) metadata is accurate relative to the cranial-caudal (CC) view.

In the global coordinate system, CC mammograms should be in the xz plane, whereas MLO mammograms should be at an oblique angle between the xy and xz planes. In both cases, however, the images are formed with coordinates local to an X-ray machine which is rotated to the desired angle, resulting in images in the xy plane. Thus, both views required rotation to regain their positions the world coordinate system.

To properly align the images and skin contours, the CC data were rotated  $+90^{\circ}$ about the *x* axis, ensuring that direction vectors were conserved (i.e. increasing values of *y* became increasing values of *z*). The MLO data were rotated to their oblique angle by applying a rotation of  $\gamma = \theta - 90^{\circ}$  for the right breast, or  $\gamma = \theta + 90^{\circ}$  for the left breast, where  $\theta$  is the angle of the X-ray beam vector (figure 4.2).



Figure 4.2: Acquisition angles of MLO mammograms

After rotation, the images were in the correct orientation relative to the breast, but not aligned to each other. To achieve this, several more assumptions were required:

- The MLO data provide an accurate nipple location in the y axis.
- The CC data provide an accurate nipple location in the z axis.
- The nipples are in the same location along the x axis.
- The midpoint between the two chest wall landmarks lies on the x axis for both data sets.
- The contour of the mammogram, describing the largest edge or shadow of the breast, is located at the centre of the volume along the X-ray beam vector.

Using these assumptions as guidelines for alignment, the midpoint between the two chest wall landmarks was computed and the CC data were shifted such that the midpoint was set to the origin at (0,0,0). This placed the CC nipple point close to but not necessarily on the x axis, as the nipple is not usually precisely in the middle of the breast.

The MLO data were also shifted to place the midpoint at the origin, but this resulted in nipple misalignment relative to the CC landmark as the two midpoints are not guaranteed to be exactly the same location. To correct this, the CC data were first rotated around the z axis to bring the nipple in line with the MLO landmark in the xy plane. Next, the MLO data were shifted so that the nipple point was at the same Euclidean distance from the origin as the CC nipple. Finally, the MLO data were rotated around the y axis, bringing the two nipples in line.

The end result is a representation of the two imaging planes as they were acquired, with the contours from each image forming a sparse wireframe of the 3D breast shape. An example is shown in figure 4.3.



Figure 4.3: Sparse wireframe of the 3D breast shape

### 4.2 Skin Surface Estimation

The breast contours of the wireframe model of figure 4.3 intersect a given zy or coronal plane between the chest wall and nipple at four points. These points were determined by finding the locations along the contour splines that corresponded to the desired xcoordinate, then used to construct an estimate of the coronal slice contour.

Coronal slices of the breast are roughly elliptical in shape. However, four points are insufficient data for ellipse fitting, as a uniquely defined ellipse requires at least five data points [46]. Furthermore, these four points are not orthogonal to each other, resulting in an uneven distribution around the edge of the ellipse. To overcome these issues, three assumptions were made:

- 1. The ellipse is centred on the centroid of the four points.
- 2. The eccentric anomaly can be approximated by the angle formed between z and y.
- 3. The ellipse is angled at no more than  $30^{\circ}$ .

Ellipses can be described in parametric form as:

$$z(t) = z_c + a\cos t\cos\phi - b\sin t\sin\phi$$
$$y(t) = y_c + a\cos t\sin\phi + b\sin t\cos\phi$$
(4.1)

where the parameter t is the eccentric anomaly at location z(t), y(t) illustrated in figure 4.4a;  $z_c$  and  $y_c$  are the coordinates of the centre of the ellipse; and  $\phi$  is the angle of the major axis (at the contact point of the auxiliary) of the ellipse relative to the z axis. The ellipse of figure 4.4 is in canonical position, with  $\phi = 0$ .

To completely describe the ellipse, the parameters  $a, b, z_c, y_c$  and  $\phi$  must be determined from the four known points. As previously mentioned, the centre point is assumed to be equal to the centroid, computed by summing each coordinate of the four points and dividing by four.



Figure 4.4: Illustration of the eccentric anomaly (t) of an ellipse

The canonical form of an ellipse has a rotation  $\phi$  of zero and is centred at the origin. This reduces equation 4.1 to:

$$z(t) = a \cos t$$
$$y(t) = b \sin t \tag{4.2}$$

which can in turn be written in the linear form y = mx, with x being replaced by  $\cos t$  or  $\sin t$ . This results in eight equations of the four known points  $(z_i, y_i)$ :  $z_i = a \cos(t_i), y_i = b \sin(t_i)$ , for i = [0, 3]. This allows for the a and b parameters to be solved by the method of least squares, which reduces to:

$$m = \frac{\sum_{i=0}^{3} x_i y_i}{\sum_{i=0}^{3} x_i^2} \tag{4.3}$$

where m = a is found by setting  $x_i = \cos t_i$  and m = b is found with  $x_i = \sin t_i$ .

For the ellipse with  $\phi = 0^{\circ}$ , the parameter  $t_i$  is estimated as the angle formed by z and y:

$$t_{i} = \tan^{-1} \left( \frac{|y_{i} - y_{c}|}{|z_{i} - z_{c}|} \right)$$
(4.4)

where the absolute value is used to constrain the angle to the first quadrant. The angles resulting from equation 4.4 are then used in equation 4.3 to determine the best fit a and b parameters for the ellipse.

Equation 4.4 uses assumption 2 to estimate the  $t_i$  parameter. This is a reasonable assumption, as the ellipses that best fit the shape of the breast are close to being circular.

The final parameter  $\phi$  cannot be calculated directly using least squares minimization, as it is not possible to rewrite equation 4.1 as a linear function of  $\phi$ . Therefore, iterative methods are used.

For every angle  $\alpha$  of integer spacing from  $-30^{\circ}$  to  $+30^{\circ}$ , a rotation around the centroid is applied to the four coordinates. This effectively rotates the coordinate system by  $\alpha$ , allowing for the parameters a and b to be computed as if the ellipse were in canonical form.

For each desired x location, the sum of squared errors between the actual coordinates of the four points and the coordinates resulting from the least squares fit a and b parameters was calculated for each angle  $\alpha$ . The angle  $\alpha$  resulting in the smallest average sum of squared errors was then set as  $\phi$ .

With all the parameters of equation 4.1 determined, an ellipse was created and displayed at the specified x location. This process was repeated for a total of twenty of ellipses along the x axis; figure 4.5 shows a sample skin surface reconstruction.

### 4.3 Internal Feature Reconstruction

After developing an estimate of the breast skin surface, internal regions of interest must be identified and located in 3D space. While there are a number of studies on automated or assisted identification of lesions in mammograms, this work focuses on visualizing features in 3D space relative to the surface estimate. Thus, corresponding regions of



Figure 4.5: Sample skin surface reconstruction (only five coronal slices shown for clarity)



Figure 4.6: Reconstruction of internal feature points

interest on the CC and MLO views are identified manually as a proof of concept.

