Predicting Tumour Location by Simulating the Deformation of the Breast using Nonlinear Elasticity and the Finite Element Method

Pras Pathmanathan
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2006

To my parents

Abstract

Breast cancer is one of the biggest killers in the western world, and early diagnosis is essential for improved prognosis. The shape of the breast varies hugely between the scenarios of magnetic resonance (MR) imaging (patient lies prone, breast hanging down under gravity), X-ray mammography (breast strongly compressed) and surgery (patient lies supine), rendering image fusion an extremely difficult task. This thesis is concerned with the use of the finite element method and nonlinear elasticity to build a 3D, patient-specific, anatomically-accurate model of the breast. The model is constructed from MR images and can be virtually deformed to simulate breast shape and predict tumour location during mammography or surgery. Two extensions of the standard elasticity problem need to be solved: an inverse elasticity problem (arising from the fact that only a deformed, stressed, state is known initially), and the contact problem of modelling compression. Both modelling and numerical analysis issues are addressed for each of these types of problem. In the latter part of this thesis, we consider how to model skin in an efficient and realistic manner as a pre-stressed membrane.

Acknowledgements

Firstly, I would like to thank Professor David Gavaghan for all his help and advice over the course of my PhD, and for being such a relaxed but motivating supervisor. I would also like to thank Professor Jon Chapman for providing so much help with the mathematical chapters in this thesis, and I also need to thank Dr Jon Whiteley, for being the person to see on any question about finite elements, or for a discussion about football.

In New Zealand, I am indebted to Dr Martyn Nash, for allowing me to spend the formative months of the PhD in Auckland and spending a lot of his time helping with the modelling during an important period. I am also grateful to Vijay Rajapogal there, for his assistance throughout.

Back in Oxford, I must also thank Professor Mike Brady, Mike Kadour, Dr Andy Wathen, Professor Endre Süli and Ralph Highnam for their help at various times during the PhD. Finally, I would like to express my thanks to my family and friends, who have made my time in Oxford such an enjoyable one.

Contents

1 Introduction **1**

2 Motivation and Overview **4**

2.1 Breast Anatomy 4
2.2 Imaging Modalities for the Breast 5
  2.2.1 Mammography 6
  2.2.2 Magnetic Resonance Imaging 7
  2.2.3 Ultrasound 9
2.3 Motivation for a Deformable Model of the Breast 10
  2.3.1 The Prone-to-Supine Problem 10
  2.3.2 Matching of MR Images with Mammograms 11
  2.3.3 CC-MLO Matching of Mammograms 12
  2.3.4 Other Uses 14
2.4 Modelling the Deformation of the Breast and Previous Efforts 14
  2.4.1 Segmentation and Edge Detection 15
  2.4.2 Mesh Generation 15
  2.4.3 Elasticity Theory 16
  2.4.4 Determining Material Laws 16
  2.4.5 Removing Gravity to Obtain the Reference State (the Backward Problem) . . 17
  2.4.6 Simulations (the Forward and Contact Problems) 17

3 Nonlinear Elasticity Theory **19**

3.1 Stress and Strain 20
  3.1.1 The Deformation Gradient, F 21
  3.1.2 Measures of Strain 23
  3.1.3 Stress 25
3.2 The Equation of State and Weak Formulation 27

3.2.1 Equations of Motion 28

3.2.2 Boundary Conditions 30

3.2.3 The Weak Formulation 31

3.3 The Constitutive Law and Strain Energy Functions 32

3.3.1 The General Form of the Constitutive Relation 34

3.3.2 The Strain Energy Function 34

3.4 Incompressibility and Pressure 38

3.5 Linear Elasticity and Linear Material Laws 39

4 The Finite Element Method for Solving Problems in Nonlinear Elasticity 42

4.1 Introduction to the Finite Element Method 44

4.1.1 Nodes, Elements and Basis Functions 44

4.1.2 Weak Formulations and the Finite Element Formulation 49

4.1.3 Linear System Assembly for Linear FE Problems 51

4.1.4 Nonlinear FE Problems and Newton's Method 52

4.1.5 Numerical Integration and Gaussian Quadrature 54

4.1.6 Solution of Linear Equations 54

4.2 The FE Method for Nonlinear Elasticity 56

4.2.1 FE Assembly for Nonlinear Elasticity Problems 56

4.2.2 Methods of Improving Newton-Convergence 60

4.2.3 Preconditioning by Non-dimensionalisation 61

4.3 Results 63

5 The Backward Problem for Computing the Reference State 71

5.1 Analysis and Reformulation 72

5.1.1 The Reference State and the full Backward Problem 72

5.1.2 Potential Numerical Methods for Solving the Backward Problem 73

5.1.3 The Weak Form and Finite Element Equations using the Direct Approach . . 75

5.2 Numerical Results and Discussion 77

5.2.1 The Backward Problem with a Transversely-Isotropic Material Law 78

5.3 GMRES on the Forward and Backward Problems 80

6 The Contact Problem 88

6.1 Introduction to the Contact Problem 88

6.2 Mathematical Analysis 90

6.3 Numerical Algorithms for Contact 94

6.3.1 The Penalty Method 96

6.3.2 The Augmented Lagrangian Method 97
6.4 Results and Discussion 98
6.4.1 Comparison of the Penalty Method and the Augmented Lagrangian Method on a Simple Problem 98
6.4.2 Results with 3D Finite Elasticity and Contact 102

7 Implementation of the Model of the Breast and Numerical Experiments 106
7.1 Pre-Computational Steps 107
7.1.1 Image Segmentation 107
7.1.2 Mesh Generation 112
7.2 Implementation of the Full Model 120
7.2.1 Modelling Assumptions 120
7.2.2 Mechanical Properties of Breast Tissue 123
7.3 Simulations using The Full Model 126
7.3.1 Computing the Reference State and the Prone-to-Supine Problem 127
7.3.2 CC and MLO Simulation and Mammographic Matching 128
7.4 Numerical Experiments 136
7.4.1 Sensitivity of the Model to Material Laws and Parameters 136
7.4.2 Effect of Skin on the Model 139
7.4.3 Effect of a Hard Tumour on the Model 141
7.5 Discussion and Further Work 141

8 Modelling Skin as an Elastic Membrane 146
8.1 Introduction 146
8.2 Membranes and Thin Elastic Solids in 2D 147
8.2.1 The Equations of a 1D String 147
8.2.2 The Thin Elastic Solid in Two Dimensions 151
8.3 The Thin Elastic Solid in Three Dimensions 158
8.3.1 Curvilinear Coordinate Systems 158
8.3.2 Stress Analysis via Energy Minimisation 162
8.3.3 Analysis of Specific Problems 165
8.3.4 Analysis of the Governing Equations 167

9 Membrane Results and the Backward Problem for Skin 169
9.1 Numerical Results 170
9.1.1 The Finite Element Implementation 170
9.1.2 Results 173

9.2 The Backward Problem for Skin . . . 176
  9.2.1 Introduction . . . 176
  9.2.2 The Backward Problem for Skin using the Fluid Filled Membrane Model . . 178
  9.2.3 Existence of an Undeformed State . . . 181
  9.2.4 Classification . . . 182
  9.2.5 The Axi-symmetric Case . . . 184
  9.2.6 Discussion on the Backward Problem using a Fluid-Filled Membrane . . . 188

10 Conclusions **192**
10.1 Summary . . . 192
10.2 Further Work . . . 194
10.3 Conclusion . . . 196

Bibliography **197**

Glossary of Symbols

Elasticity

$\mathbf{x}, x^i$	deformed position
$\mathbf{X}, X^M$	undeformed position
Ω_0	undeformed body
Ω	deformed body
$\mathbf{u}$	displacement, defined as $\mathbf{x} - \mathbf{X}$
F, F^i_M	deformation gradient, defined as $\dfrac{\partial x^i}{\partial X^M}$
J	deformation Jacobian, defined as $\det(F)$
C, C_{MN}	Green deformation tensor, defined as $F^i_M F^i_N$
E, E_{MN}	Green strain tensor, defined as $\dfrac{1}{2}(C_{MN} - \delta_{MN})$
I_1, I_2, I_3	principal invariants of C
σ, σ^{ij}	Cauchy stress tensor
S, S^{Mi}	1st Piola-Kirchoff stress tensor
T^{MN}	2nd Piola-Kirchoff stress tensor
$\mathbf{N}, \mathbf{n}$	undeformed and deformed normals
ρ_0, ρ	undeformed and deformed densities
$\mathbf{b}$	prescribed body force
$\mathbf{s}$	prescribed surface tractions
$\mathbf{u}^{\text{disp}}$	prescribed surface displacement
$\partial\Omega_0^{\text{disp}}$	part of boundary with displacement boundary conditions
$\partial\Omega_0^{\text{trac}}$	part of boundary with traction boundary conditions
$W, \bar{W}$	strain energy function and material-only part of strain energy
$\mathcal{E}$	total energy
p	internal pressure
$\mathbf{M}, \mathbf{m}$	undeformed and deformed fibre directions
K_1, K_2	strain measures in fibre direction
ϵ_{ij}	linear elasticity infinitesimal strain tensor, defined as $\dfrac{1}{2}(u_{i,j} + u_{j,i})$
λ, μ	linear elasticity Lamé coefficients
E, ν	linear elasticity Young's modulus and Poisson's ratio

Finite Element Method

$\mathbf{X}_i$	nodes
K_i	elements
$\mathcal{N}$	number of nodes (or number of unfixed nodes)
$\mathcal{M}$	number of elements
Ω_0^h	computation domain
$\boldsymbol{\xi}, (\xi, \eta, \nu)$	a point in the canonical element
N_i	basis functions on the canonical element
ϕ_i	basis functions in physical space
$\mathbb{1}_{K_i}$	indicator function (i.e. piecewise constant basis function) for element K_i
$\mathbf{U}$	vector of unknown x-displacements at nodes
$\mathbf{V}$	vector of unknown y-displacements at nodes
$\mathbf{W}$	vector of unknown z-displacements at nodes
$\mathbf{P}$	vector of unknown pressures at elements
$\mathbf{a}$	vector of finite element unknowns, defined as $(\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{P})$
$\mathcal{F}$, LoadVec	finite element residual
$\mathcal{J}$, Jac	finite element Jacobian

Contact Problem

$\mathbf{t}$	surface force
t_N	magnitude of normal surface force, defined as $\|\hat{\mathbf{n}}^{\mathrm{T}}\sigma\hat{\mathbf{n}}\|$
$\Omega_0^{(1)}, \Omega_0^{(2)}$	undeformed slave and master bodies
$\Gamma^{(1)}, \Gamma^{(2)}$	deformed slave and master contact surfaces
$\mathbf{x}^{(1)}$	deformed point in the slave (usually on the surface)
$\bar{\mathbf{x}}^{(2)}$	nearest point on master surface to $\mathbf{x}^{(1)}$
$\mathbf{d}_N$	distance vector, defined as $\mathbf{x}^{(1)} - \bar{\mathbf{x}}^{(2)}$
d_N	gap function, the signed norm of $\mathbf{d}_N$
λ_N	constraint Lagrange multipliers, equal to t_N
Σ	constraint penalisation function
P	penalty parameter
$\lambda_N^{(k)}$	estimates of λ_N in Augmented Lagrangian method

Membrane Theory

Ω_0, Ω	undeformed and deformed skin/thin solid
$\Omega^{(f)}$	fluid interior
L_0	length of the string/typical large length-scale
s	undeformed arclength in the string
τ	deformed arclength in the string
$\mathbf{t}$	non-normalised tangent vector of string/mid-surface
$\mathbf{n}$	non-normalised outward normal
W	strain energy of the string/thin solid
$\mathcal{K}$	remaining energy in system excl. strain energy
T	tension in the string
κ	curvature of a 1D curve in 2D space
λ	Lagrange multiplier corresponding to fluid incompressibility
P	pressure corresponding to fluid incompressibility, equal to $-\lambda$
$\rho^{(s)}, \rho^{(f)}$	density of skin and of fluid interior
ϵ	small parameter, thickness divided by typical length
h_0, h	initial and deformed thickness of membrane
$\mathbf{h}$	vector measuring normal deformation of membrane
$\mathbf{x}^{(k)}$	coefficient of ϵ^k in expansion of $\mathbf{x}$
$\tilde{\Theta}$	parameter space of the 3D thin solid
Θ	parameter space for the mid-surface of the thin solid
$\boldsymbol{\theta}$	a member of the parameter space
Γ_0, Γ	undeformed and deformed mid-surfaces of the thin solid
$\tilde{\mathbf{A}}_\alpha, \tilde{\mathbf{a}}_\alpha$	undeformed and deformed volume basis vectors
$\mathbf{A}_\alpha, \mathbf{a}_\alpha$	undeformed and deformed surface basis vectors
$\tilde{A}_{\alpha\beta}, \tilde{a}_{\alpha\beta}$	undeformed and deformed 3D volume metrics
$A_{\alpha\beta}, a_{\alpha\beta}$	undeformed and deformed 2D surface metrics
$\tilde{A}^{\alpha\beta}, \tilde{a}^{\alpha\beta}, A^{\alpha\beta}, a^{\alpha\beta}$	inverses of corresponding metrics
$\tilde{A}, \tilde{a}, A, a$	determinants of corresponding metrics
$\tilde{\mathbf{A}}^\alpha, \tilde{\mathbf{a}}^\alpha, \mathbf{A}^\alpha, \mathbf{a}^\alpha$	dual basis vectors to corresponding basis

$\Gamma^{\gamma}_{\alpha\beta}$	Christoffel symbols (on the deformed mid-surface), equal to $\dfrac{\partial \mathbf{a}_\alpha}{\partial \theta^\beta} \cdot \mathbf{a}^\gamma$
$b_{\alpha\beta}$	curvature tensor, defined as $\dfrac{\partial \mathbf{a}_\alpha}{\partial \theta^\beta} \cdot \hat{\mathbf{n}}^{(\text{inward})}$
κ_1, κ_2	principal curvatures (eigenvalues of b^{α}_{β})
$\tilde{T}^{\alpha\beta}$	3D generalised stress tensor
$T^{\alpha\beta}$	2D generalised in-plane stress tensor
W^{mem}	2D membrane strain energy corresponding to 3D strain energy
$F^{\alpha\beta}$	in-plane force tensor (stress residual)

Notation

- Unless explicitly stated otherwise, summation convention (where repeated indices imply summation over the index, i.e. $a_i b_i \equiv \sum_{i=1}^{N} a_i b_i$) is assumed throughout this thesis.
- Derivatives with respect to vectors or matrices represent the vector or matrix of componentwise derivatives, i.e. if $\mathbf{t}$ is a vector and W a scalar function of $\mathbf{t}$, $\frac{\partial W}{\partial \mathbf{t}}$ represents the vector whose i-th component is $\frac{\partial W}{\partial t_i}$.
- Vectors with a $\hat{}$ are assumed to be the normalised version of the original vector, for example if $\mathbf{t}$ is a vector, $\hat{\mathbf{t}}$ is assumed to be $\frac{\mathbf{t}}{\|\mathbf{t}\|}$.
- Commas represent derivatives, for example, if $x \equiv x(s)$, then $x_{,s} = \frac{\mathrm{d}x}{\mathrm{d}s}$.

Chapter 1

Introduction

Breast cancer is one of the biggest killers of women in the Western world, killing about 400,000 and affecting over a million people each year worldwide. The second most common cancer in women, and the second leading cause of cancer-related deaths in women[1], breast cancer will be diagnosed in an estimated one in seven women at some point in their lifetimes. As with all cancers, treatment can be lengthy, extremely unpleasant, and ultimately fruitless. Early diagnosis is the most effective way to improve prognosis, and for this reason screening programmes have been established in many western countries.

Currently, the imaging modalities used to image the breast are mammography, magnetic resonance (MR) imaging, and ultrasound. Mammograms (X-rays) are most commonly used, and the basis of most national screening programmes. They are used to obtain projected 2D, high resolution images of the compressed breast. MR imaging can be used to produce a set of 2D images of the different 'slices' of the breast, enabling a 3D picture to be formed, although at a lower resolution than mammography.

There has been a great deal of recent interest in the use of mathematics and computing to analyse medical images, the aim being to produce software that will aid detection or diagnosis, or for use as surgical guides. One such area is the modelling of the deformation of the breast. Due to the highly deformable nature of the breast, and the difference in body position and external forces applied between different imaging procedures, the shape of the breast varies massively between images of different types, and between the images and the shape of the breast during surgery. Our goal is to use data from MR images to build a patient-specific anatomically-accurate model of the breast, which can be accurately deformed under prescribed body forces or loading conditions, and predicts the shape and internal structure of the breast during surgery or mammography. Such a model would

---

[1] The most common cancer in women is nonmelanoma skin cancer, whereas the leading killer is lung cancer.

have a plethora of clinical uses, ranging from image-registration to surgical guidance to simple 3D visualisation. Two of the most important applications are that a model of breast-deformation would allow radiologists to better compare two different types of mammogram, and act as a guide to aid minimally-invasive surgery and biopsy.

Various methods have been proposed in the literature to model breast deformation. Most are either based on linear elasticity theory, the theory used for small strains (derivatives of displacement), or on non-physical models. However, the deformations undergone by the breast do not involve just small strains, especially when the breast is compressed during mammography, so the full nonlinear theory of elasticity is required to model the deformation accurately. The equations of nonlinear elasticity over irregular three-dimensional domains are generally too complex to be solved analytically, instead requiring the use of a numerical method for computing approximate solutions, of which the finite element (FE) method is especially well-suited. In this thesis, we describe the development of a finite element model of the breast based on nonlinear elasticity theory. The FE mesh has a geometry determined from MR images of the subject, with anatomical structure that is determined from segmented MR data. As such, we propose a patient-specific, not generic, model. Such a method models the deformation in a physically realistic manner, hence has a significantly higher likelihood of accurately simulating breast deformation than a linearised or non-physical method. Modelling large deformations with nonlinear elasticity has previously been considered computationally intractable; however, as we shall demonstrate, simulations on a moderately high-resolution mesh can be computed in a reasonable length of time on a standard modern PC.

The modelling of the deformation using nonlinear elasticity and the FE method forms the heart of the project, but there are a number of other important stages in this problem, which include image segmentation, mesh-generation, two re-workings of the standard elasticity problem, and biomechanical experimentation. There are two aims of this thesis. The first is to present an overview of all stages which need to be completed, so that a full idea of the complexity of the project is gained. We will consider, to lesser and greater extents, all the sub-problems that must be completed for a clinically-useful model of the breast to be possible. The second aim is to perform an in-depth study of the mathematical modelling and numerical analysis stages, and it will be in these types of problems that most of our work will be concentrated.

We begin in Chapter 2 by providing an introduction to breast anatomy and the imaging modalities used for the breast, as well as motivating the need for a deformable model. We also survey previous work modelling breast deformation, and give an introduction to each of the stages that will be encountered in this thesis.

The foundation for the modelling is laid down in Chapters 3 and 4, where we formulate the theory of nonlinear elasticity (Chapter 3) and describe how the finite element method can be used to solve the resulting equations (Chapter 4).

In Chapter 5, we consider how to solve an inverse elasticity problem (which we dub the *backward problem*), one which arises immediately when the deformation of the breast is tackled. The backward problem is fundamentally different to standard nonlinear elasticity problems (which we call *forward problems*), and is rarely encountered in the literature. This chapter deals with the mathematical formulation of the backward problem, numerical methods to solve it, and numerical analysis issues that arise from it.

Chapter 6 deals with changes that need to be made when the compression of the breast during mammography is modelled, which, again, is a fundamentally different problem to the forward problem. This is known as the *contact problem*, because of the unknown region of contact between the skin and the compression plates, and requires further mathematical reformulations to be stated and more sophisticated numerical techniques to be solved computationally.

In Chapter 7, the major chapter in this thesis, we bring together the theory and numerical methods from the preceeding chapters to simulate breast deformation. The first part of Chapter 7 is concerned with the building of the mesh of the breast from MR images, which involves image segmentation (assigning a tissue to each pixel in the set of images), and mesh generation. Here, we also present a survey of bio-mechanical experimental work carried out in the literature on mechanical properties of breast tissues. The second half of the chapter involves simulations using the model, where we simulate the supine and compressed breast shape, use the model to match different types of mammograms, and run a number of numerical experiments using the model.

In Chapters 8 and 9 we move away from the full problem and focus on one particular aspect of the model that is a source of problems, the manner in which the skin is modelled. Initially, skin will be modelled as a thin elastic solid, whereas in Chapter 8 we perform a rigourous mathematical analysis to derive equations for modelling the skin as a membrane. Finally, Chapter 9 deals with results using membrane skin, and a new backward problem of determining stresses in the skin so that they may be included in the full model.

Chapter 2

Motivation and Overview

In this chapter we set up the problem which is to be investigated in this thesis. We first describe the structure of the breast. Our model will be an anatomically-accurate model of the breast, and the anatomical structure which we choose to include and that which we choose to ignore can potentially have far-reaching implications on the model's performance in accurately predicting deformations. In Section 2.2, we survey the three imaging modalities used for the breast, paying particular attention to mammography and MR imaging. We will see that the uses of a deformation model are both obvious and potentially extremely valuable, and these are discussed in Section 2.3. Finally, in Section 2.4, we consider the different stages involved in the problem, whilst discussing previous efforts at modelling the deformation of the breast in the literature.

2.1 Breast Anatomy

Figure 2.1 illustrates the structure of the adult female breast. The interior of the breast is made up of three major tissue types, parenchymal (functional) or glandular tissue, fibrous tissue and fat. The parenchymal tissue of the breast is compartmentalised into 15-20 lobes, which is where milk-production occurs. The lobes are separated by fibrous walls and can be further divided into smaller compartments called lobules. The lobules are connected to the nipple by a network of ducts. The ducts begin as tiny ductules at the lobules, which merge together as they near the nipple into larger and larger ducts. The ducts and lobules are contained and supported by fibrous or connective tissue. We can group all this glandular (lobes and ducts) and fibrous tissue together as *fibroglandular tissue*.

Outside the fibroglandular tissue is a layer of *adipose* or *fatty tissue*. The proportion of fibroglandular to adipose tissue varies between individuals and with age. There is always a significant proportion of fatty tissue, with even small breasts rarely being composed of less than a third fat.

After menopause, the fibroglandular tissue atrophies, resulting in a much higher proportion of fat to fibroglandular tissue.

Beneath the breast is a large muscle, the *pectoralis major* (pectoral muscle), which separates the breast from the rest of the body. The pectoral muscle is the muscle that is contracted when the hands are placed on the hips and the elbows thrown back. The pectoral minor is found between the pectoral major and the ribs. The whole of the breast is supported by fibrous strands called *Cooper's ligaments*, which are attached to the muscle wall. Cooper's Ligaments strengthen the breast significantly, and weakening of these ligaments can cause the breast to sag. The whole of the breast is also supplied with blood vessels. In this thesis we assume that the breast has not been augmented with silicon implants.

There are two other major tissue types which we will need to consider. The *skin* is a layer of highly elastic tissue and is generally 1-2 millimetres thick. The final possible tissue type is *tumour tissue.*

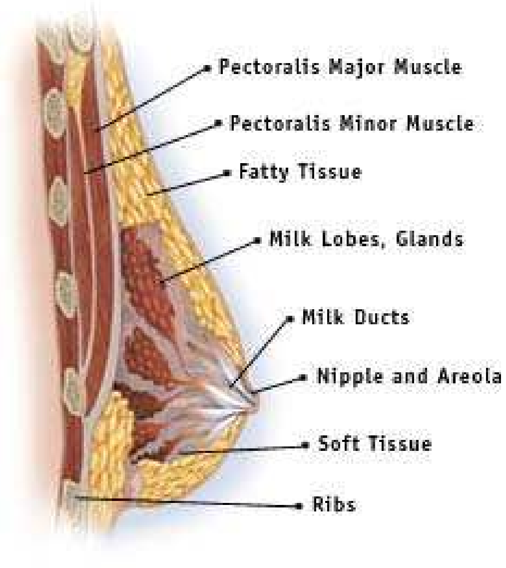

Figure 2.1: The breast structure.

2.2 Imaging Modalities for the Breast

Each of the three main imaging modalities used for breast cancer have their own particular advantages and disadvantages. We now describe the scanning procedures of the three modalities, and discuss reasons for their use.

2.2.1 Mammography

In mammography, a 2D image of the breast is formed by passing X-rays through the breast. Absorption of X-ray photons varies between tissues, so the grey-level of a pixel in a mammogram gives an indication of the proportion of different tissues in the column of the breast that the X-rays corresponding to that pixel passed through. Therefore a mammogram is a *projected* 2D image of the breast, rather than a true 2D image of a slice of the breast (compared with an MR image, discussed shortly, which is a true 2D image of a slice of the breast).

Where no absorption of X-ray photons has occurred (such as outside the breast) the mammogram appears black, and where total absorption has occurred it will appear white. Fibroglandular tissue is highly X-ray absorbent, whereas fat is not, so the fibroglandular region of the breast will appear brighter in mammograms. Because of this, mammography as a method of screening is only performed on older women, younger women tend to have dense breasts with a lot of fibroglandular tissue, and X-ray images of their breasts can sometimes turn out to be completely white. Typically, the lower age limit chosen for screening programs is about 50 years, an estimate for the onset of menopause. After menopause fibroglandular tissue is replaced by fatty tissue, resulting in better contrast mammograms.

For women for whom good quality mammograms can be taken, useful information can be gathered using mammography. Mammograms have the highest resolution of the three imaging modalities, and the small-scale structure of the breast can be visualised. In particular, this means that *micro-calcifications* can sometimes be spotted. These are small deposits of calcium salts that are often the earliest indicators of the presence of a tumour. If unobscured, they appear in X-ray images as small regions of very high contrast.

There are various angles from which the X-ray image can be taken, the most common being *cranio-caudal* (CC), where the direction of the X-ray photon travel is head-to-toe, and *medio-lateral oblique* (MLO), which is shoulder-to-opposite-hip. Less common views are medio-lateral (ML) (centre of the chest outwards to the side), and latero-medial (LM) (the opposite direction to ML). These four views are illustrated in Figure 2.2. Screening mammography typically involves two views of the breast, usually CC and MLO. The use of a combination of views allows the radiologist to study the breast anatomy from different angles and come to a more informed judgement.

When mammograms are taken the breast is compressed heavily by two plates with normals in the direction on the X-ray photon travel. For example, in CC mammography, the breast is compressed in the vertical direction. Compression is carried out for two reasons: firstly, it spreads out the tissue, allowing the radiologist to see breast structures more clearly, and secondly, it reduces the distance between X-ray source and receptor, so that a smaller dose can be used.

The imaging procedure for CC mammography is shown in Figure 2.3. Note how the amount of compression is quite substantial, and how the compressed breast shape is very different to the undeformed shape. Example CC and MLO mammograms are shown in Figure 2.4.

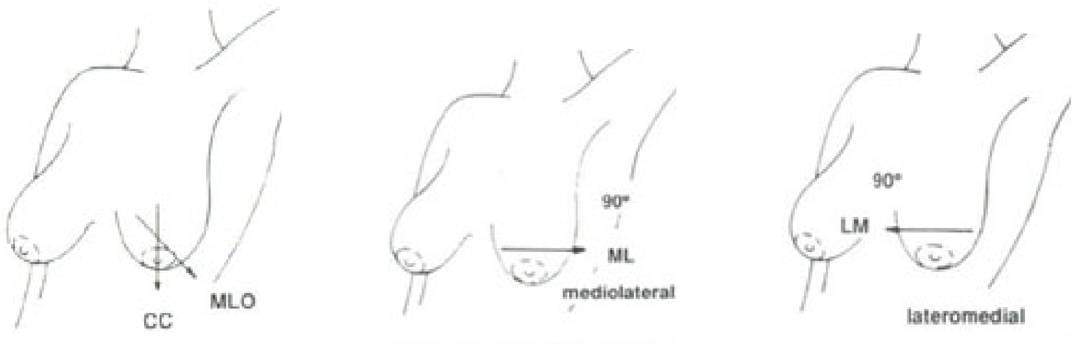

Figure 2.2: The CC, MLO, ML and LM views in mammography

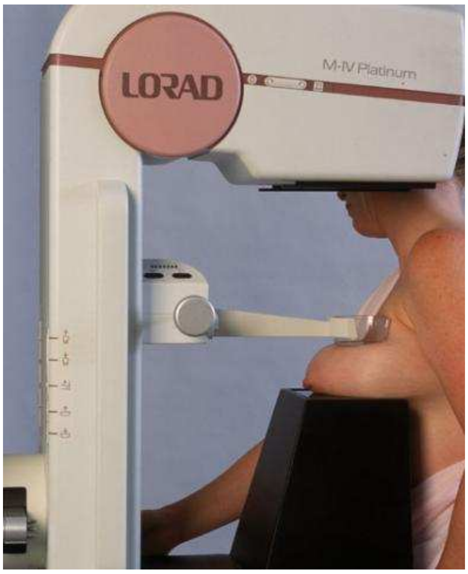

Figure 2.3: The CC mammographic imaging procedure.

2.2.2 Magnetic Resonance Imaging

The magnetic resonance phenomenon was discovered independently by Edward Purcell and Felix Bloch in 1946, and resulted in the pair sharing a Nobel Prize in 1952. Atoms with nuclei with odd numbers of protons and neutrons were known to, in the presence of a static magnetic field, tilt at a certain angle and precess at a certain speed; they discovered that the precession will intensify under

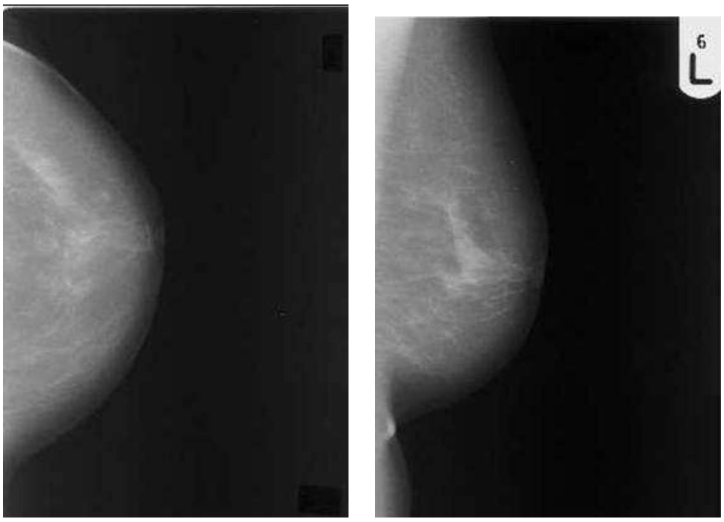

(a) (b)

Figure 2.4: (a) A CC mammogram, (b) an MLO mammogram.

an electromagnetic field at a precise frequency, known as the resonance frequency. The response to the field varies between atoms and is affected by other atoms in the same molecule, which can be exploited to obtain structural information on a molecule. These discoveries were originally put to use in analytic chemistry and biochemistry, in studies determining the composition of chemical compounds, and was known as NMR (nuclear magnetic resonance) spectroscopy. In the 1970s it was discovered that with the inclusion of a weak gradient magnetic field, together with the strong magnetic and electromagnetic fields, the MR phenomenon could provide spatial information, which eventually led to its use as a means for non-invasive tissue imaging. Clinical use began in the 1980s, for which the procedure was renamed magnetic resonance imaging, the word ‘nuclear’ being dropped because of its obvious negative connotations.

The MR machine is a large machine with a cylindrical hole to contain the patient. For breast MR imaging, the patient lies prone (face down) on the scanning table, with the breasts hanging pendulously in depressions that contain the breast coils (the equipment used to detect the resonance signal). The table is moved into the MR machine, which houses the magnet, and imaging can then takes place, a procedure which can last 10 to 30 minutes. 2D images can be taken of different ‘slices’ of the breast, and a 3D picture obtained. The images can be taken in any desired plane. Figure 2.5 displays the MR machine and the position of the patient, and some sample MR images can be seen in Figures 2.6 and 2.7.

There are two methods which can be used to search for tumours. One is to simply search for

Figure 2.5: A patient lying on the scanning table outside the magnet.

suspicious masses in the images, by eye or using image analysis algorithms. The second is a technique that can be used to search for vascular tumours.[1] If a contrast-enhancing agent (a chemical with significantly different MR properties to normal tissue, which appears very bright in the images) is injected into the blood and MR images taken at different times, blood vessels will appear to get lighter and darker as blood density (and thus contrast-enhancing agent density) varies. Similarly, masses with their own blood supply, such as vascular tumours, will also change intensity. Vascular tumours often show up extremely clearly in contrast-enhanced MRI.

The main obvious advantage of MR scans is that they provide three-dimensional data on the breast. Each image is a true 2D image of that slice, not a projected 3D onto 2D image as with mammograms. It can be used equally effectively for younger (pre-menopausal) women as with older women. However the resolution is a lot lower than mammograms, and structures such as microcalcifications cannot be resolved. There are also practical disadvantages: MR imaging is expensive, and there are relatively few MRI machines in the world. Also, the process can be mildly unpleasant, and due to the huge magnetic fields involved, not possible on patients with, for example, pacemakers or metallic implants. Since the imaging procedure takes some time, MR images are highly affected by blurring and other motion-related artifacts, motion due to both chest movement for breathing and shifts in position by the patient.

2.2.3 Ultrasound

Ultrasound imaging is performed with the breast in another new position to that of MRI or mammography. Here, the patient sits up, with the breast falling to their side. The ultrasound probe is

---

[1] When tumours first form, they do not have their own blood supply and are known as avascular tumours. As they grow, they release chemical agents which cause blood vessels to branch toward to the tumour, eventually leading to the tumour having its own capillary network, in which case it is known as a vascular tumour. The transition process is known as angiogenesis.

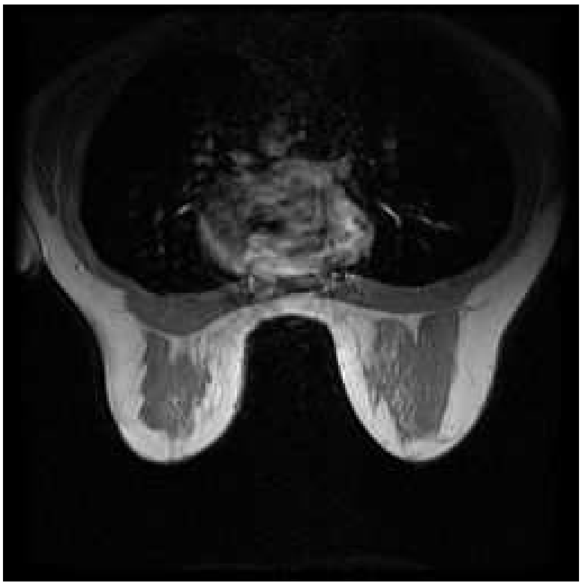

Figure 2.6: An MR image of a prone patient, the image taken in the horizontal (transverse) plane (normal in the head-to-toe direction). The fat (light coloured), fibroglandular tissue (darker coloured) and muscle are clearly visible, as is the heart.

pushed gently against the breast and dragged along the skin.

Ultrasound is the least useful of the three imaging techniques as ultrasound images are very poor quality, and is not suitable as a screening tool. It can be used as an alternative method for pre-menopausal women if MR imaging is not possible. However, there is one advantage unique to ultrasound: since the probe is pressed against the skin, the effect of different levels of compression on the breast can be seen. Tumours are generally much harder than surrounding tissue, and a method of identifying tumours is to look for masses that retain their shape under compression. Benign masses tend to be softer, so ultrasound is a useful diagnostic tool.

2.3 Motivation for a Deformable Model of the Breast

Having looked at the different imaging modalities, we are now ready to describe the possible uses of a deformable model of the breast.

2.3.1 The Prone-to-Supine Problem

The simplest example of a use of the model is what we will call the *prone-to-supine problem.* When MR images are taken of a patient, they are generally carried out with the patient prone (i.e. lying on their front), with the breast hanging down under gravity. During surgery (in particularly during

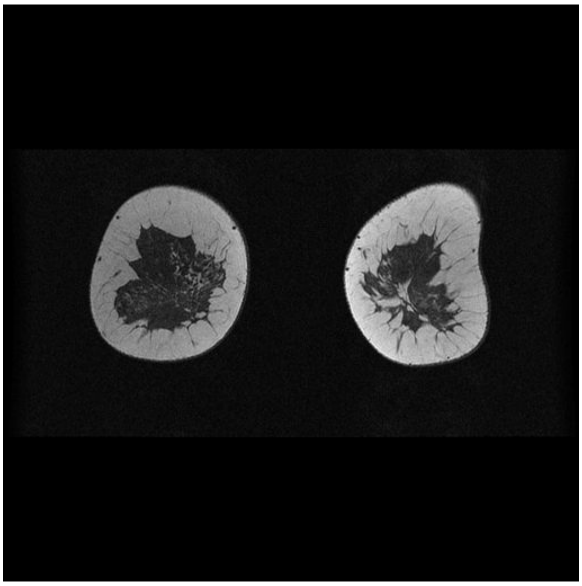

Figure 2.7: An MR image of a prone patient, the image taken in the coronal plane (normal in the front of body to back of body direction)

a lumpectomy procedure, where a tumour is removed by the surgeon), or during open biopsy (where a sample of a lump is removed by the surgeon for examination by a pathologist) the patient will be in the supine position (lying on their back). A surgeon using MR images or mammograms as a guide will have to use some degree of guesswork when locating the tumour, which hinders minimally-invasive surgery. A deformable model of the breast could be used as guide for the surgeon: the model would be built from the original MR images, and the direction of gravity reversed to simulate the breast during surgery.

Similarly, the model can be used as a guide during fine-needle or core biopsy (where a small sample of the tumour is removed for examination using a needle, a thin needle in the case of fine-needle biopsy and a larger needle in core biopsy). Needle biopsy is a very accurate test provided enough correct material is removed, but since the biopsy has to be guided by the images, it is not always easy to get a good sample. A model of the breast could be used as a guide during biopsy, and in fact if the model was highly accurate it could be used to automate the biopsy procedure.

2.3.2 Matching of MR Images with Mammograms

MR data is very useful because it gives an accurate picture of the 3D geometry and structure of the breast, but mammograms are also vitally important, since they can be used to study the fine structure of the breast, especially when used to spot small objects like micro-calcifications which cannot be seen in MR images. Therefore there is great deal of value in being able to merge the data together.

Current methods to do this rely on direct *image registration*, where structures from one image are identified with the same structures in another. For example, in [BMA$^+$00] the following procedure is used: first a 2D image is formed from MR data that is the (approximate) equivalent image of a mammogram, by projecting the MR image in the direction of the X-ray photon travel. The registration is then performed in two parts: a partial registration where the boundaries of the mammogram and the projected MR image are registered, and then the complete registration where the interiors of the images are registered. In both stages, a selection of landmarks (objects that are identifiable in both images) are chosen, and a function computed which maps the landmarks from one image onto the other. The problem with this method is that for high accuracy in matching an arbitrary part of an image, a large number of landmarks need to be found. Apart from the nipple, obvious choices of landmark are not readily available. The huge difference between the breast shape in MR images and mammograms means the projected MR image will not match the mammograms closely, so the registration will be a ‘major’ registration (the mapping from one image to another will have to move points by substantial amounts).

An alternative and potentially better method would be to use a deformable 3D model of the breast to deform the MR images of the uncompressed breast to that of the compressed breast. If this was calculated exactly, then by projecting the (compressed) MR images in the mammographic direction we would obtain an image which precisely matches the mammogram, for which no registration is needed. Even if the model of compression was not perfectly accurate, the projected compressed MR image would still more closely match the mammogram than the projected uncompressed MR image, so the registration required would then only be a minor registration, and therefore more likely to be accurate.

2.3.3 CC-MLO Matching of Mammograms

When comparing two mammograms taken from different angles, the radiologist needs to be able to say which part of the image in one mammogram corresponds to a given point in the other mammogram. A point in one mammogram, the CC image say, will have an equivalent curve in the other mammogram, the MLO image say, that is the image of the same part of the breast. Radiologists currently use very simple methods to estimate this equivalent curve, for example by guessing that the curve which corresponds to a given point $\mathbf{x_0}$ in the CC image is the set of points in the MLO image whose distance from the nipple is the same as the distance from $\mathbf{x_0}$ to the nipple in the CC image.

Given only the two mammograms, automatic determination of the equivalent curve to a given point involves 5 stages, which are illustrated in Figure 2.8, and are described below:

1. Un-projection: finding the line in the (compressed) breast that mapped onto that point in the CC image when the X-ray image was taken.
2. Un-compression: the formation of the 3D uncompressed breast from the 3D compressed breast, in which case the line will deform to a curve.
3. Rotation: rotation of the 3D breast to the MLO angle.
4. Compression: compression of the 3D breast in the new angle, so that the curve deforms again.
5. Projection: projection of the 3D compressed breast to obtain the MLO image, to get the solution curve.

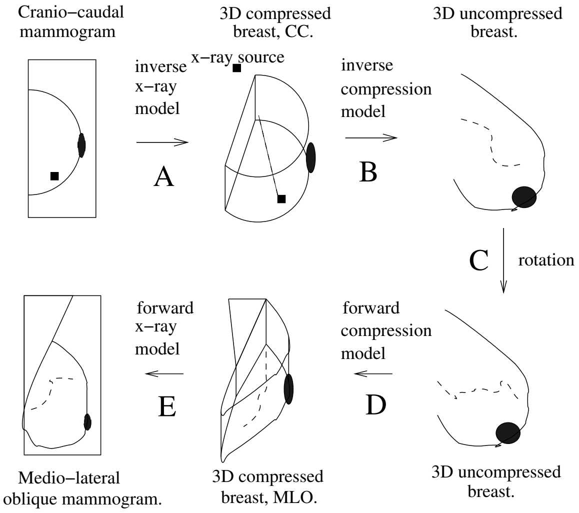

Figure 2.8: The stages needed in matching a point in the CC image to a curve in the MLO image.

Step 3 is obviously trivial. Stage 5 is also trivial, we just project the curve from the X-ray source onto the MLO image. If we just have CC and MLO data, carrying out Stage 1 is made difficult by the fact that some sort of breast shape has to assumed. Stages 2 and 4 involve modelling the compression of the breast and are clearly non-trivial.

Current methods used to tackle this problem are wholly geometric, and not patient specific. For example, in [KHB01] assumptions are made on the general shape of the breast, and how lines map under compression. The breast edge is assumed to take the shape of a quadratic curve (one quadratic curve above the nipple and another quadratic curve below the nipple), and these curves are assumed to map to new quadratic curves under compression.

To use the deformable model of the breast on this problem the procedure would be as follows. We begin with MR data of the patient, from which we build the 3D model of the breast. We then simulate the compression the breast undergoes for both CC and MLO mammography, i.e. we compress the model in the CC direction, giving a mapping, $f_{\text{MRI-to-CC}}$ say, from points in the breast in the MRI position to points in the breast in the CC imaging position, and also compress the model in the MLO direction, giving a mapping, $f_{\text{MRI-to-MLO}}$ from points in the breast in the MRI position to points in the breast in the MLO imaging position. Stage 1 as described above is now as simple as Stage 5, since we now have the breast shape in the CC position and we simply 'un-project' the point of interest in the CC image ($\mathbf{x_0}$ say) up to the X-ray source to get a straight line, $\mathcal{L}_{\mathbf{x_0}}$ say. To find the region in the MLO image which corresponds to $\mathbf{x_0}$, we have to form the curve $f_{\text{MRI-to-MLO}}((f_{\text{MRI-to-CC}})^{-1}(\mathcal{L}_{\mathbf{x_0}}))$.

2.3.4 Other Uses

The model can also be of use in situations where the deformation is not large, such as temporal matching of images of the same modality. Comparing different images of the same patient over a period in time can be very important for early diagnosis of cancers, since a suspicious mass can be identified on the basis of it being present in one image but absent in images taken in previous years. To compare quantitatively such images the radiologist has to account for the patient's position being slightly different in the different images, and a deformable model of the breast could be used to deal with these issues.

There are various other uses of the model. It could simply be used as a visualisation tool, allowing a doctor to show the patient the position and size of the tumour in 3D. It could be used as a teaching tool for radiologists, since it would enable a user to study and visualise the internal structure of the breast. The model could also be used as a tool for simulation of the breast after reconstructive (or cosmetic) surgery. Breast reconstruction involves the use of prostheses or tissue from other parts of the body to replace tissue that has had to be removed during mastectomy, and as (an extension of) the model could be used to simulate the new shape of the breast following the replacement procedure, the model would be an invaluable tool for planning reconstructive surgery.