For the purpose of identifying the general region of interest in space, lesions were modelled as simple spheres. Corresponding points on the CC and MLO views were identified and appropriate radii for 2D circles determined through an interactive display. These points were then located in 3D space by calculating the intersection of the two lines orthogonal to the imaging planes (figure 4.6).

Mathematically, this is computed as follows:

$$x = \frac{x_{cc} + x_{mlo}}{2}$$

$$y = y_{mlo} + \frac{z_{mlo} - z_{cc}}{\tan^{-1} \gamma}$$

$$z = z_{cc}$$
(4.5)

The radius of the sphere representing the 3D region of interest was calculated as the average of the radii of the two lesions identified on the mammogram views.

### 4.4 Reconstruction Evaluation

As mentioned in the previous chapter, quantification of results requires a ground truth. In the case of the reconstruction problem, the MR image volume is the only available 3D ground truth; therefore, errors are calculated with respect to features visible in the MR image.

Correct alignment of the MR image relative to the mammogram data was subject to errors, as creation of the wireframe of figure 4.3 involved independent modification of the CC and MLO data sets as well as several assumptions (described in §4.1). As a result, only the shifting that was performed to align the centre points of the chest wall landmarks was applied to the MR image.

To quantify the accuracy of the ellipse surface, the location of the skin surface on the MR image must first be determined. This was accomplished by applying a 3D version the fuzzy connected region growing described in appendix B.1 to the background area of the MRI, resulting in segmentation regions representing air and breast. A binary median smoothing algorithm as previously described in §3.1.2 was used to remove outliers from these regions.

To determine trends in errors around the contour of the breast, the error between the actual skin edge and the ellipse estimate was computed in the  $\pm z$  and  $\pm y$  directions at each of the ellipse contour x locations. Errors above and below the nipple  $(\pm y)$  were labelled cranial and caudal respectively, whereas errors to the left and right  $(\pm z)$  were labelled medial and lateral for the right breast and vice versa for the left. These directions are indicated in figure 2.6.

The errors were calculated according to the following algorithm, recalling the definitions of the a and b parameters of an ellipse (§4.2):

- 1. Begin at the centre of the ellipse  $(z_c, y_c)$ .
- 2. Iterate through each pixel in the cranial direction (-y) until a pixel in the air region is reached at point  $p = (z_p, y_p)$ .
- 3. Calculate the error as  $e = y_p (y_c b)$ .
- 4. Return to the centre of the ellipse and iterate in the caudal (+y) direction.
- 5. When an "air" pixel is reached, compute the error as  $e = (y_c + b) y_p$ .
- 6. Repeat steps 1-6 for the medial and lateral directions, substituting z for y and a for b.

In this manner, negative errors are always representative of underestimation of the skin contour, and positive errors are always representative of overestimation, regardless of the direction of the axes. For the purposes of evaluating results, the original MR image used to create the MRPI of chapter 3 was loaded, shifted, and displayed concurrently with the 3D surface and lesion reconstructions.

Figure 4.7 shows a comparison between the sample data set (091208R) skin surface estimation and slices from the original MR image at two coronal locations.

The maximum absolute errors for all data sets are tabulated in table 4.1, while the average absolute errors are presented in table 4.2. For these data, the maximum and average of the absolute values of the errors were taken to avoid averaging between underand over-estimation errors. The absolute errors were presented instead of relative errors in order to compare between nipple and chest wall regions, as the small ellipses near the nipple would appear to have exaggerated errors.

Error values are plotted on two graphs relative to location along the x axis, shown in figure 4.8. Best fit quadratic curves for these data as well as the combined data along each axis are also shown.



(a) Near Chest Wall



(b) Near Breast Centre

Figure 4.7: Skin outline comparison at two slice locations



(b) Medial and Lateral

Figure 4.8: Skin Reconstruction Errors vs. x location (nipple at -x, chest wall at +x)

Subject	Breast	Cranial	Caudal	Medial	Lateral	Average
091208	L	23.0	10.4	4.9	12.0	12.6
091208	R	23.7	15.4	13.3	9.9	15.6
091210	L	27.4	21.5	12.1	9.1	17.5
091210	R	23.1	17.6	15.5	13.7	17.5
100201	L	16.3	15.8	16.9	14.5	15.9
100201	R	16.0	15.4	10.3	12.4	13.5
100704	L	29.4	8.2	18.7	21.8	19.5
100704	R	11.8	22.7	8.7	4.5	11.9
Average		21.3	15.9	12.6	12.2	15.5

Table 4.1: Maximum absolute errors in mm for each data set and direction

Subject	Breast	Cranial	Caudal	Medial	Lateral	Average
091208	L	13.1	11.7	9.4	4.2	9.6
091208	R	16.5	6.6	3.0	8.9	8.7
091210	L	14.7	9.9	6.3	5.4	9.1
091210	R	17.4	14.4	8.4	9.9	12.5
100201	L	12.5	7.1	11.3	4.7	8.9
100201	R	10.2	6.8	3.3	3.9	6.0
100704	L	21.7	4.6	8.7	7.3	10.6
100704	R	5.5	13.9	6.5	2.8	7.2
Average		14.0	9.4	7.1	5.9	9.1

Table 4.2: Average absolute errors in mm for each data set and direction

Comparison slices in the sagittal direction for all data sets, along with plots showing the trend of errors along the x axis and a short description, are presented in §4.5.1-§4.5.4.

#### 4.5.1 Subject 091208

While the left breast of this data set appears to be a better fit than the right by visual inspection, the error graphs indicate that the fit is close to equal for both sets. For the right breast, a caudal shift is apparent, which is reflected in the error plot; this may be a result of misalignment of the ground truth MR image rather than a true underestimation.



Figure 4.9: Result of reconstructing data set 091208

Like subject 091208 of  $\S4.5.1$ , the left breast of subject 091210 appears to match the ground truth better than the right at the slice shown. However, the largest error of the left breast exceeds that of the right by almost five mm, and both are errors in the cranial region near the chest wall. Both breast surfaces underestimate the ground truth with no overestimation errors; this can also be observed in the original registration of  $\S3.5.3$  and  $\S3.5.4$ .



Figure 4.10: Result of reconstructing data set 091210

### 4.5.3 Subject 100201

Both breasts in this example show a caudal shift in the surface estimation similar to that of figure 4.9c. However, the error plots are unusual in that all anatomical directions follow a similar trend of shifting from negative to positive along the x axis. Again, this may be indicative of errors resulting from the alignment assumptions of section 4.4.



Figure 4.11: Result of reconstructing data set 100201

#### 4.5.4 Subject 100704

The final data set tested in this work exhibits a caudal shift relative to the ground truth in the left breast (4.12a) and a cranial shift in the right (4.12c). This is reflected in the error plots, where the caudal surface is significantly better than the cranial for the left breast, and the inverse is true for the right. The error plot for the left breast shows a positive shift in error at approximately x = -35, possibly corresponding to the distortion of the breast observed at the edge of the MR RF coil.