2.4 Modelling the Deformation of the Breast and Previous Efforts

There are several stages involved in finite element modelling of the breast, covering a range of fields which include image analysis, applied mathematics, numerical analysis and bio-mechanical experimentation. In this final section of this chapter, we briefly introduce each stage. Where applicable, we also discuss previous work carried out by groups modelling the deformation of the breast. Works which use linear elasticity or a non-physical method will be mentioned as appropriate, but we focus mainly on the methods used by two particular groups which have modelled deformations of the breast using a non-linear elasticity theory: *Azar et al.* (see [Aza01] or [AMS01]), and *Samani et al* (see [SBYP01], [SBRP99] or [PBSS00]). These two groups are amongst a very small number of groups to have extensively studied breast deformation without using linear elasticity[2].

2.4.1 Segmentation and Edge Detection

The first stage that has to be performed is edge detection and segmentation of MR images. Here, we have to detect the breast edge in each of the 2D MR images, and obtain everything else that is needed for the model. The exact requirements depend on how the anatomy is modelled. For example, one method of modelling the internal anatomy is to consider there to be three internal tissue types (fat, fibroglandular and tumour tissue), and in this case we need to be able to say which tissue a given point in the interior of the breast corresponds to. This could be obtained by carrying out *full voxel segmentation*, where an algorithm is devised to assign a tissue (from those three possibilities) to each voxel in the MR set of images, a voxel being the 3D equivalent of a pixel.

Although the segmentation process depends on the nature of the model, in that we need to obtain only the tissues which will be needed for the model, the opposite statement is also true: the nature of the model depends on the segmentation, in that we can only include tissues in the model that can be ‘seen’ in the MR images (unless we wish to assume some generic structure). It is for this reason that previous work in deformable modelling of the breast have limited the number of internal tissues to fat, fibroglandular and tumour. Fatty tissue can be distinguished from non-fatty tissue in MR images relatively easily, but it is often very difficult to make out other structures. Both *Azar et al.* [AMS01] and *Samani et al.* [SBYP01] segment the tissues in this manner, using full voxel segmentation. We will also restrict the internal tissues to these three types, and use voxel segmentation. Image segmentation is discussed in Section 7.1.1.

2.4.2 Mesh Generation

The next stage is mesh-generation, where the breast is decomposed into a set of regular shaped elements. Because of the complicated nature of the geometry being handled, this is a non-trivial problem, and made more so by the fact that the type of element used, for example tetrahedral or hexahedral, affects the accuracy of the solution. The techniques used by *Azar et al.* [AMS01] and *Samani et al.* [SBYP01] both involve meshes comprised of hexahedral elements in the interior of the breast, although they are generated in very different ways. *Azar et al.* mesh each image slice with N by N quadrilaterals, for some fixed N, and stack the quadrilaterals to obtain the hexahedral

---

[2] The only others who have studied breast deformation using finite elasticity are, as far as we are aware, [TDS$^+$01] and [Rui03]. We do not mention their models in this survey because both consider a number of different models and study differences in the results, rather than develop their own particular model.

elements. *Samani et al.* consider two methods of mesh generation, one where voxels are grouped into elements, and a second method where the geometry is transformed onto the unit cube and a regular mesh on the unit cube mapped back into physical space. They both use three-dimensional hexahedral elements for the interior of the breast, but two dimensional elements on the surface to model the skin. The other possibility to model the skin is to have an outer layer of thin, but three-dimensional, elements. Methods of mesh generation are described in Section 7.1.2, and we will consider both thin skin elements (Section 7.1.2) and 2D membrane skin elements (Chapters 8 and 9).

2.4.3 Elasticity Theory

The breast is one of the most deformable organs in the human body, and the deformations which we wish to simulate are large deformations, especially compression for mammography. Therefore, when forming the equations that govern the deformations, we need to use the full theory of large elastic deformations, *nonlinear elasticity theory*, also known as *finite elasticity theory.* We need to make various assumptions on the properties of breast tissues when setting up the equations which determine the deformation. For example, we have to decide to what extent we model anisotropy and heterogeneity of breast tissues, and also whether to use compressible or incompressible elasticity theory.

Both *Samani et al.* and *Azar et al.* assume isotropy and incompressibility. *Samani et al.* use the full finite elasticity theory in their model, whereas *Azar et al.* consider the large deformation as a sequence of small deformations, and use linear elasticity on each of the small deformations. Works which use linear elasticity to model breast-deformation include: [DDP$^+$04] (with pseudo-incompressibility and isotropy also assumed, in order to simulate Near Infrared Tomography), [BABM02] (where an extremely complex, but generic, anatomical model is built), [WM04] (where linear elasticity is used to model compression of 2D MR images), and [IJBM06]. Heuristic methods include [KHB01], where the breast shape is assumed to be given by a quadratic curve, as was discussed in Section 2.3.3.

In Chapter 3 we develop the theory of finite deformation elasticity. We discuss the modelling assumptions made in Section 7.2.1. Further mathematical modelling is carried out in Chapter 8.

2.4.4 Determining Material Laws

The equations of finite elasticity require knowledge of the relationship between stress and strain for each tissue in the model. These are known as the *material law* or the *constitutive law* of the tissue. Stress and strain are 2nd-order tensor functions of space, and the general relationship between them can be extremely complex. The constitutive law for materials can only be determined by experiment.

Functional forms need to proposed for experimental data, and parameters estimated by fitting. This is a difficult task, especially for biological tissues. Obtaining accurate and reliable experimental biological data is often the biggest stumbling block in biological modelling.

In image analysis problems, we obviously only have information about the starting and the ending states, so we need to formulate a model which ignores viscous (time-dependent) effects in the material laws. In Chapter 3, we define the strain-energy function, which is a function of strain only, and from which stress is computed. Tissues for which a strain-energy function exists are known as *hyperelastic*. Hyperelasticity is one of our modelling assumptions and is discussed together with the other assumptions in Section 7.2.1. The material laws we use for breast tissues are discussed in Section 7.2.2.

2.4.5 Removing Gravity to Obtain the Reference State (the Backward Problem)

As we shall see in Chapter 3, the equations of finite elasticity assume that the initial shape of the object under consideration (the *reference state*, or *reference configuration*) is stress free, which means there can be no body or surface forces acting on the object. Clearly, in our case, where we are building the model using MR data, the breast is under the influence of gravity. Before any clinical simulations can be performed, we must remove the effect of gravity from the model to obtain the true reference state, which can then be used as the initial state in deformation calculations. To do this we must either invert the equations of finite elasticity, or develop an iterative procedure that converges to the true reference state. This is discussed in Chapter 5. We shall see in Chapter 9 that a separate backward problem for skin arises, which also needs to be studied, and we discuss methods of tackling this new problem in that chapter.

Despite the fact that the breast is clearly gravity-loaded in the MR images, neither *Azar et al.* or *Samani et al.* consider the backward problems. Instead both take the MR state as their starting, and supposedly unstressed, state. However, both simulate mammographic compression, and it could be argued that the stresses and displacements in the MR state are very small compared to the stresses and displacements in the compressed state. This might be true in some cases, but it is easy to imagine scenarios with older women with breasts containing a high proportion of fat which deform significantly when hanging down under gravity.

2.4.6 Simulations (the Forward and Contact Problems)

Once we have built up the model from the segmented MR images, assigned the different tissues a material law, and obtained the reference state, we will be ready to simulate the breast position during surgery or mammography. Calculating the breast position for the patient in the supine position is essentially just a standard finite elasticity computation, which we call the forward problem. Example simulations are given in Section 7.3.1. Simulation of mammography is an example of a contact problem. We discuss how to state and solve the contact problem in Chapter 6, and sample simulations are carried out in Section 7.3.2.

Chapter 3

Nonlinear Elasticity Theory

In this chapter we derive the equations of the *theory of nonlinear elasticity*, also known as *finite elasticity*[1]. Elasticity theory is founded upon the concepts of *stress*, which is a measure of internal pressure, and *strain*, a measure of internal stretching. We begin in Section 3.1 by defining the deformation gradient tensor and the strain. The difference between quantities defined on the undeformed and deformed states of an elastic body is important in finite elasticity, and we will emphasise the distinction between the different types of strain that can be defined. Similarly, stress can be defined in several ways, and this is discussed in Section 3.1.3. The equations of state that determine the shape of a body given the stresses within it are then derived in Section 3.2, as is the weak form of the problem, and later, the energy minimisation form, in Section 3.3.2. We will see that both of these alternative formulations are more useful in computation, and sometimes also in analysis, than the direct formulation.

To relate stress to strain, we need a material-specific expression connecting them, known as the *constitutive law*. In Section 3.3 we discuss the general form of the constitutive law, before considering the specific case of *hyperelasticity*, when there exists a strain energy whose derivative with respect to strain gives the stress. The form of the strain energy can be simplified if the material in question has the property of isotropy (identical material properties in all directions) or transverse isotropy (identical material properties in a plane at each point), and these simplifications are discussed. Also, the form of the constitutive relation and the equations to be solved are different when constraints, such as incompressibility, are imposed, and this is considered in Section 3.4.

Finally, in Section 3.5, we examine the case of small displacements, which leads to the equations of *linear elasticity*.

---

[1] Finite in this context refers to the fact that the strains are of a magnitude comparable to the overall size of the elastic body (in contrast with infinitesimal or linear elasticity, where the magnitudes of strain is assumed to be small or tiny in comparison to the size of the object).

3.1 Stress and Strain

The important quantities used in characterising the principles of nonlinear elasticity are stress and strain. Stress describes the force per unit area acting on small surfaces throughout a deformable body, and strains are dimensionless measures of the stretching the body has undergone. In building the theory of nonlinear elasticity, the first step is to consider the kinematics of elastic deformations (i.e. the study of the motion of a body exclusive of mass or force), and define quantities such as displacement, deformation gradient and strain. Next, various measures of the stress in the body can be defined. Equations of motion for the deformation of the body can then be derived by considering Newton's Second Law.

Consider an undeformed elastic body $\Omega_0 \subset \mathbb{R}^3$, as shown in Figure 3.1. Given a point P in Ω_0, let its position be given by $\mathbf{X} = X^M \mathbf{e}_M$, where X^M are the (contravariant)[2] components of the position with respect to the standard orthonormal basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ for $\mathbb{R}^3$. Ω_0 is known as the *reference configuration* and the coordinates X^M are known as *Lagrangian coordinates.* The reference configuration is the known starting state of a body in nonlinear elasticity.

A general deformation of the body is given by a mapping $\mathbf{\Phi} : \Omega_0 \to \mathbb{R}^3$. We define $\mathbf{x} = x^i \mathbf{e}_i$ to be the mapped position of the point P, i.e.

$$\mathbf{x} = \mathbf{\Phi}(\mathbf{X}),$$

and define the deformed body Ω to be

$$\Omega = \mathbf{\Phi}(\Omega_0) = \{\mathbf{\Phi}(\mathbf{X}) : \mathbf{X} \in \Omega_0\}.$$

The coordinates[3] x^i are known as *Eulerian coordinates.* The distinction between the Lagrangian coordinate system and the Eulerian coordinate system is important. We will see that it is possible to

---

[2] Recall that tensors are quantities which transform according to certain *tensor transformation laws.* Tensors provide a natural and concise mathematical framework for formulating and solving problems in areas of physics, whereas non-tensor quantities cannot usually be used to characterise the behaviour of a physical system because they depend on the arbitrary choice of basis. There are three types of tensor. A covariant tensor $T_{ab\ldots c}(x)$ satisfies the transformation law

$$T'_{ab\ldots c}(x') = \frac{\partial x^p}{\partial x'^a} \frac{\partial x^q}{\partial x'^b} \cdots \frac{\partial x^r}{\partial x'^c} T_{pq\ldots r},$$

where x and x' are the old and new coordinates, and a contravariant tensor $T^{ab\ldots c}$ satisfies

$$T'^{ab\ldots c} = \frac{\partial x'^a}{\partial x^p} \frac{\partial x'^b}{\partial x^q} \cdots \frac{\partial x'^c}{\partial x^r} T^{pq\ldots r}.$$

Gradients of scalar potentials are examples of covariant tensors, and differentials examples of contravariant tensors. The third type are mixed tensors, e.g. tensors of the form $T^{ab}{}_c$, which transform according to the obvious extension of the co and contravariant laws. For an introduction to tensors, see [Nio85].

[3] We will use the following convention to distinguish between quantities in undeformed and deformed space. If possible, upper case letters will be used for the undeformed quantity and lower case letters for the deformed quantity (e.g. the positions $\mathbf{X}$ and $\mathbf{x}$, or the indices M and i in X^M and x^i). For Greek characters such as Ω this is not especially clear, so the subscript zero will instead be used for the undeformed quantity, with no subscript for the deformed quantity.

formulate Lagrangian and Eulerian descriptions of elastic deformations, the former being a description in terms of fields defined on Ω_0, the latter being in terms of fields over Ω. We assume that the reference configuration is the known configuration, with the deformed configuration being the unknown configuration which we wish to determine, so the aim is to formulate Lagrangian descriptions of elastic problems.

Given the old position $\mathbf{X}$ and the new position $\mathbf{x}$, the displacement $\mathbf{u}$ is given by their difference:

$$\mathbf{u}(\mathbf{X}) = \mathbf{x}(\mathbf{X}) - \mathbf{X}.$$

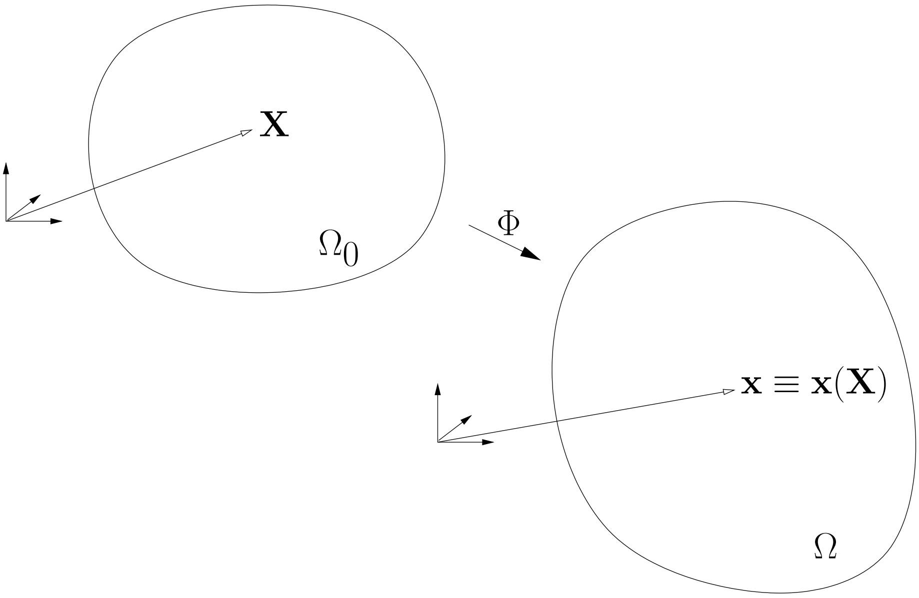

Figure 3.1: The undeformed state Ω_0 is mapped to the deformed state Ω by the deformation Φ.

3.1.1 The Deformation Gradient, F

The *deformation gradient tensor* is defined as the partial derivative of new position with respect to old position

$$F^i_M = \frac{\partial x^i}{\partial X^M}, \tag{3.1}$$

and describes the nature of the deformation. F describes the transformation between differentials, since, by the chain rule

$$\mathrm{d}\mathbf{x} = F\mathrm{d}\mathbf{X}.$$

F^i_M is not a true mixed tensor *per se*, but a two point tensor field defined on two configurations with possibly two bases (although in the current setup we have used the same (orthogonal) basis for both configurations).

Now, any deformation can be split into two parts, a stretch and a rigid body deformation (i.e. a rigid body rotation and a translation). Correspondingly, the deformation gradient tensor F can be factorised as $F = RU$, where R is an orthogonal tensor representing the rotation (clearly F will be independent of any translation, being a derivative), and U is a symmetric positive-definite tensor representing the stretching. This is known as the *polar decomposition theorem*. In fact, the polar decomposition theorem is more general than this application to elasticity. It states that for any real matrix $A \in \mathbb{R}^{n\times n}$, there is an orthogonal matrix R and a positive semi-definite matrix U such that $A = RU$, and that R and U are unique if A is invertible[4] [HJ91]. Thus the decomposition exists for all real matrices, not just those which are deformation gradients. We will show shortly that F must be non-singular for valid deformations, so this implies that R and U are unique for a given deformation.

The determinant of F^i_M, which is the Jacobian for the deformation, is an important quantity:

$$J = \det(F).$$

Note that $J = \det(RU) = \det(R)\det(U) = \det(U)$, so is only dependent on the elastic part of the deformation. Since $\mathrm{d}V = J\mathrm{d}V_0$, J is the factor by which volumes transform. In particular, this means that J must be non-zero (everywhere in Ω_0), since otherwise there would be a small region of non-zero volume that deforms to a region with zero volume, which is physically unrealistic. Hence F is invertible, and it is easy to see that the inverse of the deformation gradient is the deformation gradient of the inverse map

$$\left(F^{-1}\right)^M_i = \frac{\partial X^M}{\partial x^i}.$$

Also, it is not physical for the orientation of line elements to change during a deformation, and this provides us with the constraint that J is positive.

Since J is the factor by which volumes scale, in particular it must be equal to one everywhere if the material is incompressible. This can be seen by considering the volume change of a small region of the body. Let $\delta\Omega_0 \subset \Omega_0$, and let $\delta\Omega = \mathbf{\Phi}(\delta\Omega_0)$. Then

$$\int_{\delta\Omega_0} \mathrm{d}V_0 = \int_{\delta\Omega} \mathrm{d}V = \int_{\delta\Omega_0} J\,\mathrm{d}V_0,$$

the first equality holding by the assumption of incompressibility. Since $\delta\Omega_0$ is arbitrary, $J = 1$ everywhere. We will discuss the case of incompressible materials further in Section 3.4. A deformation for which $J = 1$ everywhere is said to be *isochoric* (volume-preserving). Similarly, J is the factor by which density scales. This can be seen by considering conservation of mass. If the density fields in the undeformed and deformed bodies are ρ_0 and ρ respectively, then

$$\int_{\delta\Omega_0} \rho_0\,\mathrm{d}V_0 = \int_{\delta\Omega} \rho\,\mathrm{d}V = \int_{\delta\Omega_0} \rho J\,\mathrm{d}V_0,$$

---

[4] For any real matrix there also exists a decomposition of the form $A = U'R'$, where R' is orthogonal and U' is positive definite and symmetric. This is clear from considering the polar decomposition of A^{T}.

the first equality holding by the principle of conservation of mass, so that $\rho_0 = \rho J$.

F can also be used to map between surfaces in the two configurations. If $\mathbf{N}$ is the normal to an infinitesimal surface element $\mathrm{d}S_0$ in the undeformed configuration, with $\mathbf{n}$ and $\mathrm{d}S = \mathbf{\Phi}(\mathrm{d}S_0)$ the corresponding normal and surface element in the deformed body (note that the normal is not embedded in the material, and does not transform like line elements), then

$$n_i \mathrm{d}S = J N_M (F^{-1})_i^M \mathrm{d}S_0,$$

or

$$\mathbf{ds} = J(F^{-1})^{\mathrm{T}} \mathbf{dS}. \tag{3.2}$$

This is known as Nanson's formula. For a proof refer to [Ogd97].

In general, F is not symmetric, and may not be diagonalisable. However, the stretching component of F is the positive-definite symmetric matrix U, which must have real strictly positive eigenvalues. These eigenvalues of U, which are denoted by λ_1, λ_2 and λ_3, are important quantities, since they represent the *principal stretches*. The eigenvectors of U are the *principal stretch directions*. Note that the eigenvalues of U' when F has the polar decomposition of the form $F = U'R'$ are also λ_1, λ_2 and λ_3.

3.1.2 Measures of Strain

Strain is a dimensionless measure of the length change gradient. For a one-dimensional object, one standard form of strain used is the displacement divided by the initial length, so that if a 1D rod is stretched uniformly from length L_0 to length L, the strain is $\frac{L-L_0}{L_0}$ everywhere. In higher dimensions strain has a tensor form. The concept of strain arises because in elasticity theory we wish to ignore the rotational component of the deformation (which is generally fixed by boundary conditions), and study the stretching (elastic) component. Since the rotational component at a point is given by the orthogonal tensor R in $F = RU$, it is natural to define the following tensor, known as the *right Cauchy-Green deformation tensor*

$$C_{MN} = (F^{\mathrm{T}} F)_{MN} = \frac{\partial x^i}{\partial X^M} \frac{\partial x^i}{\partial X^N}. \tag{3.3}$$

Then $C = U^{\mathrm{T}} R^{\mathrm{T}} R U = U^{\mathrm{T}} I U = U^2$, which is independent of any rotations. Similarly, the *left Cauchy-Green tensor* (or *Finger tensor*) is defined to be $B = FF^{\mathrm{T}}$. Note that C is a *Lagrangian tensor*, in that it acts on vectors in the undeformed space. C is symmetric and positive definite, and describes how lengths change, since $\mathrm{d}\mathbf{x}^2 = (F\mathrm{d}\mathbf{X})^{\mathrm{T}} F \mathrm{d}\mathbf{X} = \mathrm{d}\mathbf{X}^{\mathrm{T}} C \mathrm{d}\mathbf{X}$. If the deformation is a rigid body deformation, F will be an orthogonal tensor, and C will reduce to the identity.

A point is said to be *unstrained* if the deformation at that point is purely rigid (in other words, if $F = I$ at that point). This motivates the following definition of strain

$$E_{MN} = \frac{1}{2}\left(C_{MN} - \delta_{MN}\right), \tag{3.4}$$

which is known as the the *Green-Lagrange strain tensor*, or sometimes just as the *Green strain tensor*. The factor of $\frac{1}{2}$ is a normalisation factor which can be understood when the linearisation of the strain tensor is considered in linear elasticity theory (Section 3.5). It is easy to see that the Green strain tensor satisfies

$$\mathrm{d}\mathbf{x}^2 - \mathrm{d}\mathbf{X}^2 = 2\mathrm{d}\mathbf{X}^{\mathrm{T}} E \mathrm{d}\mathbf{X}.$$

Similarly, we could have defined a strain to be the tensor e satisfying

$$\mathrm{d}\mathbf{x}^2 - \mathrm{d}\mathbf{X}^2 = 2\mathrm{d}\mathbf{x}^{\mathrm{T}} e \mathrm{d}\mathbf{x}.$$

It is straightforward to show that

$$e = \frac{1}{2}\left(I - B^{-1}\right) \qquad \text{or} \qquad e_{ij} = \frac{1}{2}\left(\delta_{ij} - \frac{\partial X^M}{\partial x^i}\frac{\partial X^M}{\partial x^j}\right),$$

in contrast to

$$E_{MN} = \frac{1}{2}\left(\frac{\partial x^i}{\partial X^M}\frac{\partial x^i}{\partial X^N} - \delta_{MN}\right).$$

e_{ij} is known as the *Eulerian* or *Almansi strain tensor*.

Note that the Green strain is defined over the initial configuration, whereas the Eulerian strain is defined over the deformed configuration. In Section 3.1.3 we will define the Cauchy and 2nd Piola-Kirchoff stresses, where the latter is defined on the initial configuration, the former defined on the deformed configuration, and we will show that the Eulerian strain is the appropriate strain for the Cauchy stress, and the Green strain the appropriate strain for the 2nd Piola-Kirchoff stress. Note also that $E = 0 \Leftrightarrow e = 0 \Leftrightarrow \mathrm{d}\mathbf{x}^2 - \mathrm{d}\mathbf{X}^2 = 0 \Leftrightarrow$ all line segments have unchanged length $\Leftrightarrow$ the body undergoes a rigid body deformation.

In fact, various strain tensors can be defined, and even families of strain tensors. $C = U^2$, so $E = \frac{1}{2}(U^2 - I)$, and E could be considered to be the second member of the following family of strain tensors[5]:

$$E^{(m)} = \frac{1}{m}(U^m - I) \qquad \text{for } m \in \mathbb{N}.$$

Different strain measures can be useful in different problems, although in practice the Green strain and the Almansi strain are most commonly used. In Section 3.1.3 we will describe how stress and strain tensors can be seen to be conjugate quantities when mechanical work is considered.

---

[5] As it turns outs, this family can be extended to all integers, by using inverses of U for negative m and using $E^{(0)} = \log U$ (which exists due to the positive-definitiveness of U) for $m = 0$. It is noted in [Ogd97] that this family could even be extended to all $m \in \mathbb{R}$.

In our work we will need to only consider the Green strain tensor E_{MN}, although in practise it often turns out easier and more natural to work with the Green deformation tensor C_{MN} rather than E_{MN}. As will be seen in later sections, for some materials the stress in a material depends not on all the components of the strain tensor (or equivalently, not on all the components of C_{MN}), but only on the *principal invariants*[6], of C_{MN}. Since U has eigenvalues λ_1, λ_2, λ_3, C_{MN} has eigenvalues λ_1^2, λ_2^2, λ_3^2, and the principal invariants of C_{MN} in terms of the stretch ratios are

$$\begin{aligned}
I_1 &= \mathrm{tr}(C) = \lambda_1^2 + \lambda_2^2 + \lambda_3^2 \\
I_2 &= \frac{1}{2}((\mathrm{tr}(C))^2 - \mathrm{tr}(C^2)) = \frac{1}{2}(I_1^2 - C_{MN}C_{MN}) = \lambda_1^2\lambda_2^2 + \lambda_2^2\lambda_3^2 + \lambda_3^2\lambda_1^2 \\
I_3 &= \det(C) = J^2 = \lambda_1^2\lambda_2^2\lambda_3^2.
\end{aligned}$$

3.1.3 Stress

Stress is defined as the force per unit area acting on an infinitesimally small surface element. For any surface element, the force on that surface element will have a component (possibly equal to zero) acting parallel to the normal of the surface, which is known as the *axial* or *normal stress*, and a component acting in a direction tangential to the surface, known as the *shear stress*. This suggests that stress has a tensor form, which is known as Cauchy's Theorem. For the standard proof of Cauchy's theorem using a tetrahedron aligned with the coordinate system, see [Ogd97].

There are several formulations the stress tensor can take, with different formulations being appropriate for different situations. The most obvious formulation is the Cauchy stress tensor, defined shortly, which is a measure of the 'true' stress; however, we will see that this tensor is not the most appropriate stress tensor to use to formulate large deformation problems, and therefore we will also define the Piola-Kirchoff stress tensors.

The Cauchy Stress Tensor, σ^{ij}

The Cauchy stress tensor represents the force measured per unit *deformed* area acting on the *deformed* body. The superscript i represents the direction of the normal of the surface, and j represents the direction of the force. The Cauchy stress is an Eulerian tensor field.

For non-polar materials (materials for which there are no couple stresses or body couples present), it can be shown [Ogd97] that conservation of angular momentum implies that σ^{ij} is symmetric. We will only ever be considering non-polar materials[7] so henceforth the symmetry of the Cauchy stress tensor will be assumed.

---

[6] The principal invariants of a 3 by 3 matrix A are: $I_1(A) = \mathrm{tr}(A)$, $I_2(A) = \frac{1}{2}((\mathrm{tr}(A))^2 - \mathrm{tr}(A^2))$ and $I_3(A) = \det(A)$, three quantities which are unchanged under an orthonormal change of basis. Note that they are the coefficients in the characteristic polynomial of A, ie $\chi_A(\lambda) = \det(A - \lambda I) = -\lambda^3 + I_1\lambda^2 - I_2\lambda + I_3$.

[7] In polar materials small elements of the material exert moments on neighbouring elements, which is an extremely unusual microscopic property. Liquid crystals are examples of polar materials, biological tissues certainly are not.

If the basis for the deformed configuration is orthonormal, the Cauchy stress tensor gives the true stress (the stress in the deformed body), and it is the stress that has to be used when forming equilibrium equations. However, the problem with the Cauchy stress tensor for analysing large deformations is that we generally do not know the area in the deformed configuration. We have defined the (Green) strain tensor in the undeformed configuration, and to relate stress and strain we therefore need a stress tensor which is defined in the undeformed configuration. The first step to obtain this is to define the first Piola-Kirchoff tensor.

The 1st Piola-Kirchoff Stress Tensor, S^{Mj}

The 1st Piola-Kirchoff stress tensor represents the force measured per unit *undeformed* area, acting on the *deformed* body. The first index indicates the normal direction of the surface in the reference configuration, and is therefore a Lagrangian (i.e. upper case) index. The second index indicates the direction of the force in the deformed body, and is a spatial (i.e. lower case) index.

The 1st Piola-Kirchoff stress is useful in experiments, since undeformed areas are more easily measured than deformed areas, but the forces can only be measured in the deformed body. Since the indices represent different states, the 1st Piola-Kirchoff tensor is not in general symmetric, and it is not convenient to use it in the formulation of problems. For this reason the 2nd Piola-Kirchoff tensor is defined.

The 2nd Piola-Kirchoff Stress Tensor, T^{MN}

The 2nd Piola-Kirchoff tensor represents the force measured per unit *undeformed* area, acting on a surface in the *undeformed* body, and as such is a Lagrangian tensor field. The first index represents the normal direction, the second the direction of the force. Since the forces act only on the deformed body, the 2nd Piola-Kirchoff stress has little physical meaning, in the sense that it cannot actually be measured, and is just a mathematical convenience. However, we can obtain the (theoretical) force acting on the undeformed area, $\mathbf{T}$, from that, acting on the deformed area, $\mathbf{t}$, by using $\mathbf{T} = F^{-1}\mathbf{t}$ (just as line segments transform as $\mathrm{d}\mathbf{X} = F^{-1}\mathrm{d}\mathbf{x}$).

It will be shown in the following section that the 2nd Piola-Kirchoff stress tensor is symmetric if the Cauchy stress tensor is symmetric, which is the case for all materials considered here.

Relationships Between the Stress Tensors

It it straightforward to derive relationships held between the stress tensors. From the definition of the Cauchy and 1st Piola-Kirchoff stresses, the force, $\mathrm{d}\mathbf{t}$, acting on a surface $\mathrm{d}S_0$ in the undeformed body (with normal $\mathbf{N}$), which is the surface $\mathrm{d}S$ in the deformed configuration (with normal $\mathbf{n}$) is given by

$$\mathrm{d}t^i = \sigma^{ji} n_j \mathrm{d}S = S^{Mi} N_M \mathrm{d}S_0.$$

By using Nanson's Theorem (3.2) to map from the undeformed space to the deformed space, we obtain

$$(S^{Mi} - J(F^{-1})^M_j \sigma^{ji}) N_M \mathrm{d}S_0 = 0$$

which implies that

$$S^{Mi} = J(F^{-1})^M_j \sigma^{ji} \qquad \text{or} \qquad S^{Mi} = J \frac{\partial X^M}{\partial x^j} \sigma^{ji} \qquad \text{or} \qquad S = JF^{-1}\sigma. \tag{3.5}$$

Similarly, it can easily be shown that[8]

$$T^{MN} = S^{Mj} \frac{\partial X^N}{\partial x^j} \qquad \text{or} \qquad T = SF^{\text{-T}}. \tag{3.6}$$

Putting these relations together gives the relationship between the Cauchy stress and 2nd Piola-Kirchoff stress:

$$\sigma^{ij} = \frac{1}{J} T^{MN} \frac{\partial x^i}{\partial X^M} \frac{\partial x^j}{\partial X^N} \qquad \text{or} \qquad \sigma = \frac{1}{J} FTF^{\mathrm{T}}. \tag{3.7}$$

We see that, as claimed, the 2nd Piola-Kirchoff stress tensor is symmetric if and only if the Cauchy stress tensor is symmetric.

The different stresses can be considered conjugate to appropriate strains by considering the concept of work-conjugacy. A work-conjugate pair of stress and strain tensors is defined as a pair for which the product[9] of stress with strain rate gives the rate of work per unit volume. It can be shown [HKS05] that the rate of work is given by

$$\int_\Omega \sigma^{ij} \dot{e}_{ij} \, \mathrm{d}V,$$

and that this is equal to

$$\int_{\Omega_0} T^{MN} \dot{E}_{MN} \, \mathrm{d}V_0.$$

Hence the Cauchy stress is work-conjugate to the Almansi strain, and the 2nd Piola-Kirchoff stress is conjugate to the Green strain. The 1st Piola-Kirchoff stress is in fact conjugate to the deformation gradient.

3.2 The Equation of State and Weak Formulation

We can now derive the governing equations determining the equilibrium state in static problems using nonlinear elasticity, and the equation of motion for dynamic problems. These equations are known as the strong forms of the problems. We then reformulate the strong form into the equivalent and more useful weak form.

---

[8] Here we use the notation $G^{\text{-T}}$ for the inverse transpose, $(G^{\mathrm{T}})^{-1} \equiv (G^{-1})^{\mathrm{T}}$, of a matrix G.

[9] By 'product' of two tensors A and B, we mean $\mathrm{tr}(AB) = A_{MN}B_{MN}$, often denoted as $A : B$.

3.2.1 Equations of Motion

The equations governing elastic deformations can be derived from the principle of conservation of linear momentum. In this section we will consider the time-dependent case for completeness, although the situations modelled throughout this thesis will only ever be static time-independent problems. Time-independent steady-state equations can be obtained from time-dependent equations of motion by setting all time derivatives to zero.

Suppose the reference state of a deformable body is given by Ω_0, a compact connected subset of $\mathbb{R}^3$ and let Ω_t now represent the region of space occupied by the body at time $t \geq 0$. We suppose the body is acted on by two types of applied force. A *body force* is a force acting throughout the body, and is specified through a *body force density*, $\mathbf{b}(\mathbf{x}, t)$. Given mass density $\rho(\mathbf{x}, t)$, the total force acting on a region $\delta\Omega_t \subset \Omega_t$ is

$$\int_{\delta\Omega_t} \rho(\mathbf{x}, t)\mathbf{b}(\mathbf{x}, t)\,\mathrm{d}V.$$

Gravity is an example of a body force, in which $\mathbf{b} = \mathbf{g} = (0, 0, -9.81)^{\mathrm{T}}$, the standard acceleration due to gravity. The second possible applied force is a *surface force*, or a contact force. These are forces which act on some part of $\partial\Omega_t$ (the boundary of Ω_t) and are specified by providing the *surface traction*, or force per unit area (i.e. the surface stress), denoted by $\mathbf{s}$, on $\partial\Omega_t$. If $\delta S \subset \partial\Omega_t$, the contact force is given by

$$\int_{\delta S} \mathbf{s}(\mathbf{x}, t)\,\mathrm{d}S.$$

Surface tractions can be split into two components, the normal component to the surface, acting like a pressure force, and a tangential component, acting like friction.

Let us define the velocity vector

$$\mathbf{v} = \frac{\partial \mathbf{x}(\mathbf{X}, t)}{\partial t} = \frac{\partial \mathbf{u}(\mathbf{X}, t)}{\partial t}$$

as the time-derivative of the displacement, and consider again a small region $\delta\Omega_t \subset \Omega_t$. Define $\mathbf{f}$ to be the force acting on the boundary of $\delta\Omega_t$. The *linear momentum* of the region $\delta\Omega_t$ is

$$\int_{\delta\Omega_t} \rho(\mathbf{x}, t)\mathbf{v}(\mathbf{x}, t)\,\mathrm{d}V.$$

The principal of linear momentum states that the rate of change of the momentum must be equal to the applied forces, and hence

$$\int_{\partial(\delta\Omega_t)} \mathbf{f}\,\mathrm{d}S + \int_{\delta\Omega_t} \rho\mathbf{b}\,\mathrm{d}V = \frac{\mathrm{d}}{\mathrm{d}t} \int_{\delta\Omega_t} \rho\mathbf{v}\,\mathrm{d}V,$$

where $\partial(\delta\Omega_t)$ is the boundary of $\delta\Omega_t$. The force $\mathbf{f}$ on a surface element $\mathrm{d}S$ is given by

$$f^i \mathrm{d}S = \sigma^{ji} n_j \mathrm{d}S.$$

Integrating over the small region and applying the divergence theorem, we see that

$$\int_{\partial(\delta\Omega_t)} f^i \,\mathrm{d}S = \int_{\partial(\delta\Omega_t)} \sigma^{ji} n_j \,\mathrm{d}S = \int_{\delta\Omega_t} \frac{\partial \sigma^{ji}}{\partial x^j} \,\mathrm{d}V.$$

Thus

$$\int_{\delta\Omega_t} \frac{\partial \sigma^{ji}}{\partial x^j} \,\mathrm{d}V + \int_{\delta\Omega_t} \rho b^i \,\mathrm{d}V = \frac{\mathrm{d}}{\mathrm{d}t} \int_{\delta\Omega_t} \rho v^i \,\mathrm{d}V. \tag{3.8}$$

This is true for any arbitrary volume $\delta\Omega_t \subset \Omega_t$, so (assuming that $\mathbf{b}$ is continuous, ρ is constant, and using the symmetry of σ), our equation of state is

$$\frac{\partial \sigma^{ij}}{\partial x^j} + \rho b^i = \rho \frac{\mathrm{D} v^i}{\mathrm{D}t} \qquad \text{in } \Omega_t, \tag{3.9}$$

where $\frac{\mathrm{D}}{\mathrm{D}t} \equiv \frac{\mathrm{d}}{\mathrm{d}t} + \mathbf{v} \cdot \boldsymbol{\nabla}$ is the material derivative[10]. This equation is known as (the time-dependent form of) *Cauchy's Law*. The time independent form of Cauchy's Law is given by setting the velocity vector to zero.

$$\frac{\partial \sigma^{ij}}{\partial x^j} + \rho b^i = 0 \qquad \text{in } \Omega. \tag{3.10}$$

Henceforth we will consider only the time-independent version.

Cauchy's law is not useful for computations, since it is formed in terms of derivatives with respect to new position x^i, which are precisely the unknowns which we wish to compute. In other words it is an Eulerian formulation, and we need a Lagrangian formulation. We need an equation of state on our undeformed domain Ω_0, and the first step towards one involves making use of the 1st Piola-Kirchoff stress tensor to map to Ω_0. Let $\delta\Omega_0 \subset \Omega_0$, with $\delta\Omega = \boldsymbol{\Phi}(\delta\Omega_0)$. Now,

$$\sigma^{ji} n_j \mathrm{d}S = S^{Mi} N_M \mathrm{d}S_0$$

(by definition), so

$$\int_{\partial(\delta\Omega)} \sigma^{ji} n_j \,\mathrm{d}S = \int_{\partial(\delta\Omega_0)} S^{Mi} N_M \,\mathrm{d}S_0 = \int_{\delta\Omega_0} \frac{\partial S^{Mi}}{\partial X^M} \,\mathrm{d}V_0,$$

by again using the divergence theorem. Also,

$$\int_{\delta\Omega} \rho \mathbf{b} \,\mathrm{d}V = \int_{\delta\Omega_0} \rho \mathbf{b} J \,\mathrm{d}V_0 = \int_{V_0} \rho_0 \mathbf{b} \,\mathrm{d}V_0,$$

which means the transformed momentum equation is

$$\int_{\delta\Omega_0} \frac{\partial S^{Mi}}{\partial X^M} \,\mathrm{d}V_0 + \int_{\delta\Omega_0} \rho_0 b^i \,\mathrm{d}V_0 = 0.$$

---

[10] The material derivative is a measure of the rate of change of some property (here velocity) of a small parcel of a solid or fluid. The $\mathbf{v} \cdot \boldsymbol{\nabla}$ term takes into account the fact that the parcel moves with time. The material derivative allows a time derivative to be taken inside an integral over a time-varying domain, i.e. $\frac{\mathrm{d}}{\mathrm{d}t} \int_{\delta\Omega_t} \mathbf{v} \,\mathrm{d}V = \int_{\delta\Omega_t} \frac{\mathrm{D}\mathbf{v}}{\mathrm{D}t} \,\mathrm{d}V$.

Since $\delta\Omega_0$ is arbitrary, the Lagrangian equation of motion is

$$\frac{\partial S^{Mi}}{\partial X^M} + \rho_0 b^i = 0 \qquad \text{in } \Omega_0. \tag{3.11}$$

These equations can be written as

$$\begin{aligned} \operatorname{div}(\sigma) + \rho\mathbf{b} &= 0 \qquad \text{in } \Omega \\ \operatorname{Div}(S) + \rho_0\mathbf{b} &= 0 \qquad \text{in } \Omega_0 \end{aligned}$$

where Div represents divergence in undeformed space, and div represents divergence in deformed space. Finally, we replace the 1st Piola-Kirchoff stress with the 2nd Piola-Kirchoff stress using (3.6), so our final equation of state is

$$\frac{\partial}{\partial X^M}\left(T^{MN}\frac{\partial x^i}{\partial X^N}\right) + \rho_0 b^i = 0 \quad \text{in } \Omega_0. \tag{3.12}$$

3.2.2 Boundary Conditions

There are two types of boundary condition that can be specified: displacement boundary conditions on one part of the boundary (Dirichlet boundary conditions), and traction (i.e. force per unit area) boundary conditions on another part (Neumann boundary conditions). We define $\partial\Omega_0^{\text{disp}} \subset \partial\Omega_0$ and $\partial\Omega_0^{\text{trac}} \subset \partial\Omega_0$ as these two regions (and assume $\partial\Omega_0^{\text{disp}} \cup \partial\Omega_0^{\text{trac}} = \partial\Omega_0$ and $\partial\Omega_0^{\text{disp}} \cap \partial\Omega_0^{\text{trac}} = \emptyset$). The boundary conditions are

$$\begin{aligned} \mathbf{u} &= \mathbf{u}^{\text{disp}} \qquad \text{on } \partial\Omega_0^{\text{disp}} \\ S^{Mi}N_M &= s^i(\mathbf{X}, F) \qquad \text{on } \partial\Omega_0^{\text{trac}} \end{aligned}$$

where $\mathbf{u}^{\text{disp}}$ and $\mathbf{s}$ are the prescribed displacement and traction respectively.

The boundary conditions are therefore mixed Dirichlet-Neumann boundary conditions. For the problem to be well-posed, $\partial\Omega_0^{\text{disp}}$ needs to be non-empty[11], as otherwise the solution will only be defined up to a rigid body displacement. Note that $\mathbf{s}$ can be dependent on the deformation as well as position, which means we can essentially specify tractions on the edge of the deformed body (i.e. on $\partial\Omega^{\text{trac}} = \mathbf{\Phi}(\partial\Omega^{\text{trac}})$). If $\mathbf{s}$ is not a function of F, $\mathbf{s} \equiv \mathbf{s}(\mathbf{X})$ only, then the tractions are known as *deadload tractions*. An example of tractions that do depend on the deformation are those that arise from an external pressure, P say, in which case the traction must act only in the deformed normal direction, so that $\sigma\mathbf{n} = -P\mathbf{n}$ on $\partial\Omega^{\text{trac}}$. The traction boundary condition in this case is $\mathbf{s} = -JPF^{-\mathrm{T}}\mathbf{N}$ on $\partial\Omega_0^{\text{trac}}$.

---

[11] Or more precisely, have non-zero area (lower-dimensional measure).

3.2.3 The Weak Formulation

Force balance formulations such as Cauchy's Law (3.10), which give differential equations determining the new position, are one of three common mathematical descriptions of physical systems. For static conservative systems, if the total energy of the system can be represented in functional form, the system can be stated as an equivalent minimisation problem: minimise the total energy. The third common type are weak formulations, in which a functional $\mathcal{F}(u, v)$ of two variables is defined and the problem stated as: find $u \in V$ such that $\mathcal{F}(u, v) = 0$ for all $v \in W$, with appropriate function spaces V and W. By analogy to the weak form, the force balance form is known as the strong form of the problem. The important point about the weak form is that the differential operator $\mathcal{F}$ will involve lower order derivatives of u than the strong form, which can be analytically and computationally easier to deal with[12].

Minimisation problems can be converted to weak formulations by taking directional derivatives (which we discuss in Section 3.3.2). However, even when energy formulations are not available, the weak formulation of a problem exists, and can be determined by multiplying the strong form by an arbitrary function (known in this context as a test function), and then integrating over the domain. Alternatively, the weak form can be derived using the principal of virtual work [Spe04]. Consider a virtual displacement $\delta\mathbf{x}$, which is defined as an infinitesimal displacement compatible with the boundary conditions (so that $\mathbf{x} + \delta\mathbf{x} = \mathbf{x}^{\text{disp}}$ on $\partial\Omega_0^{\text{disp}}$, ie $\delta\mathbf{x} = 0$ on $\partial\Omega_0^{\text{disp}}$). The principal of virtual work states that the total work done by any virtual displacement is zero, which means that

$$\int_{\partial\Omega} \sigma^{ij} n_i \delta x_j \, \mathrm{d}S = \int_{\partial\Omega^{\text{trac}}} s^i \delta x_i \, \mathrm{d}S.$$

The divergence theorem can be used to expand the left-hand side of the equation:

$$\int_{\Omega} \frac{\partial \sigma^{ji}}{\partial x^j} \delta x_i + \sigma^{ji} \frac{\partial(\delta x_i)}{\partial x^j} \, \mathrm{d}V = \int_{\partial\Omega^{\text{trac}}} s^i \delta x_i \, \mathrm{d}S.$$

Cauchy's law of motion (3.10) can now be used to eliminate the stress derivative from this equation, leading to

$$\int_{\Omega} \sigma^{ij} \frac{\partial(\delta x_j)}{\partial x^i} \, \mathrm{d}V = \int_{\Omega} \rho b^i \delta x_i \, \mathrm{d}V + \int_{\partial\Omega^{\text{trac}}} s^i \delta x_i \, \mathrm{d}S. \tag{3.13}$$

Similar analysis could have been carried out using the undeformed body and the 1st Piola-Kirchoff stress tensor:

$$\int_{\Omega_0} S^{Mi} \frac{\partial(\delta x_i)}{\partial X^M} \, \mathrm{d}V_0 = \int_{\Omega_0} \rho_0 b^i \delta x_i \, \mathrm{d}V_0 + \int_{\partial\Omega_0^{\text{trac}}} s^i \delta x_i \, \mathrm{d}S_0.$$

---

[12] Since the weak form involves lower order derivatives, requirements on the solution are loosened (hence the term 'weak'), another reason why the weak form is useful. For example, the ODE $u'' = g$ requires u to have a second derivative, and a solution may not strictly exist if, for example, g is not continuous. However, the weak form of this ODE is $\int u'v' - gv = 0$ for all $v \in W$ (for some suitable space W), which admits solutions that do not have a strict second derivative. The weak form is actually the natural statement of the problem in physical situations where the forcing terms or boundary conditions are not continuous.

Finally, we replace the 1st Piola-Kirchoff stress with the second Piola-Kirchoff stress:

$$\int_{\Omega_0} T^{MN} \frac{\partial x^i}{\partial X^M} \frac{\partial(\delta x_i)}{\partial X^N} \, \mathrm{d}V_0 = \int_{\Omega_0} \rho_0 b^i \delta x_i \, \mathrm{d}V_0 + \int_{\partial\Omega_0^{\mathrm{trac}}} s^i \delta x_i \, \mathrm{d}S_0. \tag{3.14}$$

This equation has to be true for any virtual displacement $\delta\mathbf{x}(\mathbf{X})$ satisfying $\delta\mathbf{x} = 0$ on $\partial\Omega_0^{\mathrm{disp}}$, and is the weak form of the problem, essentially an orthogonality relation which the solution must satisfy. We will use the weak formulation to set up equations to be solved using the finite element method.

For certain types of material, we can also derive an energy minimisation formulation of large deformation problems. To do so, we need to define the *strain energy*, which is carried out in the next section, in the context of the constitutive relation.

3.3 The Constitutive Law and Strain Energy Functions

To relate the kinematic analysis to the conservation laws we need a relationship between stress and the kinematic variables. This relationship is known as the *constitutive relation.* It is a material dependent law, the exact form of which can only be determined experimentally. The stress analysis carried out so far could apply equally to fluids (liquid or gas) or solids; it is only in the relation between stress and the kinematic variables that the mechanics of these states differ. For idealised fluids, functional forms of constitutive laws can be derived, for example, a non-viscous fluid has an isotropic stress of the form $\sigma^{ij} = p\delta^{ij}$, where p is an internal pressure, whereas a Newtonian viscous fluid has a term which is linear in the strain rate: $\sigma^{ij} = -p\delta^{ij} + D^{ijKL}V_{KL}$, where $V_{KL} = \frac{\mathrm{d}}{\mathrm{d}t}(E_{KL})$ is the strain rate [Fun94] and D^{ijKL} a fourth-order tensor of viscosity coefficients.

Functional forms for constitutive laws can be proposed for different types of solid, as we shall see, but determining the parameters of the law experimentally is usually a difficult task. This is especially true for biological tissues, since ideally experiments should be carried out in *in vivo.* Experiments are generally carried out *ex vivo* on tissue samples, but it is difficult to keep excised tissue in normal living conditions when analysing stress-strain properties. Experimental data is often unreliable, and can vary by orders of magnitude in certain cases [Fun94].

The constitutive relation is often extremely complex for real materials. Simple relationships can be formed for idealised materials, such as Hooke's Law (see for example [Gou99]), which states that strain is proportional to stress. This may be approximately true for a given material for small ranges of strain, stress, temperature etc., but can probably never be true in general - suggested by the fact that virtually every known solid material can be broken by sufficiently large stresses. Still, Hooke's Law is useful, because most structural materials obey it in their normal range of stresses and strains. However, soft biological materials rarely obey Hooke's Law in their typical strain ranges;

the stress-strain relationships for living tissue is usually nonlinear. Tissues are also anisotropic in general, further complicating the stress-strain relationship, as well as heterogeneous, in which case the constitutive law is a function of position.

There is also more complex mechanical behaviour that can be exhibited by materials. When a body is suddenly strained and the strain maintained constant afterward, the corresponding stresses induced in the body can decrease with time, a phenomenon known as *stress relaxation*. If the body is suddenly stressed and the stress maintained constant, the body can continue to deform, a phenomenon known as *creep*. If the body is subjected to cyclic loading, the stress-strain relationship for the loading process is usually different to that for the unloading process, a phenomenon known as *hysteresis*. This behaviour is demonstrated in Figure 3.2. If the body is loaded for a second time (following the same procedure), the loading curve is often different to the first loading curve. However, if the loading-unloading process is repeated many times, the difference in the stress-strain curves between cycles eventually disappear, a phenomenon known as *preconditioning*. The reason for this is that the internal structure of the material has changed with the cycling. By repeated cycling, a steady state will eventually be reached. However, changing the cycling procedure (e.g. by loading to a different maximum stress) will change the internal structure again, and the material will no longer be preconditioned. In general, the constitutive relation of a material should be a function not only of strain, but of strain, history of strain and strain rate. Materials displaying phenomena such as hysteresis are known as *viscoelastic*. *Elastic* materials are defined as materials which return to their original state when loaded and then unloaded, in other words, materials for which there is a one-to-one relationship between stress and strain (at a constant temperature).

In static problems, we obviously only have information about the starting and ending states. Throughout this work, we neglect viscoelastic effects and assume pure elasticity for all tissues, so stresses are functions of strain only, not strain rate or history of strain. This can be partially justified. Preconditioned tissues can be considered as *pseudo-elastic*, a concept introduced in [Fun94] to simplify viscoelastic constitutive laws dramatically. In preconditioned samples hysteresis type effects can be ignored, and there are effectively two separate one-to-one stress-strain laws that need to be considered: the loading stress-strain relationship and the unloading relationship. Thus, a preconditioned viscoelastic material can be considered as one pure elastic material during loading, and a different elastic material during unloading. Although tissue *in vivo* is not preconditioned, samples for which experimental laws have been found usually are. Also biological tissues have been shown to have elastic properties that are largely insensitive to strain rate, to several orders of magnitude [Fun94], and we can ignore the history of strain since we are only looking at steady-state solutions. The simplification of pseudo-elasticity to pure elasticity is a necessary first assumption.

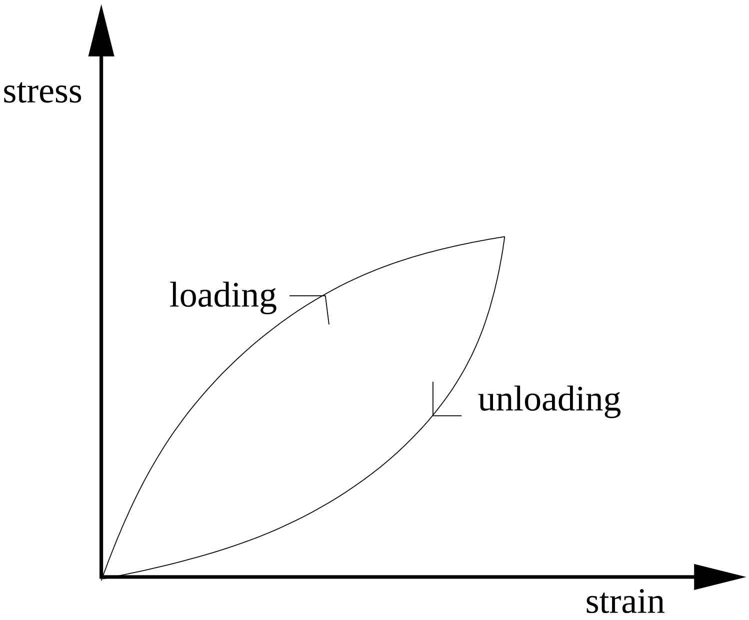

Figure 3.2: A sample hysteresis curve, demonstrating the path-dependent relationship between stress and strain.

3.3.1 The General Form of the Constitutive Relation

The constitutive relation can be expressed in the form

$$\sigma = \mathcal{G}(F)$$

for some symmetric tensor-valued function $\mathcal{G}$. $\mathcal{G}$ is known as the *response function.* Materials for which such a function exists are called *Cauchy elastic materials.* We assume that $\mathcal{G}(I) = 0$, i.e. that the stress vanishes throughout the body in the undeformed state (so that the undeformed state is the *natural configuration* or *reference configuration*).

We will formulate our constitutive relation using T and E rather than σ and F. This is equivalent because we have a relation between σ and T, and because $\mathcal{G}$ must only be a function of U (where $F = RU$)[13], and E is certainly an invertible function of U (as $E = \frac{1}{2}(U^2 - I)$ with U symmetric and positive definite). To formulate the constitutive relationship between T and E, we will make use of the strain energy function.

3.3.2 The Strain Energy Function

The strain energy function is defined as a scalar potential, W, which is a function of the components of the Green's strain tensor, and which relates the 2nd Piola-Kirchoff stress tensor to the strain

---

[13] Due to the axiom of objectivity for constitutive laws, see [Ogd97] for details.

tensor is the following way[14]:

$$T^{MN} = \frac{1}{2}\left(\frac{\partial W}{\partial E_{MN}} + \frac{\partial W}{\partial E_{NM}}\right).$$

Alternatively, if we demand that the strain energy is written in a way that is symmetric in the (equal) strains E_{MN} and E_{NM}, then the relationship between the stress and the strain energy function is simply

$$T^{MN} = \frac{\partial W}{\partial E_{MN}},$$

(or $T^{MN} = 2\frac{\partial W}{\partial C_{MN}}$ if W is written as a function of C).

Materials for which a strain energy function exists are known as *hyperelastic.* For perfectly elastic materials, the existence of the strain energy function in the case of either an isoentropic or an isothermal deformation can be justified using a thermodynamic argument [Fun94].

The strain energy can also be used to obtain other stress tensors, by taking derivatives with respect to other variables. For example, consider $\frac{\partial W}{\partial F^i_M}$:

$$\begin{aligned}\frac{\partial W}{\partial F^i_M} &= \frac{\partial W}{\partial E_{PQ}}\frac{\partial E_{PQ}}{\partial F^i_M}\\ &= T^{PQ}\frac{\partial E_{PQ}}{\partial F^i_M}.\end{aligned}$$

Using $E_{PQ} = \frac{1}{2}(F^j_P F^j_Q - \delta_{PQ})$ it is straightforward to show that it follows that

$$\begin{aligned}\frac{\partial W}{\partial F^i_M} &= T^{PM}F^i_P\\ &= S^{Mi}.\end{aligned} \tag{3.15}$$

If a deformable object is hyperelastic, the strong and weak forms can be re-stated as an energy minimisation problem. Suppose the body is subjected to body forces $\mathbf{b}$ and surface tractions $\mathbf{s}$ as in Section 3.2.1. The total energy as a function of $\mathbf{x}$ is given by

$$\mathcal{E}(\mathbf{x}) = \int_{\Omega_0} W(E_{MN}(\mathbf{x})) - \rho_0\mathbf{b}\cdot\mathbf{x}\,\mathrm{d}V_0 - \int_{\partial\Omega_0^{\mathrm{trac}}} \mathbf{s}\cdot\mathbf{x}\,\mathrm{d}S_0 \tag{3.16}$$

and the minimisation problem is

$$\underset{\mathbf{x}\,:\,\mathbf{x}(\partial\Omega_0^{\mathrm{disp}})=\mathbf{x}^{\mathrm{disp}}}{\text{minimise}} \int_{\Omega_0} W\left(E_{MN}\right) - \rho_0\mathbf{b}\cdot\mathbf{x}\,\mathrm{d}V_0 - \int_{\partial\Omega_0^{\mathrm{trac}}} \mathbf{s}\cdot\mathbf{x}\,\mathrm{d}S_0. \tag{3.17}$$

To see that this is equivalent to the weak formulation (3.14), we need to make use of *directional derivatives.* Let $\mathbf{x}$ be a function and $f(\mathbf{x})$ a functional. The directional derivatives of f in the direction $\mathbf{x}^*$ at $\mathbf{x}$ is defined to be the scalar

$$D_{\mathbf{x}^*}[f(\mathbf{x})] = \lim_{\epsilon\longrightarrow 0}\frac{f(\mathbf{x}+\epsilon\mathbf{x}^*) - f(\mathbf{x})}{\epsilon},$$

---

[14] Note that this definition implies that the strain energy satisfies $\frac{\mathrm{d}W}{\mathrm{d}t} = T^{MN}\frac{\mathrm{d}}{\mathrm{d}t}\left(E_{MN}\right)$. This can be used as the definition of strain energy.

if the limit exists. (Note that we will generally denote the directions by $\delta\mathbf{x}$ rather than $\mathbf{x}^*$). Now, suppose the functional f is minimised by $\mathbf{x}_{\text{min}}$. This is equivalent to

$$D_{\delta\mathbf{x}}[f(\mathbf{x}_{\text{min}})] = 0 \qquad \forall\, \delta\mathbf{x}.$$

Consider the total energy given by (3.16). Clearly

$$D_{\delta\mathbf{x}}\left[\int_{\Omega_0} \rho_0 \mathbf{b}\cdot\mathbf{x}\,\mathrm{d}V_0\right] = \int_{\Omega_0} \rho_0 \mathbf{b}\cdot\delta\mathbf{x}\,\mathrm{d}V_0,$$

and similarly with the traction term, so in order to show that the directional derivative of the total energy is the weak integral (3.14) (and hence that the weak formulation is equivalent to the energy minimisation problem), it is required to prove that

$$D_{\delta\mathbf{x}}\left[\int_{\Omega_0} W(E_{MN})\,\mathrm{d}V_0\right] = \int_{\Omega_0} S^{Mi}\frac{\partial \delta x^i}{\partial X^M}\,\mathrm{d}V_0.$$

This follows easily by considering W as a function of F^i_M, and using the chain rule and (3.15).

Since the Green-Lagrange strain E is symmetric, the strain energy function is essentially a function of 6 independent variables. However, for certain types of material, the dependency on E can be simplified. For *isotropic materials*, which are materials for which the material properties are identical in all directions from a point, the strain energy should be invariant under orthogonal changes of the coordinate axes. It can be shown that this is equivalent to the strain energy being a function of the invariants $I_1(C)$, $I_2(C)$, and $I_3(C)$ [Ogd97]:

$$W \equiv W(I_1, I_2, I_3).$$

Transversely isotropic materials are materials for which there is a preferred direction (which can be a function of position), known as the *fibre direction*, with material properties identical in all directions in the plane orthogonal to the fibre direction. Hyperelastic transversely isotropic materials can be shown [Ogd97] to have strain energy functions of the form

$$W \equiv W(I_1, I_2, I_3, K_1, K_2)$$

where, if $\hat{\mathbf{m}} \equiv \hat{\mathbf{m}}(\mathbf{X})$ is the unit fibre direction,

$$K_1 = \hat{\mathbf{m}}^{\mathrm{T}} E \hat{\mathbf{m}},$$

i.e. the strain in the fibre direction, and

$$K_2 = \hat{\mathbf{m}}^{\mathrm{T}} E^2 \hat{\mathbf{m}} - K_1^2 = \|E\hat{\mathbf{m}}\|^2 - K_1^2.$$

Materials which are neither isotropic nor transversely isotropic are known as *anisotropic*.

Some Example Strain Energy Functions

For illustrative purposes, we now describe some sample isotropic, transverse isotropic and anisotropic strain energy functions.

It is common to model isotropic incompressible materials as *Mooney-Rivlin* [Moo40], where the (material part of the) strain energy function has the form

$$W = c_1(I_1 - 3) + c_2(I_2 - 3) \tag{3.18}$$

for some material constants c_1 and c_2 which have to be determined by experiment. (We have stated that this is only the material part of the strain energy function because incompressible materials have strain energies with a non-material-dependent term, which we will derive in Section 3.4). *Neo-Hookean* materials are Mooney-Rivlin materials for which $c_2 = 0$. Ogden [Ogd72] proposed a law which is a series expansion in the principle stretches

$$W = \sum_n c_n(\lambda_1^{\alpha_n} + \lambda_2^{\alpha_n} + \lambda_3^{\alpha_n} - 3),$$

where c_n and α_n are constants, with α_n potentially any real number.

Mooney-Rivlin strain energy functions do not model body tissues well. One strain energy function often used for body tissues, for example in [HFS03], is the Veronda exponential material law [Ver70], taking the form

$$W = \alpha\left(e^{\beta(I_1-3)} - 1\right) + \gamma(I_2 - 3), \tag{3.19}$$

or

$$W = \alpha\left(e^{\beta(I_1-3)} - 1\right). \tag{3.20}$$

An adaption of the exponential law for a transversely isotropic material is used in [Hir02], where a term of the form

$$W^{\text{fibre}} = \alpha'\left(e^{\beta'(\lambda_f^2-1)}\right) - \alpha'\beta'\lambda_f^2, \tag{3.21}$$

(where $\lambda_f^2 = \hat{\mathbf{m}}^{\mathrm{T}}C\hat{\mathbf{m}}$ is the square of the fibre stretch) is added to the strain energy.

A strain energy function for an anisotropic material which models the fact that for biological tissue the stress often tends to infinity as the strain approaches some maximal value, used in [Nas98], is the *pole-zero* strain energy:

$$W \quad = \quad \sum_{N,M=1}^{3} \frac{k_{MN}E_{MN}^2}{(E_{MN} - A_{MN})^{\beta_{MN}}} \qquad \text{(no summation convention)}$$

for some constants $k_{MN}(= k_{NM})$ which are measures of stiffness, $A_{MN}(= A_{MN})$ which are the maximal strains, and $\beta_{MN}(= \beta_{NM})$. β_{MN} is often assumed to be equal to 2, and the 12 parameters for stiffness and maximal strain have to be determined experimentally.

3.4 Incompressibility and Pressure

In modelling the deformation of biological tissues, we may wish to impose the constraint of incompressibility. In this case the stress-strain relationship must be treated differently, as can be seen by considering a cube of an incompressible material that has equal forces applied on each face. Since the material is incompressible, there will be no deformation, i.e. no strain, but internal stresses must exist to balance the applied surface tractions. Thus the stress in general for an incompressible material cannot be determined only from the strain - there is an internal pressure which is independent of the strain, and which is a new variable that may need to be computed.

In order to deal with incompressibility we summarise the general theory of constraints. Recall that the constitutive relation can be written in the form

$$\sigma = \mathcal{G}(F).$$

Any constraint can be expressed as

$$\mathcal{K}(F) = 0,$$

for some scalar function $\mathcal{K}$. For incompressibility, the constraint is

$$\mathcal{K}(F) \equiv \det(F) - 1 = 0.$$

Another constraint that is often required is that of inextensibility. Here there is a direction, given by a unit vector $\hat{\mathbf{m}} \equiv \hat{\mathbf{m}}(\mathbf{X})$, along which no extension is allowed. In this case $\mathcal{K}$ is given by

$$\mathcal{K}(F) = \hat{\mathbf{m}} \cdot (C\hat{\mathbf{m}}) - 1,$$

since $\hat{\mathbf{m}} \cdot (C\hat{\mathbf{m}})$ is equal to $\|F\hat{\mathbf{m}}\|^2$.

The constraint $\mathcal{K} = 0$ introduces a new term into the governing equations. Consider the weak formulation

$$\int_{\Omega_0} S^{Mi} \frac{\partial(\delta x_i)}{\partial X^M} \, \mathrm{d}V_0 = \int_{\Omega_0} \rho_0 b^i \delta x_i \, \mathrm{d}V_0 + \int_{\partial\Omega_0^{\mathrm{trac}}} s^i \delta x_i \, \mathrm{d}S_0 \qquad \forall \; \delta\mathbf{x}.$$

Since $0 = \delta\mathcal{K} = \frac{\partial\mathcal{K}}{\partial F^i_M} \frac{\partial \delta x_i}{\partial X^M}$, we can impose the constraint by adding this term multiplied by an arbitrary Lagrange multiplier, q, into the weak formulation:

$$\int_{\Omega_0} \left(S^{Mi} + q \frac{\partial\mathcal{K}}{\partial F^i_M} \right) \frac{\partial(\delta x_i)}{\partial X^M} \, \mathrm{d}V_0 = \int_{\Omega_0} \rho_0 b^i \delta x_i \, \mathrm{d}V_0 + \int_{\partial\Omega_0^{\mathrm{trac}}} s^i \delta x_i \, \mathrm{d}S_0 \qquad \forall \; \delta\mathbf{x}.$$

It follows that we should replace the old 1st Piola-Kirchoff stress for an unconstrained material with

$$S^{\mathrm{constr}} = S^{\mathrm{unconstr}} + q \frac{\partial\mathcal{K}}{\partial F^{\mathrm{T}}}$$

for a constrained material. Since $S = JF^{-1}\sigma$, this means that the Cauchy stress is replaced by

$$\sigma^{\mathrm{constr}} = \sigma^{\mathrm{unconstr}} + qJ^{-1}F \frac{\partial\mathcal{K}}{\partial F^{\mathrm{T}}}$$

or (dropping the superscript)

$$\sigma = \mathcal{G}(F) + qJ^{-1}F\frac{\partial \mathcal{K}}{\partial F^{\mathrm{T}}}.$$

The second term is known as the constraint stress and is not otherwise derivable from material properties.

For incompressibility we have[15] $\frac{\partial \mathcal{K}}{\partial F^{\mathrm{T}}} = \frac{\partial \det(F)}{\partial F^{\mathrm{T}}} = \det(F)F^{-1} = JF^{-1}$, so

$$\sigma = \mathcal{G}(F) + qI.$$

We now define the *pressure* (sometimes referred to as the hydrostatic pressure), p, by $p = -q$, so

$$\sigma = \mathcal{G}(F) - pI.$$

The equivalent form of the 2nd Piola-Kirchoff stress is obtained by using the relationship between the Cauchy stress and the 2nd Piola-Kirchoff stress (3.7); it is

$$T^{MN} = \frac{\partial \overline{W}(I_1, I_2)}{\partial E_{MN}} - p(C^{-1})^{MN}, \tag{3.22}$$

where $\overline{W}(I_1, I_2)$ is the material-dependent part of the strain energy. The pressure $p \equiv p(\mathbf{X})$ is a scalar field and an unknown which needs to be computed in calculating the new position. The four equations which determine the four unknowns $\mathbf{x}$ and p are Cauchy's equations (3.10) (or (3.11)) and the constraint $\det(F) = 1$. The full weak form in this situation is obtaining by summing each of Cauchy's equations, multiplied by δx_i, and the constraint multiplied by a variation in the pressure, δp, resulting in:

$$\begin{aligned} 0 = \quad & \int_{\Omega_0} T^{MN}\frac{\partial x^i}{\partial X^M}\frac{\partial(\delta x_i)}{\partial X^N}\,\mathrm{d}V_0 - \int_{\Omega_0} \rho_0 b^i \delta x_i\,\mathrm{d}V_0 \\ & + \int_{\partial\Omega_0^{\mathrm{trac}}} s^i \delta x_i\,\mathrm{d}S_0 + \int_{\Omega_0} \delta p(\det F - 1)\,\mathrm{d}V = 0 \qquad \forall \delta\mathbf{x}, \delta p. \end{aligned} \tag{3.23}$$

Finally, because $\frac{\partial I_3}{\partial E_{MN}} = 2I_3(C^{-1})^{MN}$, and using the definition of the strain energy function $T^{MN} = \frac{\partial W}{\partial E_{MN}}$, we see that the full strain energy function for incompressible materials has the form

$$W^{\mathrm{incomp}} = \overline{W}(I_1, I_2) - \frac{p}{2}(I_3 - 1).$$

3.5 Linear Elasticity and Linear Material Laws

The theory of linear elasticity is the theory of small-strain deformations, in which the strain is linearised in the displacement. This type of linearity is known as *geometric linearity.* There is a second type of linearity encountered in elasticity: *material linearity*, which refers to constitutive

---

[15] Here we use the result that for a matrix A, $\frac{\partial(\det(A))}{\partial A_{pq}} = \det(A)(A^{-1})_{qp}$.

laws which state that stress is linearly related to strain. Linear elastic problems are generally both geometrically and materially linear, however nonlinear elasticity problems can still involve linear materials.

In linear elasticity we assume the displacement is small and linearise quantities accordingly. Consider the Lagrangian strain given by (3.4).

$$\begin{aligned} E_{MN} &= \frac{1}{2}(C_{MN} - \delta_{MN}) \\ &= \frac{1}{2}\left(\frac{\partial(X^i + u^i)}{\partial X^M}\frac{\partial(X^i + u^i)}{\partial X^N} - \delta_{MN}\right) \\ &= \frac{1}{2}\left(\frac{\partial X^i}{\partial X^M}\frac{\partial X^i}{\partial X^N} + \frac{\partial u^M}{\partial X^N} + \frac{\partial u^N}{\partial X^M} + \frac{\partial u^i}{\partial X^M}\frac{\partial u^i}{\partial X^N} - \delta_{MN}\right) \\ &= \frac{1}{2}\left(\frac{\partial u^M}{\partial X^N} + \frac{\partial u^N}{\partial X^M}\right) + \text{terms quadratic in } u. \end{aligned}$$

Hence, the *infinitesimal strain tensor* is defined to be

$$\epsilon_{ij} = \frac{1}{2}\left(\frac{\partial u^i}{\partial X^j} + \frac{\partial u^j}{\partial X^i}\right).$$

Assuming material linearity, so the Cauchy stress is linearly related to the infinitesimal strain, we have

$$\sigma^{ij} = C^{ijkl}\epsilon_{kl},$$

for some fourth-order tensor C^{ijkl}. If the material is isotropic, C must be an isotropic fourth-order tensor. It can be shown that any isotropic fourth-order tensor has to have the form

$$C^{ijkl} = \alpha\delta^{ij}\delta^{kl} + \beta\delta^{ik}\delta^{jl} + \gamma\delta^{il}\delta^{jk},$$

for some functions of space α, β and γ. Substituting, and defining $\lambda = \alpha(\mathbf{x})$ and $2\mu = \beta(\mathbf{x}) + \gamma(\mathbf{x})$, it can easily be seen that

$$\sigma_{ij} = \lambda u_{k,k}\delta_{ij} + \mu(u_{i,j} + u_{j,i}),$$

or

$$\sigma_{ij} = \lambda\epsilon_{kk}\delta_{ij} + 2\mu\epsilon_{ij}. \tag{3.24}$$

We see that any linear isotropic material has its elastic properties determined by two parameters. λ and μ are known as the *Lamé coefficients*.

Now, let us define

$$E = \frac{\mu(3\lambda + 2\mu)}{\lambda + \mu}, \tag{3.25}$$

and

$$\nu = \frac{\lambda}{2(\lambda + \mu)}. \tag{3.26}$$

E is known as the *Young's modulus* for the material, and determines the strain in the same direction as stress. It takes values in $(0, \infty)$, larger values representing stiffer materials. ν is known as *Poisson's ratio*, and determines strain in orthogonal directions to stress. It measures a material's ability to resist compression, and is usually positive[16]. It is always less than $\frac{1}{2}$, with values near $\frac{1}{2}$ representing highly incompressible materials. Note $\nu = \frac{1}{2}$ is equivalent to λ being infinite and $E = 3\mu$. Now, consider the incompressibility constraint $\det(F) = 1$. Since

$$\det(F) = \begin{bmatrix} 1 + \frac{\partial u^1}{\partial X^1} & \frac{\partial u^1}{\partial X^2} & \frac{\partial u^1}{\partial X^3} \\ \frac{\partial u^2}{\partial X^1} & 1 + \frac{\partial u^2}{\partial X^2} & \frac{\partial u^2}{\partial X^3} \\ \frac{\partial u^3}{\partial X^1} & \frac{\partial u^2}{\partial X^1} & 1 + \frac{\partial u^3}{\partial X^3} \end{bmatrix},$$

to highest order $\det(F) - 1 = \frac{\partial u^1}{\partial X^1} + \frac{\partial u^2}{\partial X^2} + \frac{\partial u^3}{\partial X^3} = \boldsymbol{\nabla} \cdot \mathbf{u}$, which we note is also the trace of the strain. Therefore the linearised constraint of incompressibility is $\epsilon_{kk} = 0$. One possibility of implementing incompressibility indirectly, used throughout the literature, is to set ν to be a value close to 0.5, such as 0.49 (see, for example, [JP99] (foot modelling) or [RGC01] (tongue modelling)), 0.495 (e.g. [DDP$^+$04] (breast modelling)), or 0.499 (e.g [GLLG$^+$05] (brain modelling)). However this suffers from the user having to make a somewhat arbitrary choice of the value of ν. An alternative possibility to impose incompressibility is to use the same approach as in the nonlinear theory and introduce a pressure field, p, as an lagrange multiplier. $-p$ replaces $\lambda\epsilon_{kk}$ in the stress, which means the stress in the incompressible case is

$$\sigma^{ij} = -p\delta^{ij} + \mu\epsilon^{ij}, \tag{3.27}$$

subject to the constraint $\epsilon_{kk} = 0$.

Finally, we need to derive the linear equation of state. Consider Cauchy's equation written in terms of the undeformed configuration (3.11). The 1st Piola-Kirchoff stress is related to the Cauchy stress through (3.5), and this is expanded in terms of displacement, we see that, since $J = 1$ and $F^i_M = \delta^i_M$ to first order, the 1st Piola-Kirchoff stress is identical to the Cauchy stress to first order. Thus, the linearised form of (3.11) is

$$\frac{\partial \sigma^{ij}}{\partial X^j} + \rho_0 b^i = 0, \tag{3.28}$$

where σ_{ij} is given by (3.24). Essentially this means that in linear elasticity the distinction between derivatives with respect to deformed and undeformed coordinates does not need to be considered.

In this chapter we have derived the equations of finite elasticity, discovering that they are highly nonlinear. In order to solve finite elasticity problems, a numerical scheme is necessary. In the next chapter, we discuss the most common computational method for solving solid mechanics problems, the finite element method.

---

[16] But not always: some strange materials have negative Poisson's ratio. A bar of such material will expand in cross-sectional radius when stretched.

Chapter 4

The Finite Element Method for Solving Problems in Nonlinear Elasticity

Mathematical modelling can be used to study and predict behaviour in a huge range of physical and biological processes. The aim of mathematical modelling is to develop systems of equations that describe the behaviour of the process or a simplified version of the process. These systems of equations will often take the form of coupled ordinary or partial differential equations (ODEs or PDEs), which, in general, cannot be solved analytically. Analytical solutions may exist for simple cases, and analytical methods may sometimes be applied to simplify or reduce problems, but mathematical models generally become intractable as they get more complex, especially when nonlinear effects are included, or when the number of spatial dimensions increases, or crucially, when the domain is allowed to take arbitrary shapes. The equations of nonlinear elasticity developed in Chapter 3 are highly nonlinear and over three dimensions, and we wish to solve them over non-regular shaped domains, which necessitates the use of a numerical scheme to compute approximate solutions to these equations.

Various methods have been developed for numerically integrating ordinary and partial differential equations, such as finite difference methods [FW04], finite volume methods [VM96], finite element methods [Red93], and spectral methods [Tre01]. Of these, the finite difference method is the most elementary, being founded upon Taylor's Theorem, and involving the approximation of derivatives of a variable with linear combinations of values of that variable at nearby points. The finite volume method is used extensively in computational fluid dynamics applications. Whereas in the finite difference method the PDE is approximated only at a grid-point (a 'node'), in the finite volume method the PDE is approximated at 'control regions', which are volumes that enclose a single node. The PDE is integrated over the control region, allowing boundary conditions to be applied through use of the divergence theorem. The finite element method also involves integrating the PDE over small volumes, here called 'elements', but is more complex than the finite volume method as quantities of interest are allowed to take various possible functional forms, depending on the types of 'basis functions' (defined shortly) used. In the finite element method, the basis functions are generally piecewise polynomials and are non-zero only in small regions (i.e. have local compact support). Spectral methods can be considered as an analogous method in which globally smooth basis functions are used, either trigonometric polynomials or orthogonal polynomials, together with specific quadrature rules. They work incredibly effectively on problems involving highly smooth functions and simple domains [Tre01].

Of the numerical methods available, the best-suited for computation of solutions of solid mechanics problems is by far the finite element (FE) method. The FE method was first proposed in 1943 by Richard Courant [Cou43], although most of the essential mathematical analysis behind the method was not carried out until the 1960s. In the FE method, differential equations are reformulated as equations in terms of integrals over the domain, involving functions from infinite-dimensional function spaces. By approximating these functions by functions from *finite-dimensional* function spaces we obtain the finite element problem, an integral problem over a finite-dimensional space that can be rewritten as a linear matrix problem (or a finite set of nonlinear equations) and solved using numerical techniques in linear algebra. Since the finite element method involves integral equations and the weak formulation of a problem, it has several advantages over other numerical schemes for solving partial differential equations. The finite element mesh is a set of regular-shaped volume elements which together comprise the whole domain. Since numerical integration over any of these elements is a straightforward procedure, the difficulty in dealing with irregularly shaped domains in removed from the numerical method (although moved to the meshing stage). Also, Neumann boundary conditions are handled through assimilating them into the weak formulation, so, unlike the finite-difference method, the finite element method provides a natural and consistent way to deal with boundary conditions. Similarly, the finite element method is good for problems involving distinct regions of different properties, since the equations are written as integral equations and the domain of integration can be split into different regions of different properties.

4.1 Introduction to the Finite Element Method

4.1.1 Nodes, Elements and Basis Functions

The first stage in using the finite element method to solve a problem numerically is to create a *mesh* which encompasses the domain of the problem. A mesh is a collection of *nodes*, which are points in the domain, and *elements*, which are regular-shaped subsets of the domain, generally tetrahedra or hexahedra (or their lower dimensional equivalents in 1D or 2D), and having nodes for vertices. Let us formalise this. Let Ω_0 be the domain over which the problem has been posed. Define the nodes as a set of points $\{\mathbf{X}_i \in \Omega_0,\, i = 1, \ldots, \mathcal{N}\}$, and let $K_i \subset \Omega_0$, for $i = 1, \ldots, \mathcal{M}$, be the regular-shaped subsets satisfying $K_i \cap K_j = \emptyset$ if $i \neq j$, and containing only nodes at vertices. A node may not necessarily be a vertex of an element however. Define the *computational domain*, Ω_0^h, to be the union of the closures[1] of the elements

$$\Omega_0^h = \bigcup_{i=1}^{\mathcal{M}} \bar{K}_i.$$

We demand that if a node is a vertex for any element, it is also a vertex of any neighbouring element. A node which does not satisfy this property is known as a *hanging node*, in which case the mesh is known as *inadmissable*. A mesh of a two-dimensional computational domain with 10 nodes and 10 elements is shown in Figure 4.1(a), and an element with a hanging node is shown in Figure 4.1(b). The computational domain is the domain over which integrands which appear in the weak formulation will be integrated, and elements are regions over which integration will be considered separately.

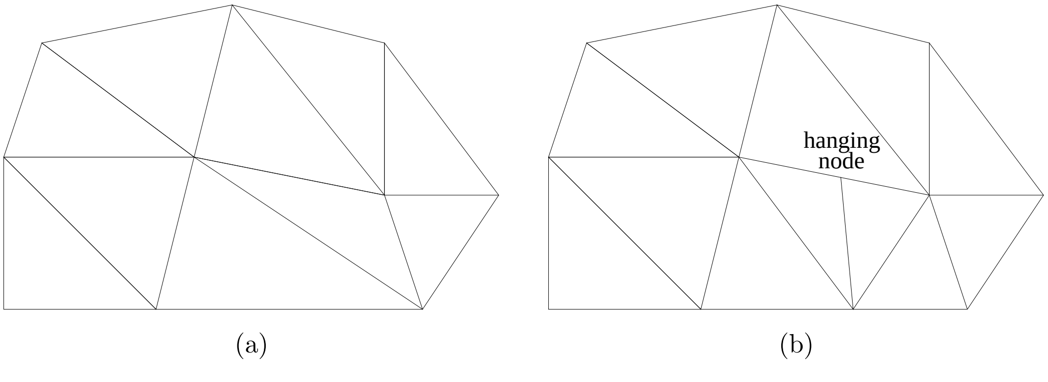

(a) (b)

Figure 4.1: (a) A sample computational domain, with 10 nodes and 10 elements. (b) A similar but inadmissable computational domain with a hanging node.

In the finite element method, unlike the finite difference method, values of the numerical solution

---

[1] By convention elements are taken to be open.

can be calculated at any point in the mesh, because the numerical solution is assumed to take a known *functional form*, and not just consist of a set nodal values. The numerical solution *is* computed at the nodes however, like in finite difference methods. To interpolate values of the numerical solution at non-nodal points from their nodal values, we need to define *basis functions.* These functions take the form of polynomials (in 1, 2, or 3 variables) over any particular element, and (apart from the case of piecewise constant basis functions), are continuous throughout the computational domain and satisfy the important property that

$$\phi_i(\mathbf{X}_j) = \delta_{ij}, \tag{4.1}$$

where ϕ_i is the basis function associated with the node $\mathbf{X}_i$. As well as the unknown quantity being numerically computed, basis functions are used to interpolate any quantity of interest in the computational domain. Different orders of basis function can (and will) be used to interpolate different quantities. Suppose we have a quantity u whose nodal values are u_i at the nodes $\mathbf{X}_i$. Given a set of basis functions $\{\phi_j\}$ associated with each node, the functional form of u when interpolated by these basis functions is given by

$$u = \sum_{i=1}^{\mathcal{N}} u_i \phi_i \qquad \text{or} \qquad u(\mathbf{X}) = \sum_{i=1}^{\mathcal{N}} u_i \phi_i(\mathbf{X}).$$

Note that this then satisfies $u(\mathbf{X}_j) = u_j$, because of (4.1).

Piecewise Constant Basis Functions

The lowest order bases are *piecewise constant basis functions*, which are, as noted above, slightly different from the other basis functions. In this case there is one basis function associated with each element, which takes the value one inside the element, and zero on every other element. In other words, the piecewise constant basis function for an element, K_i say, is the indicator function for K_i. We will explicitly use the standard notation for indicator functions, $\mathbb{1}_{K_i}$, to denote piecewise constant basis functions.

$$\mathbb{1}_{K_i}(\mathbf{X}) = \begin{cases} 1 & \text{if } \mathbf{X} \in K_i \\ 0 & \text{otherwise.} \end{cases} \tag{4.2}$$

A function interpolated with piecewise constant basis functions will be of the form

$$p = \sum_{i=1}^{\mathcal{M}} p_i \mathbb{1}_{K_i}$$

where p_i is the value of the function in element K_i.

Piecewise (Bi/Tri-)Linear Basis Functions

The next highest order basis functions are those which are linear in each element. The constraint (4.1) and the condition of linearity in each element uniquely fixes the set of basis functions. In one dimension they are easily seen to be the so-called 'hat' functions as shown in Figure 4.2. In higher dimensions we have to consider different types of elements separately. If the element is triangular or tetrahedral the basis functions are linear, if the element is rectangular or hexahedral the basis functions are bilinear or trilinear.

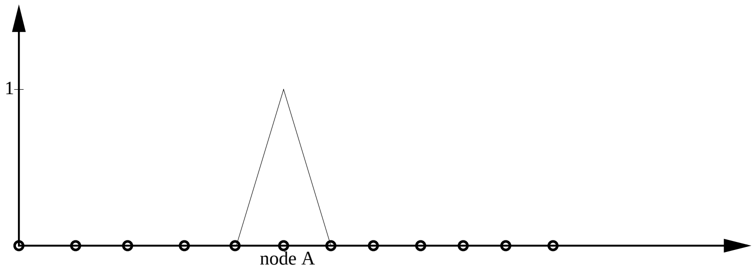

Figure 4.2: The 1D 'hat' functions are piecewise linear functions taking the value 1 at one node and zero at all other nodes. Here the hat function for the node labelled 'A' is shown.

To specify the basis functions, it is easiest to just consider the bases possible on one single element. Consider the different *canonical elements.* The canonical triangle (Figure 4.3(a)) is the triangle on the plane with vertices $(0,0),(1,0),(0,1)$ labelled 1, 2, and 3 respectively, and the canonical tetrahedron (Figure 4.3(b)) is the tetrahedron in 3D space with vertices $(0,0,0),(1,0,0),(0,1,0),(0,0,1)$ labelled 1, 2, 3 and 4 respectively. The linear basis functions on these elements are the linear functions which take the value one on one node and zero on the others. If the coordinates of the canonical element are (ξ,η) and (ξ,η,ν) in two- and three-dimensional space respectively, the linear basis functions are

$$\begin{aligned} N_1(\xi,\eta) &= 1-\xi-\eta \\ N_2(\xi,\eta) &= \xi \\ N_3(\xi,\eta) &= \eta \end{aligned}$$

on the canonical triangle, and

$$\begin{aligned} N_1(\xi,\eta,\nu) &= 1-\xi-\eta-\nu \\ N_2(\xi,\eta,\nu) &= \xi \\ N_3(\xi,\eta,\nu) &= \eta \\ N_4(\xi,\eta,\nu) &= \nu \end{aligned}$$

on the canonical tetrahedron.

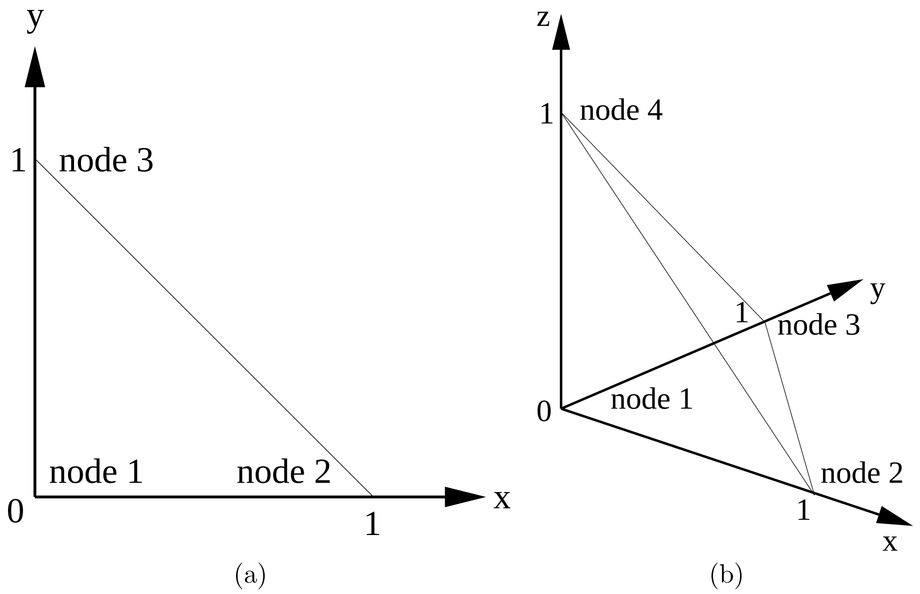

Figure 4.3: (a) The canonical triangle and (b) the canonical tetrahedron, with node orderings.