Figure 4.12: Result of reconstructing data set 100704

Figure 4.13 shows a comparison between a reconstructed feature and the MR volume of the only data set containing a discrete feature visible in both modalities (100704-R). The red sphere marks the estimated region of the lesion resulting from backprojection, and the corresponding lesion on the MR is clearly visible as an opaque mass.



(a) Coronal View

(b) Sagittal View



### 4.6 Discussion

The method used to generate 3D estimates of the skin surface was capable of producing reasonable breast shapes such as the example shown in figure 4.5 for all data sets. This indicates that the ellipse fitting algorithm is robust and unlikely to result in extreme shapes, such as ellipses with the long axis at the MLO projection angle. In addition, the average absolute error between the estimated and ground truth skin contours was less than 10 mm (table 4.2), suggesting an excellent approximation given the limited data available.

The skin reconstruction tended to underestimate the surface of the breast in all direc-

tions. This can be seen visually by examining the sagittal comparison results of figures 4.9-4.12, and is also evident numerically in the error graphs. The errors in the mediallateral (ML) axis were lower on average than those in the cranial-caudal (CC) axis; this can be seen in the graphs of the overall trends (figures 4.8a and 4.8b).

Cranial errors showed a trend of increasing underestimation approaching the chest wall, while caudal errors exhibited the opposite trend (figure 4.8a). In addition, caudal errors tended to be lower than cranial, with average maxima of 15.9 mm and 21.3 mm respectively. The combined trend of both CC directions remained at a near constant 10 mm underestimation, suggesting that an improved estimate in one direction resulted in increased error in the other.

Medial and lateral errors tended to be roughly equal, following a trend of decreasing in error from the nipple to the chest wall (increasing x values in figure 4.8b).

These trends are an expected result, given the limitations and assumptions of the processing methods. In the ML axis, information is obtained from all four *a priori* points, whereas CC information is only available from the MLO data; thus, it is logical that errors in the CC axis are greater. Furthermore, the cranial region tends to be the most asymmetric.

The overall underestimation is largely a result of the tendency during registration (chapter 3) for the mammograms to underestimate the area of the MRPIs. The relative lack of data in the CC axis is also a factor, as it tends to be the longer axis of the elliptical cross-section.

Of the eight data sets used in this work, only one had a lesion visible in both the mammograms and the MR image. This lesion was used as a test case for internal feature estimation. Visual inspection of the comparison between the 3D estimation and the corresponding region on the MR image (figure 4.13) shows that the feature reconstruction algorithm was successful in locating the feature within the same quadrant of the breast.

The centroid of the red sphere of figure 4.13 is misaligned by approximately 15 mm in both the coronal and sagittal planes. This is well within the targeted quadrant-level accuracy, and is comparable to the centimetre level resolution of TSAR imaging.

The misalignment error of the lesion can be attributed to several factors. First, the lesion of the MR image is not seen on either MRPI (figure 3.24), and thus did not have an effect in driving the registration. Secondly, inspection of figure 4.13b reveals that the landmark used to identify the chest wall of the MLO view is significantly misplaced, indicating that all the data may be somewhat shifted. Finally, as the lesion is close to the skin, any inaccuracies in the orthogonal projection scheme (e.g. slight variations in acquisition angle) are magnified.

While the first factor could be improved by modifying the method used to generate the MRPI, it is important to note that the accuracy is still quite good despite the fact that the lesion itself did not influence the registration process. This indicates that the registration algorithm provides a means of comparing regions of interest even where corresponding features are not obvious.

For all of the error measurements and visual comparisons, error introduced by inexact alignment of the reference MR image may be a factor due to the independent rotations required of the two mammograms (as described in §4.4). It is important to note that this does not affect the accuracy of the reconstruction; indeed, it indicates that the actual errors are likely to be smaller overall.

Based on the literature review provided in §2.5, this work is the first to quantify the accuracy of mammographic skin surface estimation. While Kita and Yam *et al.* used a similar surface fitting technique, they did not have a ground truth skin surface available for comparison [30, 29].

As the internal feature localization accuracy is affected by the registration accuracy, this value can be used to compare the registration techniques used in this work to previous methods. A major limitation of this work is that only one data set was available to test internal lesion localization, so it is this single result that is compared to average values presented in literature.

As discussed in §2.4, the method developed by Behrenbruch *et al.* resulted in localization errors averaging 20 % local deformation, or approximately 10 mm [25]. While this is an improvement over the 15 mm error resulting from this work, Behrenbruch points out that lesions were often used as landmarks for registration, while the intensity-based registration method described in §3.2 assumes no such correlation.

Localization errors presented by Ruiter *et al.* are approximately three times smaller than the errors resulting from the single lesion localization result of this work [27]. However, the patient-specific FE model technique employed by Ruiter is more complex and requires more user interaction than was desired for the applications of this work. Furthermore, the model-based technique was limited to mammograms undergoing 21% strain or less, excluding the majority of clinical mammograms.

In summary, the average 10 mm error between estimated and true skin represents an excellent reconstruction from the two mammographic views. While previous work obtained internal feature localization accuracy in the 5-10 mm range, the estimated 15 mm error of this work exceeds the goal of achieving quadrant-level accuracy and is therefore acceptable. Furthermore, the reconstruction method is fully automated and extremely fast, requiring only one or two seconds to reconstruct a complete data model.

Contributions of this work to the literature are:

- Development of an ellipse fitting algorithm from only four unequally spaced data points.
- Quantification of breast surface estimation error as compared to a ground truth.

## Chapter 5

## Conclusion

The need for improved breast cancer diagnostics has driven research in many areas of breast image acquisition and analysis. In this thesis, two specific aims were achieved: registration of mammograms and magnetic resonance breast images, and estimation of three dimensional features from mammograms.

An efficient and automated algorithm to register 2D mammograms to projection images created from MR image volumes was presented. This method was validated through the use of simulated images deformed via finite element modelling, and was shown to be robust to rotational errors of up to 10°. Results of registering sixteen pairs of images showed excellent external alignment with an average Dice's coefficient of 0.93, and mutual information metrics showed improvement over similar work found in literature.

Following registration to MR projection images, each of the two mammogram views was considered undistorted; that is, the distortion introduced by mammographic compression was undone. These images were treated as projections through an uncompressed breast and a skin surface estimate was computed by fitting ellipses to twenty equally spaced coronal slices. This reconstructed surface was then compared to the 3D MR surface and was found to have an average error of 10 mm in all orthogonal directions.

The 3D location of internal lesions was estimated from the undistorted mammograms by backprojecting along an orthogonal trajectory from each of the two mammographic views and finding the intersection in 3D space. Only one data set with a visible lesion on both mammograms and the MR image was available, and the lesion localization error was determined to be approximately 15 mm. This is slightly worse than comparable methods presented in literature, but is well within the targeted goal of obtaining an estimate of lesion location with quadrant-level accuracy. Furthermore, this degree of error is comparable to TSAR's 1 cm resolving power, and therefore provides a reasonable estimate for TSAR purposes.