The canonical quadrilateral is the unit square $[0,1]\times[0,1]$, with nodes labelled as shown in Figure 4.4(a). The bilinear basis functions on the square are

$$\begin{aligned} N_1(\xi,\eta) &= (1-\xi)(1-\eta) \\ N_2(\xi,\eta) &= \xi(1-\eta) \\ N_3(\xi,\eta) &= (1-\xi)\eta \\ N_4(\xi,\eta) &= \xi\eta. \end{aligned}$$

Similarly, the canonical hexahedron is the unit cube $[0,1]\times[0,1]\times[0,1]$, with nodes labelled as in Figure 4.4(b), with trilinear basis functions

$$\begin{aligned}
N_1(\xi,\eta,\nu) &= (1-\xi)(1-\eta)(1-\nu) \\
N_2(\xi,\eta,\nu) &= \xi(1-\eta)(1-\nu) \\
N_3(\xi,\eta,\nu) &= (1-\xi)\eta(1-\nu) \\
N_4(\xi,\eta,\nu) &= \xi\eta(1-\nu) \\
N_5(\xi,\eta,\nu) &= (1-\xi)(1-\eta)\nu \\
N_6(\xi,\eta,\nu) &= \xi(1-\eta)\nu \\
N_7(\xi,\eta,\nu) &= (1-\xi)\eta\nu \\
N_8(\xi,\eta,\nu) &= \xi\eta\nu.
\end{aligned} \tag{4.3}$$

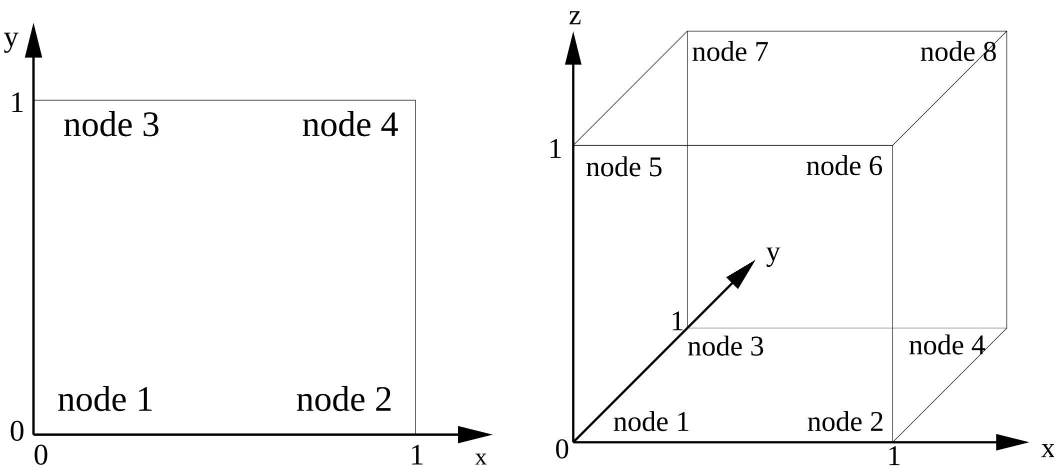

(a) (b)

Figure 4.4: (a) The canonical quadrilateral (unit square) and (b) the canonical hexahedron (unit cube), with node orderings.

Other Basis Functions

There are a number of other types of basis function that can be chosen to interpolate a field. We mention some here so they can be included in discussions later on. *Quadratic* basis functions are piecewise quadratic functions over the domain. The extra degrees of freedom with quadratic functions compared to linear means that each element must have extra nodes at which the field values are given. For example, in 1D, each element has to have three nodes (the two end-points, and one internal node). The three quadratic basis functions on the element are the three quadratic polynomials which take the value 1 at one node and zero at the other two. A field, u say, is interpolated on the element as $u = \Sigma_1^3 u_i \phi_i$ (where u_i are the three nodal values and ϕ_i the three bases), which is the unique quadratic to take the values u_i at node i. If an unknown is interpolated with quadratic basis functions, the extra degrees of freedom lead to increased numerical complexity. Quadratic basis functions are good at interpolating curved boundaries, and their use can lead to increased accuracy in the FE method.

The basis functions described thus far are all known as *Lagrange* basis functions; there is another, very different, class known as *Hermite* basis functions. These are basis functions that are used to interpolate both the field and the spatial derivative of the field over the element, which requires both nodal values and nodal derivatives to be known (or computed). An example is *cubic-Hermite* bases. In 1D, a cubic polynomial has four degrees of freedom, and the four cubic basis functions on a 2-noded 1D element are: $\{\phi_i \,:\, i = 1, 2\}$, which take the value 1 at node i, 0 at the other node, and have derivative 0 at both nodes, and $\{\psi_i \,:\, i = 1, 2\}$, which take the value 0 at both nodes, and have derivative 1 at node i and 0 at the other. If the nodal values of a field u are known, u_i say, and the nodal derivatives of u also known, t_i say, the interpolation is $u = u_1\phi_1 + u_2\phi_2 + t_1\psi_1 + t_2\psi_2$, which is the unique cubic to have value u_i and derivative t_i at node i. Cubic-Hermite bases can also lead to increased accuracy, but at the cost of a massively increased number of unknowns to be computed.

Finally, we briefly mention the *discontinuous Galerkin* (DG) finite element method [RH73]. This also uses interpolation with piecewise polynomials, but with no kind of continuity constraint on the field across element boundaries. Continuity is instead enforced 'weakly', by adding a suitable function to the weak formulation. The DG method introduces a huge number of new unknowns (as the nodal value of a field has to be computed for each element containing that node[2]), but has many advantages over the standard FE method, which include increased accuracy (despite, rather counter-intuitively, the lack of continuity), easy mesh refinement, and the ability to handle hanging nodes.

4.1.2 Weak Formulations and the Finite Element Formulation

Let $\Omega_0 \subset \mathbb{R}^n$ be compact. The first Sobolev space H^1 on Ω_0 is defined as the set of functions on Ω_0 which are square-integrable and whose first (weak) derivatives are also square-integrable. (A weak derivative is a generalisation of classical derivative useful for functions which do not have a strict classical derivative but have a classical derivative almost everywhere[3]. For precise definitions

---

[2] For example, an internal node in a 3D mesh made up of hexahedral elements could be expected to be contained in 8 elements. Therefore there would be approximately 8 times as many unknowns in a DG FEM implementation compared to a continuous FEM implementation.

[3] Almost everywhere means everywhere in the space Ω_0 except for a subset with measure zero (zero integral). For example, if $\Omega_0 = \mathbb{R}$, let f be defined by $f(x) = 1$ if $x \in \{0, 1, 2, \ldots, 10^{10^{10^{10}}}\}$ and $f(x) = 2$ otherwise. f is equal to 2 almost everywhere.

of weak derivatives and Sobolev spaces refer to [Sul02]). Let $\Gamma \subset \partial\Omega_0$, and define $H^1_0 = \{v \in H^1 : v = 0 \text{ on } \partial\Omega_0\}$, and $H^1_{0,\Gamma} = \{v \in H^1 : v = 0 \text{ on } \Gamma\}$.

The general weak problem for one unknown u subject to zero-Dirichlet boundary conditions on Γ is: find $u \in H^1_{0,\Gamma}$ such that

$$\int_{\Omega_0} f_1(u, v, \boldsymbol{\nabla} u, \boldsymbol{\nabla} v)\,\mathrm{d}V_0 + \int_{\partial\Omega_0} f_2(\boldsymbol{\nabla} u \cdot \hat{\mathbf{n}}, v)\,\mathrm{d}S_0 = \int_{\Omega_0} g(v)\,\mathrm{d}V_0 \qquad \forall\, v \in H^1_0$$

where f_1, f_2 and g are some known functionals. In this context the functions $v \in H^1_0$ are known as *test functions*. Existence and uniqueness of solutions of weak problems in the linear elliptic case can be established if the weak form satisfies the conditions of the Lax-Milgram lemma, see [Cia78] for details.

Defining

$$\begin{aligned} a(u, v) &= \int_{\Omega_0} f_1(u, v, \boldsymbol{\nabla} u, \boldsymbol{\nabla} v)\,\mathrm{d}V_0 \\ b(v) &= \int_{\partial\Omega_0} f_2(\boldsymbol{\nabla} u \cdot \hat{\mathbf{n}}, v)\,\mathrm{d}S_0 + \int_{\Omega_0} g(v)\,\mathrm{d}V_0, \end{aligned}$$

we can write this more succinctly as: find $u \in H^1_{0,\Gamma}$ such that

$$a(u, v) = b(v) \qquad \forall\, v \in H^1_0.$$

The finite element problem is obtained by replacing the domain with the computational domain, and by solving the weak formulation over different function spaces. Given basis functions on Ω^h_0 with which to interpolate u (and which are zero on Γ), we define V^h to be the set spanned by these basis functions, and $V^h_0 = \{v^h \in V^h : v^h = 0 \text{ on } \partial\Omega^h_0\}$. The finite element problem is then: find $u^h \in V^h$ such that

$$a^h(u^h, v^h) = b^h(v^h) \qquad \forall\, v^h \in V^h_0, \tag{4.4}$$

where[4]

$$\begin{aligned} a^h(u^h, v^h) &= \int_{\Omega^h_0} f_1(u^h, v^h, \boldsymbol{\nabla} u^h, \boldsymbol{\nabla} v^h)\,\mathrm{d}V_0 \\ b^h(v^h) &= \int_{\partial\Omega^h_0} f_2(\boldsymbol{\nabla} u^h \cdot \hat{\mathbf{n}}, v^h)\,\mathrm{d}S_0 + \int_{\Omega^h_0} g(v^h)\,\mathrm{d}V_0 \end{aligned}$$

The existence and uniqueness of a solution to the finite element problem can be determined for some problems using the Lax-Milgram lemma, similarly to the initial weak problem. The question of under which conditions the finite element solution approximates the solution to the full problem is a fundamental one to finite element analysis. Problems and conditions for convergence of the FE solution to the true solution as a FE mesh is refined is discussed in detail in [Sul02].

---

[4] Note that often the distinction between Ω_0 and Ω^h_0 is not made, in which case we could just have written the FE problem as: find $u^h \in V^h$ such that $a(u^h, v^h) = b(v^h)$ for all $v^h \in V^h_0$.

Suppose we wish to solve the finite element problem for values of the unknown u at the nodes, excluding those nodes which are on the Dirichlet boundary, where we know the value of u. The number of nodes is $\mathcal{N}$, so supposing the number of nodes on the Dirichlet boundary is $\mathcal{N}_d$, the number of unknowns is $\mathcal{N} - \mathcal{N}_d$. The dimension of the space V_0^h is also $\mathcal{N} - \mathcal{N}_d$, and we obtain the $\mathcal{N} - \mathcal{N}_d$ equations for the $\mathcal{N} - \mathcal{N}_d$ unknowns by considering (4.4) and setting v^h to be the basis functions corresponding to non-Dirichlet nodes (i.e. the basis functions which span V_0^h). Thus the $\mathcal{N} - \mathcal{N}_d$ equations determining the unknown nodal values of u are

$$\mathcal{F}_i(\{u_j\}) \equiv a^h(u^h, \phi_i) - b^h(\phi_i) = 0.$$

The vector whose components are $\mathcal{F}_i$ is known as the *residual vector*. For linear problems, it will be a linear vector function of the unknowns u_j. In this case it can be written in the form

$$\mathcal{F} = \mathcal{A}\mathcal{U} - \mathcal{B} \tag{4.5}$$

where $\mathcal{U}$ is the vector of unknowns u_j, $\mathcal{U} = (u_1, u_2, \ldots, u_{\mathcal{N}})^{\mathrm{T}}$. It is straightforward to show that the matrix $\mathcal{A}$ is given by

$$\mathcal{A}_{ij} = a^h(\phi_i, \phi_j),$$

and the vector $\mathcal{B}$ is given by

$$\mathcal{B}_i = b^h(\phi_i).$$

$\mathcal{A}$ is known as the *global stiffness matrix* (GSM) and $\mathcal{B}$ is known as the global load vector (GLV). For nonlinear problems $\mathcal{F}$ will be a nonlinear function of the unknowns, in which case a numerical scheme for solving finite-dimensional nonlinear systems must be used. This is discussed further in Section 4.1.4.

4.1.3 Linear System Assembly for Linear FE Problems

Setting up the matrix equation (4.5) involves the computation of integrals over the computational domain. To do this we split the integral into a sum of integrals over elements. The idea is to loop over different elements, calculating the contributions to the global stiffness matrix and global load vector from that element, and adding these contributions to the main matrix and vector. These contributions take the form of *element stiffness matrices* (ESM) and *element load vectors* (ELV). The element stiffness matrix for the element K is the matrix $\mathcal{A}^{(K)}$ which satisfies

$$\int_K f_1(u^h, \phi_j^{(K)}, \boldsymbol{\nabla} u^h, \boldsymbol{\nabla}\phi_j^{(K)})\,\mathrm{d}V_0 = \left((\mathcal{U}^{(K)})^{\mathrm{T}}\mathcal{A}^{(K)}\right)_j$$

where $\mathcal{U}^{(K)}$ is a small vector containing the nodal values of u for *only* the nodes contained in the element K, and $\phi_j^{(K)}$ is the j-th basis function for the element K. In other words, the element stiffness matrix comes from restricting all integrals to just the element in consideration. The element stiffness vector is defined similarly, it is the vector $\mathcal{B}^{(K)}$ satisfying

$$\int_{K\cap\partial\Omega_0^h} f_2(\boldsymbol{\nabla} u^h\cdot\hat{\mathbf{n}},\phi_j^{(K)})\,\mathrm{d}S_0+\int_K g(\phi_j^{(K)})\,\mathrm{d}V_0=\left(\mathcal{B}^{(K)}\right)_j,$$

(which is independent of u^h as either $u=0$ or $\boldsymbol{\nabla} u\cdot\hat{\mathbf{n}}$ is given on the boundary).

Hence, for triangular elements with linear basis functions, the ESM is a 3 by 3 matrix, since there are three nodes in a triangular element, and the ELV is a 3 by 1 vector. Once computed, the components of the ESM and the ELV are added to the appropriate entries of the global stiffness matrix and load vector. For example, if the first node in the triangle has global number 10, and the second node global number 24, the (1,1) component of the ESM is added to the (10,10) component of the GSM, the (1,2) component of the ESM is added to the (10,24) component of the GSM, the first component of the ELV is added to the tenth component of the GLV, and so on.

4.1.4 Nonlinear FE Problems and Newton's Method

For problems which are inherently nonlinear in u, such as nonlinear elasticity, the residual vector is a nonlinear function of the unknown coefficients u_j, and a *nonlinear root-finding* method is required to solve the finite element formulation. Finding roots of nonlinear equations is fundamental to solving a huge range of mathematical models of physical phenomena. These types of problems often arise from maximisation or minimisation problems, by taking the gradient of the functional being optimised, so root-finding methods can be considered to be a subclass of optimisation methods.

The *Newton method* (or Newton-Raphson method) is one common method for numerically computing the zero of a function, where the nonlinear problem is tackled by computing the solution to a sequence of linear problems. It is motivated in one-dimensional problems by considering the tangent of a function $f(x)$ at a point x_0 which is taken to be a first guess of the zero of f. The zero of the tangent will generally be a better approximation to the zero of f than x_0, and is easily computed, being the solution of a linear equation, so this is set to be the new approximation. This process can be repeated, leading to the following algorithm: given an initial guess x_0 of the zero of f, let

$$x_{n+1}=x_n+\Delta x_n,\qquad\qquad n=0,1,2,\ldots$$

where the increment Δx_n is given by

$$\Delta x_n=-\frac{f(x_n)}{f'(x_n)}.$$

It can be shown that for initial guesses sufficiently close to a root, and for suitably smooth functions f, the Newton method converges to the root, with quadratic rate of convergence [NW99]. However, there are several problems with the Newton method, the major one being that the method is not guaranteed to converge. If the initial guess is not close to the true root, the method may lead to a divergent sequence. In certain cases the method can lead to an oscillatory sequence. Another problem is that the derivative has to be known.

The Newton method can be easily extended to problems in more than one-dimension. Suppose we wish to find the root of a nonlinear vector function $\mathcal{F}(\mathbf{a})$, where $\mathbf{a}$ is the vector of unknowns. Given an estimate $\mathbf{a}_n$ the new approximation $\mathbf{a}_{n+1}$ is set to be $\mathbf{a}_n + \Delta\mathbf{a}_n$, where the increment is given by

$$\Delta\mathbf{a}_n = -\left(\mathcal{J}(\mathbf{a}_n)\right)^{-1}\mathcal{F}(\mathbf{a}_n) \tag{4.6}$$

Here $\mathcal{J}$ is the *Jacobian* of the problem $\mathcal{J}_{ij} = \frac{\partial \mathcal{F}_i}{\partial a_j}$, evaluated at the old guess. Therefore, for nonlinear finite element problems, we deal with a sequence of Jacobians and residual vectors instead of global stiffness matrices and global load vectors. These Jacobians and residuals can be assembled element-by-element is an analogous way to that described in Section 4.1.3.

Damping

In order to reduce the risk of the Newton method diverging, it is important to make a small modification to the method and introduce some *damping.* Here we replace the update equation with

$$\mathbf{a}_n = \mathbf{a}_{n-1} + s\Delta\mathbf{a}_{n-1}, \tag{4.7}$$

for some $0 < s \leq 1$, known as the damping factor. This cautious updating is sometimes necessary because the direction $\Delta\mathbf{a}$ is only guaranteed to decrease the (norm of the) residual *initially.* A full Newton step ($s = 1$) may increase the residual, and in fact may cause immediate divergence in cases where the current guess is far from the final solution. We will see that in cases where there is a very large elastic deformation damping is essential. The best value of s at any Newton step could be computed from a 1D minimisation problem: given $\mathbf{a}_{n-1}$ and $\Delta\mathbf{a}_{n-1}$, compute s by minimising the 1D-function

$$f(s) = \|\mathcal{F}(\mathbf{a}_{n-1} + s\Delta\mathbf{a}_{n-1})\|_2.$$

However, in a given Newton step, high accuracy in computing the best choice of s is not usually necessary (in that it will not generally lead to large gains in decreasing the residual), so it is enough to simply test a number of different values for s and choose the best one. (In fact, we can test values of s which are larger than 1 as well, since for some Newton steps the residual may continue to decrease further in the Newton direction). Only near the solution is the best damping factor guaranteed to be one, and once in this region convergence to the true solution is quadratic.

4.1.5 Numerical Integration and Gaussian Quadrature

The entries of Jacobian matrices and residual vectors are given by integrals of known quantities over the computational domain. In assembly, these are reduced to a sum of integrals of known quantities over elements. These quantities, being interpolated over the element by polynomial basis functions, will be continuous and sometimes even infinitely differentiable, but in general the integrals cannot be computed analytically, and a numerical integration scheme is necessary. To numerically approximate an integral, we transform the integral to the canonical element, $\hat{K}$, and use *Gaussian quadrature*, where the integral of a function is approximated by a weighted sum of function evaluations. Gaussian quadrature involves the use of a fixed number of known *quadrature points*, $\boldsymbol{\xi}_g$, $g = 1, \ldots, G$, in the canonical element, and known weights ω_g, with an integral of the form $\int_{\hat{K}} F(\boldsymbol{\xi})\,\mathrm{d}V$ being approximated by

$$\int_{\hat{K}} F(\boldsymbol{\xi})\,\mathrm{d}V \approx \sum_{g=1}^{G} \omega_g F(\boldsymbol{\xi}_g)$$

Quadrature points and weights can be chosen so that Gaussian quadrature is exact on all polynomials on the element of degree $\leq 2G$. Details of Gaussian quadrature schemes can be found in [Eng80].

4.1.6 Solution of Linear Equations

In each Newton step we need to solve the linear system $-\mathcal{J}\Delta\mathbf{a} = \mathcal{F}$. These systems of linear equations are generally too large in dimension to be solved using a direct approach such as the Gaussian Elimination algorithm [Dem97]. Instead, it is necessary to employ an iterative solution method. The Jacobian matrices that need to be inverted are in general non-symmetric, but do have the property that they are *sparse*, i.e. contain only a small number of non-zero entries in each row. This is because the basis functions chosen are non-zero only in small regions of the computational domain, and therefore the equations at a given node can only be affected by nearby nodes.

Let us use the standard notation for linear matrix problems and suppose we wish to solve the equation $A\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$ where A is a large sparse n by n matrix, and $\mathbf{b}$ an n-dimensional vector. The idea of an iterative method is to start from an initial guess (often the zero vector), and update the guess in such a way as to converge to the true solution $\mathbf{x}$. To measure the extent that a guess $\mathbf{x}^{(k)}$ satisfies the matrix equation, we define the residual vector (for the linear matrix problem, not to be confused with the residual vector for the nonlinear finite element problem), $\mathbf{r}^{(k)} = \mathbf{b} - A\mathbf{x}^{(k)}$. Clearly we wish $\mathbf{r}^{(k)}$ to be as close to zero as possible. A suitable stopping criterion for an iterative method is for the scaled 2-norm of the residual, $\frac{1}{n}\|\mathbf{r}^{(k)}\|_2$ to drop below a fixed tolerance, or for the ratio of the 2-norms of the current residual to initial residual, $\frac{\|\mathbf{r}^{(k)}\|}{\|\mathbf{r}^{(0)}\|}$, to drop below a given tolerance.

In the case of symmetric positive-definite matrices, the choice of iterative method is straightforward: usually the conjugate gradients algorithm [HE52] is suitable. The matrices arising from the finite element method in incompressible finite elasticity are not symmetric or positive definite however, due to the nonlinearities and the pressure Lagrange multipliers. The choice of iterative scheme for general nonsymmetric matrices is not so simple, with different schemes working better than others in different situations. A suitable iterative method for solving large sparse non-symmetric matrix equations is the restarted *generalised minimum residual* method (GMRES) [SS86], which is used for computing simulations in this thesis. GMRES is an example of a Krylov subspace method, in which the current iterate, $\mathbf{x}^{(k)}$, is formed by minimising the residual over the *Krylov subspace*, defined as

$$\mathcal{K}_k(A, \mathbf{r}^{(0)}) = \operatorname{span}\{\mathbf{r}^{(0)}, A\mathbf{r}^{(0)}, A^2\mathbf{r}^{(0)}, \ldots, A^{k-1}\mathbf{r}^{(0)}\}.$$

It is clear from this that GMRES residuals must monotonically decrease with k (although not necessarily strictly), and GMRES is guaranteed to converge in n steps (assuming exact arithmetic), as the n-th Krylov subspace must be $\mathbb{R}^n$. Of course, the latter observation is impractical for very large systems, where we require a good approximation in much less iterations than n. Unfortunately, very little can be said in general about the convergence: matrices (including well-conditioned matrices) can be constructed for which GMRES does not converge, or even reduce the residual, until the n-th step [vdV03]. The GMRES algorithm makes use of an orthogonal set of vectors, one computed each iteration, which must be stored. This leads to prohibitively large storage costs and a decrease in speed as the number of iterations increases, and to avoid this GMRES is often restarted after m iterations, for some user-chosen restart parameter m. The resulting algorithm is known as *restarted GMRES* and often denoted as GMRES(m). GMRES(m) is used throughout this thesis. We will discuss GMRES and GMRES(m) further in Section 5.3.

Preconditioning

Preconditioning is often a vital stage when efficiently solving linear systems. It involves replacing the linear system with an equivalent system which has favourable properties so that an iterative linear solver converges in fewer iterations. In *left-preconditioning*, the system $A\mathbf{x} = \mathbf{b}$ is replaced by

$$LA\mathbf{x} = L\mathbf{b}, \tag{4.8}$$

where the matrix L is known as the (left-)preconditioner. The idea is to pick L so that the solver converges faster on the matrix LA than A. The optimal choice is $L = A^{-1}$, which of course is not available, but good left-preconditioners in a sense approximate the inverse of A. In *right preconditioning*, we define $\mathbf{y} = R^{-1}\mathbf{x}$, for some (right-)preconditioner R, and solve

$$AR\mathbf{y} = \mathbf{b}. \tag{4.9}$$

There are a few generic preconditioners, such as the *diagonal preconditioner*, where L is diagonal with elements

$$L_{ii} = \max_j |A_{ij}| \qquad \text{(no summation convention)}, \tag{4.10}$$

or preconditioners based on the LU factorisation of A [vdV03]. However, preconditioners that exploit specific features of the physical problem are often more effective, and the development of good preconditioning methods is an active area of research for a range of mechanical problems, mostly commonly in fluid dynamics (see for example [Tur99]). There has been some effort in developing preconditioners for linear elasticity (e.g. [MMRT01]), but such research for finite elasticity problems appears very rare.

In this thesis we do attempt to develop specialised preconditioners for elasticity, and do not precondition the Jacobians matrices with anything more than the diagonal preconditioner described above. However, we will make two hugely important modifications prior to forming the Jacobian that effectively act as preconditioners (although they are not implemented explicitly as such), and massively reduce GMRES computation time. The first is described in Section 4.2.3, and the second, which is important for the backward problem, in Section 5.3.

4.2 The FE Method for Nonlinear Elasticity

The finite element method has been the primary numerical tool for problems in solid mechanics since the 1970s, both in structural mechanical problems where the linear theory of elasticity can usually be used and in bio-mechanical problems with the nonlinear theory. Currently, it is used by thousands of researchers and engineers across the world, modelling elastic deformations of a huge number of materials in a wide range of scenarios. These applications range from the elastic modelling of buildings, to vehicles, to earthquake modelling [HH04], to polymer-bounded explosives [BA02], with biological applications including the modelling of: whole organs such as the breast, heart [Nas98], liver, brain; as well as smaller systems, for example arterial modelling; and larger systems, for example the whole face [VSCS03].

In this section we describe how the finite element method can be used to solve problems in finite elasticity.

4.2.1 FE Assembly for Nonlinear Elasticity Problems

Let us assume the material being modelled is homogeneous, isotropic and incompressible. This means the unknowns to be computed are both the displacement (or position) and the pressure field. In the following we assume the material is hyperelastic, and explain how the finite element method can be used for a general isotropic strain energy function $\overline{W}(I_1, I_2) - \frac{1}{2}p(I_3 - 1)$. We will also assume that there are no surface tractions, and the only body force is gravity.

We will use trilinear interpolation for the displacement. It is well-accepted that triangular and tetrahedral elements perform poorly in comparison with quadrilateral and hexahedral elements when low-order schemes are used in elasticity problems (see [BPM$^+$95] for a discussion), and for this reason all meshes in this thesis are composed solely of hexahedral elements.

This also dictates that we should use piecewise constant interpolation for the pressure field. It has been shown that the pressure field should be interpolated by polynomials of degree one less than those used for the displacement, and in this case the accuracy of the FE method actually increases compared to the same degree polynomials being used for both. This is essentially because if linear basis functions are used for position and displacement, the strain, and therefore the material component of the stress, will be constant in any element (because strains are derivatives of displacement), so the second term in the stress $-p\delta^{ij}$ should ideally also be constant in the element. This is an example of a *mixed method.*

Suppose we have $\mathcal{N}$ unfixed nodes (note the change of notation to Section 4.1.1, where $\mathcal{N}$ was the total number of nodes) and $\mathcal{M}$ elements. Since we are using trilinear interpolation for displacement and piecewise constant interpolation for pressure, the finite element unknowns will be the nodal displacements and the elemental pressures, so we have $3\mathcal{N} + \mathcal{M}$ unknowns. Let $\{\phi_i\}$ be the standard trilinear basis functions, and define

$$\mathbf{u}^h = \begin{bmatrix} \sum_{i=1}^{\mathcal{N}} U_i\phi_i \\ \sum_{i=1}^{\mathcal{N}} V_i\phi_i \\ \sum_{i=1}^{\mathcal{N}} W_i\phi_i \end{bmatrix} \tag{4.11}$$

and

$$p^h = \sum_{j=1}^{\mathcal{M}} P^j \mathbb{1}_{K_j}$$

to be the finite element solutions for the displacement and pressure.

Let

$$\mathbf{U} = \begin{bmatrix} U_1 \\ \vdots \\ U_{\mathcal{N}} \end{bmatrix} \qquad \mathbf{V} = \begin{bmatrix} V_1 \\ \vdots \\ V_{\mathcal{N}} \end{bmatrix} \qquad \mathbf{W} = \begin{bmatrix} W_1 \\ \vdots \\ W_{\mathcal{N}} \end{bmatrix} \qquad \mathbf{P} = \begin{bmatrix} P_1 \\ \vdots \\ P_{\mathcal{M}} \end{bmatrix}$$

and let the vector of unknowns (of size $3\mathcal{N} + \mathcal{M}$) be

$$\mathbf{a} = \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \\ \mathbf{W} \\ \mathbf{P} \end{bmatrix}. \tag{4.12}$$

The weak formulation of the incompressible problem is given by (3.23)

$$\int_{\Omega_0} T^{MN}(\mathbf{u},p)F^k_M(\mathbf{u})F^k_N(\delta\mathbf{u})\,\mathrm{d}V_0 - \int_{\Omega_0}\rho_0\mathbf{g}\cdot\delta\mathbf{u}\,\mathrm{d}V_0 + \int_{\Omega_0}(\det F(\mathbf{u})-1)(\delta p)\,\mathrm{d}V_0 = 0 \qquad \forall\,\delta\mathbf{u},\ \delta p \tag{4.13}$$

(with $T^{MN} = \frac{\partial\overline{W}(I_1,I_2)}{\partial E_{MN}} - p(C^{-1})^{MN}$). The equivalent weak finite element problem is

$$\int_{\Omega_0^h} T^{MN}(\mathbf{u}^h,p^h)F^k_M(\mathbf{u}^h)F^k_N(\delta\mathbf{u}^h)\,\mathrm{d}V_0 - \int_{\Omega_0^h}\rho_0\mathbf{g}\cdot\delta\mathbf{u}^h\,\mathrm{d}V_0 + \int_{\Omega_0^h}(\det F(\mathbf{u}^h)-1)(\delta p^h)\,\mathrm{d}V_0 \;=\; 0 \qquad \forall\,\delta\mathbf{u}^h,\ \delta p^h \tag{4.14}$$

where $\delta\mathbf{u}^h$ is allowed to be any linear combination of the displacement basis vectors, and δp^h any linear combination of the pressure vectors. To this end, we set $\delta\mathbf{u}^h = \begin{bmatrix}\phi_i\\0\\0\end{bmatrix}, \begin{bmatrix}0\\\phi_i\\0\end{bmatrix}$ and $\begin{bmatrix}0\\0\\\phi_i\end{bmatrix}$ $(i = 1,\ldots,\mathcal{N})$ and $\delta p^h = \mathbb{1}_j$ $(j = 1,\ldots,\mathcal{M})$ to obtain the $3\mathcal{N}+\mathcal{M}$ nonlinear equations which need to be solved:

equations $i = 1$ to $\mathcal{N}$:

$$\mathcal{F}_i \equiv \int_{\Omega_0} T^{MN}(\mathbf{U},\mathbf{V},\mathbf{W},\mathbf{P})F^1_M(\mathbf{U},\mathbf{V},\mathbf{W})\frac{\partial\phi_i}{\partial X^N}\,\mathrm{d}V_0 - \int_{\Omega_0}\rho_0 g_1\phi_i\,\mathrm{d}V_0 = 0,$$

equations $i = \mathcal{N}+1$ to $2\mathcal{N}$

$$\mathcal{F}_i \equiv \int_{\Omega_0} T^{MN}(\mathbf{U},\mathbf{V},\mathbf{W},\mathbf{P})F^2_M(\mathbf{U},\mathbf{V},\mathbf{W})\frac{\partial\phi_{i-\mathcal{N}}}{\partial X^N}\,\mathrm{d}V_0 - \int_{\Omega_0}\rho_0 g_2\phi_{i-\mathcal{N}}\,\mathrm{d}V_0 = 0,$$

equations $i = 2\mathcal{N}+1$ to $3\mathcal{N}$:

$$\mathcal{F}_i \equiv \int_{\Omega_0} T^{MN}(\mathbf{U},\mathbf{V},\mathbf{W},\mathbf{P})F^3_M(\mathbf{U},\mathbf{V},\mathbf{W})\frac{\partial\phi_{i-2\mathcal{N}}}{\partial X^N}\,\mathrm{d}V_0 - \int_{\Omega_0}\rho_0 g_3\phi_{i-2\mathcal{N}}\,\mathrm{d}V_0 = 0,$$

equations $i = 3\mathcal{N}+1$ to $3\mathcal{N}+\mathcal{M}$:

$$\mathcal{F}_i \equiv \int_{\Omega_0}(\det F(\mathbf{U},\mathbf{V},\mathbf{W})-1)\,\mathbb{1}_{K_{i-3\mathcal{N}}}\,\mathrm{d}V_0 = 0 \Rightarrow \int_{\text{element } i-3\mathcal{N}}(\det F - 1)\,\mathrm{d}V_0 = 0.$$

These are the equations specifying the components of the residual vector $\mathcal{F}(\mathbf{a})$. To solve $\mathcal{F}(\mathbf{a}) = 0$ with Newton's method, we need to be able to compute the Jacobian $\mathcal{J}$, which means we need to calculate terms of the form $\frac{\partial\mathcal{F}_i}{\partial U^\alpha}, \frac{\partial\mathcal{F}_i}{\partial V^\alpha}, \frac{\partial\mathcal{F}_i}{\partial W^\alpha}$ and $\frac{\partial\mathcal{F}_i}{\partial P^\alpha}$, for each of the four different forms of $\mathcal{F}_i$. (The matrix $\mathcal{J}_{i\alpha}$ therefore has a natural 4-by-4 block matrix structure). Suppose $i \in \{1,\ldots,\mathcal{N}\}$. Then

$$\begin{aligned}\mathcal{J}_{i\alpha} = \frac{\partial\mathcal{F}_i}{\partial U^\alpha} &= \frac{\partial}{\partial U^\alpha}\left(\int_{\Omega_0} T^{MN}(\mathbf{U},\mathbf{V},\mathbf{W},\mathbf{P})F^1_M(\mathbf{U},\mathbf{V},\mathbf{W})\frac{\partial\phi_i}{\partial X^N}\,\mathrm{d}V_0\right)\\ &= \int_{\Omega_0}\frac{\partial T^{MN}}{\partial U_\alpha}F^1_M\frac{\partial\phi_i}{\partial X^N} + T^{MN}\frac{\partial F^1_M}{\partial U^\alpha}\frac{\partial\phi_i}{\partial X^N}\,\mathrm{d}V_0\\ &= \int_{\Omega_0}\frac{\partial T^{MN}}{\partial E_{PQ}}\frac{\partial E_{PQ}}{\partial U^\alpha}F^1_M\frac{\partial\phi_i}{\partial X^N} + T^{MN}\frac{\partial F^1_M}{\partial U^\alpha}\frac{\partial\phi_i}{\partial X^N}\,\mathrm{d}V_0.\end{aligned} \tag{4.15}$$

To evaluate this we need to use

$$\frac{\partial F^1_M}{\partial U^\alpha} = \frac{\partial}{\partial U^\alpha}\left(\delta^1_M + \frac{\partial\left(\phi_k U^k\right)}{\partial X^M}\right) = \frac{\partial \phi_\alpha}{\partial X^M}$$

and therefore

$$\frac{\partial E_{PQ}}{\partial U^\alpha} = \frac{\partial}{\partial U^\alpha}\left(\frac{1}{2}\left(F^i_P F^i_Q - \delta_{PQ}\right)\right) = \frac{1}{2}\left(\frac{\partial \phi_\alpha}{\partial X^P}F^1_Q + \frac{\partial \phi_\alpha}{\partial X^Q}F^1_P\right).$$

(Computing T^{MN} and $\frac{\partial T^{MN}}{\partial E_{PQ}}$ is described shortly). Similar expressions can be derived for $\frac{\partial \mathcal{F}_i}{\partial U^\alpha}$, $\frac{\partial \mathcal{F}_i}{\partial V^\alpha}$, $\frac{\partial \mathcal{F}_i}{\partial W^\alpha}$, for any $i \leq 3\mathcal{N}$.

Suppose again $i \in \{1, \dots, \mathcal{N}\}$. Then

$$\mathcal{J}_{i,3\mathcal{N}+\alpha} = \frac{\partial \mathcal{F}_i}{\partial P^\alpha} = \frac{\partial}{\partial P^\alpha}\left(\int_{\Omega_0} T^{MN} F^1_M \frac{\partial \phi_i}{\partial X^M}\,\mathrm{d}V_0\right) = \int_{\text{element } i} -(C^{-1})^{MN} F^1_M \frac{\partial \phi_i}{\partial X^M}\,\mathrm{d}V_0,$$

and again similar expressions hold for $i \leq 3\mathcal{N}$.

The final group of calculations we need to make are those of the form $\frac{\partial \mathcal{F}_i}{\partial U^\alpha}$ for $i \in \{3\mathcal{N} + 1, \dots, 3\mathcal{N} + \mathcal{M}\}$:

$$\frac{\partial \mathcal{F}_i}{\partial U^\alpha} = \frac{\partial}{\partial U^\alpha}\left(\int_{\text{element } i - 3\mathcal{N}} (\det(F) - 1)\,\mathrm{d}V_0\right) = \int_{\text{element } i - 3\mathcal{N}} \frac{\partial \det(F)}{\partial U^\alpha}\,\mathrm{d}V_0.$$

Here we make use of the fact that $\frac{\partial \det(F)}{\partial U^\alpha} = \frac{\partial \det(F)}{\partial F^i_M}\frac{\partial F^i_M}{\partial U^\alpha} = \det(F)(F^{-1})^M_i \frac{\partial \phi_\alpha}{\partial X^M}\delta^i_1$.

Computing T^{MN} and $\frac{\partial T^{MN}}{\partial E_{PQ}}$ in the isotropic case

Finally, we complete the description of the computations that have to be performed by describing how to compute the stress T^{MN} and the stress derivative $\frac{\partial T^{MN}}{\partial E_{PQ}}$. We consider only the isotropic case here. We are assuming a 2nd Piola-Kirchoff stress of the form

$$T^{MN} = \frac{\partial \overline{W}(I_1, I_2)}{\partial E_{MN}} - p(C^{-1})^{MN},$$

so that

$$T^{MN} = \frac{\partial \overline{W}}{\partial I_1}\frac{\partial I_1}{\partial E_{MN}} + \frac{\partial \overline{W}}{\partial I_2}\frac{\partial I_2}{\partial E_{MN}} - p(C^{-1})^{MN}.$$

We need to be able to compute T^{MN} and $\frac{\partial T^{MN}}{\partial E_{PQ}}$. Now

$$\begin{aligned}
\frac{\partial T^{MN}}{\partial E_{PQ}} = & \quad \frac{\partial^2 \overline{W}}{\partial {I_1}^2}\frac{\partial I_1}{\partial E_{MN}}\frac{\partial I_1}{\partial E_{PQ}} + \frac{\partial \overline{W}}{\partial I_1}\frac{\partial^2 I_1}{\partial E_{MN}\partial E_{PQ}} \\
& + \frac{\partial^2 \overline{W}}{\partial {I_2}^2}\frac{\partial I_2}{\partial E_{MN}}\frac{\partial I_2}{\partial E_{PQ}} + \frac{\partial \overline{W}}{\partial I_2}\frac{\partial^2 I_2}{\partial E_{MN}\partial E_{PQ}} \\
& + \frac{\partial^2 \overline{W}}{\partial I_1 \partial I_2}\left(\frac{\partial I_1}{\partial E_{MN}}\frac{\partial I_2}{\partial E_{PQ}} + \frac{\partial I_2}{\partial E_{MN}}\frac{\partial I_1}{\partial E_{PQ}}\right) \\
& - p\frac{\partial (C^{-1})^{MN}}{\partial E_{PQ}},
\end{aligned}$$

and since $\frac{\partial}{\partial E_{MN}} = 2\frac{\partial}{\partial C_{MN}}$, we need to construct the following matrices and fourth order tensors:

$$\frac{\partial I_1}{\partial C_{MN}} \qquad \frac{\partial^2 I_1}{\partial C_{MN}\partial C_{PQ}} \qquad \frac{\partial I_2}{\partial C_{MN}} \qquad \frac{\partial^2 I_2}{\partial C_{MN}\partial C_{PQ}} \qquad \frac{\partial (C^{-1})^{MN}}{\partial C_{PQ}}.$$

Firstly, $I_1 = \mathrm{tr}(C)$, so

$$\frac{\partial I_1}{\partial C_{MN}} = \delta^{MN} \quad \text{and} \quad \frac{\partial^2 I_1}{\partial C_{MN}\partial C_{PQ}} = 0.$$

$I_2 = \frac{1}{2}(\mathrm{tr}(C)^2 - \mathrm{tr}(C^2)) = \frac{1}{2}(I_1^2 - C_{PQ}C_{PQ})$, so

$$\frac{\partial I_2}{\partial C_{MN}} = I_1\delta_{MN} - C_{MN},$$

and it follows that

$$\frac{\partial^2 I_2}{\partial C_{MN}\partial C_{PQ}} = \delta_{MN}\delta_{PQ} - \delta_{MP}\delta_{NQ}.$$

Finally, we need $\frac{\partial (C^{-1})^{MN}}{\partial C_{PQ}}$. This can be obtained by differentiating the equality $(C^{-1})^{MN}C_{NP} \equiv \delta^M_P$; we find that:

$$\frac{\partial (C^{-1})^{MN}}{\partial C_{PQ}} = -(C^{-1})^{MP}(C^{-1})^{QN}.$$

4.2.2 Methods of Improving Newton-Convergence

The damped Newton method as described in Section 4.1.4 is not guaranteed to converge, and to ensure the best chance of Newton convergence, a couple of modifications are required.

Firstly, instead of using an initial Newton guess of zero (i.e. taking $\mathbf{a}_0 = 0$), it is more sensible to use the solution of the zero-displacement problem. This naturally means the displacement components of **a** are zero, but the hydrostatic pressure is not also zero. It is straightforward to compute what the zero-displacement pressures should be. If there is no displacement the body remains in the reference state and therefore the stress T^{MN} must be equal to zero. T^{MN} is given by

$$\begin{aligned} T^{MN} &= \frac{\partial \overline{W}(I_1, I_2)}{\partial E_{MN}} - p(C^{-1})^{MN} \\ &= 2\frac{\partial \overline{W}}{\partial I_1}\delta^{MN} + 2\frac{\partial \overline{W}}{\partial I_2}(I_1\delta_{MN} - C_{MN}) - p(C^{-1})^{MN} \end{aligned}$$

If we set C and C^{-1} to be the identity, I_1 to be three (the trace of the identity), and equate this with zero, we obtain

$$0 = \left(2\frac{\partial \overline{W}}{\partial I_1} + 4\frac{\partial \overline{W}}{\partial I_2} - p\right)\delta^{MN}$$

Hence the zero-displacement pressure is the constant $p = 2\frac{\partial \overline{W}}{\partial I_1} + 4\frac{\partial \overline{W}}{\partial I_2}$, evaluated at $I_1 = I_2 = 3$. For example, if we had a homogeneous material with exponential strain energy $W = \alpha(e^{\beta(I_1-3)} - 1) - \frac{p}{2}(I_3 - 1)$, the zero-displacement pressure is $p = 2\alpha\beta$, and we would take

$$\mathbf{a} = (0, \ldots, 0, 0, \ldots, 0, 0, \ldots, 0, 2\alpha\beta, \ldots, 2\alpha\beta)^{\mathrm{T}}$$

as our initial Newton guess.

Choosing an initial guess which is as close as possible to the solution is important because the Newton method is only guaranteed to converge given an initial guess in some neighbourhood near the solution. The second modification is related to this, and involves breaking down the original problem into a sequence of subproblems, and using the solution of one problem as the initial guess of the next. If the solution to each subproblem is a relatively small change in displacement, Newton convergence is far more likely. In the case of a body deforming under gravity, we can accomplish this by solving a number of problems with gravity applied at increasing fractions of its true value. For example, we use applied body forces of $\frac{\mathbf{g}}{N}, \frac{2\mathbf{g}}{N}, \ldots, \frac{(N-1)\mathbf{g}}{n}, \mathbf{g}$, for some small integer N, and perform a few Newton steps on each subproblem, using the final solution of one subproblem as the starting guess for the next. Of course, it is not necessary to solve each subproblem to high accuracy, only the final one. This principal of altering the problem to reduce the change in displacement in any given Newton step and decrease the chance of divergence can be applied to other types of problem. In Chapter 6 we will discuss the contact problem, where an elastic body is compressed by rigid plates. To prevent divergence in this case we can solve a sequence of subproblems in which the plate position is incremented, so that full compression is only applied gradually.

4.2.3 Preconditioning by Non-dimensionalisation

We stated in Section 4.1.6 that we will make two modifications which effectively act as preconditioners, and are vital to fast GMRES convergence. The first modification is to nondimensionalise all variables described in Chapter 3 prior to computing residual vectors or Jacobians. We are only concerned with static problems, for which time does not figure, so we have two dimensions, length and mass, which may be rescaled. We rescale lengths by a typical length scale l_{typ}, and masses by a typical mass parameter m_{typ}. Hence

$$\mathbf{X}' = \frac{\mathbf{X}}{l_{\text{typ}}} \qquad \mathbf{x}' = \frac{\mathbf{x}}{l_{\text{typ}}} \qquad \mathbf{u}' = \frac{\mathbf{u}}{l_{\text{typ}}} \qquad \mathrm{d}V_0' = \frac{\mathrm{d}V_0}{l_{\text{typ}}^3}.$$

The quantities F, E and C are dimensionless, so $F' = F$, $E' = E$ and $C' = C$. Now[5], [density] = $[M][L]^{-3}$ and [acceleration] = $[L][T]^{-2}$, so

$$\begin{aligned} \rho_0' &= \frac{l_{\text{typ}}^3}{m_{\text{typ}}}\rho_0 \\ \mathbf{g}' &= \frac{\mathbf{g}}{l_{\text{typ}}} \end{aligned}$$

---

[5] Here we use the standard notation in dimensional analysis: square brackets represent the dimension of a quantity, and $[L]$, $[M]$ and $[T]$ denote dimensions of length, mass and time respectively.

Finally, the dimensions of pressure, and of the strain energy function, are $[M][L]^{-1}[T]^{-2}$, so the rescaled p and W are

$$\begin{aligned} p' &= \frac{l_{\mathrm{typ}}}{m_{\mathrm{typ}}} p \\ W' &= \frac{l_{\mathrm{typ}}}{m_{\mathrm{typ}}} W \end{aligned}$$

Since W is a function of the dimensionless strain (or dimensionless strain invariants $I_j^{'}$), rescaling W amounts to rescaling the material parameters, for example if $W = a(e^{b(I_1-3)} - 1)$ or $W = c_1(I_1 - 3) + c_2(I_2 - 3)$ then

$$a' = \frac{l_{\mathrm{typ}}}{m_{\mathrm{typ}}} a \qquad b' = b \qquad c_1^{'} = \frac{l_{\mathrm{typ}}}{m_{\mathrm{typ}}} c_1 \qquad c_2^{'} = \frac{l_{\mathrm{typ}}}{m_{\mathrm{typ}}} c_2.$$

We should choose the length parameter l_{typ} so that the new lengths are $O(1)$, and the mass parameter m_{typ} so that the material parameters are $O(1)$.

It is obvious that this rescaling implicitly preconditions all linear systems that are formed for the Newton method. It is also readily observed that it is equivalent to both a right and a left preconditioner. Since the unknown vector (4.12) is replaced by

$$\mathbf{a}' = \begin{bmatrix} \frac{\mathbf{U}}{l_{\mathrm{typ}}} \\ \frac{\mathbf{V}}{l_{\mathrm{typ}}} \\ \frac{\mathbf{W}}{l_{\mathrm{typ}}} \\ \frac{l_{\mathrm{typ}}}{m_{\mathrm{typ}}} \mathbf{P} \end{bmatrix},$$

the right-preconditioner can be seen to be

$$R = \begin{pmatrix} l_{\mathrm{typ}} I_{3\mathcal{N}} & 0 \\ 0 & \frac{m_{\mathrm{typ}}}{l_{\mathrm{typ}}} I_{\mathcal{M}} \end{pmatrix}$$

(where I_P is the P by P identity matrix). By choosing appropriate l_{typ} and m_{typ}, this right preconditioner has the effect of rescaling the components of the unknown vector to be approximately the same magnitude, in particular so that the unknown displacements are of a similar magnitude to the unknown pressures. We do not construct the left-preconditioner explicitly here.

In Section 4.3 we will display some sample computations on the deformation of an elastic cube. A comparison of the rate of convergence of GMRES on dimensionalised and non-dimensionalised linear systems arising from these simulations are shown in Figures 4.5 and 4.6. In all simulations in this thesis, the restarted version of GMRES is used rather than full GMRES, since memory-requirements are prohibitive in the latter, and therefore we only look at the efficiency of GMRES(m) on these matrices. Figure 4.5 shows the convergence histories using GMRES(200)[6]. The relative residual, is

---

[6] Note that 200 is a relatively high choice of restart parameter.

plotted against the number of outer iterations, where by 'outer iteration' we mean a complete of set of m inner iterations (200 inner iterations, in this case) (in other words, the number of restarts). This figure clearly displays the effect of non-dimensionalisation preconditioning. The residual for the dimensionalised linear system decays approximately exponentially, but at a slow rate, whereas the residual for the non-dimensionalised linear system decays virtually immediately to as low as it will go.

Figure 4.6 displays even more dramatic results. Here, convergence histories are shown using GMRES(50). The non-dimensionalised residual decreases somewhat less quickly, as would be expected, but now the dimensionalised residual does not even follow an (approximately-)exponential pattern, and GMRES(50) *stagnates* on the dimensionalised system. Clearly a higher restart parameter is required in order to have any hope of convergence.

In general when using GMRES(m), it is preferable to use as low a value of the restart parameter as possible, both in terms of storage and speed. The decrease in speed as m increases is due to computation becoming more and more expensive towards the end of any GMRES(m) outer iteration, and is quite considerable. Convergence will often occur quicker (in time) using a small restart parameter, even though more total inner iterations may be necessary. For example, to compute the solution of the non-dimensionalised linear system, GMRES(200) takes 169 total iterations and 2.36 seconds, whereas with GMRES(50) 342 total iterations are required but only 1.97 seconds. The difference in time increases significantly as the size of the linear system increases. With this in mind, we see that preconditioning by non-dimensionalisation is vital, as without it not only is computation time increased, but small restart parameters become unusable.

4.3 Results

Figures 4.7 to 4.11 illustrate some sample finite element simulations of an incompressible elastic cube deforming under gravity. We use a cube rather than a mesh of the breast geometry here to study the performance of the method on a well-shaped regular mesh, and to separate the performance of the method from issues that may arise using the mesh of the breast. Similarly, in Chapters 5 and 6 we will initially test the new methods on cubic meshes. The meshes used here are obviously relatively coarse.

For Figures 4.7, 4.8 and 4.9, the cube is taken to be homogeneous and isotropic, and made of a hyperelastic material with Veronda exponential material law (3.20). Material parameters were chosen to give significant deformations. Figure 4.10 displays the result of a simulation with a heterogeneous cube where one part of the cube is significantly stiffer (α in (3.20) twice as large) than the rest of the cube. The difference in the deformation is clearly seen: the stiffer material deforms far less and causes the total deformation to be much smaller than in Figure 4.9.

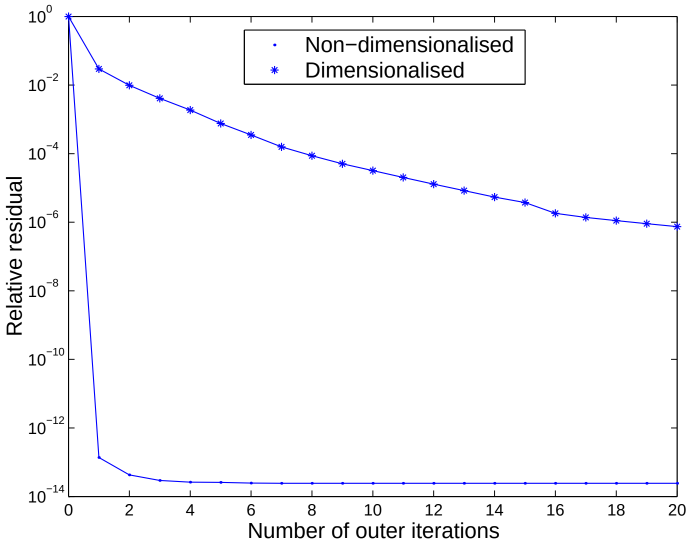

Figure 4.5: Convergence history for GMRES(200) on dimensionalised and non-dimensionalised linear systems.

In order to verify that the numerical method has been implemented correctly, we appeal to a principal that has been discussed by Peter Lax, who has said that for very complex numerical algorithms, a suitable method of verification is to compare two independent implementations and check that they agree. The results in Figures 4.7 to 4.10 have been compared to the results of simulations of the same problems with the same meshes carried out with CMISS [Hun], an academic finite element package developed in the University of Auckland, and have been seen to produce identical results in displacement and internal pressures.

In Table 4.2, the necessity of the damping factor s in (4.7) is illustrated. The table gives the norm of the residual $\mathcal{F}$ for various values of s $(0 \leq s \leq 1)$ after the first Newton iteration of the sheared cube simulation (Figure 4.9). The best choice of s from those tested is in fact very small, only 0.04. Taking a full Newton step $(s = 1)$ would not just lead to the norm of the residual increasing, it would cause the norm to increase by 38 orders of magnitude. In other words, an absence of the damping factor would bring about immediate divergence and failure of the numerical method.

Figures 4.11 display the results of simulations using a (homogeneous) fibre-reinforced material and a transversely isotropic material law. Fibres were chosen to be in the vertical direction, with a Veronda material law (3.20) with an exponential term for the fibre direction component (3.21)

$$W(I_1, I_2, I_3, K_1) = a(e^{b(I_1-3)} - 1) + a'e^{2b'K_1} - 2a'b'K_1 - \frac{p}{2}(I_3 - 1). \tag{4.16}$$

a was set to be 80Pa, and a' allowed to vary between 0 (no fibre-reinforcement) to 800 (very stiff fibres reinforcing the object in the vertical direction). Gravity was set to act upwards. The simulations

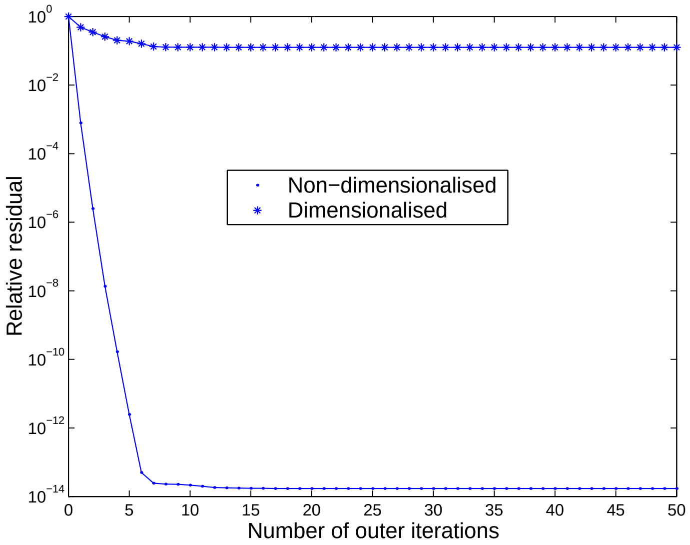

Figure 4.6: Convergence history for GMRES(50) on dimensionalised and non-dimensionalised linear systems.

Number of nodes	x-displacement	y-displacement	z-displacement
5^3	2.93	2.93	15.74
10^3	1.94	1.94	15.53
15^3	1.70	1.70	15.56
20^3	1.60	1.60	15.57

Table 4.1: Displacement of a top corner node for each of the simulations shown in Figure 4.12. (The length of the cube is 100).

clearly reproduce the expected behaviour: as the fibre stiffness increases the elastic deformation decreases.

Finally, Figure 4.12 shows that the numerical simulations converge to a limit solution as the mesh is refined. In these figures the squashed cube simulation is carried out with meshes with 5^3, 10^3, 15^3 and 20^3 nodes. Quantitative verification of this is given in Table 4.1.

s	$\lVert\mathcal{F}(\mathbf{a}_0 + s\Delta\mathbf{a}_0)\rVert_2/\text{numNodes}$
0.0	6.59952e-07
0.01	6.53363e-07
0.02	6.47094e-07
0.03	6.42239e-07
0.04	6.40938e-07
0.05	6.46682e-07
0.06	6.64629e-07
0.07	7.01719e-07
0.08	7.66421e-07
0.09	8.68168e-07
0.1	1.01708e-06
0.2	1.22511e-05
0.3	0.000820427
0.4	0.383896
0.5	992.434
0.6	1.42047e+07
0.7	1.18698e+12
0.8	5.93934e+17
0.9	1.74384e+24
1	2.92631e+31

Table 4.2: The norm of the residual vector for a range of damping factors, for the first Newton iteration of the simulation used to produce Fig 4.9.

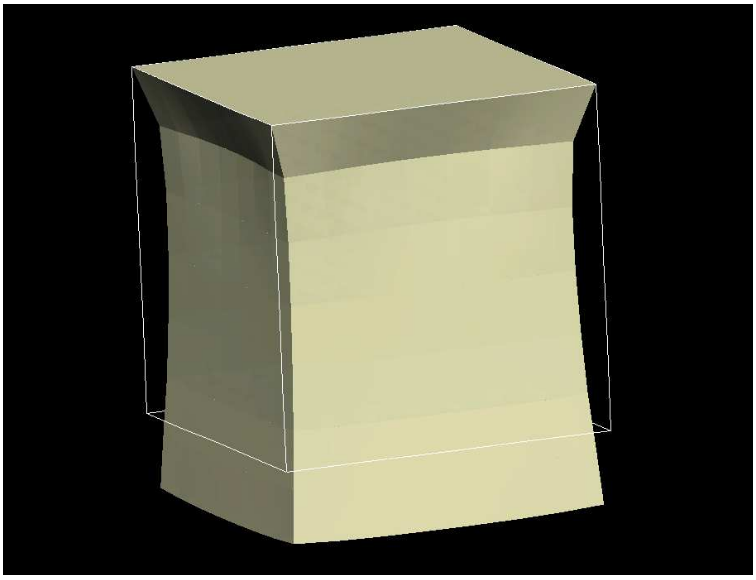

Figure 4.7: Finite element simulation of a homogeneous isotropic elastic cube stretched under gravity (top surface fixed, gravity acting downwards). The white lines represent the undeformed state.

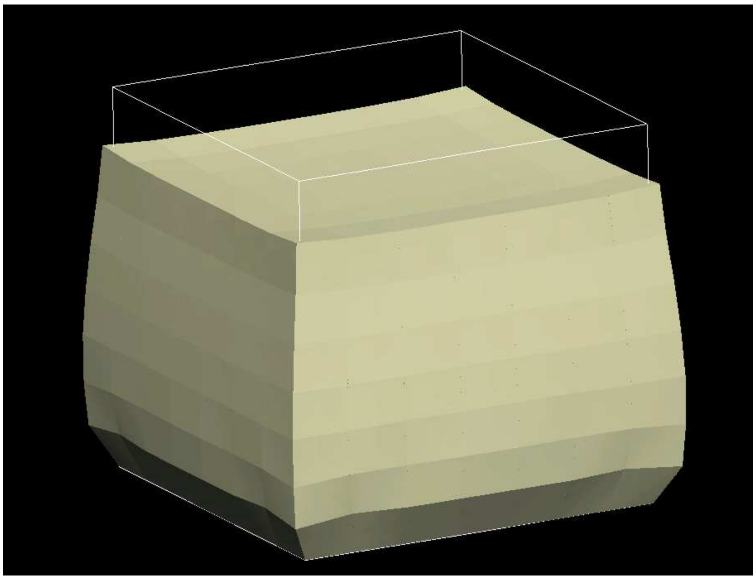

Figure 4.8: Finite element simulation of a homogeneous isotropic elastic cube squashed under gravity (bottom surface fixed, gravity acting downwards). The white lines represent the undeformed state.

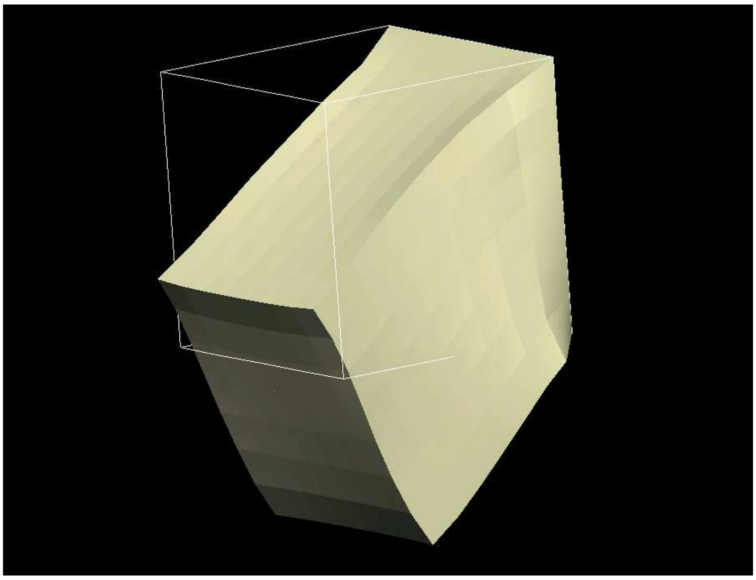

Figure 4.9: Finite element simulation of a homogeneous isotropic elastic cube sheared under gravity (side surface fixed, gravity acting downwards). The white lines represent the undeformed state.

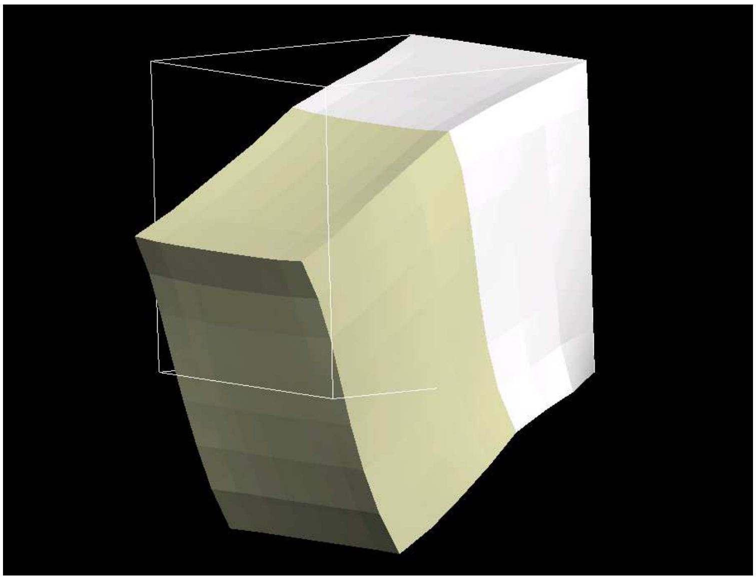

Figure 4.10: Simulation of a heterogeneous cube sheared under gravity. The white elements correspond to material which is significantly stiffer than the remainder of the cube.

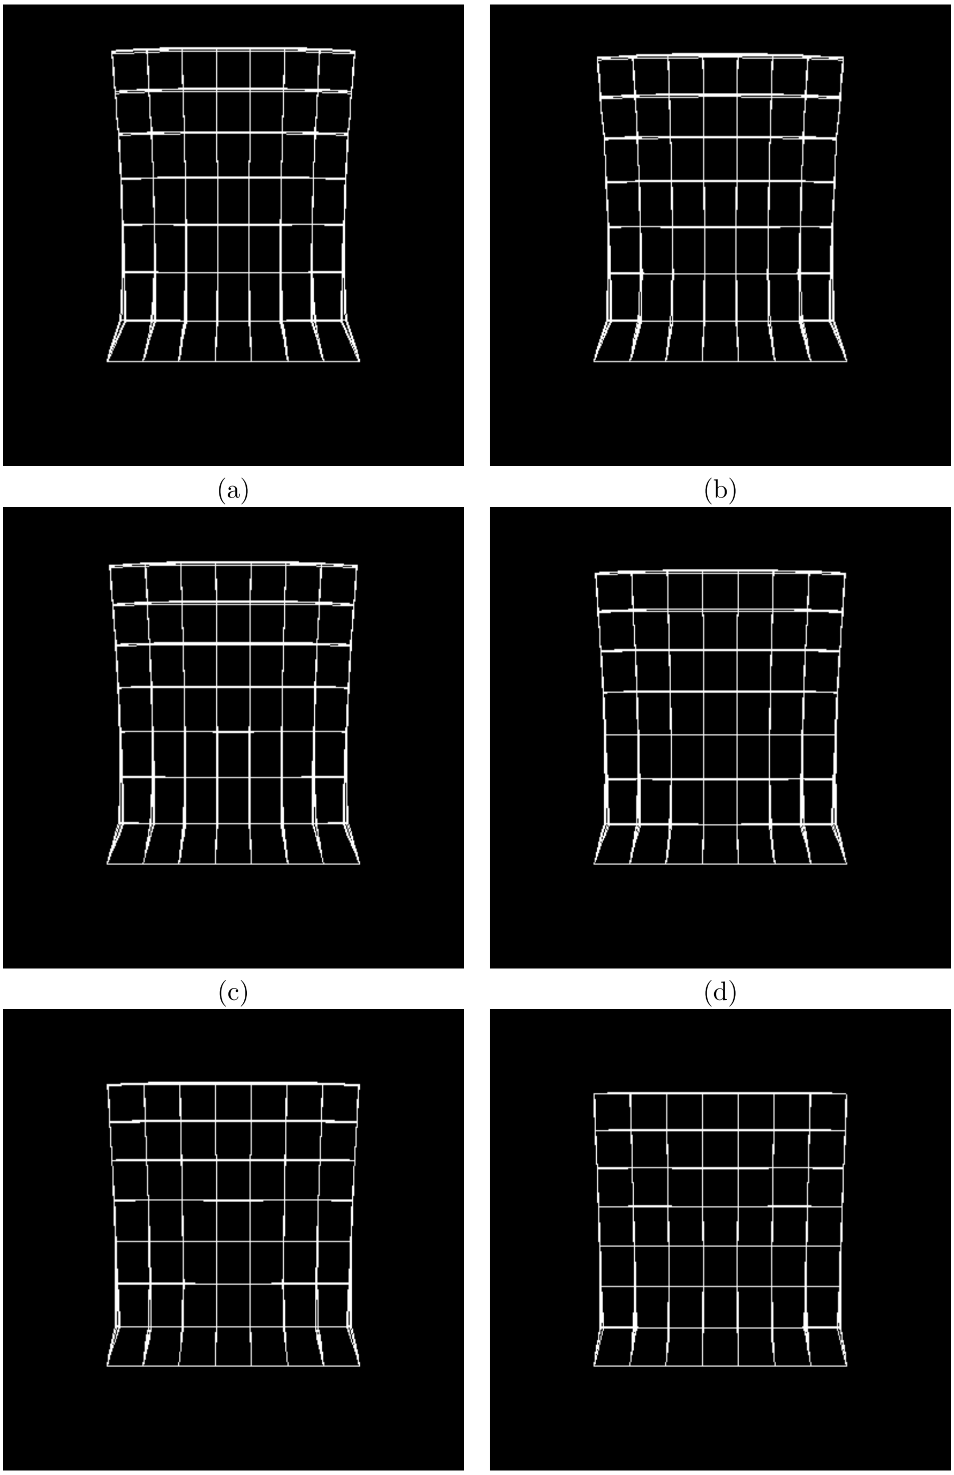

(a) (b) (c) (d) (e) (f)

Figure 4.11: Results of simulations with a fibre-reinforced material and transversely-isotropic material law, a' (from equation (4.16)) given by: (a) 0 (b) 50 (c) 100 (d) 200 (e) 400 (f) 800.

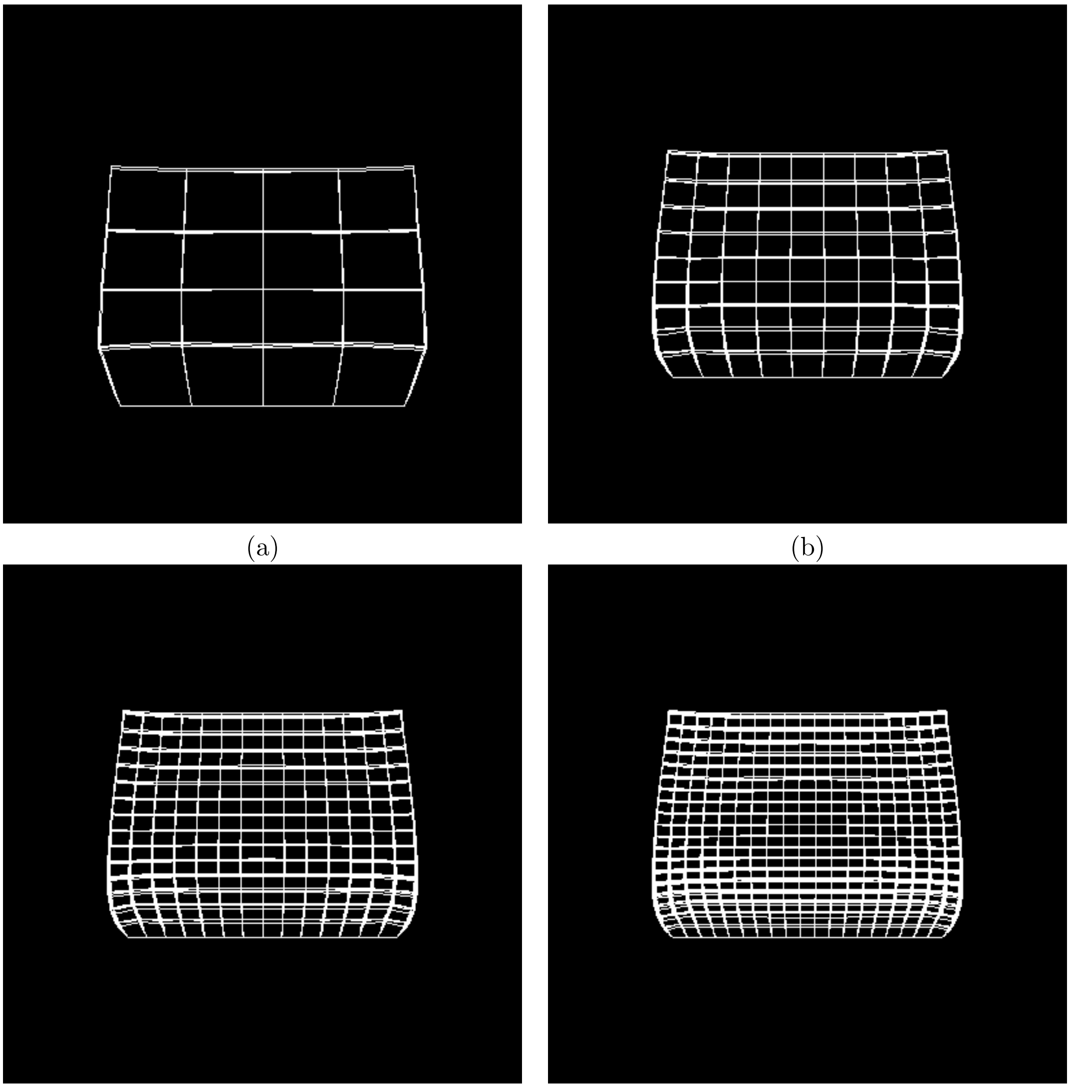

(a) (b) (c) (d)

Figure 4.12: Verification that the solution converges as the mesh is refined. Meshes have (a) 5^3, (b) 10^3, (c) 15^3, (d) 20^3 nodes.

Chapter 5

The Backward Problem for Computing the Reference State

In this chapter we discuss a new type of deformation simulation we have to perform, the *backward problem*. The backward problem[1] is the problem of determining the reference state of a deformable body, given a deformed configuration with known body forces. In other words, it is an inverse nonlinear elasticity problem, we are given $\mathbf{x}$ and wish to calculate $\mathbf{X}$. We do not use the terminology 'inverse problem' however, as inverse problems generally refer to the determination of material parameters given a known deformation (for example, [KS05]). Instead we define the problems discussed in Chapter 3, for example (3.23), as *forward problems*, where $\mathbf{X}$ is the known and $\mathbf{x}$ the unknown, and define the opposite as backward problems.

As we saw in Chapter 3, the equations of nonlinear elasticity assume the initial shape of the object under consideration is stress-free, i.e. there are no body or surface forces acting on the object, and the internal Cauchy stress is zero. In other words, it is assumed that the initial state is the true reference configuration. However, our aim is to build a mesh of the breast from MR data, in which case the breast will clearly be under the influence of gravity. Any initial mesh will be of the breast in a *gravity-loaded* state, not of the reference state. Before any other deformation simulations can be performed, we must remove the effect of gravity from the model to obtain the true reference state, which can then be used as the initial state in subsequent deformation calculations.

An important question which needs to be addressed is whether a true reference state exists. In general the answer is no, and that this is the case for the breast is evident by considering the skin. Skin *in vivo* is known to be under tension, which means that there is residual stress in the skin which will be present even in the absence of gravity. Similarly, it could be argued that most organs in the

---

[1] Terminology invented here.

body have no true reference state, as they will often carry internal stresses produced during growth. We begin in Section 5.1.1 by discussing the full backward problem, before stating our backward problem as one in which residual stress effects have been neglected. The effect of tension in the skin will be discussed further in Chapter 8.

In linear elasticity, if we wish to compute the undeformed shape given a deformed shape and the direction of gravity, we could solve this problem simply by computing a forward linear elastic simulation with gravity applied in the opposite direction. Unfortunately, in nonlinear elasticity we cannot write a backward problem as a forward problem with reversed body forces. We show how the equations of finite elasticity can be reformulated to solve the backward problem in Section 5.1.3, where we also describe how the new equations can be solved with the finite element method. Results of simulations are given in Section 5.2, along with verification of their accuracy and a discussion of the method. We will see that there is initially a massive difference in computation time between the forward and backward problems, and the final part of the chapter, Section 5.3, is concerned with determining why this arises, and motivating a very simple change which removes the time difference.

5.1 Analysis and Reformulation

5.1.1 The Reference State and the full Backward Problem

To state the full backward problem let us first introduce the concept of *compatibility*. Suppose we know the deformed shape of an elastic object, and have measured the stress field throughout the body. This means that we can compute the strain, using the strain energy function. The inverse problem for determining the reference state is then to determine an inverse deformation $\mathbf{X}(\mathbf{x})$ which produces that strain, which involves integrating the strain for the displacement. However, strain has 6 independent components (in 3D, and because of symmetry), but there are only 3 unknowns (the 3 components of displacement, or of old position), so such a problem is over-constrained, and there must be some additional conditions satisfied by the strain components. These conditions are known as *compatibility conditions*. A specified hypothetical or experimental strain (or a specified deformation gradient) is compatible if there exists a deformation giving rise to that strain or deformation gradient. In linear elasticity the compatibility question has a relatively simple answer. The strain is related to displacement by

$$\epsilon_{ij} = \frac{1}{2}(u_{i,j} + u_{j,i}),$$

and it can be seen after a little algebra that any such tensor defined in this way must satisfy

$$\nabla \times \nabla \times \epsilon = 0 \qquad \text{i.e.} \qquad \epsilon^{\text{ALT}}_{ikl}\,\epsilon^{\text{ALT}}_{jmn}\,\epsilon_{ln,mn} = 0$$

where $\epsilon^{\mathrm{ALT}}_{ijk}$ is the alternating tensor[2]. Given a symmetric tensor ϵ_{ij}, $\nabla \times \nabla \times \epsilon = 0$ is therefore a necessary condition for there to exist a displacement $\mathbf{u}$ from which ϵ_{ij} arises as a strain. Assuming the deformed state is a simply-connected[3] region, it can also be shown [SZJ$^+$96] that this is a sufficient condition as well. In other words, a strain ϵ_{ij} is compatible in linear elasticity if and only if $\nabla \times \nabla \times \epsilon$ vanishes. If the strain is compatible, it is possible to determine $\mathbf{u}$ by integrating ϵ_{ij}, see for example [SZJ$^+$96]. In finite elasticity the compatibility condition is more difficult to state, involving the Riemannian curvature tensor. We will discuss this further in Chapter 9.

We have introduced the concepts of compatibility and integration of strains to emphasise that the backward problem which we solve in this chapter is a simplification to the full problem of determining the true reference state, if one exists. The full problem cannot be stated in terms of deformed positions alone, instead strains in the deformed configuration also need to specified. These strains would then have to be analysed to determine if they were compatible, and if so, an inverse displacement computed by integrating the strains. If the strains are not compatible, the reference state can be considered to be a discontinuous state which is locally stress-free everywhere (see [Ols04]). For the breast, this could involve treating the skin and the interior as separate entities in a discontinuous reference state.

In this chapter we neglect any residual stresses in the breast (and in particular, in the skin), and make the assumption that the deformed breast arises from a stress-less continuous reference state loaded by gravity. Thus we implicitly make the assumption that the strains in the MR-imaged breast are compatible, because we assume that an inverse displacement exists, which we use to compute strains. In order to verify the validity of this statement, the approximate reference state which we will shortly show how to form could be used to compute the stresses in the MRI gravity-loaded breast. These stresses would then have to be compared to experimental data on the true stresses in the prone breast.

5.1.2 Potential Numerical Methods for Solving the Backward Problem

We now describe some possible methods that could be used to obtain the reference state, assuming that one exists and the known state is the deformed state obtained from that reference state under known body forces and tractions.

One very simple method which may be used in the case of zero-tractions to obtain a potential approximation is to solve a forward problem with the direction of the body force reversed. In linear elasticity, this method would give the exact reference state, precisely because the equations are

---

[2] The alternating tensor $\epsilon^{\mathrm{ALT}}_{ijk}$ satisfies: $\epsilon^{\mathrm{ALT}}_{ijk} = 1$ if (i, j, k) is an even permutation of $(1, 2, 3)$; $\epsilon^{\mathrm{ALT}}_{ijk} = -1$ if (i, j, k) is an odd permutation of $(1, 2, 3)$; and $\epsilon^{\mathrm{ALT}}_{ijk} = 0$ if any two of i, j or k are equal. We use the notation $\epsilon^{\mathrm{ALT}}_{ijk}$ instead of the more common notation ϵ_{ijk} to avoid confusion with the infinitesimal strain ϵ_{ij} in this equation.

[3] A domain is simply-connected if, loosely-speaking, it contains no holes.

linear, and because $\frac{\partial}{\partial x^i}$ is assumed to be equal to $\frac{\partial}{\partial X^M}$. Of course, in finite elasticity, the result is not guaranteed to even be similar to the true reference state.

A more sophisticated method is to solve a *sequence* of forward deformations, and use a non-linear optimisation procedure to try to converge to the true reference state. Suppose the initial known deformed state is Ω, and we wish to find the reference state Ω_0. Write $\mathcal{P}_g$ for the operation of computing a forward deformation with gravity acting in the 'correct' direction, and $\mathcal{P}_{-g}$ for the operation of computing the forward deformation with gravity acting in the reversed direction. Using the first method described above, we can obtain a first approximation to the reference state, $\Omega_0^{(0)}$ say, by computing $\Omega_0^{(0)} = \mathcal{P}_{-g}(\Omega)$. The important observation is that we can test how well $\Omega_0^{(0)}$ approximates the unknown Ω_0 by computing $\mathcal{P}_g(\Omega_0^{(0)})$ and comparing this with Ω. Given a suitable norm (which would be some function of nodal values), we can quantify this, and define an error $E_m = \|\mathcal{P}_g(\Omega_0^{(m)}) - \Omega\|$. Numerical optimisation algorithms are available which do not rely on the availability of the exact functional form for calculating a residual (in this case E_m) or its derivative. Such black-box algorithms instead only require a routine for computing the residual and could in principal be used to solve this nonlinear optimisation problem. However, this method would clearly suffer from being computationally expensive, with each run requiring the computation of potentially a large number of forward problems. This method does have one very large advantage though, which we will return to after describing the final method presented here.

The third method is essentially a mathematical implementation of the second method, and involves direct inversion of the equations of finite elasticity. Here, we set $\mathbf{x}$ as the known variable and $\mathbf{X}$ as the unknown, and make the observation that the Eulerian formulations that were derived in Chapter 3, which at the time were of little use and had to be converted to Lagrangian formulations, are exactly what we require now. Since this method effectively involves directly inverting the elasticity equations, we will refer to it as *the direct approach.* It is the method we choose to use to compute the reference state, with the details described in Section 5.1.3. This method has the advantage of being a lot faster than the second method in theory, and being an analytic inversion, more likely to be able to produce arbitrarily accurate results. It also has the advantage that analytic functional forms for residuals and Jacobians can be calculated by hand, so there is no need for numerical differentiation. However, the big drawback of the direct approach is that it is model-dependent. If a new or modified model of the deformable body was proposed, a new inversion would have to be carried out, and as we shall see in Section 5.1.3, the inversion and finite element implementation can be an algebraically complex procedure. The big advantage of the second method was that it was completely model-independent. The model of the forward deformation could be changed without having to make any alterations to the second method.

5.1.3 The Weak Form and Finite Element Equations using the Direct Approach

We now derive the finite element equations that are used to solve the backward problem by the direct approach. We assume no surface tractions, gravity as the body force, and incompressibility.

Let us define

$$G_i^M = \frac{\partial X^M}{\partial x^i}$$

as the deformation gradient of the inverse (backward) map, which is the same as the inverse of the standard (forward) deformation gradient. Incompressibility implies that $\det(F) = 1$, so that $\det(G) = 1$. We use G because it is explicitly a function of the known variable $\mathbf{x}$.

We begin from the weak formulation (3.13), amended for the incompressibility constraint

$$\int_\Omega \sigma^{ij}(\mathbf{u}(\mathbf{x}), p(\mathbf{x})) \frac{\partial(\delta u_j)}{\partial x^i}\,\mathrm{d}V - \int_\Omega \rho g^i \delta u_i\,\mathrm{d}V + \int_\Omega (\det(G(\mathbf{u})) - 1)\delta p\,\mathrm{d}V = 0 \qquad \forall \delta\mathbf{u},\, \delta p \tag{5.1}$$

where $\delta\mathbf{u}$ and δp are test functions.

This problem can be solved with the finite element method, using the same techniques as before. To evaluate σ we have to use the relationship between σ and T, and the definition of T for incompressible materials, giving

$$\sigma^{ij} = F_M^i \left(2\frac{\partial \overline{W}}{\partial C_{MN}} - p(C^{-1})^{MN}\right) F_N^j$$

or, in terms of G,

$$\begin{aligned} \sigma^{ij} &= (G^{-1})_M^i \left(2\frac{\partial \overline{W}}{\partial C_{MN}} - p(C^{-1})^{MN}\right) (G^{-1})_N^j \\ &= 2(G^{-1})_M^i \frac{\partial \overline{W}}{\partial C_{MN}} (G^{-1})_N^j - p\delta^{ij}, \end{aligned}$$

with $C_{MN} = (G^{-1})_M^i (G^{-1})_N^i$.

Defining $\mathcal{N}$ and $\mathcal{M}$ as in Section 4.2.1, we set $\delta\mathbf{u}^h = \begin{bmatrix} \phi_i \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ \phi_i \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 0 \\ \phi_i \end{bmatrix}$ $(i = 1, \ldots, \mathcal{N})$ and $\delta p^h = \mathbb{1}_{K_j}$ $(j = 1, \ldots, \mathcal{M})$ to obtain the $3\mathcal{N} + \mathcal{M}$ nonlinear equations which need to be solved:

$$\begin{aligned} \text{equations } 1 \text{ to } \mathcal{N}: &\quad \int_\Omega F_M^1 T^{MN} F_N^j \frac{\partial \phi_k}{\partial x^j}\,\mathrm{d}V - \int_\Omega \rho g_1 \phi_k\,\mathrm{d}V = 0, \\ \text{equations } \mathcal{N}+1 \text{ to } 2\mathcal{N}: &\quad \int_\Omega F_M^2 T^{MN} F_N^j \frac{\partial \phi_{k-\mathcal{N}}}{\partial x^j}\,\mathrm{d}V - \int_\Omega \rho g_2 \phi_{k-\mathcal{N}}\,\mathrm{d}V = 0, \\ \text{equations } 2\mathcal{N}+1 \text{ to } 3\mathcal{N}: &\quad \int_\Omega F_M^3 T^{MN} F_N^j \frac{\partial \phi_{k-2\mathcal{N}}}{\partial x^j}\,\mathrm{d}V - \int_\Omega \rho g_3 \phi_{k-2\mathcal{N}}\,\mathrm{d}V = 0, \\ \text{equations } 3\mathcal{N}+1 \text{ to } 3\mathcal{N}+\mathcal{M}: &\quad \int_\Omega (\det G - 1)\, \mathbb{1}_{K_{i-3\mathcal{N}}}\,\mathrm{d}V = 0 \Rightarrow \int_{\text{element } i - 3\mathcal{N}} (\det G - 1)\,\mathrm{d}V = 0. \end{aligned}$$

Defining the vector of unknowns $\mathbf{a}$ as in (4.12), we again write these equations as

$$\mathcal{F}(\mathbf{a}) = 0.$$

To reduce the number of cases that need to be considered and simplify the notation, let us define μ to be 1 if $0 < \alpha \leq \mathcal{N}$, 2 if $\mathcal{N} < \alpha \leq 2\mathcal{N}$ and 3 if $2\mathcal{N} < \alpha \leq 3\mathcal{N}$. Also, let $\gamma = \alpha - (\mu - 1)\mathcal{N}$. All of the first $3\mathcal{N}$ components of the residual vector can then be written as

$$\mathcal{F}_\alpha = \int_\Omega F^\mu_M T^{MN} F^j_N \frac{\partial \phi_\gamma}{\partial x^j} \, \mathrm{d}V - \int_\Omega \rho g^\mu \phi_\gamma \, \mathrm{d}V = 0.$$

We wish to express $\mathcal{J}_{\alpha\beta}$, the Jacobian used in Newton's method. Given β, let K and δ be defined as μ and γ were for α. Also, let $U^{K\gamma}$ represent U^γ if $K = 1$, V^γ if $K = 2$, and W^γ if $K = 3$. We can then express the components of the Jacobian, for $\alpha, \beta \leq 3\mathcal{N}$, as

$$\mathcal{J}_{\alpha\beta} = \int_\Omega \frac{\partial F^\mu_M}{\partial U^{K\delta}} T^{MN} F^j_N \frac{\partial \phi_\gamma}{\partial x^j} + F^\mu_M T^{MN} \frac{\partial F^j_N}{\partial U^{K\delta}} \frac{\partial \phi_\gamma}{\partial x^j} + \frac{\partial T^{MN}}{\partial E_{PQ}} \frac{\partial E_{PQ}}{\partial U^{K\delta}} F^\mu_M F^j_N \frac{\partial \phi_\gamma}{\partial x^j} \, \mathrm{d}V.$$

We need to calculate $\frac{\partial F^i_M}{\partial U^{K\delta}}$. Using the chain rule and calculating $\frac{\partial F}{\partial G}\frac{\partial G}{\partial U}$, we eventually find that

$$\frac{\partial F^i_M}{\partial U^{K\delta}} = F^i_K F^j_M \frac{\partial \phi_\delta}{\partial x^j}.$$

From this, we have

$$\begin{aligned} \frac{\partial E_{PQ}}{\partial U^{K\delta}} &= \frac{\partial}{\partial U^{K\delta}} \left(\frac{1}{2} \left(F^l_P F^l_Q - \delta_{PQ} \right) \right) \\ &= \frac{1}{2} \left(\frac{\partial F^l_P}{\partial U^{K\delta}} F^l_Q + \frac{\partial F^l_Q}{\partial U^{K\delta}} F^l_P \right) \\ &= \frac{1}{2} \left(\left(F^l_K F^k_P \frac{\partial \phi_\delta}{\partial x^k} \right) F^l_Q + \left(F^l_K F^k_Q \frac{\partial \phi_\delta}{\partial x^k} \right) F^l_P \right). \end{aligned}$$

Putting all this together we obtain

$$\begin{aligned} \mathcal{J}_{\alpha\beta} = \int_\Omega \quad & F^\mu_K F^k_M F^j_N T^{MN} \frac{\partial \phi_\gamma}{\partial x^j} \frac{\partial \phi_\delta}{\partial x^k} + F^\mu_M F^j_K F^k_N T^{MN} \frac{\partial \phi_\gamma}{\partial x^j} \frac{\partial \phi_\delta}{\partial x^k} \\ + \quad & \frac{1}{2} \frac{\partial T^{MN}}{\partial E_{PQ}} F^\mu_M F^j_N F^l_K (F^l_P F^k_Q + F^l_Q F^k_P) \frac{\partial \phi_\gamma}{\partial x^j} \frac{\partial \phi_\delta}{\partial x^k} \, \mathrm{d}V. \end{aligned} \tag{5.2}$$

The other block entries of the Jacobian are easier to compute. If $\alpha \leq 3\mathcal{N}$, $\beta > 3\mathcal{N}$, then

$$\begin{aligned} \mathcal{J}_{\alpha\beta} = \frac{\partial \mathcal{F}_\alpha}{\partial P_{\beta - 3\mathcal{N}}} &= \frac{\partial}{\partial P_{\beta - 3\mathcal{N}}} \left(\int_\Omega \sigma^{\mu j} \frac{\partial \phi_\gamma}{\partial x^j} - \rho g^\mu \phi_\gamma \, \mathrm{d}V \right) \\ &= -\int_{\text{element } \beta - 3\mathcal{N}} \frac{\partial \phi_\gamma}{\partial x^\mu} \, \mathrm{d}V. \end{aligned}$$

For $\alpha > 3\mathcal{N}$, $\beta \leq 3\mathcal{N}$

$$\begin{aligned}
\mathcal{J}_{\alpha\beta} = \frac{\partial \mathcal{F}_\alpha}{\partial U^{K\delta}} &= \frac{\partial}{\partial U^{K\delta}} \left(\int_{K_{\alpha-3\mathcal{N}}} (\det G - 1) \, \mathrm{d}V \right) \\
&= \int_{K_{\alpha-3\mathcal{N}}} \frac{\partial(\det G)}{\partial G_i^M} \frac{\partial G_i^M}{\partial U^{K\delta}} \, \mathrm{d}V \\
&= \int_{K_{\alpha-3\mathcal{N}}} -(\det G) F_K^i \frac{\partial \phi_\delta}{\partial x^i} \, \mathrm{d}V.
\end{aligned}$$

Finally, for $\alpha > 3\mathcal{N}$, $\beta > 3\mathcal{N}$,

$$\mathcal{J}_{\alpha\beta} = \frac{\partial \mathcal{F}_\alpha}{\partial P_{\beta-3\mathcal{N}}} = \frac{\partial}{\partial P_{\beta-3\mathcal{N}}} \left(\int_{K_{\alpha-3\mathcal{N}}} (\det G - 1) \, \mathrm{d}V \right) = 0.$$

Using these results, we can assemble the residual and Jacobian for the backward problem and solve using Newton's method as described in Chapter 4.