## 5.1 Summary of Contributions

The contributions of this work arising from the first specific aim are as follows:

- The incorporation of precise mammographic acquisition parameters in the generation of synthetic mammograms for the purposes of image registration.
- The ability to manually override landmark detection in difficult cases.
- The use of a robust and automated intensity-based mutual information technique to refine registration of breast images.
- A registration validation technique incorporating a simulated image and an FE model.

Together, these three contributions suggest development of the most accurate, robust, and automated registration of mammograms and MR projection images reported to date. This observation is supported by quantification using Dice's coefficient as well as mutual information.

The second specific aim resulted in the following contributions:

- Development of an ellipse fitting algorithm from only four unequally spaced data points, permitting computationally efficient estimation of the skin surface from two non-orthogonal 2D mammograms.
- Quantification of breast surface estimation error as compared to a ground truth.

These two contributions suggest that a rapid method to estimate the skin surface has been developed. This method is accurate to an average of 10 mm, and errors are traceable to limitations in image registration and ground truth alignment.

Achievement of the two specific aims resulted in the ability to track lesions in 3D space with errors on the centimetre scale.

### 5.2 Limitations

This work is limited in several respects. As mentioned in chapter 3, breast images exhibiting significant deformation due to contact with the MR coil were rejected, indicating that this work only applies to breasts below a certain size threshold. Similarly, preliminary landmark alignment was not always successful, as the landmarks were not clearly visible on all images.

The reconstruction algorithm of chapter 4 was also not flawless. The ellipse fitting method to estimate the skin surface resulted in consistent errors in the cranial portion of the breast, indicating that improvements could by made by anticipating this error. Furthermore, alignment of the ground truth MR data was a challenge, and improvements to this procedure would likely improve validation results.

The largest limitation of this work was the small number of data sets used for algorithm development and testing. In particular, the internal feature reconstruction method could only be tested on a single data set. While the results from this example were satisfactory, more data sets are required to prove robustness.

### 5.3 Future Work and Potential Applications

Registration of mammograms to MR images is a useful technique not only for undistorting mammograms but also for improved multimodal diagnostic power. Estimation of the 3D

location of features seen in mammograms also contributes to this end as well as providing a method of visualizing mammograms in 3D. However, perhaps the most significant contribution of this work is the information generated, which has the potential to assist in a number of other research and clinical applications.

#### 5.3.1 Statistical Deformation Model

It may seem counter intuitive to attempt to estimate the surface of the breast based on two mammographic views, as the MR image provides this data with no estimation or processing necessary. The reconstruction methods of chapter 4 deliberately do not depend on MR data, as it is anticipated that in the future the use of MR information will be optional.

MR imaging is not currently a standard procedure in Canadian breast cancer care [2]. While MR has the advantage of providing 3D information through the use of nonionizing radiation, it is considerably slower and more expensive than mammography, and mammography will likely remain the gold standard for the near future [2]. Researchers and clinicians alike stand to benefit from 3D information provided by breast imaging techniques, but if this information can be obtained from mammograms, the time and financial burdens on both the health care system and the patient can be reduced.

The EBS (§3.2.1) and B-spline (§3.2.2) transforms used to deform the mammogram are defined by the translations of specific nodes: the landmark positions (EBS) and the regular warping grid (B-spline). Given a data set of sufficient size, it is conceivable that a statistical deformation model could be developed by determining average deformations for different breast types. This model could then be used to undistort mammograms without the need to acquire MR images, allowing for 3D estimation of mammographic features directly.

#### 5.3.2 Applications to TSAR

As described in the introduction to this thesis, TSAR is an emerging microwave-based breast imaging technique [3]. While undergoing a TSAR scan, the patient lies prone on a table with her breast pendant into a tank of immersion medium such as canola oil [3]. An antenna is then placed close to the skin surface and scanned in a three dimensional pattern around the breast while emitting ultrawideband electromagnetic signals and measuring the reflections. As malignancies have been shown to affect these reflections, TSAR has the potential to detect cancer based on tissue properties in addition to structure [3].

Scanning the TSAR antenna around the breast requires prior knowledge of the breast size, shape, and location as the antenna must be in close proximity to the skin surface [12]. This can be estimated through the use of a "pre-scan" using either a laser or the TSAR antenna, but this introduces additional imaging and processing time, resulting in patient discomfort and possible motion artifacts [47]. Use of the mammographic skin surface estimate developed in this work would overcome these problems by allowing for the TSAR scanning pattern to be determined prior to patient arrival. Furthermore, image formation from the EM signals requires knowledge of the skin surface location, as assumed EM properties on the inside and outside the breast are used to calculate the distance between the antenna and the reconstructed pixel location.

Like mammography and MR imaging, acquisition of a TSAR image results in a distortion of the breast shape. While patient positioning is similar to that of MR, the use of immersion medium counteracts the effects of gravity by introducing floatation. To account for these changes, an FE model capable of predicting the deformation due to floatation has been developed [45]. In the future, the skin surface estimation obtained from this work could be used as a starting point for this FE model.

In addition to the skin surface prior knowledge, the ability to reconstruct the 3D location of internal regions of interest seen on mammograms has the potential to reduce

scanning time and improve precision of the TSAR system by targeting the microwave beam at a specific area of interest. Diagnostic power can also be increased by considering the joint information from multiple modalities. Finally, knowledge of the 3D location of features as seen in mammograms could serve as a validation tool for the TSAR system. This is illustrated in figure 5.1, where a TSAR image, the reconstructed tumour location superimposed on the corresponding MR slice, and the original mammograms from data set 100704R are shown. While it is evident that the largest feature seen in the TSAR image corresponds to the tumour as seen on the MR image and 3D estimate, the same position is difficult to locate and visualize from the 2D mammograms.

#### 5.3.3 Surgical Planning

Current clinical guidelines prescribe either core biopsy, fine needle aspiration, or excision of suspected malignant regions detected through mammography [48]. Core biopsies and fine needle aspiration are performed by inserting a biopsy needle (large or fine gauge) into the region of interest and extracting cells or tissue, typically under ultrasound guidance [48]. However, not all lesions are visible under ultrasound, and the accuracy of these procedures depends on the expertise of the physician [48].

Excision of malignant regions involves placing fine wires into the lesion under radiographic or ultrasound image guidance while consulting a radiologist and the original mammograms [48]. These wires remain in place until the lesion is surgically removed, after which another mammogram is obtained to verify removal [48].

Both lesion biopsy and excision procedures could benefit from 3D reconstructions of mammographic lesion location relative to the skin surface. By presenting this visualization in conjunction with the original mammograms, there is the potential to reduce the inter-operator variability in accuracy of image-guided wire or needle placement.





## References

- [1] Canadian Cancer Society, "Canadian Cancer Statistics 2009," 2009.
- [2] Steering Committee on Clinical Practice Guidelines for the Care and Treatment of Breast Cancer, "The palpable breast lump: information and recommendations to assist decision-making when a breast lump is detected," *Canadian Medical Association Journal*, vol. 158, Oct. 1998.
- [3] E. C. Fear, P. M. Meaney, and M. A. Stuchly, "Microwaves for breast cancer detection?" *Potentials, IEEE*, vol. 22, pp. 12–18, 2003.
- [4] M. Kriege, C. T. M. Brekelmans, I. M. Obdeijn, C. Boetes, H. M. Zonderland, S. H. Muller, T. Kok, R. A. Manoliu, A. P. E. Besnard, and M. M. A. Tilanus-Linthorst, "Factors Affecting Sensitivity and Specificity of Screening Mammography and MRI in Women with an Inherited Risk for Breast Cancer," *Breast Cancer Research and Treatment*, vol. 100, pp. 109–119, 2006.
- [5] C. K. Kuhl, S. Schrading, C. C. Leutner, N. Morakkabati-Spitz, E. Wardelmann, R. Fimmers, W. Kuhn, and H. H. Schild, "Mammography, Breast Ultrasound, and Magnetic Resonance Imaging for Surveillance of Women at High Familial Risk for Breast Cancer," *Journal of Clinical Oncology*, vol. 23, pp. 8469–8476, 2005.
- [6] Y. Guo, R. Sivaramakrishna, C.-C. Lu, J. Suri, and S. Laxminarayan, "Breast image registration techniques: a survey," *Medical and Biological Engineering and Computing*, vol. 44, pp. 15–26, 2006.
- [7] T. England and N. Sharples, "Dielectric Properties of the Human Body in the Microwave Region of the Spectrum," *Nature*, vol. 163, no. 4143, pp. 487–488, Mar. 1949.

- [8] A. J. Surowiec, S. S. Stuchly, J. R. Barr, and A. Swarup, "Dielectric properties of breast carcinoma and the surrounding tissues," *Biomedical Engineering, IEEE Transactions on*, vol. 35, pp. 257–263, 1988.
- [9] M. Lazebnik, C. Watkins, S. Hagness, J. Booske, D. Popovic, L. McCartney, M. Okoniewski, M. Lindstrom, T. Breslin, J. Harter, S. Sewall, W. Temple, D. Mew, A. Magliocco, and T. Ogilvie, "The dielectric properties of normal and malignant breast tissue at microwave frequencies: analysis, conclusions, and implications from the wisconsin/calgary study," in *Proceedings of the IEEE APS International Symposium*, 2007, pp. 2172–2175.
- [10] J. Shea, P. Kosmas, B. V. Veen, and S. Hagness, "Three-dimensional microwave tomography for breast imaging and cancer detection," in *International Conference* on Biomedical Applications of Electrical Impedance Tomography, 2009, pp. 5–6.
- [11] N. Joachimowicz, C. Pichot, and J. P. Hugonin, "Inverse Scattering: an iterative numerical method for electromagnetic imaging," *IEEE Transactions on Antennas* and Propagation, vol. 39, no. 12, pp. 1742–1753, 1991.
- [12] T. C. Williams, J. M. Sill, and E. C. Fear, "Breast Surface Estimation for Radar-Based Breast Imaging Systems," *IEEE Transactions on Biomedical Engineering*, vol. 55, pp. 1678–1686, 2008.
- [13] A. Fhager, C. Chen, and M. Persson, "On the Use of A Priori Data in Microwave Tomography," URSI General Assembly, Chicago (IL), vol. 2, no. 1, pp. 2–5, 2008.
- [14] A. H. Golnabi, P. M. Meaney, S. D. Geimer, M. W. Fanning, and K. D. Paulsen, "Microwave imaging utilizing a soft prior constraint," in *Proceedings of SPIE*. Spie, 2009.

- [15] J. L. Prince and J. M. Links, Medical Imaging Signals and Systems. Pearson Prentice Hall, 2006.
- [16] D. J. Dowsett, P. A. Kenny, and R. E. Johnston, *The Physics of Diagnostic Imaging*. Oxford University Press, 2006.
- [17] IEEE Computer Society, "IEEE Std 754-2008 IEEE Standard for Floating-Point Arithmetic," 2008.
- [18] Health Canada, "Canadian Mammography Quality Guidelines," 2002.
- [19] D. Saslow, C. Boetes, W. Burke, S. Harms, M. O. Leach, C. D. Lehman, E. Morris, E. Pisano, M. Schnall, S. Sener, and Others, "American Cancer Society Guidelines for Breast Screening with MRI as an Adjunct to Mammography," *CA: A Cancer Journal for Clinicians*, vol. 57, pp. 75–89, 2007.
- [20] W. R. Crum, "Non-rigid image registration: theory and practice," British Journal of Radiology, vol. 77, pp. S140–S153, Dec. 2004.
- [21] T. S. Yoo, Insight into images: principles and practice for segmentation, registration, and image analysis. AK Peters, Ltd., 2004.
- [22] R. Knowlton, Clinical Applications of Image Registration. San Diego, CA: Academic Press, 2000, ch. 38, p. 901.
- [23] C. Jack, B. Brinkmann, D. Hanson, M. O'Connor, T. O'Brien, B. Mullan, and E. So, "Subtraction ictal SPECT co-registered to MRI in partial epilepsy: Description and technical validation of the method with phantom and patient studies," *Nuclear Medicine Communications*, vol. 19, no. 1, pp. 31 – 45, 1998.
- [24] Y. Zhang, S. Sarkar, Y. Qiu, D. B. Goldgof, and L. Li, "3D Finite Element

Modeling of Nonrigid Breast Deformation for Feature Registration in X-ray and MR Images," in *IEEE Workshop on Applications of Computer Vision*, 2007.

- [25] C. Behrenbruch, K. Marias, P. Armitage, M. Yam, N. Moore, R. English, P. Clarke, F. Leong, and J. M. Brady, "Fusion of contrast-enhanced breast MR and mammographic imaging data," *The British Journal of Radiology*, vol. 77, pp. 311–340, Jan. 2003.
- [26] R. Marti, R. Zwiggelaar, C. Rubin, and E. Denton, "Two-Dimensional-Three-Dimensional Correspondence in Mammography," *Cybernetics and Systems*, vol. 35, pp. 85–105, 2004.
- [27] N. Ruiter, R. Stotzka, T. Muller, H. Gemmeke, J. Reichenbach, and W. Kaiser, "Model-based registration of X-ray mammograms and MR images of the female breast," *IEEE Transactions on Nuclear Science*, vol. 53, pp. 204–211, Feb. 2006.
- [28] T. Muller, R. Stotzka, A. Hochmuth, and W. Eppler, "Volume reconstruction of clustered microcalcifications in mammograms," *Computational Imaging and Vision*, vol. 13, pp. 321–328, 1998.
- [29] M. Yam, J. M. Brady, R. Highnam, C. Behrenbruch, R. English, and Y. Kita, "Three-dimensional reconstruction of microcalcification clusters from two mammographic views," *IEEE Transactions on Medical Imaging*, vol. 20, no. 6, pp. 479–89, Jun. 2001.
- [30] Y. Kita, "Correspondence between Different View Breast X Rays Using Curved Epipolar Lines," Computer Vision and Image Understanding, vol. 83, pp. 38–56, Jul. 2001.
- [31] C. Daul, P. Graebling, A. Tiedeu, and D. Wolf, "3-D Reconstruction of Microcalcification Clusters Using Stereo Imaging: Algorithm and Mammographic Unit

Calibration," *IEEE Transactions on Biomedical Engineering*, vol. 52, no. 12, pp. 2058–2073, 2005.