5.2 Numerical Results and Discussion

Some sample backward computations are displayed in Figures 5.1 and 5.2. In Figure 5.1, we take the initial deformed state to be the result of the simulation displayed in Figure 4.7 in Chapter 4 (i.e. the stretched cube). The true reference state, which this backward simulation aims to compute numerically, is therefore a cube. The result of the backward computation is shown in Figure 5.1(b), correctly taking the form of a cube. A similar computation with the sheared cube as the starting state is given in Figure 5.2. Again, the result of the backward computation is a cube. The average error (the average difference at the nodes between the result of the backward computation and the exact cube, divided by the cube's width) in this two computations are 0.00040 in Figure 5.1 and 0.00234 in Figure 5.2. These values verify that this direct method solves the backward problem to high accuracy. For a comparison, we have taken the sheared cube as the starting state, and performed a forward simulation with the direction of gravity reversed. The result, shown in Figure 5.3, is surprisingly accurate, looking similar to the cube, but has visibly curved edges. The average error here is 0.025.

The errors in these results arise from discretisation errors solving the forward problem followed by a backward problem. Although the weak form (5.1) for the backward problem is the exact mathematical inverse to the weak form of the forward problem (4.13), we would not expect the finite element versions of these weak forms to be exactly inverse. However, the error should converge to zero as the number of nodes increases. To investigate this, we have solved the same problems on several meshes of the cube, ranging from a very coarse mesh with $2^3 = 8$ nodes to a fairly dense mesh with $14^3 = 2744$ nodes. The average error per node is tabulated against mesh size in Table

Number of nodes	Scaled error
2^3	0.00682
4^3	0.00134
6^3	0.000634
8^3	0.000401
10^3	0.000276
12^3	0.000202
14^3	0.000156

Table 5.1: Scaled error (average error, per node, divided by cube width) between a cube and the result of a forward then backward computation (for the stretched cube simulation), against the number of nodes.

5.1, which verifies that the error decreases as the number of nodes increases, although it does so slowly.

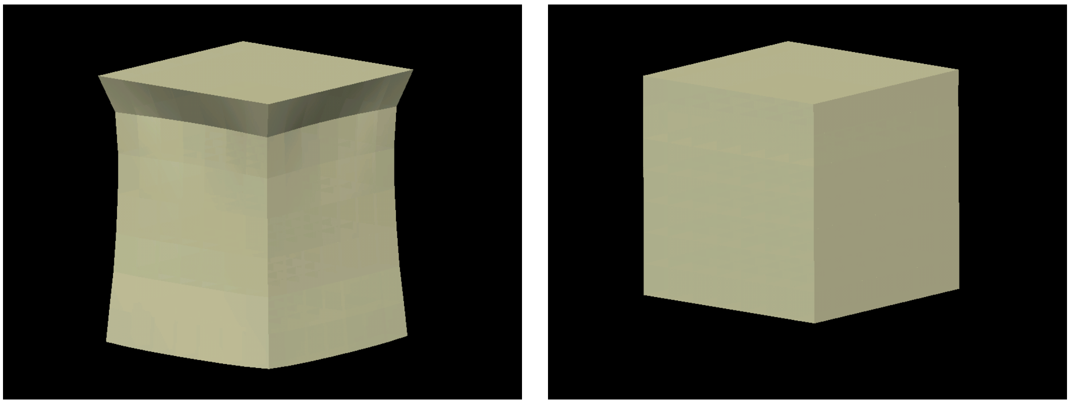

(a) (b)

Figure 5.1: Solving the backward problem directly: (a) Shape taken as the deformed starting state for a backward problem computation (obtained by solving a forward problem on a cube, see Fig 4.7), (b) result of the backward computation.

5.2.1 The Backward Problem with a Transversely-Isotropic Material Law

The backward problem using a material which is transversely-isotropic merits its own discussion, as in this situation there are two distinct types of problem possible, which differ conceptually, although are similar in implementation. We need to have a fibre direction in order to implement a transversely isotropic law, and the two scenarios arise when we consider the meaning of the fibre direction in the backward problem, and our *a priori* knowledge of it.

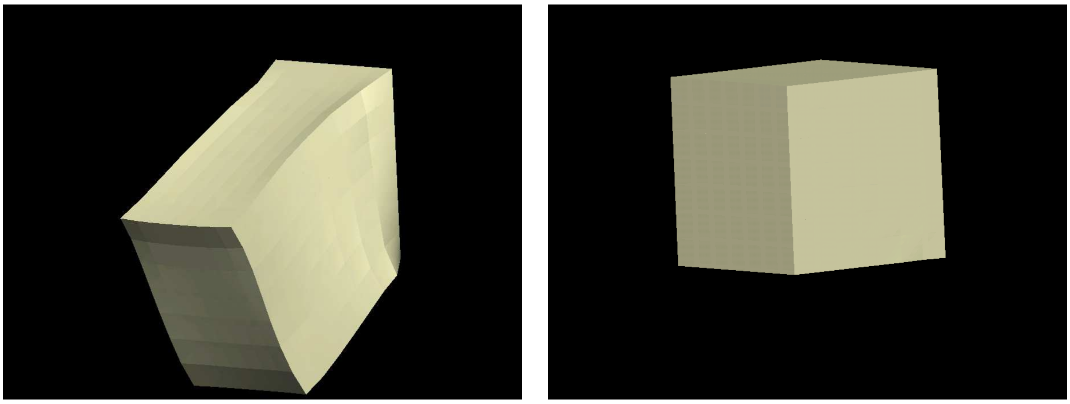
(a) (b)

Figure 5.2: Solving the backward problem directly: (a) Shape taken as the deformed starting state for a backward problem computation (obtained by solving a forward problem on a cube, see Fig 4.9), (b) result of the backward computation.

The first scenario is that we know the fibre directions in the loaded (deformed) state. This can occur if, for example, the loaded state was measured from a set of images, and it was possible to image the fibre directions as well. In breast MR images however, it is not possible to see any fibre directions, so the fibre direction must be chosen. The second scenario arises if we choose the fibre direction in the reference state, not the loaded state. This may seem an unnatural situation, since it involves assuming fibre directions in the unknown state, but it is one which may occur when we consider possible fibre directions for fatty tissue, with the most sensible choice being to take the fibres to be in the muscle-to-skin direction in the reference state.

There is very little difference between these scenarios when implementing the direct method to solve the backward problem with transversely isotropic materials, because the fibre direction in the undeformed and deformed states are simply related by $M_N = F^i_N m_i$, where $\mathbf{M}$ and $\mathbf{m}$ are the (non-normalised) fibre directions in the undeformed and deformed states respectively. To verify the results of backward computations using transversely isotropic laws, we used the direct method to retrieve the reference state using the result shown in Figure 4.11(d) as the known starting state. Again, the reference state was correctly computed as the cube, and the average error per node was 0.00025.

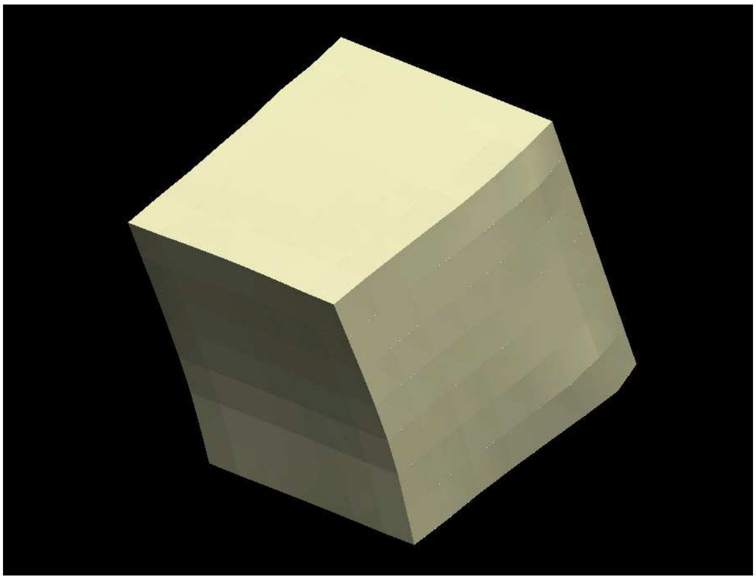

Figure 5.3: The result of attempting to solve a backward problem by computing the solution of a forward problem applying gravity in the opposite direction.

5.3 GMRES on the Forward and Backward Problems

One of the limitations encountered solving the backward problem with this direct approach is that computation time is significantly higher than for forward problems on the same sized meshes. Table 5.3 gives the time taken to solve a sample forward and backward problem, and shows a large discrepancy between the two types of problem. The total computation time grows much faster with mesh-size for the backward problem than the forward problem. This is obviously unsatisfactory, ideally we wish to be able to solve both in similar times. This time-difference would be especially highly restrictive for breast-deformation problems, as here the first calculation, obtaining the reference state from a mesh formed from MR images, would be limited by mesh-size. This would inhibit dense meshes being used in subsequent forward simulations.

The main reason for the time difference is that the solution of linear systems with GMRES(m) takes much longer for backward problem linear systems, suggesting that the preconditioning does not work as effectively with the backward problem as the forward problem. We now consider why this is the case, and identify the fundamental factor that dictates the speed of GMRES(m) on both backward and forward problems.

Figure 5.4 displays the rate of convergence of GMRES(50) on a sample (non-dimensionalised)

Number of nodes	Computation Time Forward Problem (s)	 Backward Problem (s)
2^3	0.5	0.5
4^3	15	19
6^3	88	111
8^3	204	345
10^3	451	1555
12^3	1077	7428
14^3	1884	48504

Table 5.2: Total computation times for forward and backward simulations on a number of meshes.

backward problem linear system, compared to the convergence history for an equivalent (non-dimensionalised) forward problem. The residuals for the backward problem do decrease exponentially, and do not stagnate, but the rate of decrease is much slower than with the forward problem.

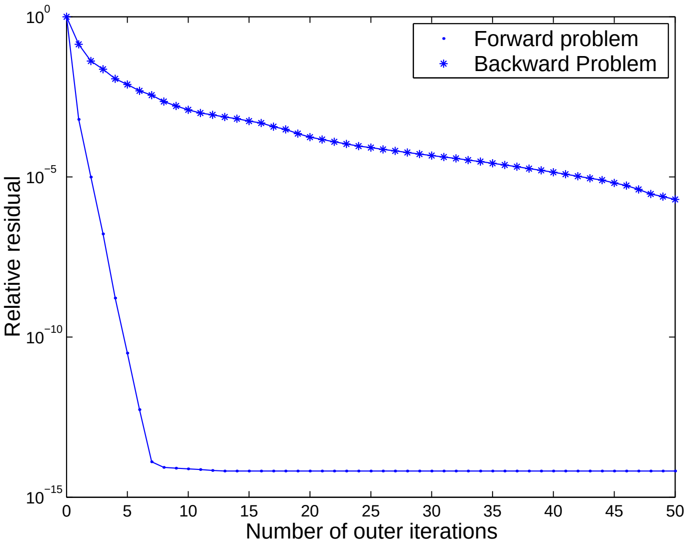

Figure 5.4: Convergence histories for GMRES(50) on linear systems arising from a forward and backward problem.

As we discussed in Section 4.1.6, positive theorems characterising GMRES-convergence are not as readily available as for other solvers, because many examples can be constructed for GMRES where convergence is slow, or stagnation occurs. In particular, it can be shown that the convergence behaviour of GMRES cannot be characterised by the spectrum[4] of the matrix (unlike, for example, conjugate gradients, where the convergence rate is related to the ratio of the largest eigenvalue

---

[4] The spectrum of a matrix is its set of eigenvalues.

to the smallest [vdV03])[5]. In addition, very little can be said about GMRES(m), not even that convergence is faster as m increases. For example, linear systems are constructed in [Emb99] for which GMRES(1) converges but GMRES(2) stagnates.

However, such constructions are in a sense un-physical, and it has been observed that linear systems arising from physical problems will not generally lead to extreme behaviour. For matrices formed from physical problems, it has been observed that GMRES will generally converge faster if the matrix has *clustered eigenvalues*, and also that convergence will generally be poorer if there are eigenvalues near the origin. Although such statements cannot be true in general for all matrices, they do suggest that the spectrum of a matrix can give important insight into the convergence behaviour of GMRES.

The spectra of some of the matrices encountered thus far are plotted in Figure 5.5. We saw in Section 4.2.3 the massive effect non-dimensionalisation had on GMRES convergence. Consider Figures 5.5(a) and (b), which show the spectra of a dimensionalised and a non-dimensionalised forward problem matrix. The former has eigenvalues which are completely real and contained in the interval $[0, 5]$, the latter has eigenvalues that have non-zero imaginary part, but are more closely clustered together, with real part $\in [0, 3.5]$ and imaginary part $\in [-1, 1]$. The spectra would appear to verify the statement that clustered eigenvalues lead to better GMRES properties. Now consider Figures 5.5(c) and (d). Figure 5.5(c) is the spectrum of a matrix arising from the first Newton iteration of a (non-dimensionalised) backward problem, 5.5(d) is the spectrum for the second Newton iteration of the same calculation. We see that backward problem spectra are fundamentally different to forward problem spectra. The spectrum for the first Newton iteration is completely real, the spectrum for the second has small non-zero imaginary components, but in both cases, the spectrum has eigenvalues which have *negative real parts*, i.e. the backward problem matrices are indefinite. All spectra for forward problems were definite[6].

Now, let $A \in \mathbb{R}^{m\times m}$ be a positive-definite symmetric square matrix, and $B \in \mathbb{R}^{m\times(n-m)}$ be a rectangular matrix with full rank ($\text{rank}(B) = n - m$), and define

$$\mathcal{A}_+ = \begin{bmatrix} A & B \\ B^{\mathrm{T}} & 0 \end{bmatrix}. \tag{5.3}$$

---

[5] The following (devastating) theorem [GPS96] shows that the convergence behaviour of GMRES cannot be characterised by the eigenvalues of the matrix, and worse. Let a_k, $k = 0, \ldots, N$ be a monotonically decreasing (not necessarily strictly) set of positive real numbers, with $a_0 = 1$ and $a_N = 0$. Then, for *any* given set of eigenvalues, a matrix A with these eigenvalues and a vector $\mathbf{b}$ can be constructed such the convergence history of GMRES applied to the matrix problem $A\mathbf{x} = \mathbf{b}$ has relative residuals satisfying $\frac{\|\mathbf{r}_k\|}{\|\mathbf{r}_0\|} = a_k$. We could take $a_k = 1$ for all $k = 0, \ldots, N-1$, and then the theorem states that for any set of eigenvalues (even all equal to one), it is possible for GMRES to stagnate completely until the final step.

[6] The terminology *definite* and *indefinite* can be used to refer to matrices with non-real spectra. Definite means the spectrum is confined to the right-half plane (all eigenvalues have positive real part), indefinite means the spectrum is not confined to either half-plane.

$\mathcal{A}_+$ is therefore an $n \times n$ symmetric matrix. Also, define

$$\mathcal{A}_- = \begin{bmatrix} A & B \\ -B^{\mathrm{T}} & 0 \end{bmatrix} \tag{5.4}$$

which is neither symmetric nor skew-symmetric, but 'block-skew-symmetric'. Since $\mathcal{A}_+$ is symmetric it has real eigenvalues. It can also be shown that $\mathcal{A}_+$ is indefinite, with n positive eigenvalues and $n-m$ negative eigenvalues. $\mathcal{A}_-$ is not symmetric so may have complex eigenvalues, but it can be shown that all eigenvalues must have positive real part, and therefore $\mathcal{A}_-$ is definite.

Matrices such as $\mathcal{A}_+$ and $\mathcal{A}_-$ arise from finite element implementations of linear mechanical problems with constraints, for example incompressible linear elasticity, or the *Stokes' problem*, for which the equations for the flow velocity $\mathbf{u}$ and hydrostatic pressure p of an incompressible fluid are given by

$$\begin{aligned} \nabla^2\mathbf{u} + \boldsymbol{\nabla} p &= \mathbf{f} \\ -\boldsymbol{\nabla}\cdot\mathbf{u} &= 0 \end{aligned}$$

where the second equation is the incompressibility constraint. The constraint is written as $-\boldsymbol{\nabla}\cdot\mathbf{u} = 0$ instead of $\boldsymbol{\nabla}\cdot\mathbf{u} = 0$ because the gradient operator is adjoint to the negative of divergence, and therefore the whole problem is self-adjoint, which leads naturally to a symmetric finite element global stiffness matrix of the form $\mathcal{A}_+$.

The matrices constructed using Newton's method in finite elasticity (both forward or backward problems) are all of the form

$$\mathcal{J} = \begin{bmatrix} P & Q \\ R^{\mathrm{T}} & 0 \end{bmatrix},$$

where $P \in \mathbb{R}^{m\times m}$ is square but not necessarily symmetric, and $Q, R \in \mathbb{R}^{m\times(n-m)}$ are rectangular but not equal. The Jacobian $\mathcal{J}$ is not symmetric nor block-symmetric nor block-skew-symmetric. However, if it was linearised about small displacements it would reduce to a linear elasticity stiffness matrix and take the form of $\mathcal{A}_+$ or $\mathcal{A}_-$. It now becomes clear, from Figure 5.5, that the eigenvalues of the finite elasticity Jacobians have distributions which are dominated by the linear part of the Jacobian, because all the spectra have forms similar to the spectra of $\mathcal{A}_+$ or $\mathcal{A}_-$. For example, the forward problem Jacobians are not block-skew-symmetric, but their eigenvalues still all have positive real part. It is clear that the forward problem Jacobians are (similar[7] to) perturbations of a block-skew-symmetric matrix of the form $\mathcal{A}_-$. Likewise, the backward problem Jacobians whose spectra are shown in Figures 5.5(c) and 5.5(d) were not symmetric, but it is clear they were (similar to) perturbations of a matrix of the form $\mathcal{A}_+$. Both have spectra which lie on or close to the real line, but are indefinite, with m positive eigenvalues and $n-m$ negative eigenvalues.

---

[7] Here we mean similar in the linear algebra sense, i.e. represent the same linear map under a different basis.

The indefinite spectra in Figure 5.5 are much less closely clustered than the definite spectra, and also have eigenvalues closer to zero, both of which are expected to slow GMRES convergence. We can attempt to remove the GMRES time-difference by converting the indefinite spectrum to a definite spectrum, by replacing a backward problem Jacobian, which is a perturbation of $\mathcal{A}_+$, with one which is a perturbation of $\mathcal{A}_-$. To do so, we simply have to change the sign of the appropriate block of the Jacobian. This could be thought of as applying a simple preconditioner. It is also identical to replacing the backward weak formulation (5.1)

$$\int_\Omega \sigma^{ij}(\mathbf{u}(\mathbf{x}), p(\mathbf{x})) \frac{\partial(\delta u_j)}{\partial x^i}\,\mathrm{d}V - \int_\Omega \rho g^i \delta u_i \,\mathrm{d}V + \int_\Omega (\det(G(\mathbf{u})) - 1)\delta p \,\mathrm{d}V = 0 \qquad \forall \delta\mathbf{u},\, \delta p$$

with the completely equivalent weak form

$$\int_\Omega \sigma^{ij}(\mathbf{u}(\mathbf{x}), p(\mathbf{x})) \frac{\partial(\delta u_j)}{\partial x^i}\,\mathrm{d}V - \int_\Omega \rho g^i \delta u_i \,\mathrm{d}V - \int_\Omega (\det(G(\mathbf{u})) - 1)\delta p \,\mathrm{d}V = 0 \qquad \forall \delta\mathbf{u},\, \delta p.$$

The spectrum of a modified backward problem Jacobian is shown in Figure 5.5(e), as hoped it lies wholly in the right-half plane. Figure 5.6 displays the convergence history of GMRES(50) on the modified backward problem Jacobian, and shows that the modification completely removes the time-difference between the two types of problem, with GMRES(50) converging at the same rate for the backward problem as the forward problem. Finally, the computations in Table 5.3 are repeated with this modification, and shown in Table 5.3. Massive gains in computation time are evident, with the time to solve backward problems now growing at the same rate as forward problems. All backward problems take approximately 1.1 to 1.2 times longer than their equivalent forward problems. (This is due to Jacobian assembly taking approximately 1.1 times longer for backward problems). Note that Table 5.3 gives the time to solve a full nonlinear problem, and the gains in time show that the modification works for linear systems in *all* the Newton iterations, even later Newton iterations where displacements will be at their highest, and when Jacobians will look least like $\mathcal{A}_+$ or $\mathcal{A}_-$.

We have shown in this section how the choice of weak formulation is crucial when it comes to solving linear systems. This applies to forward problems as well as backward problems. For example, if we had chosen the weak form

$$\int_{\Omega_0} T^{MN}(\mathbf{u}, p) F^k_M(\mathbf{u}) F^k_N(\delta\mathbf{u}) \,\mathrm{d}V_0 - \int_{\Omega_0} \rho_0 \mathbf{g} \cdot \delta\mathbf{u} \,\mathrm{d}V_0 - \int_{\Omega_0} (\det F(\mathbf{u}) - 1)(\delta p) \,\mathrm{d}V_0 = 0 \qquad \forall\, \delta\mathbf{u},\, \delta p$$

instead of

$$\int_{\Omega_0} T^{MN}(\mathbf{u}, p) F^k_M(\mathbf{u}) F^k_N(\delta\mathbf{u}) \,\mathrm{d}V_0 - \int_{\Omega_0} \rho_0 \mathbf{g} \cdot \delta\mathbf{u} \,\mathrm{d}V_0 + \int_{\Omega_0} (\det F(\mathbf{u}) - 1)(\delta p) \,\mathrm{d}V_0 = 0 \qquad \forall\, \delta\mathbf{u},\, \delta p$$

(both of which are completely equivalent), then the forward problem Jacobians would all be perturbations of $\mathcal{A}_+$ and their spectra would be indefinite. GMRES-convergence would then be as slow for the forward problem as it had been for the backward problem. The importance of using a weak form

	Computation Time		
Number of nodes	Forward (s)	Backward (s)	Modified Backward (s)
2^3	0.5	0.5	0.5
4^3	15	19	18
6^3	88	111	98
8^3	204	345	279
10^3	451	1555	535
12^3	1077	7428	1214
14^3	1884	48054	2426

Table 5.3: Total computation times for forward, original backward, and modified backward simulations on a number of meshes.

in finite elasticity which leads to a definite Jacobian does not seem to be well-known, we have found no mention of it in the literature. Studies of iterative solver performance in the context of the block structure appear to be restricted to fluid dynamics problems, for example for Stoke's flow [BBGS04] or the Navier-Stokes equations [GW98]. Note that although the Navier-Stokes equations are nonlinear, the nonlinearities are confined to the momentum equation. The constraint of incompressibility is still linear ($\boldsymbol{\nabla} \cdot \mathbf{u} = 0$), and therefore the matrices still have a block symmetric form of the form $\mathcal{A}_+$ or $\mathcal{A}_-$, just with A in (5.3) no longer symmetric and positive definite. In finite elasticity the incompressibility constraint is nonlinear ($\det(F) = 1$), which destroys block-symmetry/block-skew-symmetry. It would also appear that GMRES performance does not depend so dramatically on the block-symmetry/skew-symmetry in other applications than it apparently does for finite elasticity. It may be that our identification that block-structure dictates GMRES convergence in finite elasticity is new, and the precise explanation why this is so is an open problem.

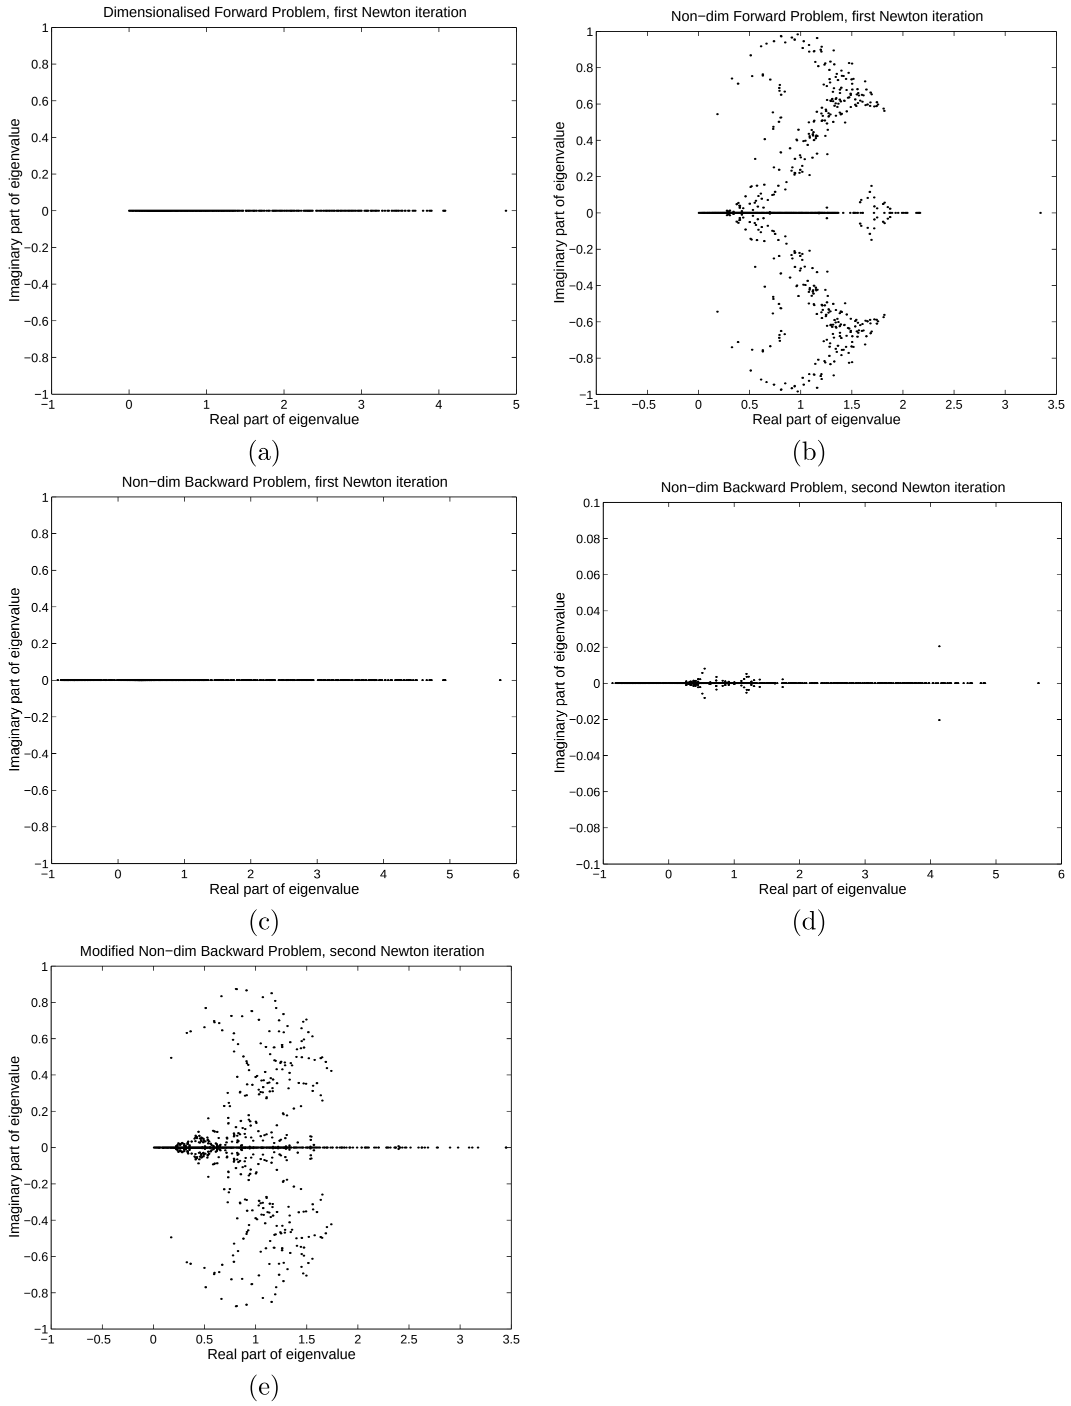

(a) (b) (c) (d) (e)

Figure 5.5: Spectra of some matrices: (a) dimensionalised forward problem, (b) non-dimensionalised forward problem, (c) non-dimensionalised backward problem, first Newton iteration (observe how this is fundamentally different to (a) as the eigenvalues do not all have positive real part), (d) non-dimensionalised backward problem, second Newton iteration, (e) modified non-dimensionalised backward problem, second Newton iteration.

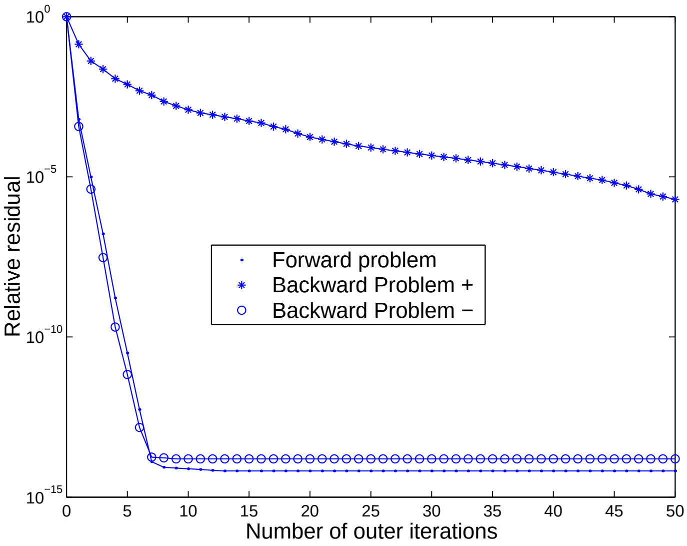

Figure 5.6: Convergence histories for GMRES(50) on the forward problem and two versions of a backward problem matrix. ‘Backward Problem +’ represents the original backward problem matrix, a perturbation of $\mathcal{A}_+$, and ‘Backward Problem -’ represents the modified backward problem matrix, a perturbation of $\mathcal{A}_-$.

Chapter 6

The Contact Problem

6.1 Introduction to the Contact Problem

In this chapter we discuss the *contact problem*, which arises when we attempt to simulate the shape of the breast when compressed for mammography. Contact problems differ from standard forward problems because the region of contact is unknown, and therefore surface tractions, even supposing they were known, cannot be applied as Neumann boundary conditions. Instead, the problem has to be formulated as two bodies deforming under their body forces, with displacement and zero traction boundary conditions, with the *constraint* that neither body is allowed to penetrate the other. This leads to different strong, weak and energy formulations, and, as we shall see, to a new class of numerical problem to be solved.

There are two types of contact, *frictionless* (or *smooth*) contact, and *frictional* contact. In the former we specify zero friction between the contact surfaces, from which it follows that the tangential tractions along the contact surfaces are zero. Given a Cauchy stress field σ, the traction on a surface is given by $\mathbf{t} = \sigma\hat{\mathbf{n}}$, where $\hat{\mathbf{n}}$ is the unit normal. Let us decompose $\mathbf{t}$ into its normal and tangential components

$$\mathbf{t} = \mathbf{t}_N + \mathbf{t}_T,$$

where the normal component $\mathbf{t}_N$ is given by $\mathbf{t}_N = \hat{\mathbf{n}} \cdot \mathbf{t} = \hat{\mathbf{n}} \cdot \sigma\hat{\mathbf{n}}$ and the tangential component $\mathbf{t}_T$ is $\mathbf{t}_T = \mathbf{t} - \mathbf{t}_N$. In frictionless problems we set $\mathbf{t}_T = 0$ and use $t_N = \|\mathbf{t}_N\|$ in the analysis. Frictional problems are most commonly modelled using the Coulomb model of friction, which states that

$$\|\mathbf{t}_T\| \leq \mu\, t_N,$$

where $\mu > 0$ is the coefficient of friction and depends on both the surfaces in contact. There are two types of frictional scenario: either $\|\mathbf{t}_T\| < \mu\, t_N$, in which case the points in contact are said to *stick*, or $\|\mathbf{t}_T\| = \mu\, t_N$, in which case the points are said to *slip*. In this work, we restrict ourselves to the frictionless case only.

When contact problems are simulated using the finite element method, immediate problems arise. One is that the finite element mesh is not smooth, so contact between surfaces cannot be modelled at every point between the finite element boundaries of two curved surfaces. Another issue is that contact mechanics involves the use of the surface normal, which does not vary continuously across the boundary of a finite element mesh.

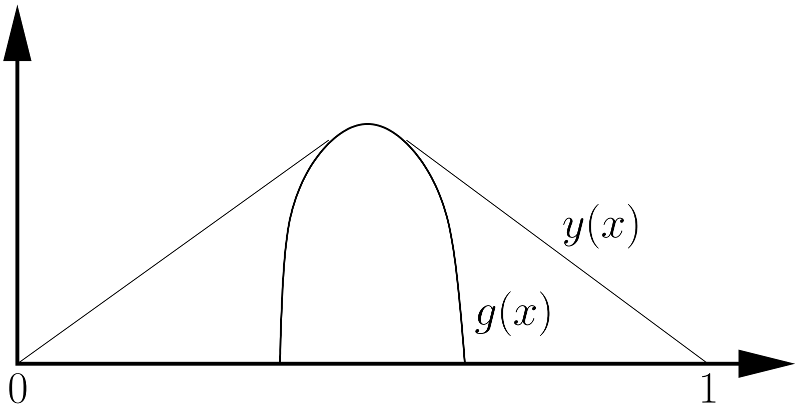

Figure 6.1: The obstacle problem: a string $y(x)$ stretched over a fixed obstacle $g(x)$. The string takes the shape that minimises the total length of y, subject to $y \geq g$

Before beginning the analysis of the full contact problem in three-dimensions, let us consider a simple 1D example to illustrate some of the main concepts. Suppose an elastic string is stretched over a smooth object with its ends fixed, as illustrated in Figure 6.1. Let the height of the string be given by $y(x)$, for $0 \leq x \leq 1$, and let the height of the obstacle be $g(x)$. This is known as the *obstacle problem.* Without worrying at the moment about the equations of state for an elastic string, we can make several observations on y. Firstly, we have the constraint on the string that $y \geq g$, or in other words $y - g \geq 0$. Also, y must have either zero or negative second derivative. The former case occurs when the string is not in contact with the obstacle and therefore stretched taut to a straight line. The fact that y must have negative second derivative when in contact is intuitively obvious: if y'' is positive at a point, then locally y looks like a concave quadratic, so the string is not stretched taut. The final observation is that either the string is taut, in which case $y'' = 0$, or the string is in contact with the obstacle, in which case $y - g = 0$. We can write this more succinctly as

$y''(y-g)=0$ everywhere. Putting these three observations together we get

$$\begin{aligned} y'' &\leq 0 \\ y-g &\geq 0 \\ y''(y-g) &= 0 \end{aligned} \tag{6.1}$$

These equations form a *linear complementarity problem*, and correspond to the strong form of an unconstrained problem. In fact, there is a formulation which corresponds to the weak form of the unconstrained problem: (6.1) can be shown to imply that $\int_0^1 y''(y-z)\,\mathrm{d}x \geq 0$ for all $z \in C^1([0,1])$ satisfying $z < g$ [KO88]. This is a *variational inequality.* Now, note that in the absence of an obstacle, the equation of state (Euler-Lagrange equation) for the string would be $y''=0$. Hence, in the region where the constraint does not hold, the unconstrained Euler-Lagrange equation does. This is a crucial observation. Finally, let us consider the gradient of the string. If the obstacle does not have a continuous first derivative, the string need not be smooth either, as can be inferred by considering a string stretched over a thin obstacle that has a thin sharp point. However if the obstacle is smooth, a geometrical argument can be used to come to the conclusion that y has continuous first derivative. Of course, (6.1) shows that y will have a discontinuous second derivative unless no contact occurs. All these conclusions can be extended to the general case. We note that in the situation which we ultimately aim to model, breast compression by mammographic plates, the constraint for, say, the upper plate will be of the form (assuming the plate is parallel to the (x,y)-plane):

$$z \leq \begin{cases} k & \text{if } y < y_0 \\ \infty & \text{otherwise.} \end{cases}$$

where (x,y,z) is the deformed new position and k and y_0 are constants, as shown in Figure 6.2. Although this is a discontinuous constraint function, it can effectively be considered to be a piecewise linear constraint, because (in physically realistic situations) it is easy to imagine replacing the constraint with a piecewise linear constraint which gives the same solution. However, whether considered discontinuous or piecewise linear, the constraint is non-smooth, and therefore we do not expect the displacement to have a continuous first derivative.

6.2 Mathematical Analysis

In this section, we set up the equations and inequalities that govern contact problems, in order to use them to motivate numerical algorithms. As in Chapter 3, the first step is to consider the kinematics of the deformation, in order to define important quantities of interest. In the general contact problem, there are assumed to be two deformable bodies, represented by $\Omega_0^{(1)}$ and $\Omega_0^{(2)}$ in

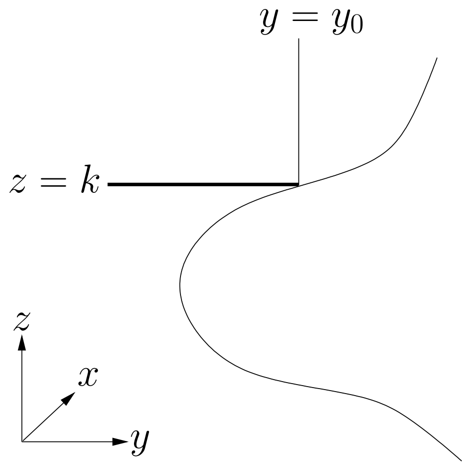

Figure 6.2: When the breast is compressed by a plate (thick line), the constraint is of the form $z < k$ if $y < y_0$. The plate, which effectively is a discontinuous constraint, can be replaced with a continuous constraint (thick line plus thin line) without changing the solution.

their undeformed configurations, and $\Omega^{(1)}$ and $\Omega^{(2)}$ in their deformed configurations. Define $\Gamma^{(1)}$ and $\Gamma^{(2)}$ as the (deformed) surfaces of contact in each body. Similarly define undeformed and deformed coordinates $\mathbf{X}^{(1)}, \mathbf{x}^{(1)}, \mathbf{X}^{(2)}$ and $\mathbf{x}^{(2)}$.

The two bodies are referred to as the master and slave bodies, with $\Omega^{(1)}$ taken to be the slave and $\Omega^{(2)}$ the master. In this chapter we will only consider the scenario where one body is rigid, and is used to compress the second deformable body, as these are the types of scenario which we wish to model. For reasons which will become clear we set the rigid plate to be the master body $\Omega^{(2)}$ and the deformable body to be the slave $\Omega^{(1)}$. Clearly, since the plate is rigid, $\Omega_0^{(2)} = \Omega^{(2)}$ and $\mathbf{X}^{(2)} = \mathbf{x}^{(2)}$.

The distinction in how the master and slave are used arises when we define the *gap function.* For any point $\mathbf{x}^{(1)}$ (in the deformed configuration) there exists a point $\bar{\mathbf{x}}^{(2)}$ which minimises the distance from $\mathbf{x}^{(1)}$ to the plate boundary:

$$\bar{\mathbf{x}}^{(2)} = \underset{\mathbf{x}^{(2)} \in \Gamma^{(2)}}{\arg\min} \|\mathbf{x}^{(1)} - \mathbf{x}^{(2)}\|.$$

It follows that $\bar{\mathbf{x}}^{(2)}$ is the orthogonal projection of $\mathbf{x}^{(2)}$ onto $\Gamma^{(2)}$, as displayed in Figure 6.3. It now becomes clear why we have chosen the rigid body as the master: in the context of the breast being compressed by plates, the plates are flat and it is trivial to compute the nearest point to a flat object. Note that, in general, the nearest point $\bar{\mathbf{x}}^{(2)}$ is not necessarily unique, in which case we are free to choose any possible nearest point.

We are now ready to define the gap function. Let the *distance vector* be

$$\mathbf{d}_N = \mathbf{x}^{(1)} - \bar{\mathbf{x}}^{(2)},$$

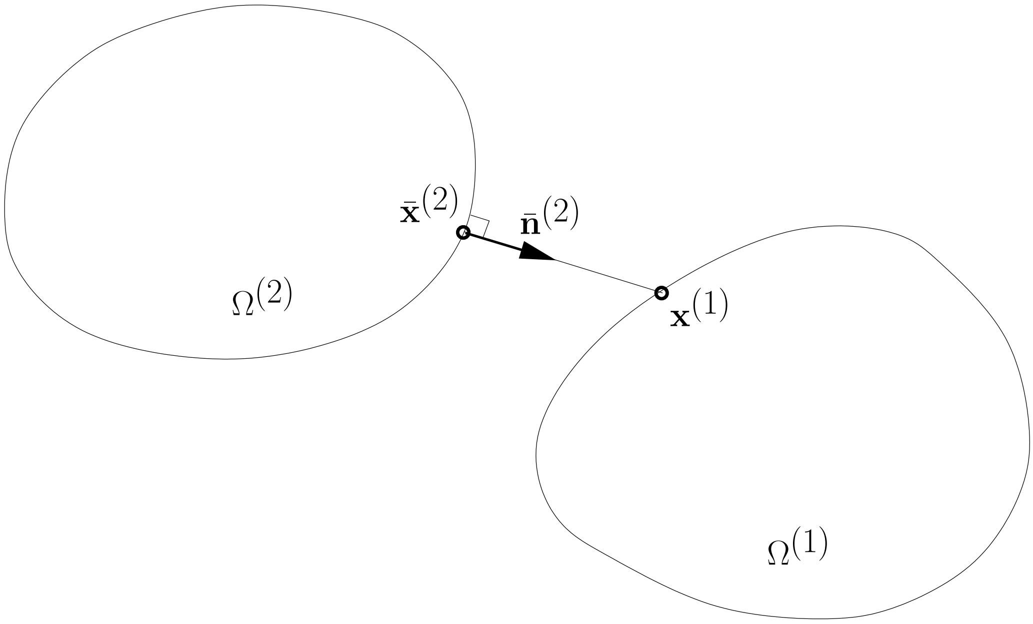

Figure 6.3: A point on the slave surface $\mathbf{x}^{(1)}$ and the nearest point $\bar{\mathbf{x}}^{(2)}$ on the master surface.

and define the gap function as

$$\begin{aligned} d_N &= \bar{\mathbf{n}}^{(2)} \cdot \mathbf{d}_N \\ &= \bar{\mathbf{n}}^{(2)} \cdot (\mathbf{x}^{(1)} - \bar{\mathbf{x}}^{(2)}) \end{aligned}$$

where $\bar{\mathbf{n}}^{(2)}$ is the unit normal of the master surface at $\bar{\mathbf{x}}^{(2)}$, which is parallel to $\mathbf{d}_N$. This is equal to $\|\mathbf{d}_N\|$ if $\mathbf{d}_N$ and $\bar{\mathbf{n}}$ point in the same direction, and equal to $-\|\mathbf{d}_N\|$ if they point in opposite directions. Therefore, the gap function is the *signed normal distance* between a point on the slave surface and the master surface. If d_N is negative at any point on the slave surface it means that *penetration* of that point into the master body has occurred, which is forbidden. We can now state the contact problem as a constrained energy minimisation problem:

$$\min_{\mathbf{x}^{(1)}} \mathcal{E}(\mathbf{x}^{(1)}) \qquad \text{subject to } d_N(\mathbf{x}^{(1)}) \geq 0 \quad \forall\, \mathbf{x}^{(1)} \in \Gamma^{(1)}, \tag{6.2}$$

where $\mathcal{E}$ is the total energy, defined in (3.16) (in the case of no applied surface tractions). To convert this constrained minimisation into a set of equalities and inequalities, we have to call upon a fundamental result from the theory of constrained optimisation. Let λ_N, a field defined on the slave contact surface, be a Lagrange multiplier corresponding to the constraint $d_N \geq 0$. (Here, the subscript N just denotes 'normal', as with d_N and t_N). Let

$$\mathcal{L}(\mathbf{x}^{(1)}, \lambda_N) = \mathcal{E}(\mathbf{x}^{(1)}) + \int_{\partial\Omega_0^{\text{trac}}} \lambda_N d_N \,\mathrm{d}S \tag{6.3}$$

The *first-order optimality conditions* for (6.2), necessary conditions that must occur at a solution of (6.2), are [NW99]

$$\begin{aligned} D_{\delta\mathbf{x}^{(1)}}(\mathcal{L}) &= 0 \qquad \forall\, \delta\mathbf{x}^{(1)} \text{ (compatible with Dirichlet b.c.s)} \\ \lambda_N &\leq 0 \\ d_N &\geq 0 \\ \lambda_N d_N &= 0. \end{aligned} \tag{6.4}$$

Here we have again used the concept of a directional derivative, $D_{\delta\mathbf{x}}$, as defined in Chapter 3. The three final equations are known as the *Karush-Kahn-Tucker (KKT) conditions.* The Lagrange multipliers $\lambda_N = \lambda_N(\mathbf{x}^{(1)})$ can be interpreted as the normal force acting between the two bodies at the point $\mathbf{x}^{(1)}$, i.e.

$$\lambda_N = t_N.$$

We can see how the KKT conditions are the equivalents of (6.1) for this problem. The interpretations of the KKT conditions are: $d_N \geq 0$ states that there is no penetration of one body into the other; $\lambda_N \leq 0$ states that the force acting from one body on the other can only be compressive, i.e. it can only push the bodies apart, not pull them together; $d_N\lambda_N = 0$ states that, at any point, either the bodies are in contact or there is no contact force there. Note that the gap d_N and the compressive force λ_N are a work-conjugate pair, in the same sense that the Cauchy stress and Almansi strain are work-conjugate, and the 2nd Piola-Kirchoff stress and Green strain are work-conjugate (as was discussed in Section 3.1.3).