- [32] "The Insight Segmentation and Registration Toolkit." [Online]. Available: www.itk.org
- [33] "The Visualization Toolkit." [Online]. Available: www.vtk.org
- [34] L. Ibanez, W. Schroeder, L. Ng, and J. Cates, "The ITK Software Guide," 2005.
- [35] J. Udupa and S. Samarasekera, "Fuzzy Connectedness and Object Definition: Theory, Algorithms, and Applications in Image Segmentation," *Graphical Models* and Image Processing, vol. 58, no. 3, pp. 246–261, May 1996.
- [36] D. Salomon, Curves and surfaces for computer graphics. Springer, 2006.
- [37] M. Davis, A. Khotanzad, D. Flamig, and S. Harms, "A physics-based coordinate transformation for 3-D image matching." *IEEE Transactions on Medical Imaging*, vol. 16, no. 3, pp. 317–28, 1997.
- [38] M. Unser, a. Aldroubi, and M. Eden, "Fast B-spline transforms for continuous image representation and interpolation," *IEEE Transactions on Pattern Analysis* and Machine Intelligence, vol. 13, no. 3, pp. 277–285, Mar. 1991.
- [39] D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank, "PET-CT image registration in the chest using free-form deformations." *IEEE Transactions* on Medical Imaging, vol. 22, no. 1, pp. 120–8, Jan. 2003.
- [40] P. Viola and W. Wells, "Alignment by maximization of mutual information," International Journal of Computer Vision, vol. 24, no. 2, pp. 137–154, 1997.

- [41] P. Thévenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," *IEEE Transactions on Image Processing*, vol. 9, no. 12, pp. 2083–99, Jan. 2000.
- [42] K. Lange, *Optimization*. Springer, 2004.
- [43] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multi-Modality Image Registration Maximization of Mutual Information," *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 187–198, 1996.
- [44] R. R. Sokal and P. H. Sneath, Principles of Numerical Taxonomy. WR Freeman and Company, 1963.
- [45] M. Kuhlmann, E. Fear, A. Ramirez-Serrano, and S. Federico, "A Coupled Eulerian-Lagrangian Mechanical Model of the Human Breast," in 11th Annual Alberta Biomedical Engineering Conference, 2010.
- [46] P. Rosin, "Further Five-Point Fit Ellipse Fitting," Graphical Models and Image Processing, vol. 61, no. 5, pp. 245–259, Sep. 1999.
- [47] D. W. Winters, J. D. Shea, E. L. Madsen, G. R. Frank, B. D. Van Veen, and S. C. Hagness, "Estimating the breast surface using UWB microwave monostatic backscatter measurements." *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 1, pp. 247–56, Jan. 2008.
- [48] Steering Committee on Clinical Practice Guidelines for the Care and Treatment of Breast Cancer, "Investigation of lesions detected by mammography," *Canadian Medical Association Journal*, vol. 158, pp. 9–14, Oct. 1998.
- [49] C. de Boor, "B(asic)-Spline Basics," p. 34, 1986.
# Appendix A

# Software Framework



Figure A.1: The main program interface in reconstruction mode

To allow for processing parameter variation and ease of use, a graphical program was developed. This software, called "Penumbra", supports image viewing, preprocessing, registration, and reconstruction modes. On the right hand side, loaded images are displayed, as shown in figure A.1. On the left hand side, controls for modifying pipeline parameters are shown. Each of the images in the image display is a possible input for the processing pipelines.

Penumbra was developed in C++ using a combination of open-source toolkits. wxWidgets 2.8 (www.wxwidgets.org) was used for the graphical interface components; the Insight toolkit (www.itk.org) was used for image processing algorithms; and the Visualization toolkit (www.vtk.org) was used for image display and 3D rendering. Figure A.2 shows an overview of the primary classes and interactions.

The following section describes the major components of Penumbra.

**PenumbraFrame** is the main user interface class, shown upon initial launch. It is the "glue" holding the program together, responsible for passing information between sub-classes and responding to program-wide (i.e. not image- or process-specific) user input.

**ImageIO** class handles reading of image files and series as well as writing of .png screenshots, vtkPolyData for curves and surfaces, and DICOM images. Only one instance of this class is available to the entire program, ensuring only a single image is read or written at a time.

WindowBase is the base class for the display windows. It contains a vtkRenderer object that allows data to be displayed, and a wxVTKRenderWindowInteractor<sup>1</sup> to interface between the user, the renderer, and the wxWidgets GUI elements. Zoom, pan, and rotate are provided by this class.

**ImageWindow** inherits from WindowBase and is responsible for displaying 2D images or slices from 3D images, as well as user interactions such as changing window/level of the display, scrolling through slices, and image saving operations. Functions specific to this work such as displaying contours and landmarks are also provided.

Window3D also inherits from WindowBase, but instead of a 2D window displaying images or slices, the render area is treated as a 3D space. This class contains specific functions for skin contour (ellipse) display, wireframe breast spline display, and internal features. Image slices can also be displayed, but they are shown as a 2D plane floating in 3D space.

ImageWindowWrapper is a helper class that handles flexible layout of images, allow-

<sup>&</sup>lt;sup>1</sup>Available from http://wxvtk.sourceforge.net/



Figure A.2: Software framework (Penumbra) overview

ing for loading and displaying as many images as can be held in memory. This class also keeps track of 3D window groups (i.e. multiple views) and performs intelligent destruction of ImageWindow objects.

**ProcessUI** is a flexible widget generator for interacting with the ITK process pipelines. The only controls explicitly defined are for input image selection and pipeline execution; all others are created through various functions at run time. Each set of controls are contained in a pane and added to a wxListBook object, which displays the list of panes and allows for parameter editing on the selected item. ProcessUI is a "dumb" class, with no knowledge of the pipeline capabilities; it simply provides an interface allowing the user to modify parameters of the pipeline.

**Preprocess, Registration, and Reconstruction Pipelines** contain the ITK pipelines implementing the algorithms described in chapters 3 and 4. For each pipeline, the list of required process objects (filters) is declared upon object instantiation, but creation (and memory allocation) does not occur until the process is enabled by the user via the corresponding **ProcessUI**. Values for processing parameters are set when the user enters a value, but connections between filters and pipeline execution is deferred until the "Run" button is pressed, thus preventing any lag while modifying parameters.

## Appendix B

## Algorithms

#### B.1 Fuzzy Connected Object Extraction

The fuzzy connectivity between two pixels is determined by considering both the locations and the intensities of each pixel. In order to describe the algorithm used to determine connectedness, several definitions must first be made.