In order to solve the contact problem, we have to solve (6.4) for the unknowns $\mathbf{x}$ and λ_N. Consider the directional derivative of the Lagrangian

$$\begin{aligned} D_{\delta\mathbf{x}^{(1)}}(\mathcal{L}) &= D_{\delta\mathbf{x}^{(1)}}(\mathcal{E}) + D_{\delta\mathbf{x}^{(1)}}\left[\int_{\Gamma^{(1)}} \lambda_N d_N \,\mathrm{d}S\right] \\ &= \mathcal{W}(\mathbf{x}^{(1)}, \delta\mathbf{x}^{(1)}) + \int_{\Gamma^{(1)}} \lambda_N \delta d_N \,\mathrm{d}S, \end{aligned} \tag{6.5}$$

where $\delta d_N = D_{\delta\mathbf{x}^{(1)}}[d_N]$ is the directional derivative of d_N, and we have introduced the notation

$$\mathcal{W}(\mathbf{x}, \delta\mathbf{x}) = D_{\delta\mathbf{x}}[\mathcal{E}(\mathbf{x})]$$

for the weak formulation function. $\mathbf{d}_N = \mathbf{x}^{(1)} - \bar{\mathbf{x}}^{(2)}$ has, in general, a complex nonlinear relationship on $\mathbf{x}^{(1)}$. A derivation of the general form of the directional derivative of d_N is given in [Klo02], but for the special case of the second body being rigid it is simply given by

$$\delta d_N = \bar{\mathbf{n}}^{(2)} \cdot \delta\mathbf{x}^{(1)},$$

which is clear when considered geometrically.

Now, consider the relationship between λ_N and d_N as given by the KKT conditions. λ_N and d_N are plotted in Figure 6.4. There is no functional relationship between λ_N and d_N, since all non-zero values of d_N occur when $\lambda_N = 0$. The numerical methods which are employed in this thesis are all essentially based on the *regularisation* of this relationship, which is discussed in the next section.

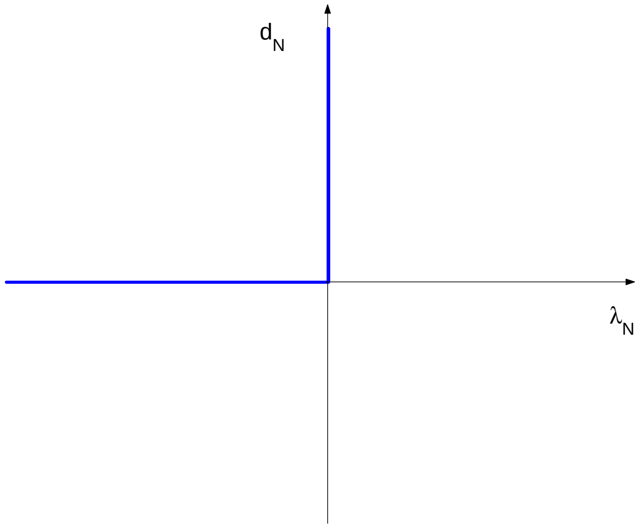

Figure 6.4: The relationship between the signed normal distance d_N between two bodies in contact and the Lagrange multiplier λ_N.

6.3 Numerical Algorithms for Contact

Numerical methods for solving optimisation problems with inequality constraints have been studied since the 1940s. One approach for developing a method for solving problems of the type (6.2) is to treat the Lagrange multipliers explicitly as unknowns to be computed. Such methods are sometimes known as mixed methods, since they involve solving for displacement and forces, as the Lagrange multipliers are tractions. One class of Lagrange multiplier methods are *active set methods*, in which the set of constraints is partitioned into two groups, those which are active (which means they attain their bound, i.e. the constraint $d_N \geq 0$ is active at a point $\mathbf{x}$ if $d_N(\mathbf{x}) = 0$), and those which are inactive ($d_N(\mathbf{x}) > 0$). Constraints which are inactive can essentially be ignored, reducing the optimisation-with-inequalities problem to an optimisation-with-equality-constraints problem. The challenge is to determine the correct set of active constraints in the solution, whilst minimising the energy functional. An example of an active set method is the gradient projection algorithm [NW99].

One drawback of methods that involve the explicit use of Lagrange multipliers is that they introduce a huge number of new variables which need to be computed. The number of iterations which need to be performed is significantly higher than with the basic Newton method for non-contact problems, so these types of solution method are just too time-consuming to be useful on large problems such as the one we are considering.

A second, more devastating, problem with the methods described above, in the context of *static* contact mechanics problems, is that they often require the initial guess to be *feasible*, which means that all the constraints must be satisfied in the initial guess. We wish to be able to simulate the equilibrium state of an object being compressed by plates, and the obvious initial guess we wish to be able to use is the initial shape of the object (the reference state), which clearly will not be feasible for non-trivial problems except in very specific scenarios[1], as if it was, it would very likely mean the solution to the constrained optimisation problem is simply that of the unconstrained optimisation problem. In general, it is very difficult to construct a deformed shape that is feasible for the initial Newton guess, without causing the body to become extremely badly-shaped, in the sense that the strains within the body are unnaturally high, leading to the initial residual being huge or essentially infinite. Note that this is a problem affecting static problems only, as in dynamic problems the solution at a previous time-step will always be a feasible and obtainable starting guess for the next time-step.

In order to avoid dealing with Lagrange multipliers explicitly, we will use a class of methods in which the original constrained problem is replaced by a sequence of unconstrained problems. In these methods, small violations of the constraints are allowed, but penalised. Instead of minimising the energy and dealing with the constraints, we add a term onto the energy which penalises the violation of the constraints and then minimise using the same numerical tools for unconstrained problems (i.e. in this thesis, the Newton method). This extra term will be a function of the gap function, and we will write it as an integral over the slave contact surface. The new unconstrained problem is then of the form

$$\min_{\mathbf{x}^{(1)}} \mathcal{E}(\mathbf{x}^{(1)}) + \int_{\Gamma^{(1)}} \Sigma(d_N; \mathbf{q})\, \mathrm{d}S, \tag{6.6}$$

where $\Sigma \equiv \Sigma(d_N) > 0$ is to be chosen. Σ will in general be dependent on a set of numerical parameters, which we denote by $\mathbf{q}$, and we express this parameter-dependence by using the semicolon in $\Sigma(d_N; \mathbf{q})$.

We need to choose Σ so that it measures violations of the constraint, and takes larger and larger values the worse the violation gets. Σ must also attain its minimum when no violation of constraints occurs. In this way, the minimisation of (6.6) discourages violation of the constraint, although it does not forbid it. The two methods of this form which we will consider are the *penalty function method* and the *Augmented Lagrangian method*, which we will describe shortly. First, let us consider how this method compares to the analytic formulation derived in Section 6.2. This will allow us to physically

---

[1] For example, such as: the unit cube, fixed at $z = 1$ with gravity applied downwards and a fixed plate at some $z = z_0 < 0$.

motivate the leap from the penalty function method to the Augmented Lagrangian method, and also allow us to numerically approximate physical quantities of interest using parameters from the numerical schemes.

Consider the weak form obtained by taking directional derivatives of the new functional

$$\mathcal{W}(\mathbf{x}^{(1)}, \delta\mathbf{x}^{(1)}) + \int_{\Gamma^{(1)}} \frac{\partial\Sigma(d_N; \mathbf{q})}{\partial d_N}\, \delta d_N \,\mathrm{d}S = 0.$$

We would like this weak form to approximate (6.5), which means we need λ_N to be equal to the derivative of Σ, or

$$t_N = \frac{\partial\Sigma(d_N; \mathbf{q})}{\partial d_N}.$$

We can use this to approximate the surface tractions. The quality of the numerical method can in some way be measured by the accuracy of this approximation. It now becomes clear why these methods are known as *regularisation methods*. As stated in Section 6.2, there is no one-to-one functional relationship between t_N $(= \lambda_N)$ and d_N. These methods effectively approximate the true non-one-to-one relation with a proper but artificial functional relationship, i.e. regularise the relationship.

6.3.1 The Penalty Method

The penalty method is the simplest method in this class, and is used commonly to solve optimisation problems with both equality or inequality constraints. An in-depth analysis of the penalty method can found in [NW99], and applications to mechanics problems can be found in [BKT02] and [KMH$^+$04].

Let us first introduce some notation: let $[.]_+$ denote the positive part of a number:

$$[x]_+ = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

Let $P \in \mathbb{R}$ be a large positive number. The function Σ for the penalty method is given by

$$\Sigma(d_N; P) = \frac{P}{2}[-d_N]_+^2,$$

so the full penalty method for solving an elasticity problem with contact is to solve the (finite element equivalent of the) following problem:

$$\min_{\mathbf{x}^{(1)}} \mathcal{E}(\mathbf{x}^{(1)}) + \int_{\Gamma^{(1)}} \frac{P}{2}[-d_N]_+^2 \,\mathrm{d}S.$$

P is known as the penalty parameter, and is used to force the penalty integral to be large if the constraint is violated. The larger the value of P, the greater the penalty integral, and the stronger violation of the constraint is penalised.

The derivative of Σ for the penalty method is easily computed; our approximation to the normal surface tractions is

$$t_N = \frac{\partial \Sigma}{\partial d_N} = P[-d_N]_+. \tag{6.7}$$

Although intuitively simple and relatively easy to implement, there are several disadvantages to the penalty method. The main problem is that the penalty method can never guarantee total fulfilment of the constraint. The penalty method has to be applied repeatedly with P tending to ∞ in order to force constraint violation down to within any given tolerance. However, as P increases, the Jacobian matrices will become increasingly ill-conditioned [NW99], leading to greater errors when the corresponding linear system is solved. Another problem with the penalty method is that the new energy being minimised is no longer twice differentiable, since the $[.]_+^2$ operator is not twice differentiable, and therefore the Newton method on the new energy loses its quadratic convergence.

6.3.2 The Augmented Lagrangian Method

The Augmented Lagrangian method, also known as *the method of multipliers*, is one of the most commonly used methods for solving optimisation problems with equality or inequality constraints. It has been used a great deal for contact mechanics problems, examples being [Che01] and [Klo02]. Analysis of the method can be found in [NW99]. This method is based on the observation that the approximation (6.7) of the surface tractions (or the equivalent approximation for Lagrange multipliers in a more general problem) given by the penalty method is not a good one and can be improved. Firstly, we define a set of *constants* $\lambda_N^{(0)}$, which are our approximations to the true Lagrange multipliers (at the surface nodes, for example). Also, let P be a penalty parameter as in the previous section. The idea is to solve a minimisation with constraints problem with a suitable function $\Sigma(d_N; \{P, \lambda_N^{(0)}\})$, and then, by using the solution to approximate the tractions/Lagrange multipliers, *update* the constants $\lambda_N^{(0)}$ and repeat the process.

In particular, the function Σ used for the Augmented Lagrangian method is

$$\Sigma(d_N; \{P, \lambda_N^{(k)}\}) = \frac{1}{2P}[-(\lambda_N^{(k)} + P d_N)]_+^2$$

It is easily seen that this function reduces to the penalty function if all the $\lambda_N^{(k)}$ are set to zero. Also, if we look at the energy if P is set to be zero (after multiplying out the bracket and dropping the $(\lambda_N^{(0)})^2$ term, which is just a constant and does not affect an energy-minimisation), we are left with the Lagrangian (6.3). This demonstrates how the Augmented Lagrangian bridges the gap between the Lagrange multiplier form and the penalty method.

The approximation to the surface tractions resulting from this choice of Σ is

$$t_N = -[-(\lambda_N^{(k)} + P d_N)]_+ \tag{6.8}$$

Now, suppose we are in the k-th iteration, with $\lambda_N^{(k)}$ the approximations of the surface tractions in that iteration. To move to the next iteration, we solve the minimisation problem with the corresponding augmented energy, and use (6.8) to obtain better estimates of the surface traction. In other words, we set

$$\lambda_N^{(k+1)} = -[-(\lambda_N^{(k)} + Pd_N)]_+$$

The Augmented Lagrangian Method, like the penalty method, involves the use of the positive-part operator $[.]_+$, and as such suffers from the same problems of non-differentiability and the resulting loss of convergence rate in the Newton method. Moreover, the Augmented Lagrangian method requires several iterations, each of which is an optimisation problem, and therefore can take considerably longer to run than the penalty method. (In practise though, it is not always necessary to solve each iteration to low tolerances, only the final one). However, this method fixes, to some extent, the two other problems with the penalty method. Firstly, P does not have to tend to infinity in order for us to obtain convergence to a solution satisfying all the constraints. It can be shown [NW99] that, under certain conditions, there exists some $P^\star \in \mathbb{R}$ such that any $P \in [P^\star, \infty)$ will suffice. Secondly, because relatively low values of P can be used, stiffness matrices arising in the Newton method are less likely to be ill-conditioned.

6.4 Results and Discussion

We now use the penalty and Augmented Lagrangian methods to simulate some contact problems, to compare their effectiveness and consider their failings. Since contact problems with exact solutions are extremely difficult to construct in three-dimensions with nonlinear elasticity, we will first investigate a simple one-dimensional linear contact problem, in order to study the accuracy of the two methods.

6.4.1 Comparison of the Penalty Method and the Augmented Lagrangian Method on a Simple Problem

Let us consider the following 1D constrained optimisation problem.

$$\min_y \int_0^1 \frac{1}{2} y'^2 \,\mathrm{d}x \qquad \text{subject to} \qquad \int_0^1 y \,\mathrm{d}x = A, \qquad y(x) \leq H(x), \qquad y(0) = y(1) = 0, \quad (6.9)$$

for some given area A and a given height function $H(x)$, an upper barrier. This is the problem of minimising the length of a linear string which is fixed at both ends, stretched over a constant area, and also compressed from above. The integral area condition is easily dealt with by introducing a

Lagrange multiplier $\bar{\lambda}$; the equivalent problem is to find the stationary point of

$$\int_0^1 \frac{1}{2}y'^2 + \bar{\lambda} y \,\mathrm{d}x \qquad \text{subject to} \qquad \int_0^1 y \,\mathrm{d}x = A, \qquad y(x) \leq H(x), \qquad y(0) = y(1) = 0. \quad (6.10)$$

This problem has been chosen because it is straightforward to calculate the analytic solution in the special case[2] of H = constant. Note first that in the absence of the compressive constraint, the solution is given by the parabola $y = -\frac{\bar{\lambda}}{2}x(1-x)$ (since the Euler-Lagrange equation would be $y'' = \bar{\lambda}$). $\bar{\lambda} = -12A$ from the area constraint, so $y = 6Ax(1-x)$. Consider the case of constant H. H must be greater than A for a solution to be possible, and H must be less than $\frac{3A}{2}$ for there to be some compression. In this case the solution must be of the form shown in Figure 6.5. There are three distinct regions: the first region is a parabola with second derivative $-\bar{\lambda}$ (since we know that the curve must satisfy the unconstrained Euler-Lagrange equations when not in contact with the barrier), the second region is where $y = H$, and the third region is a parabola symmetric to the first. We define α to be the value of x where region one ends and region two begins. By symmetry region two ends at $1 - \alpha$. The curve is completely determined by α and $\bar{\lambda}$. Once noted that the parabolas have zero gradient at α and $1 - \alpha$ (because the gradient must be continuous as we have a smooth constraint), it is elementary to compute the form of the parabolas in terms of $\bar{\lambda}$. Then the area integral condition can be used to eliminate $\bar{\lambda}$. The final relationship between the first contact point α, the area A and the height H is

$$\alpha = \frac{3}{2}\left(1 - \frac{A}{H}\right) \qquad 1 \leq \frac{H}{A} \leq \frac{3}{2},$$

with $\bar{\lambda}$ given by

$$\bar{\lambda} = -\frac{2H}{A^2}. \tag{6.11}$$

The finite element method with the penalty method and the Augmented Lagrangian method have been used to compute numerically an approximate solution to this problem. Unless stated otherwise, 100 gridpoints were used. All numerical simulations have been carried out with enough Newton iterations (or enough iterations of updating the Augmented Lagrangian method parameters $\lambda_N^{(k)}$) for the solution to have converged visually.

Verification that both methods give accurate results can be found in Figure 6.6 and Tables 6.1 and 6.2. Here, the two methods have been used to solve the minimisation problem with a constant height function. Since in this case we can calculate the analytic solution, we can compute the average error at the gridpoints, by taking the 2-norm of the difference between the true and computed solution (considering these as vectors), and scaling by the number of gridpoints. We can also compare the value of the analytic solution for $\bar{\lambda}$ given by (6.11) with that computed by the finite

---

[2] In fact, it is possible to write down algebraic equations determining the analytic solution for (at least) all height functions where the solution comes into contact with the compressing surface in one region only.

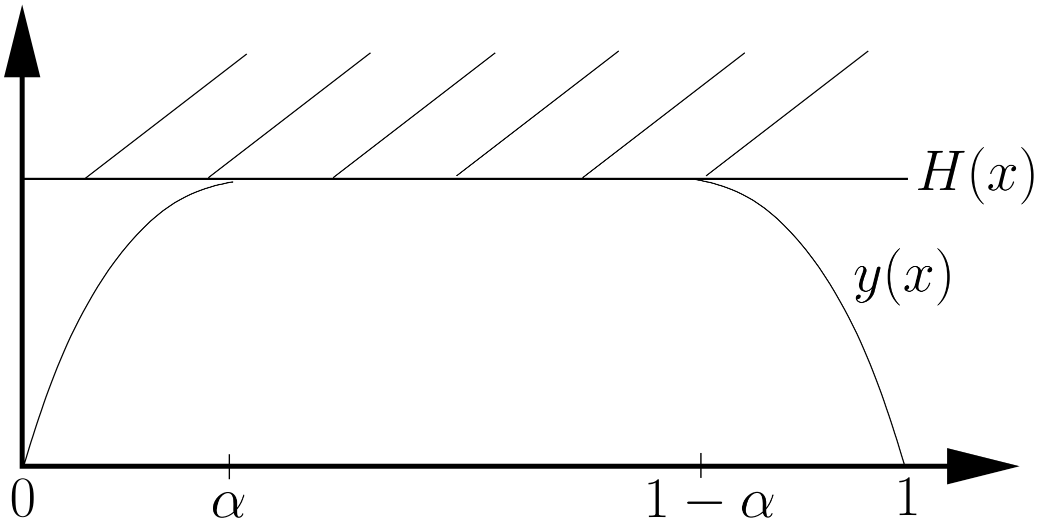

Figure 6.5: Form the solution of (6.10) must take, for constant height function H, with two parabolic regions around the central contact region.

element method. These results are displayed in Table 6.1. We see that while both methods provide accurate results, the Augmented Lagrangian Method is significantly better in terms of requiring a smaller value of the penalty parameter P. P needs to be several orders of magnitude greater in the penalty method for accuracy to be similar to that of the Augmented Lagrangian method. The error does not decrease to zero as the penalty parameter continues to increase however, and at this point the accuracy is limited by numerical error in the finite element method. Further improvements can be obtained by increasing the number of gridpoints, as displayed in Table 6.2, where the errors using the Augmented Lagrangian method with $P = 1000$ and various grids are given. As the number of gridpoints increases, the error decreases.

Figures 6.7 and 6.8 display similar results graphically with a varying height function, $H(x) = 1 + 0.1x + 0.2\sin(30x)$. Figures 6.7(a) and 6.7(b) show the final solutions obtained by the penalty method using 100 gridpoints and a range of values of P, and Figure 6.8(a) shows the same for the Augmented Lagrangian method. We see that the penalty method requires P to be of the order of 10^5 for the constraint $y \leq H$ to be satisfied visually, whereas $P = 10^3$ will do for the Augmented Lagrangian method. Figure 6.8(b) displays the computed solution y and the height function H together with the profile of the parameters λ_N computed in the Augmented Lagrangian method (with λ_N multiplied by 1/500 so as to be on the same scale). We stated that the λ_N approximate the true Lagrange multipliers which can be interpreted as compressive forces by the barrier, and the displayed profile of λ_N is consistent with this: λ_N is zero where there is no contact, and at its greatest in magnitude at the troughs of the H, i.e. where the barrier would be expected to be

P	$\|y_{\text{pen}} - y\|_2/100$	$\|y_{\text{aug}} - y\|_2/100$	$\|\bar{\lambda}_{\text{pen}} - \bar{\lambda}\|$	$\|\bar{\lambda}_{\text{aug}} - \bar{\lambda}\|$
10^1	0.014062	0.0014152	14.8561	2.2752
10^2	0.0087985	2.7151e-005	9.9727	0.056958
10^3	0.0019785	2.1934e-005	2.6041	0.044697
10^4	0.00022495	2.0449e-005	0.30441	0.035493
10^5	1.9143e-005	2.0028e-005	0.010959	0.031823
10^6	1.7952e-005	1.9461e-005	0.016566	0.022612
10^7	1.9293e-005	1.9492e-005	0.021995	0.023125
10^8	1.9444e-005	1.9461e-005	0.02255	0.022612
10^9	1.946e-005	1.9461e-005	0.022605	0.022612
10^{10}	1.9461e-005	1.9461e-005	0.022611	0.022612

Table 6.1: The differences between the true and computed solutions, and between the true and computed Lagrange multiplier associated with the area constraint, for the penalty and Augmented Lagrangian methods, for a range of penalty parameters.

number of gridpoints, N	$\|y_{\text{aug}} - y\|_2/N$	$\|\bar{\lambda}_{\text{aug}} - \bar{\lambda}\|$
4	0.088388	8110.7222
8	0.016443	8.9722
16	0.0020918	1.1003
32	0.00034248	0.25307
64	6.6099e-005	0.1062
128	1.2304e-005	0.030364
256	2.7325e-006	0.010678
512	9.4384e-007	0.0051769
1024	5.1178e-007	0.0034203

Table 6.2: The differences between the true and computed solutions, and between the true and computed Lagrange multiplier associated with the area constraint, with the Augmented Lagrangian method, for a range grid sizes.

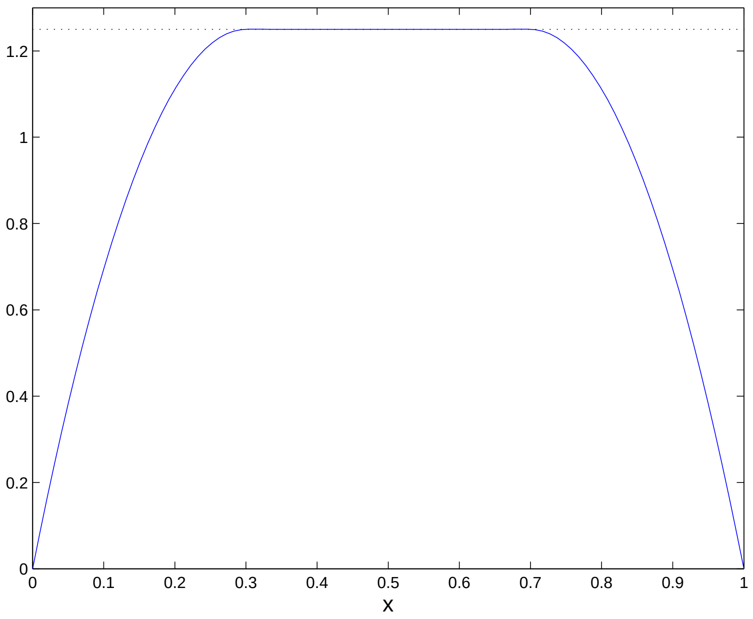

Figure 6.6: Numerical solution of (6.10) with $H = 1.25$ obtained by the Augmented Lagrangian method with $P = 100$.

pushing down the hardest.

6.4.2 Results with 3D Finite Elasticity and Contact

We now present some results of simulations on the full three-dimensional constrained finite elasticity problem. Since the problem is nonlinear, analytic solutions of model problems are not readily available, but we can still measure the extent to which the constraint is satisfied in the solution of the numerical scheme. Figure 6.9 displays the result of a simulation in which an isotropic hyperelastic cube is squashed by a plate to three-quarters of its original height. We have used the Augmented Lagrangian method with $P = 1000$. Clearly every point on the top surface of the cube should be in contact with the plate. Let E_P denote the average distance per node between the nodes on the top surface and the plate, divided by the initial length of the cube so as to be dimensionless. At the end of this simulation $E_{1000} = 4.4 \times 10^{-12}$, verification that the Augmented Lagrangian method can enforce the constraint to extremely high accuracy. For comparison, other values of the average error in the gap are: $E_{100} = 1.6 \times 10^{-10}$, $E_{10} = 5.6 \times 10^{-8}$ and $E_1 = 8.2 \times 10^{-7}$, each of which would be well within any practical tolerance. Figure 6.10 displays the λ_N profile for the numerical solution (in dimensionless units). The profile fits with the expected force distribution on the top surface of the cube. The force is smallest in magnitude at the corners, as here the cube has the most freedom to move. The force is greatest in magnitude in the centre, where the reverse is true.

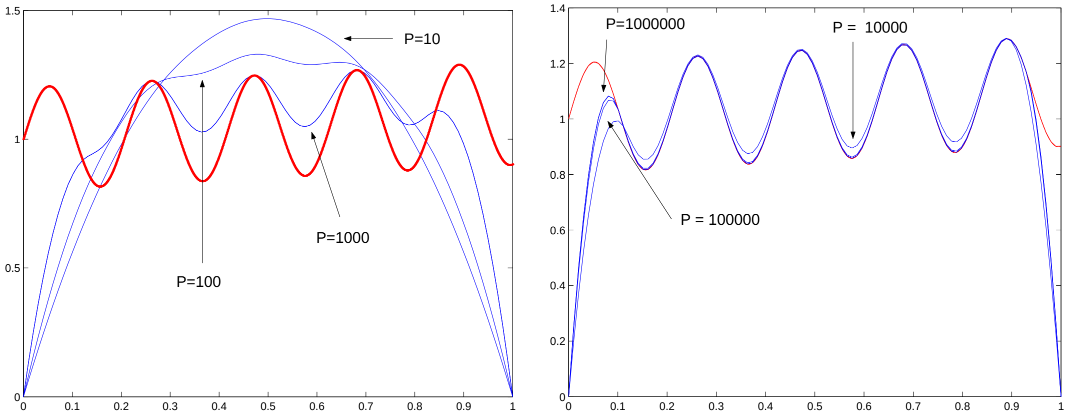

(a) (b)

Figure 6.7: (a) Height function $H(x)$ in bold red, together with solutions obtained using the penalty method with $P = 10, 100$ and 1000. (b) Solutions obtained using the penalty method with $P = 10000, 100000$ and 1000000.

If Figure 6.9 illustrates how the Augmented Lagrangian method can give incredibly accurate results, Figure 6.11 shows what can go wrong. Here, we attempt to simulate a cube half compressed by a plate. Since we are using linear interpolation in the elements, unless an edge of an element happens to line up with the edge of the plate, the geometry of the mesh does not allow the constraint to be enforced properly. Ideally the mesh should be highly refined in the region where contact with the plate ends. The average dimensionless error per node in this case is just 2.6×10^{-3}, and in fact, although the Newton method on this problem initially decreases the residual, it eventually stagnates.

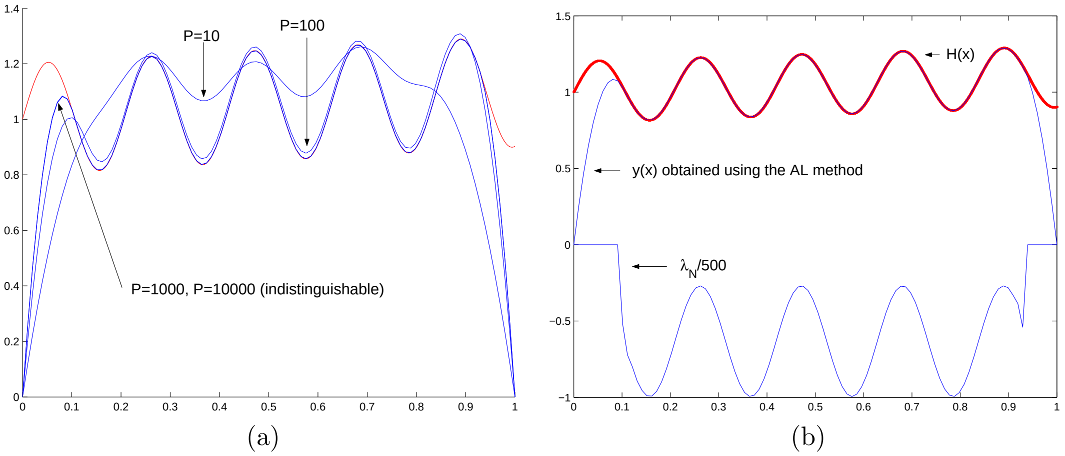

(a) (b)

Figure 6.8: (a) Height function $H(x)$ (red), together with solutions obtained using the Augmented Lagrangian method with $P = 10, 100, 1000$ and 10000. (b) Computed solution y, height function H and profile of the (scaled) Augmented Lagrangian parameters λ_N.

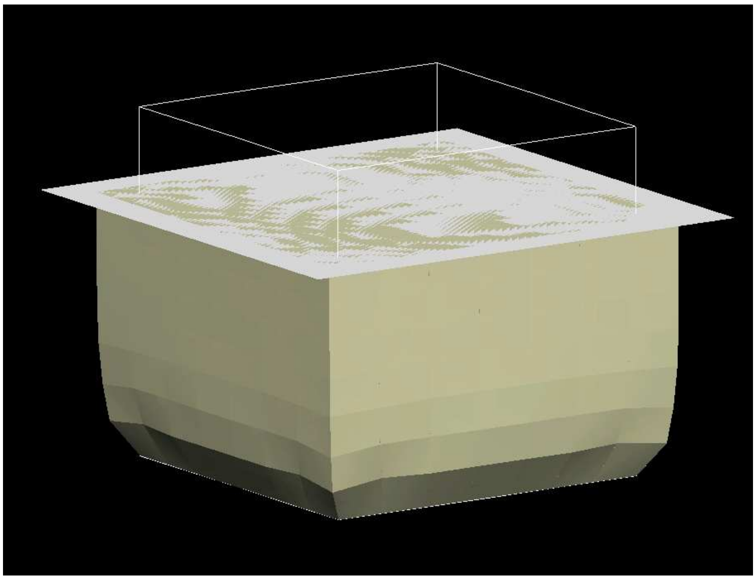

Figure 6.9: Simulation using the AL method, $P = 1000$, of an elastic cube compressed by a solid plate to three-quarters of its height. The white lines represent the initial shape. The compression plate is shown in a lighter colour than the cube.

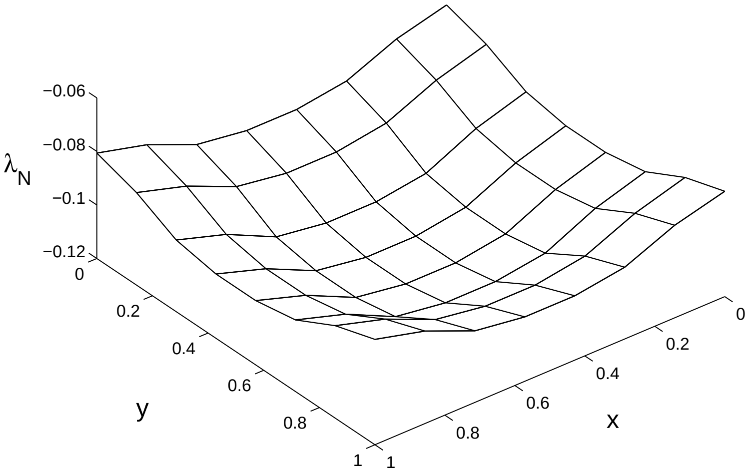

Figure 6.10: Profile of λ_N against X and Y for the fully compressed cube (units are dimensionless).

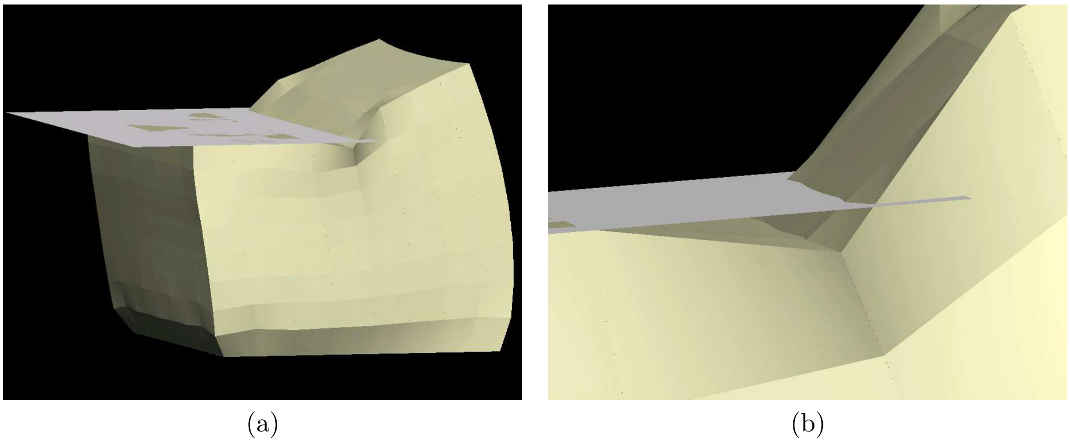
(a) (b)

Figure 6.11: (a) Simulation of a finite elastic cube half-compressed by a solid plate (lighter coloured), with (b) a close-up of the region where the plate ends.

Chapter 7

Implementation of the Model of the Breast and Numerical Experiments

In this chapter we use the numerical techniques developed in Chapters 4, 5 and 6 to model large deformations of the breast. The development of a deformable model of the breast from MR images initially involves a number of pre-computation stages, and we will begin by discussing them in turn. First, we consider the problem of *segmentation* of a set of MR images in Section 7.1.1. The result of the segmentation is information on the boundary of the breast and the distribution of tissues within the breast. The boundary data is then passed to the *mesh generation* stage, discussed in Section 7.1.2. Mesh generation of irregular 3D domains is a highly non-trivial task, especially in the case of meshes composed entirely of 3D hexahedral elements. However, the quality of the mesh can have long-reaching consequences on accuracy and computation time in finite element calculations. Elements which are not well-shaped can lead to slow and/or inaccurate numerical simulations, a phenomenon which we will discuss in order to make observations on modelling results later.

In Section 7.2.1 we state and analyse a number of assumptions which have been made, in particular debating the validity of heterogeneity, (transverse) isotropy and incompressibility, in the context of breast anatomy and physiology. We then discuss experimental results on mechanical properties of the breast tissues in question (Section 7.2.2), and consider which material laws it is appropriate to use in the model, and the degree of confidence (if any) we can have in experimental material parameters.

Once all these steps have been completed, we are able to perform large deformation simulations of the breast, and attempt to recreate the breast shape during the scenarios of surgery and mammography. This is carried out in Section 7.3, where we consider the prone-to-supine problem and CC and MLO compression. We then use the latter simulations to create simulated CC and MLO mammograms, and estimate the equivalent curves in the MLO image of various points in the CC image (Section 7.3.2). Finally, in Section 7.4, we use the model to carry out some numerical experiments, where we aim to identify important properties of the model which govern the deformation. Finally, we end in Section 7.5 with a discussion of the results in this chapter.

7.1 Pre-Computational Steps

7.1.1 Image Segmentation

Image segmentation is the common term used for a number of image processing operations, involving the labelling or numbering of constituent parts of an image. An image in general may be two or three-dimensional, with the term ‘voxel’ used as the three-dimensional equivalent of a pixel. We also use the term ‘pixel’ to denote both pixel or voxel, in the general n-dimensional context. For the application in mind, MR images can be considered as either a set of two-dimensional images, each of which needs to be segmented, or as a single three-dimensional image. The choice of image segmentation method dictates which interpretation is appropriate.

The simplest segmentation algorithm is the conversion of a monochrome image into a binary (black/white) image based on the value of intensity of individual pixels in the original image, where pixels are assigned a value of black or white based solely on whether their intensity value falls above or below a chosen threshold value. The resultant binary image then represents a partitioning of the monochrome image into two regions. This is the simplest example of a *thresholding* segmentation procedure. Thresholding is clearly highly susceptible to the effect of noise in the image, and more advanced segmentation procedures use data available at neighbouring pixels when assigning a region to a given pixel. Segmentation algorithms can be loosely grouped into two types. Region-based methods (for example region-growing or Markov Random Field methods, see [PXP00]) group pixels together using similarity properties of neighbouring pixels. Boundary detection methods (for example deformable models, see [PXP00]) use gradient information to attempt to separate regions of different properties. They can perform well on structures with well-defined boundaries (for example, segmentation of the brain cortical surface or heart ventricle wall), but are more susceptible to noise than region-based methods, since gradients are highly noise-sensitive. In recent years there has been a drive towards hybrid algorithms which take into account the complementary nature of both gradient and similarity data.

Image segmentation methods of medical images have a variety of applications, their use in providing heterogeneity data for finite element modelling being a minor one. More common applications include direct aid in diagnosis, information for image registration, tissue volume estimation, study of anatomical structure, treatment planning, and computer aided surgery. A major issue in the development of segmentation methods for medical images is the difficulty in validation of a method's accuracy. Ideally, pixel-wise data on the true distribution of tissues in the image, known as the 'gold standard', would be available, for comparison with the method's computed distribution. However, the gold standard is rarely available for images based on clinical data. For this reason, segmentation methods are often tested on phantom models with known structure, and great efforts have been put into the development of phantoms which mimic imaging properties of tissue. An alternative method of validation is to compare the automated result with a manually obtained segmentation, a method which suffers from the obvious drawback of comparison with an imperfect truth model. Statistical methods of verification provide another method of validation (see, for example, [ZWB$^+$04]). Another issue is that medical images tend to suffer from imaging artifacts such as noise, partial volume effects (the effect of different tissues on a single pixel), and motion. The various imaging modalities each have their own particular artifacts and idiosyncrasies, rendering reliable segmentation a challenging task. MR images tend to be especially noisy, are taken at relatively low resolutions, and suffer from blurring due to motion of the heart, lungs and whole body.

Segmentation of breast MR images is an important part of deformation modelling of the breast. Since, however, the aim of this thesis is to perform a mathematical study of the use of nonlinear elasticity and the finite element method in modelling breast deformation, segmentation is not considered an integral part of this thesis, and we will not devote too much attention to building or finding a segmentation algorithm which is accurate and reliable enough to be provably clinically viable. Still, in the ultimate aim of this line of research - clinically-useful software - MR segmentation would be a fundamental stage which needs to be performed automatically and robustly, so we will discuss the specific tasks involved with breast MR segmentation, look at their details and issues, and suggest some possible methods for performing these tasks.

Research into segmentation of medical images of the breast has been performed since the late 1970s [ROM$^+$05]. It has primarily been focused on segmentation of mammograms, e.g. [LKC$^+$95], [MJ95], as mammograms are generally considered the most effective imaging technique, but there has also been significant study into MR and ultrasound segmentation. Lesion segmentation, e.g. [LKC$^+$95], is the most common aim of segmentation of breast images, especially so for mammograms, but also for ultrasound and dynamic contrast-enhanced MR images. Automatic segmentation of a lesion or tumour enables the radiologist to study the size and shape of the mass and aids diagnosis. This application requires a segmentation algorithm with very high accuracy, with factors such as partial volume effects needing to be addressed. Powerful segmentation techniques for such problems have been proposed, which can segment lesions on sub-voxel levels robustly and accurately. Micro-calcification segmentation [MJ95] is another application which requires extremely high precision. Other applications include breast contour (skin-air boundary) segmentation of mammograms[1] ([ROM$^+$05], [WS03]) or MR images, for registration or as a first stage of automated diagnosis. Applications of healthy tissue segmentation of breast MR images include volume or density measurements, 3D visualisation of the breast-interior, and informing finite element models.

Figure 7.1(a) displays a sample MR image of the breast. Fibroglandular tissue appears dark gray, as does muscle. Fatty tissue appears white or light gray. The skin also appears dark gray, although generally the spatial resolution of MR images is too small to distinguish the skin/fat boundary. For our purposes we need to complete four tasks: locate the skin boundary, locate the pectoral muscle boundary, locate pixels corresponding to fatty tissue, and locate pixels corresponding to fibroglandular tissue.

Edge-Detection of Skin and Muscle in the Breast

Edge-detection, the detection of boundaries of regions characterised by sharp changes in intensity, is a common segmentation task. Edges in images can be of several different types, and there are a variety of different edge-detection methods suitable for locating different types of edge.

Edge-detection methods for locating the skin surface of breast MR images are discussed in a number of articles in the literature. Since air produces virtually zero MR signal, the background in MR images is virtually black, and direct thresholding of background from breast is a viable method for obtaining the breast edge. Thresholding has the disadvantages of being noise-sensitive, may require the user to input threshold parameters, and tends to produce jagged contours. Alternative methods used in the literature for improved breast contour segmentation include *active contour* fitting and *dynamic programming*[2] approaches. An active contour [KWT88] (also known as a snake) is a curve overlaid on the image which finds object boundaries by moving dynamically according to given rules, for instance, minimisation of some given energy. It is used in [WS03] to segment the breast contour in mammograms. In [Arb] and [Hay98] dynamic programming algorithms are used to obtain the breast edge in MR images.

In this thesis, direct thresholding is suitable for our purposes. (Actually, we use connected thresholding, discussed shortly). The small losses in accuracy and the jagged nature of the contour extracted by thresholding is insignificant compared to errors in fitting a mesh surface to the contour data (which we shall discuss in Section 7.1.2). Figure 7.1 displays the results of thresholding to

---

[1] Breast contour segmentation for mammograms is a non-trivial task, due to the tapered nature of the contour in the image, caused by the decrease in the volume of tissue that X-rays pass through near the breast edge.

[2] Dynamic programming is an optimisation approach based on recursion and dynamically filling in tables.

extract the breast surface.

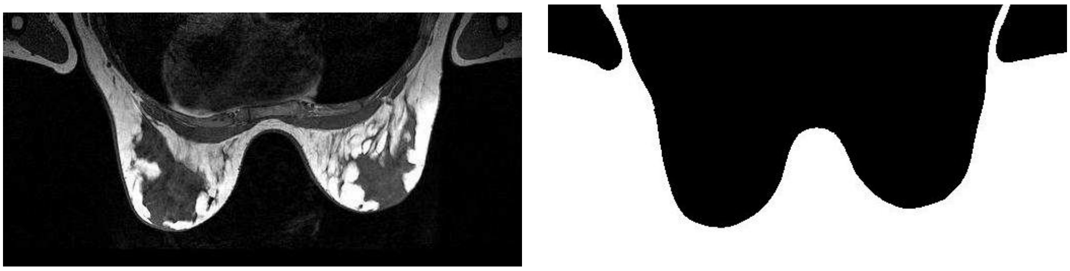

(a) (b)

Figure 7.1: Breast contour detection by direct thresholding: (a) initial MR image, (b) segmented image.