For a given set of elements X, a fuzzy subset  $\mathcal{A}$  of X is a set of ordered pairs:

$$\mathcal{A} = \{x, \mu_{\mathcal{A}}(x) | x \in X\}, \text{where } \mu_{\mathcal{A}}(x) : X \mapsto [0, 1]$$
(B.1)

and  $\mu_{\mathcal{A}}(x)$  is the *membership function* of x in X. Similarly, a fuzzy 2-ary relation  $\rho$  in X can be defined as a subset of  $X \times X$ :

$$\rho = \{(x, y), \mu_{\rho}(x, y) | (x, y) \in X \times X\}, \text{ where } \mu_{\rho}(x, y) : X \times X \mapsto [0, 1]$$
(B.2)

A digital image in Euclidean space  $\mathbb{R}^n$  is divided into discrete units, or *spels*. The coordinates of the centres of each spel c are defined by an array of length n, where the j'th coordinate of c is written as  $c_j$  and  $1 \leq j \leq n$ . The set of all spels in the image is referred to as the discrete space  $\mathbb{Z}^n$ .

With this definition of a digital image, relations between spels can be found by treating  $Z^n$  as the set  $X \times X$  in equation B.2. One such relation,  $\alpha$ , is defined as the fuzzy spel adjacency. While many different functions could be used, an example of the fuzzy adjacency between spels c and d is:

$$\mu_{\omega}(c,d) = \begin{cases} \frac{1}{1 + k_1(\sqrt{\sum_{i=1}^{n} (c_i - d_i)^2})}, & \text{if } \sum_{i=1}^{n} |c_i - d_i| \le n \\ 0 & \text{otherwise} \end{cases}$$
(B.3)

where  $k_1$  is a non negative constant. A commonly used simple form of equation B.3 is obtained by setting  $k_1 = 0$ , resulting in a hard 8-connected adjacency.

The pair  $(Z^n, \alpha)$ , where  $\alpha$  is an adjacency relation such as  $\omega$ , is referred to as a *fuzzy* digital space.

A scene over  $(Z^n, \alpha)$  is a pair  $\mathcal{C} = (C, f)$ , where the scene domain  $C = \{c | -b_j \leq c_j \leq b_j \text{ for some } b \in Z_+^n\}$ .  $Z_+^n$  is the set of positive *n*-tuples; thus, C is the set of spels in  $(Z^n, \alpha)$  between plus and minus b. f is some function of c, and if the values of f are on the range [0, 1] it is referred to as a membership function. When the range of f is  $\{0, 1\}$ ,  $\mathcal{C}$  is a binary scene, or segmentation.

The fuzzy affinity  $\kappa$  between spels c and d in a scene C can be defined by any fuzzy relation, provided it is reflexive and symmetric (see Udupa *et al.* for definitions [35]). For breast background segmentation, a Gaussian function was used:

$$\mu_{\kappa}(c,d) = e^{\frac{-\left(\frac{f(c)+f(d)}{2}+\mu\right)^2}{2\sigma^2}}$$
(B.4)

where  $\mu$  and  $\sigma^2$  are the mean and variances of the object being segmented and f(c) and f(d) are the pixel intensity values.  $\mu$  and  $\sigma$  are determined by the average values of the neighbours of the seed pixel and the user at run-time, respectively.

The final concept to be introduced prior to outlining the segmentation algorithm is that of the connectivity path. A path  $p_{cd}$  from c to d is represented by a sequence of spels  $\langle c^{(1)}, c^{(2)}, \ldots c^{(m)} \rangle$  where  $m \geq 2$  spels and  $c^{(1)} = c$ ,  $c^{(m)} = d$ . The set of all possible paths from c to d is represented by  $P_{cd}$ , and the set of all paths in C is  $P_C$ .

The strength of connectivity along a given path p is defined as the minimum affinity value between all spels along the path. Combining all strengths gives the  $\kappa$ -net  $\mathcal{N}$ , defined for each p as:

$$\mu_{\mathcal{N}}(p) = \min[\mu_{\kappa}(c^{(1)}, c^{(2)}), \mu_{\kappa}(c^{(2)}, c^{(3)}), \dots \mu_{\kappa}(c^{(m-1)}, c^{(m)})]$$
(B.5)

From this concept of path strength we arrive at a measure of connectivity  $\mathcal{K}$  between

spels c and d, simply defined as the strength of the strongest path:

$$\mu_{\mathcal{K}}(c,d) = \max_{p \in P_{cd}} [\mu_{\mathcal{N}}(p)]$$
(B.6)

The algorithm to extract the background region, given the seed spel o (set to the top-left corner), is as follows:

- Inputs:  $C, o, \sigma^2$  and  $\mathcal{K}$
- Output: K<sub>o</sub>-scene  $C_o = (C_o, f_o)$  of C; *i.e.* the scene describing the connectivity of all spels c in C to spel o

Data Structures: an nD array to hold the result  $C_o$  and a temporary queue Q

begin

calculate  $\mu$  of equation B.4; set all elements of  $C_o$  to 0, except for o which is set to 1; push all spels c of  $C_o$  that satisfy  $\mu_{\kappa}(o, c) > 0$  to Q; while Q is not empty do: remove a spel c from Q; find  $f_{max} = \max_{d \in C_o} [\min(f_o(d), \mu_{\kappa}(c, d))];$ if  $f_{max} > f_o(c)$  then: set  $f_o(c)$  to  $f_{max};$ push all spels e such that  $\mu_{\kappa}(c, e) > 0$  to Q; end if

end while

end

The resulting  $K_o$  scene is then thresholded based on a connectivity value, set by the user, to obtain a binary scene representing the background region of the image.

### B.2 Hermite Interpolation

Hermite interpolation is a method of interpolating a polynomial function between two data points. This method requires knowledge of adjacent data points, as curve segments (Hermite polynomials) that share a data point also share tangent vectors.

Given a set of points  $[P_1(t), P_2(t)] = [P(0), P(1)]$  and a pair of tangents  $[P_1^t, P_2^t]$ , the Hermite polynomial P(t) is computed as follows:

$$P(t) = at^{3} + bt^{2} + ct + d, \text{ where}$$

$$a = 2P_{1} - 2P_{2} + P_{1}^{t} + P_{2}^{t}$$

$$b = -3P_{1} + 3P_{2} - 2P_{1}^{t} - P_{2}^{t}$$

$$c = P_{1}^{t}$$

$$d = P_{1}$$

or, in matrix form:

$$P(t) = (t^{3}, t^{2}, t, 1) \begin{pmatrix} 2 & -2 & 1 & 1 \\ -3 & 3 & -2 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} P_{1} \\ P_{2} \\ P_{1}^{t} \\ P_{2}^{t} \end{pmatrix}$$
(B.7)

### B.3 The Elastic Body Spline Transformation

The EBS transform is represented in two dimensions as:

$$\vec{d}(\vec{x}) = [d_1(\vec{x}) \ d_2(\vec{x})]^T$$
 (B.8)

where  $\vec{x} = [x_1 \ x_2]^T$ . At N landmarks or control points, equation B.8 must exactly equal the displacement from the landmark point  $p_i$  to its corresponding landmark  $q_i$  on the other image; elsewhere in the image, the displacement is interpolated. The underlying physical model for the EBS is the set of Navier partial differential equations (PDEs) describing the equilibrium displacements of a homogeneous and isotropic elastic material under load. The PDEs are as follows:

$$\mu \nabla^2 \vec{u}(\vec{x}) + (\mu + \lambda) \nabla [\nabla \cdot \vec{u}(\vec{x})] = \vec{f}(\vec{x})$$
(B.9)

where  $\vec{u}(\vec{x})$  is the displacement of a point from the unloaded to the loaded position;  $\vec{f}(\vec{x})$  is the force field; and  $\mu$  and  $\lambda$  are the Lamé constants describing the properties of the material.

The EBS parameters are determined by calculating an analytical solution to the PDEs of equation B.9 and imposing the constraint that the transform must relax to affine as the distance from the landmarks approaches infinity. To do this, a smooth deformation must be enforced, as a singularity could ensue otherwise. A method of enforcing this constraint is to assume a uniform force field such as:

$$\vec{f}(\vec{x}) = \vec{c}|\vec{x}| \tag{B.10}$$

where the coefficients  $\vec{c} = [c_1 \ c_2 \ c_3]^T$  are the strengths of the force field components.

Solving equation B.9 using the forces of B.10 and representing as a matrix results in a solution to the displacements  $\vec{u}(\vec{x})$ :

$$\vec{u}(\vec{x}) = \mathbf{G}(\vec{x})\vec{c}$$
$$\mathbf{G}(\vec{x}) = [\alpha |\vec{x}|^2 \mathbf{I} - n \vec{x}\vec{x}]|\vec{x}|$$
(B.11)

where *n* is the number of dimensions,  $\alpha = 12(1 - \nu) - 1$ ,  $\nu = \lambda/[2\lambda + \mu]$  is Poisson's ratio. This allows for the Lamé constants to be replaced by a single constant, assumed to be 0.25 for an elastic material.

Equation B.11 was obtained using the Galerkin vector method, conveniently decoupling the Navier PDEs of equation B.9. As a result, the 3D model can be applied to the 2D task at hand. The EBS of equation B.8 is then found by applying the solution B.11 to each landmark point and summing as follows:

$$\vec{d}(\vec{x}) = \sum_{i=0}^{N} \mathbf{G}(\vec{x} - \vec{p_i})\vec{c_i} + \mathbf{A}\vec{x} + \vec{b}$$
 (B.12)

where **G** is now a 2 × 2 matrix,  $\vec{c_i}$  are the spline coefficients, and  $\mathbf{A}\vec{x} + \vec{b}$  is an affine transform accounting for the bulk displacement, rotation, and scaling of the image.

Given the displacements between landmark points, the spline coefficient values  $\vec{c_i}$ and the matrix **A** can be solved computationally at any arbitrary location  $\vec{x}$ , yielding a displacement value for each pixel of the image. This is explained in more detail by Davis *et al* [37].

#### B.4 B-Spline Interpolation

B-splines or "basis-splines" of degree n form the basis of the subspace of all piecewise polynomial functions of degree n and of class  $C^{n-1}$  [38]. B-splines may be computed for any non decreasing sequence of knots  $t = t_i$  [49]; for this work, only the case of equally spaced knots will be addressed.

To construct B-splines of order n, n+2 knots are required. The normalized B-spline functions are then defined as:

$$\beta^{n}(x) = \sum_{j=0}^{n+1} \frac{(-1)^{j}}{n!} {n+1 \choose j} (x-j)^{n} \mu(x-j)$$
(B.13)

where  $\binom{n+1}{j}$  are the binomial coefficients:

$$\binom{n+1}{j} = \frac{(n+1)!}{(n+1-j)!j!}$$

and  $\mu(x)$  is the step function:

$$\mu(x) = \begin{cases} 1 & \text{for } x \ge 0\\ 0 & \text{for } x < 0 \end{cases}$$

Thus, the zero-order B-spline is simply the boxcar function  $\beta^0(x) = \mu(x) - \mu(x-1)$ .

Discrete B-splines are obtained by sampling continuous spline functions with an expansion factor or step size m:

$$b_m^n(k) = \beta^n\left(\frac{k}{m}\right) = \frac{1}{m^n} \sum_{k=0}^{n+1} \frac{(-1)^j}{n!} \binom{n+1}{j} (k-jm)^n \mu(k-jm)$$
(B.14)

Taking the z-transform of equation B.14 yields the expression:

$$B_m^n(z) = \frac{1}{m^n} B_1^n(x) (B_m^0(z))^{n+1}, \text{ where}$$

$$B_m^0(z) = \sum_{k=0}^{n+1} b_1^n(k) z^{-k}$$
(B.15)

is the z-transform of the discrete signal obtained by sampling the B-spline at its knots, and where

$$B_m^0(z) = \sum_{k=0}^{m-1} z^{-k}$$

is the z-transform of a boxcar function of length m.

A discrete signal  $\{f(k)\}$  defined on  $k = -\infty, \ldots, +\infty$  can be represented using B-splines as a weighted sum or convolution:

$$f(k) = \phi^n(k) = \sum_{i=-\infty}^{+\infty} c(i)b_1^n(k-i) = b_1^n * c(k)$$
(B.16)

where n is both the degree of the polynomial functions connected at the knot points and the order of the B-spline.

The goal of B-spline interpolation is to determine the coefficients c(i) of equation B.16 such that  $\phi^n(x)$  matches of the values of  $\{f(k)\}$  at the knot points. Taking the z-transform of equation B.16 results in:

$$F(z) = B_1^n(z)C(z),$$
 (B.17)

implying that c(k) can be obtained from inverse filtering. The corresponding linear space operator  $\{s^n(k)\}$  to obtain these coefficients is called the direct spline filter of order n, and has the transfer function:

$$S^{n}(z) = B_{1}^{n}(z)^{-1} = \frac{1}{\sum_{k=0}^{n+1} b_{1}^{n}(k) z^{-k}},$$
(B.18)

which has an infinite impulse response.

Equation B.18 provides a direct method of obtaining filter coefficients for m = 1. However, to interpolate the signal  $\{f(k)\}$ , upsampling by a factor m is required, producing the new sequence:

$$[f]_{\uparrow m}(k') = \begin{cases} f(k) & \text{for } k' = mk \\ 0 & \text{otherwise} \end{cases}$$
(B.19)

Using a similar representation for c(k) yields the equivalent convolution for equation B.16:

$$f_m(k') = b_m^n * [c]_{\uparrow m}(k')$$
 (B.20)

and finally, the *z*-transform:

$$F_m(z) = B_1^n(z) \frac{1}{m^n} (B_m^0(z))^{n+1} C(z^m)$$
(B.21)

Unser *et al.* show that from equation B.21, signal interpolation can be achieved from a cascade of n + 1 moving average filters of size m and an indirect spline filter  $\{b_1^n(k)\}$  [38]. This has the advantage of having a finite impulse response, and is easily implemented in the spatial domain.