Automatic segmentation of the muscle surface is a much more challenging procedure. The image gradient at the muscle-breast boundary is not as high as that at the skin-air boundary, and this region of the MR image is strongly affected by motion of the heart and lungs. It is often difficult to manually segment the pectoral muscle in MR images, let alone automate this task. Consequently, methods discussed in the literature for chest wall segmentation of breast-MRI are less common and less successful. There are a number of articles looking at pectoral muscle segmentation in mammograms, in particular of MLO mammograms (for example [KCAR04], [ROM$^+$05]) where the muscle shows up relatively clearly as a large bright region (see the top corner of Figure 2.4(b)), but examples of segmentation of the muscle in MRI is rare. In [Arb], a similar dynamic programming approach to that used to segment the breast surface is used to segment the chest wall, but with very mixed results. A complicated dynamic programming method is used in [Hay98] for chest wall segmentation, with better results.

With the difficulty of this task in mind, and the relative lack of success in the literature, we have resorted simply to manually segmenting the muscle surface for the data used in this thesis. An example segmentation, which takes the form of a set of points on the muscle surface, is shown in Figure 7.2. Clearly, the difficulty in developing an accurate and reliable method for extracting the back of the breast from MR images is one of the major hurdles which needs to be overcome in finite element modelling of the breast.

Segmentation of Breast MR Images

Segmentation of breast MR images is also relatively uncommon, when compared to segmentation of mammograms, and segmentation of healthy tissue is again much rarer than that of lesions. Examples

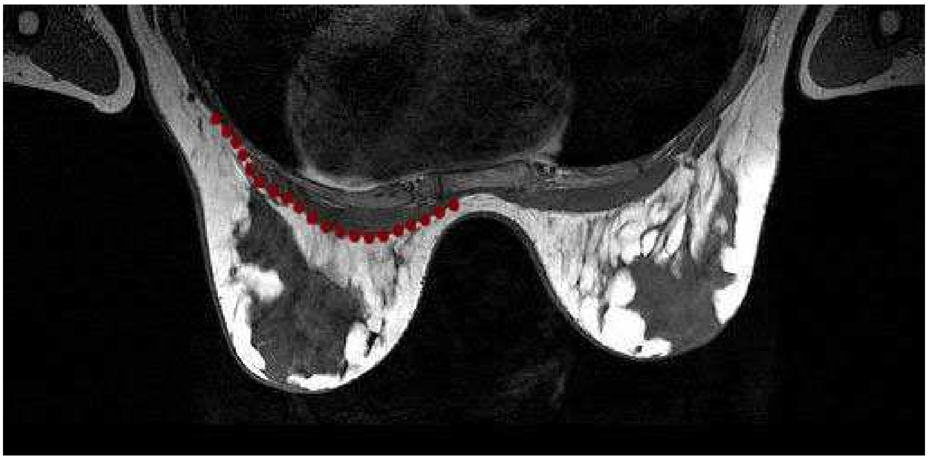

Figure 7.2: MR image with pectoral muscle surface manually segmented.

of breast MR segmentation include: [Arb], where a difference-image, obtained by subtracting a pre-contrast MR image from a post-contrast image, is segmented using two thresholding filters in order to extract a lesion; [KB04], where the whole breast is segmented into a number of regions using a Markov Random Field; and [LCSC02], where the fibroglandular tissue in the breast is segmented for visualisation.

The fibroglandular region has a complex 3D structure, rendering boundary-finding methods unsuitable and forcing the use of a region-based method. The fact that, despite its complexity, the fibroglandular tissue is necessarily one connected region (in 3D space) suggests the use of a three-dimensional *region-growing* method. These methods require the user to manually choose one or more *seed points*, defined to be voxels which are known to be within the region to be segmented. Starting from the seeds, neighbouring voxels are added to the region if they satisfy certain criteria. Region-growing methods differ in the exact criteria they use. As an example, a very simple region-growing method can be defined as the following. Given constant parameters v_1 and v_2, and a seed voxel whose intensity value v satisfies $v_1 \leq v \leq v_2$, any neighbouring voxel is included in the region if its value is also in the range $[v_1, v_2]$. This method is clearly a thresholding method that only segments one connected region, and is known as *connected thresholding.* More sophisticated methods use similarity information when assigning pixels. The *confidence connected* method [ISNC], which we have used, is an extension of such simple region-growing methods. This method goes through a user-defined number of iterations. In each iteration, the mean m and standard deviation σ of the intensity of the current segmented region is computed, and neighbouring voxels are added to the region if their intensity falls within the range $[m - C\sigma, m + C\sigma]$, for some user-defined constant C. When no more neighbouring voxels satisfy the criterion, an iteration ends and the new mean and standard deviation of the current region is computed for the next iteration. (In the first iteration, if there is only one seed point, the method automatically includes a small neighbourhood of the seed so the initial standard deviation is non-zero). The choice of the constant C and the number of iterations greatly affect the segmentation, and the user must choose values which perform well for the particular image being used.

Figure 7.3 displays some sample results using the confidence connected method to segment fibroglandular tissue, with semi-successful results.


(a) (b)

Figure 7.3: Segmentation with the confidence-connected image filter: (a) initial MR image, (b) segmented image.

7.1.2 Mesh Generation

Mesh generation is the process of breaking up the physical domain into a set of elements, each of which has a simple shape. Meshes are required for a variety of applications, the two major uses being graphics visualisation and numerical simulation of PDEs with the finite element method. Meshes can be n-dimensional in n-D space, they can be two-dimensional in three-dimensional space (surface meshes), or one-dimensional in two/three-dimensional space (curve meshes). Meshes can be composed of tetrahedral or hexahedral elements (in 3D, or equivalent shapes in lower dimensions), or can be a mixture of both types of element. The order of the basis functions used has a direct impact on the shape of the mesh too, since basis functions are used to interpolate the spatial variable over the element. For example, a mesh with second order basis functions is not necessarily made up of straight-edged elements, some elements may have curved edges. Such elements are good at approximating curved boundaries of spatial domains.

As well as the inherent difficulty in automatically meshing any irregularly-shaped domain, mesh generation is further complicated by the need for elements to be *well-shaped.* Well-shaped elements are those which have similarly-sized angles and lengths. Elements which are strongly skewed, or are (for example) long and thin, are therefore badly-shaped. Badly-shaped elements have been observed to have a detrimental effect on both computation time and solution accuracy when the finite element method is used. They can have the dual effects [She02] of decreasing interpolation accuracy (potentially leading to more Newton iterations to be required) and increasing the condition number (see below) of finite element Jacobians (leading to more iterations be required of the linear system solver). There is no single perfect measure of the quality of a single element or mesh. Some standard measures of element quality include: aspect ratio (ratio of length to height) or 3D variations; minimum/maximum angle/solid angle; and a variety of dimensionally-consistent ratios of functions of edge-length, face-area or volume[3]. One common measure for tetrahedra or triangles with a strong mathematical foundation is the condition number [Knu01], defined by $\kappa = \|J^{-1}\|_F \|J\|_F$, where J is the Jacobian of the map from the element to the canonical element, and the norm is the Frobenius matrix norm[4]. [Knu01] describes various possible element-quality measures, and forms a mathematical theory uniting them.

Mesh generation is a challenging task, studied by a great number of mathematicians and engineers around the world. As well as the automatic meshing of complex or irregularly-shaped (or sometimes regularly-shaped) domains, automatic mesh generation is important in dynamic problems involving moving phenomena such as shocks, for which the domain has to be re-meshed at each time-step, and in adaptive finite element methods, in which areas of low solution accuracy are identified automatically and the mesh refined at those locations.

Unfortunately, most of the research on the development of fully automatic mesh generators has been concerned with triangular or tetrahedral meshes, i.e of *triangulation* of physical domains. There are plenty of triangulation algorithms and freely-available triangulation packages available to the finite element modeller. However, algorithms for building hexahedral meshes are far less common, despite the fact that hexahedral meshes often display favourable numerical properties in computation. Since we are using trilinear basis functions to interpolate position and displacement, we are forced to use a hexahedral mesh, as was discussed in Section 4.2.1. Automatic generation of hexahedral/quadrilateral meshes is generally a much harder problem than triangulation, especially in three-dimensions. Hexahedral meshes can be classified into two categories, *structured* or *unstructured*. A structured mesh is one where the number of elements containing any particular node[5] is the same for each node. Structured meshes are easier to generate, but cause difficulties when refinement of the mesh in a particular localised region is required. Consider Figure 7.4(a), and suppose we wish to refine an element, say the bottom left one. The obvious approach would be

---

[3] By dimensionally consistent we mean, for example, length squared divided by area, or area$^{3/2}$ divided by volume.

[4] The Frobenius norm of a real matrix A is $\|A\|_F = \sqrt{\Sigma_{i,j} A_{ij}^2} = \sqrt{\mathrm{tr}(A^{\mathrm{T}}A)}$.

[5] This is known as the *valency* of the node.

to convert the element into four new elements as shown in Figure 7.4(b), in which case each new element would be well-shaped assuming the original element was. However, this introduces hanging nodes[6], and in order to remove them the mesh needs further refinement. If the mesh is to remain structured, the entire mesh has to be refined, resulting in the mesh shown in 7.4(c). A method for refining a single element where the mesh becomes unstructured but does not introduce hanging nodes is shown in 7.4(d). However, the new elements are less well-shaped than the original, and elements in subsequent refinements would be even more badly-shaped, so this is not a viable method of mesh-refinement. Generally, it is much easier to refine triangles or tetrahedra without creating badly-shaped elements or introducing hanging nodes. This ease of local refinement is one of the main advantages of tetrahedral meshes over hexahedral meshes.

As with MR image segmentation, we do not consider mesh generation to be an integral component of this thesis, and generation of optimal meshes for finite element modelling of the breast is considered a separate area of research. For this reason, we have resorted to using structured hexahedral meshes, despite their problems with local refinement. Although we do not attempt to build high quality meshes of the breast, we will describe some basic approaches to breast mesh generation, and outline the procedures taken to build the meshes used in this thesis.

Surface- versus Voxel-Based Approaches

In order to construct a mesh from a set of MR images of the breast, there are two possible methods that can be used, both of which have been used in the literature, *voxel-based* methods and *surface-based* methods.

In the voxel-based method (used for breast mesh generation in [SBRP99], [PBSS00]) one hexahedral element is created for each voxel in the 3D MR image that corresponds to breast tissue. This is illustrated in Figure 7.5(a). Such an approach leads to a structured mesh with each element identical in shape. If this approach would lead to a mesh with too many elements, the image resolution can be reduced prior to meshing.

Surface-based methods (used in [SBYP01] and [Rui03] for breast mesh generation) can take two forms. Either surface meshes are used to mesh the breast boundary, and the 3D mesh formed by creating interior 3D elements in the region enclosed by the surface meshes, or an initial 3D mesh is formed over some simple or regular geometry (e.g. cuboid/sphere/hemisphere), and the mesh warped so its surfaces fit the breast boundary. This is illustrated in Figure 7.5(b).

The advantage of the voxel-based approach is the extreme ease of meshing, and the fact that every element is well-shaped. The disadvantage is that the surface is inaccurately meshed in a 'jagged'

---

[6] Recall that a hanging node is a node contained in a face of an element, but is not a vertex, and that meshes with hanging nodes are inadmissable and cannot be used with continuous finite elements. This was discussed in Section 4.1.1.

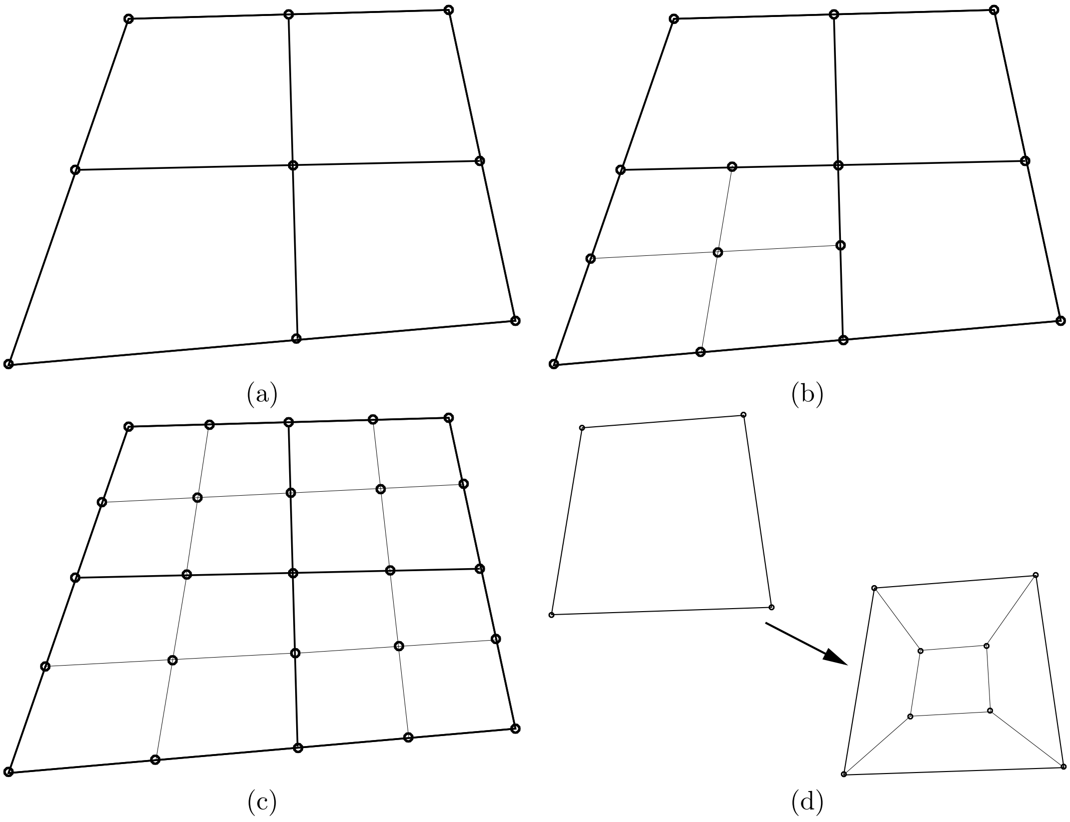
(a) (b) (c) (d)

Figure 7.4: The problem of refining a single element in a structured hexahedral mesh.

manner. Surface-based methods have the advantage of providing a more continuous surface, which would be expected to be favourable in numerical simulations. Continuous mesh surfaces are especially preferable in contact problems, as they are more likely to smoothly fit against the contact plate. However, hexahedral surface-based meshes introduce badly-shaped elements, and are significantly more difficult to construct.

In this chapter we wish to explicitly model the skin as a thin layer of elements, for which the surface-based approach would be more appropriate. More importantly, in Chapter 8 we will model the skin as a membrane, which will entail computing integrals over the skin surface. These integrals use the surface normal, which we wish to be as continuous as possible. This forces us to use the surface-based approach for mesh generation, as the surface normal in a voxel-based mesh alternates between perpendicular directions as it moves across the jagged surface. A voxel-based mesh would be expected to be completely unsuitable for membrane simulations.

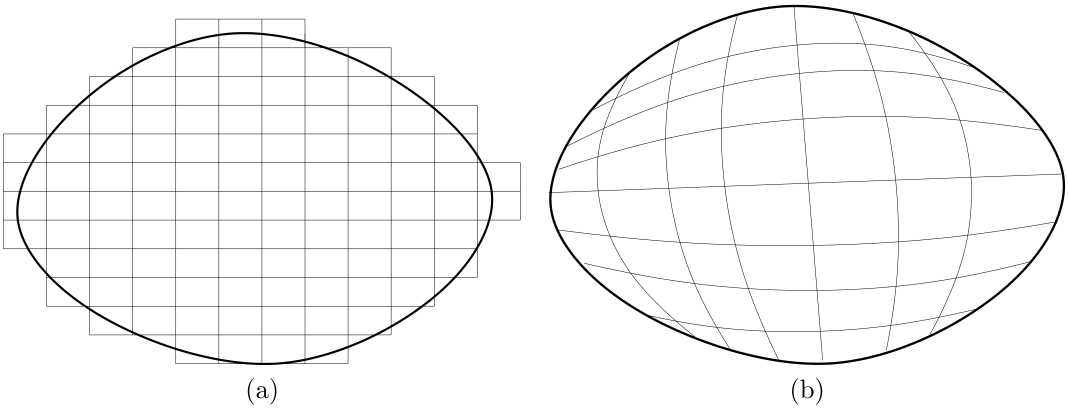
(a) (b)

Figure 7.5: (a) Voxel-based meshing. (b) Surface-based meshing.

Surface Fitting

We now suppose we have information on the location of the front skin surface and the back muscle surface of the breast from MR images, and need to construct 2D surface meshes of each of these surfaces. One approach is to take regular-shaped 2D meshes and fit nodal positions to the data, or alternatively take a regular 3D starting mesh and fit one surface to the skin surface and another to the muscle. We take the latter option. To do this we setup a simple rectangular mesh. We choose cubic-Hermite basis functions to interpolate position in this mesh. Cubic-Hermite basis functions, which were briefly introduced in Section 4.1.1, interpolate the variable (in this case spatial position) with a cubic polynomial that is a function of both the nodal values and the nodal derivatives of the variable. The surface fitting problem is to choose nodal positions and nodal position derivatives that minimise some chosen error between the resultant mesh and the data.

Figure 7.6(a) displays the skin and muscle datapoints which the mesh is to be fitted to, and the fitted cubic-Hermite mesh is shown in Figure 7.6(b). Once the mesh has been fitted it is refined and converted to a trilinear mesh for use in finite element calculations (Figure 7.6(c)). Interior elements are then assigned a tissue type based on the segmentation, with the element set to be the tissue which appears in the greatest number of voxels in the element. The fibroglandular elements are shown in Figure 7.6(d).

Note that because we have essentially fitted a regular mesh of a cuboid to the breast, with the front and back of the cuboid mapping to the skin and muscle and the sides of the cuboid mapping to the remaining (very thin) edges of the breast mesh. The final breast mesh therefore contains many badly-shaped elements, and will be inefficient for computation. Development of improved methods of meshing the breast, for example by warping a hemisphere onto the breast rather than a cuboid, is an important task, but is not pursued further in this thesis.

Generation of Skin Elements

In this chapter we model the skin by creating a layer of thin elements at the skin surface and assigning skin material properties to them, as illustrated in Figure 7.7. The skin elements are set to be constant width.

One obvious problem with this approach is that long thin elements are created, which are badly-shaped in the sense described in Section 7.1.2. The skin elements could be refined in the long direction so that they had a lower aspect ratio, but since a structured mesh is being used, the whole mesh would have to be refined to avoid hanging nodes. In Section 7.4 we will study the repercussions of modelling skin as a thin 3D elastic layer. In Chapter 8, we perform an asymptotic analysis of the equations of finite elasticity to obtain the limiting equations as the skin thickness tends to zero, in order to form a membrane model where the skin can be modelled as a 2D surface, removing the need for thin skin elements.

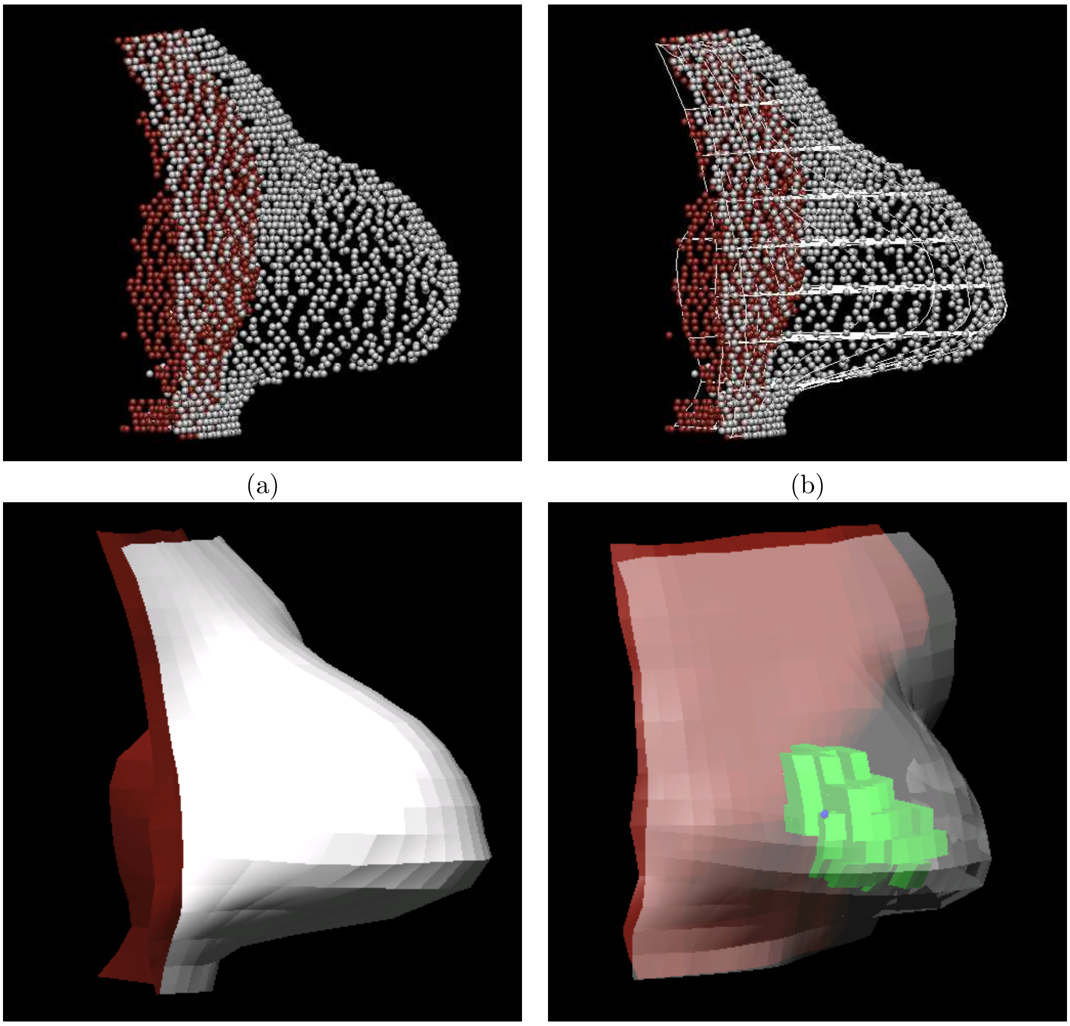

(a) (b)

(c) (d)

Figure 7.6: Mesh generation. (a) Skin (silver) and muscle (red) datapoints, (b) datapoints and a fitted cubic-Hermite mesh, (c) refined trilinear mesh (only the back surface (red) and the front surface (white) are shown), (d) refined mesh with translucent front surface so that the fibroglandular elements (green) can be seen. Fatty tissue elements are not shown. An artificial tumour is shown as the blue sphere on the edge of the fibroglandular region.

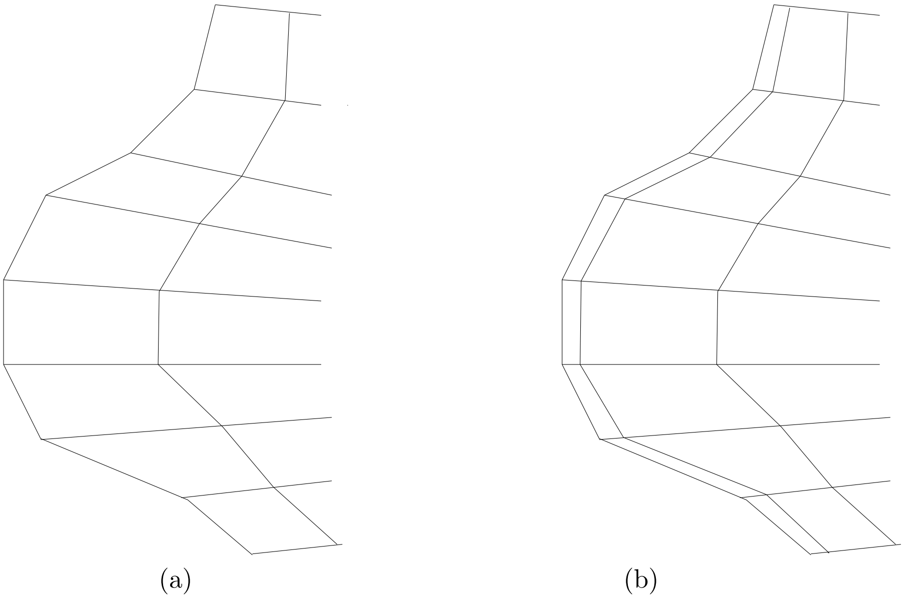

Figure 7.7: Adding thin skin elements: (a) diagram of a mesh before skin elements have been created, (b) mesh with thin skin elements.

7.2 Implementation of the Full Model

7.2.1 Modelling Assumptions

We now consider the modelling assumptions that have been made for the simulations in this chapter.

The Boundary and Boundary Conditions

We model the breast by taking the skin and the front surface of the pectoral muscle as the boundaries, and solve the equations of finite elasticity in this region. In all computations we naturally set zero-traction boundary conditions on the skin surface. We make the assumption that the shape of the pectoral muscle remains unchanged in deformations, and therefore set zero-displacement boundary conditions on this surface. The remaining surfaces of the mesh, which connect the front skin surface to the back muscle surface, are chosen semi-automatically. On these surfaces zero-displacement boundary conditions are also specified.

The assumption that the pectoral muscle does not change shape cannot be true in general. The pectoralis major is the major muscle in the chest wall and is involved in a wide range of shoulder movements. In principle, there will be little change in the muscle shape between the MR imaging scenario and the mammographic scenario, as in the former the patient's arms rest next her head, in the latter the patient's arm is lifted above her head to free space for the mammographic equipment. However, during surgery, the patient lies with their arms at their side, which necessarily involves a different shape of the pectoral muscle.

The choice of zero-displacement boundary conditions on the region connecting the pectoral muscle to the skin is motivated by the observation that skin further away from the centre of the breast is relatively unaffected by gravity or compression by plates and does not undergo deformation. If we assume that the pectoral muscle does not deform, it is sensible to assume that tissue in the relatively short region between the skin and the muscle also does not deform.

Heterogeneity

The female breast consists of parenchymal tissue, fibrous structures and fatty tissue, surrounded by skin and muscle. It is permeated by blood and lymph vessels, with the fatty tissue strengthened by Cooper's Ligaments. We model this heterogeneity with a model consisting of fatty tissue, fibroglandular tissue, skin, and potentially tumour tissue.

We do however assume each of the individual tissue types can be modelled as individually homogeneous. In reality, fatty tissue in the breast is heterogeneous due to the presence of Cooper's Ligaments, which form a 'honeycombed' structure throughout the fat. Also, subcutaneous fat is thicker than retromammary fat[7] but we neglect this difference and instead use average values for density and material parameters. Clearly fibroglandular tissue is not truly homogeneous (even at a semi-macroscopic level), consisting of two separate types of tissue, both of which are highly developed. Tumours are known to not be homogeneous but heterogeneities on such a small scale are unlikely to affect large-scale deformation.

When modelling the skin, we assume it to be completely homogeneous and isotropic. We also assume it is uniformly thick. Normal skin in the breast varies between 0.5mm and 2mm however, thickness changes with age [EdRR+89], and skin is known to be heterogeneous and anisotropic [Fun94].

We note that for some women the breast will also contain a silicon implant. Our model takes no account of this possibility. In practice, since breast cancer is more common later in life (and because breast cancer screening programmes are generally performed on the post-menopausal demographic), patients with implants are clinically uncommon in the current period of time. However, patients with implants will become more common in the future, and clinical software would ideally need to be able to handle this eventuality.

Hyperelasticity

We assume each of the breast tissues are hyperelastic, so that there exists a strain-energy function W whose derivative with respect to strain gives the 2nd Piola-Kirchoff stress. In doing so we are assuming that the constitutive relation does not depend on strain rate or history of strain. Most biological tissues display a viscous response to dynamic loading as well as an elastic response, however since we are looking at static problems and slow loading we assume the forces in the breast are dominated by the elastic response. Also, biological tissues have been shown to have elastic properties that are largely insensitive to strain rate to several orders of magnitude [Fun94]. Wellman shows in [WHDK99] that breast tissues can be modelled primarily as elastic, by comparing results of indentation tests at various strain rates.

Isotropy and Transverse Isotropy

There is no consensus in the literature on the anisotropy of healthy breast tissue. Breast tissue is often assumed to be isotropic, with anisotropies in the breast sometimes taken to be indicative of the presence of cancer, but there has also been some recent evidence that breast tissue is also inherently anisotropic [WDH+02].

In this thesis we make the assumption of isotropy for all breast tissues, although we will also consider a transversely isotropic material law for fat. The motivation for transverse isotropy is the

---

[7] The retromammary region is the area at the back of the breast, posterior to the parenchyma and anterior to the chest wall.

need to model the effect of Cooper's Ligaments on the deformation. The observation that it is Cooper's Ligaments losing their stiffness over a period of years that partially causes the breasts of older women to sag (together with the atrophy of fibroglandular tissue), shows that Cooper's Ligaments play a role in controlling the magnitude of breast-deformation. Unfortunately this observation also suggests the strength of Cooper's Ligaments is a highly age-dependent parameter.

In order to model breast tissues as transversely isotropic, a fibre direction is needed at each spatial position in the breast. Transverse isotropy means the material is isotropic in the plane normal to the fibre direction. However, such a direction may not necessarily exist physically at each point in the fatty tissue, because fatty tissue in the breast is not inherently a fibrous structure (in contrast to, for example, muscle). If we model the fatty tissue as transversely isotropic, we would be effectively aiming to average out the reinforcing effect of the Cooper's Ligaments throughout the tissue. A difficulty with using transversely isotropic material laws for fatty tissue is that the position and direction of Cooper's Ligaments are not visible in MR images and are therefore are not known *a priori.* (This can be compared to the known data in for example the finite element modelling of the heart performed in [Nas98], where highly accurate fibre directions of canine cardiac tissue were experimentally determined in order to inform the (anisotropic, in this case) material model). It is necessary for us to choose an appropriate fibre direction. For fatty tissue, the obvious choice is to set the fibres to be in posterior-to-anterior direction (the chest wall to skin direction). In the initial simulations in Section 7.3, we will take fatty tissue to be isotropic, and then consider the effect of using transverse isotropy in Section 7.4.

Skin is known to be anisotropic [Fun68], but would be expected to be well-approximated as transversely isotropic, with the 'fibre' direction in this case being the surface normal. We will use an isotropic material law for the skin, and assume that the increased stiffness in the normal direction can be neglected.

Incompressibility

In the deformation calculations carried out we will assume that all breast tissues are incompressible. It is shown in [Fun94] that biological tissue can be taken to be incompressible, since all soft tissues in the body are composed primarily of water. For this reason incompressibility is most likely a valid assumption for prone-to-supine deformations. However, volume changes may be observed during mammography, since the compression of the breast is quite substantial and could lead to a significant quantity of blood being forced from the breast, and may result in a loss of volume in the region of the breast between the plates. This may not necessarily mean the total volume of the breast decreases, as blood forced out from the region between the plates may not travel further than other parts of the breast that are still within the mesh, in which case the decrease in volume of the region being compressed will be countered by an increase in the rest of the breast. Although, if this were the case, it would not follow that incompressible material laws are suitable (as such laws enforce volume-conservation at every point, not just total volume conservation), an incompressible law would be consistent with total volume conservation, and be an approximation which averages out the dual effects of blood loss and blood gain.

Given the lack of experimental data about volume decreases in breast tissues under compression due to blood loss, and volume increases in breast tissue outside the compression region due to blood gain, we neglect the effect of blood loss and will therefore make the assumption of incompressibility everywhere. A possible area for further research in this project, given suitable experimental data, would be to alter the constraint $\det F = 1$ in such a way as to model the effect of blood movement.

Physical Assumptions

When solving the backward problem, we have assumed the breast shape obtained from the MR images is only acted upon by gravity. In particular, this means we have assumed that the breast is not in contact with the MR coils. This is a restrictive assumption: it will be true for some MR images of patients with smaller breasts but untrue for some patients with larger breasts. Modelling a scenario in which the breasts lie in contact with the coils would involve solving a backward contact problem, which would be significantly different to both the uncoupled backward and contact problems. The backward contact problem is not tackled in this thesis.

Mammographic Assumptions

In Chapter 6 we considered only the frictionless contact problem, and is this chapter we make the simplifying assumption that the effect of friction between skin and compression plate in mammography can be ignored. The magnitude of friction between skin and plate is not heavily reported in the literature, but friction might be expected to play some role in the deformation. Further work in this area of research could involve investigating the effect friction has on the deformation, and experimentally determining the coefficient of friction between skin and plate.

7.2.2 Mechanical Properties of Breast Tissue

Mechanical Properties of Internal Breast Tissues

The mechanical properties of breast tissue can only be determined by experiment. However, suitable experimental evidence is scarce. Although a number of studies have been made on breast tissue elastic properties, most experimental work is restricted to the linear regime of a simplified stress-strain law, or the assumption of a linear stress-strain relationship is explicitly made. Breast tissues *ex vivo* have been shown to be incompressible, so, under the assumption of isotropy, the problem of determining a linear material law for a given tissue reduces solely to determining the Young's modulus. The process of making these assumptions and following this line of experimentation is encouraged by the application of elastography, where tissue is characterised and imaged by its stiffness, and the theory that diagnosis can be performed on the basis of tissue stiffness measurements.

All healthy breast tissues have been shown to be nonlinear [SP04], [WHDK99]. The stiffness of breast tissue has also been shown to vary with individual and with age, and in the case of the parenchyma, over the menstrual cycle as well [LSBG03]. It is therefore only possible to determine average values of material parameters experimentally.

There are several issues hindering measurement of accurate and reliable data. Most importantly, tissue properties vary depending on whether the tissue is an excised sample or *in vivo*. *In vivo* measurements of Young's modulus is possible using elastography (e.g. [KTF$^+$98]), but *in vivo* measurements of parameters in more complex material laws is much more difficult and has not been carried out, to the best of our knowledge. Since excised tissue samples are small and must be maintained in living conditions, two and three-dimensional testing programs are difficult to perform, meaning virtually all experiments are limited to one-dimensional force application and one-dimensional stress/strain measurements. In general, experimental material laws in the literature for breast tissue are not especially reliable for use in large-deformation modelling. Precise stress-strain laws proposed for fatty and fibroglandular tissue in the literature vary hugely (see figures in [Rui03]). Even measurements of the Young's modulus of breast tissues can vary by orders of magnitude [TDS$^+$01]. Experimental Young's moduli for breast tissue can be found in [LRM$^+$99], [SGMO94] and [SB03].

One of the few nonlinear studies of the complete stress-strain relationship for breast tissue is performed by Wellman [WHDK99], with the aim of comparing the stiffnesses of cancerous and benign lumps. The tissues are assumed to be incompressible with an exponential relationship between nominal stress σ_n and nominal strain ϵ_n of the form

$$\sigma_n = \frac{b_n}{m_n}\left(e^{m_n \epsilon_n} - 1\right), \tag{7.1}$$

for a tissue labelled by the subscript n. Best-fit parameters b_n and m_n are measured for fat, fibroglandular tissue, and a variety of breast cancers.

In the finite element modelling of breast deformation in [Aza01], nonlinear elastic deformations are modelled in a pseudo-linear manner which enables the direct use of the exponential model (7.1) with Wellman's data. One major modification is made: in order to compensate for the effect of Cooper's Ligaments, which compartmentalise fatty tissue and restrict compression, it is proposed that there exists a limiting value of strain, taken to be 15%, at which fat stiffens and becomes similar to fibroglandular tissue.

	C_{10}	C_{01}	C_{20}	C_{02}	C_{11}
Fat	310	300	2250	3800	4720
Fibroglandular	330	280	4490	7720	9450

Table 7.1: Experimental material parameters (Pa) for the polynomial hyperelastic law (7.2) given in [SP04].

Wellman's data is also used in the model of the breast in [SBYP01], where a Neo-Hookean hyperelastic law $W = c_1(I_1 - 3)$ is assumed, and c_1 chosen to best fit the data for each tissue.

There are only a couple of articles in the literature, to the best of our knowledge, which treat breast tissue as hyperelastic. In [SP04] an incompressible polynomial hyperelastic material law of the form

$$W(I_1, I_2, I_3) = \sum_{0<i+j\leq 2} C_{ij}(I_1 - 3)^i (I_2 - 3)^j - \frac{p}{2}(I_3 - 1) \tag{7.2}$$

is assumed for breast tissues. This law has often been used in the modelling of rubber. No motivation is given for why this choice may be suitable for breast tissue. The law contains 5 material parameters, C_{01}, C_{10}, C_{11}, C_{20} and C_{02}, to be experimentally determined. Indentation tests were performed on recently-excised samples of fatty and fibroglandular breast tissue, with the compressive force f_{meas} measured as a function of the vertical strain E_{33}. Using a nonlinear elastic finite element model with strain energy given by (7.2), an inverse problem was solved for the 5 best material parameters which minimised the difference between the computed forces and the measured forces f_{meas}. Their results show good correspondence between the numerically and experimentally computed relationships between f_{meas} and E_{33}, suggesting that the polynomial hyperelastic truly was a suitable choice of material law. Also, it is claimed that the material parameters agree in the linear approximation to experimentally determined Young's moduli. Comparison with other functional forms are not made. The material parameters obtained are given in Table 7.1.

Experiments involving indentation tests on recently-excised samples of glandular tissue and infiltrating ductal carcinoma are described in [KJL$^+$05]. Here, a hyperelastic *Arruba-Boyce* material law (a 5th-degree polynomial in I_1) is assumed, but modified to account for residual stress due to tissue hydration and resistance to hydrostatic deformations. Material parameters are again computed by solving the inverse problem using the finite element method to simulate the indentation tests.

Finally, we note that, as far as we are aware, no work is available in the literature detailing experiments on breast tissues with a transversely isotropic or an anisotropic hyperelastic material model.

Mechanical Properties of Skin

The mechanical properties of skin have been studied much more extensively than internal breast tissues. Skin has been shown to be anisotropic [Fun68], heterogeneous, nonlinear [Ver70] and viscoelastic [AMLdR80], whose properties vary with individual, age, location on the body and other parameters [EdRR$^+$89].

Skin is subject to tension *in vivo*, and contracts upon excision [Rot54]. Excised cat skin used in experiments in [Ver70] was estimated to be stretched by a factor $\lambda \in [1.18, 1.29]$ *in vivo*. Such results show that there is no true connected reference state for the breast, as even in the absence of gravity there still exist tensile forces acting within the skin. These stresses, which act within an initial 'undeformed' body are known as *residual stresses*, and are caused by growth. In this chapter we make the assumption that residual stresses in the skin can be neglected, but we will discuss this issue further in Chapter 9.

Human skin under tension has been shown to be *biphasic*, with an initial linear extension stage which corresponds to collagen fibres straightening, followed by a stiffening stage corresponding to collagen extending [RV66]. Azar [Aza01] approximates this using a piecewise linear constitutive law when modelling breast deformation.

A hyperelastic material law for skin is proposed in [Ver70], where uniaxial experiments on small samples of excised cat skin led to both compressible and incompressible forms of an exponential hyperelastic law being suggested, the incompressible version being:

$$W = a(e^{b(I_1-3)} - 1) + c(I_2 - 3) - \frac{p}{2}(I_3 - 1) \tag{7.3}$$

with $a = 92.39\text{Pa}$, $b = 4.4$ and $c = -203.40\text{Pa}$.

7.3 Simulations using The Full Model

We can now use the model to simulate the clinical scenarios discussed in Chapter 2. We use the mesh built in Section 7.1.2 and shown in Figure 7.6, assume isotropic fibroglandular and fatty tissue material laws of the form (7.2) using material parameters estimated in [SP04] (Table 7.1), and the skin material law given in (7.3). This mesh has 4913 nodes and 4096 elements. Since the average element size is much larger than that of a typical undiagnosed tumour, and as the MR image set was taken of a healthy volunteer and contained no tumour, all elements are set to be either skin, fat or fibroglandular; no element was set to have a tumour material law. However, so that we have a landmark which we can measure the displacement of, we artificially choose a single node on the edge of a fibroglandular element to be a tumour node. This is shown in Figure 7.8(a), where the initial, gravity-loaded, MR breast in displayed. The effect the presence of a stiff tumour has on the deformation will be considered later.

7.3.1 Computing the Reference State and the Prone-to-Supine Problem

The first stage is to compute the reference state, by solving the backward problem using the techniques introduced in Chapter 5. The result of this simulation is shown in Figure 7.8(b), which we see is a relatively small deformation in this case. Without having clinical data to compare the results with, it is difficult to say whether this magnitude of deformation is realistic or not. On one hand, it is perhaps smaller than would be expected, seeing as it is clear from Figure 7.6 that the majority of this breast is fat, and reinforcement due to Cooper's Ligaments has not been modelled as yet. This would possibly suggest that the material parameters in Table 7.1 are too large. On the other hand, the initial prone mesh is one which does not appear to have been highly deformed by gravity, which would mean the size of the deformation is as expected.

Figure 7.8(c) displays the result of the forward problem using the computed reference state as the starting state, to simulate supine breast shape. This deformation is also a relatively small one. Again it is difficult to judge without experimental data whether this is due to the geometry of the particular breast modelled, or if material parameters are too stiff, or if prone-to-supine deformations are generally small. In Section 7.4, we will study the effect changes to material parameters have on the magnitude of the prone-to-supine deformation, in order to quantify the sensitivity of the model to constitutive laws.

Table 7.2 compares the total displacement (from the prone to supine position) at the nipple and tumour, with two other simulations. In the first (row 2 in Table 7.2), we solve a forward problem using the prone state as the starting state (i.e. we do not solve the backward problem prior to the forward problem). In the second (row 3), we do the same but double the magnitude of $\mathbf{g}$. In linear elasticity this would effectively solve the backward and forward problems in one computation. The results demonstrate the need for the backward computation, even for this small deformation. The displacements in the first forward-only simulation are about half of their true values, as would be expected. The displacements in the second forward-only simulation are closer to the true values, but still about 15% off. For women whose breasts deform more under gravity, the discrepancy would likely be even greater.

The computations times for the backward and forward problems are 60 minutes and 35 minutes respectively, on a 2Ghz Linux PC[8]. The mesh is relatively coarse but not excessively so, so these times demonstrate that simulating deformations of the breast with nonlinear elasticity is computationally tractable, and solutions can be computed in clinically-viable lengths of time.

---

[8] Note that this is a relatively slow desktop PC.

Method	Tumour disp.	Nipple disp.	Total mean disp.
Prone to Reference State to Supine	3.03	8.78	1.14
Prone to Supine directly, $\mathbf{g} = (0, 0, 9.81)$	1.54	4.45	0.57
Prone to Supine directly, $\mathbf{g} = (0, 0, 19.62)$	2.60	7.11	0.95

Table 7.2: Displacements (mm) in various prone-to-supine simulations.

7.3.2 CC and MLO Simulation and Mammographic Matching

Next, we simulate mammographic compression. The Augmented Lagrangian method as described in Chapter 6 is used to compress the breast in the CC and MLO directions, the results of which are shown in Figure 7.9. Figure 7.9(a) displays the initial mesh with the position of the CC plates, revealing the extent of the deformation that will be necessary to satisfy the constraint. Figures 7.9(b) and 7.9(c) are two views of the CC compressed breast, and Figures 7.9(d) and 7.9(e) are two views of the MLO compressed breast. To obtain an idea of the magnitude of the strains involved, Figure 7.9(f) displays a histogram of the largest eigenvalue (in modulus) of the strain E_{MN} sampled at the centre of each element. The histogram shows that the strains go as low as -0.5, representing compression to half the length, and as high as 1.2, representing stretching to over double the length.

Figure 7.9(a) shows that the initial state is penetrated deeply by the plates. Not surprisingly, this means the Augmented Lagrangian method will not converge if full compression is applied from the first iteration. Instead, it is vital to apply the compression incrementally (as we described in Section 4.2.2), computing a few Newton iterations for each increment. Even with increments, the method will still not converge, in the sense that the norm of the Newton residual vector does not converge to zero. Instead, the method decreases the residual norm by several orders of magnitude, after which further iterations have little effect on the residual (unlike the forward and backward problems, where the Newton residual can be reduced to any desired tolerance). This is a common phenomenon with optimisation algorithms, which often have a *domain of convergence* rather than a limit point[9]. Once the solution enters the domain of convergence, further Newton steps just move the solution by small amounts within the domain of convergence, without converging to any point.

Computing the solutions of these contact problems takes 114 minutes for the CC simulation and 123 minutes for the MLO simulation, which is longer than the forward and backward problems but still an acceptable length of time (especially given the machine used).

---

[9] By domain of convergence and limit point, we mean respectively a domain and a point in the solution space $\mathbb{R}^P$, where $P \gg 1$ is the number of unknowns, not a domain or a point in physical space.

Simulation of Mammograms

We can use the compression results to simulate CC and MLO mammograms. In order to produce a (CC) mammogram of the compressed breast, we have to carry out the following algorithm:

- Choose a mammogram image resolution RES and a sampling resolution $ZRES$.
- For each of the RES^2 pixels in the mammogram, given by coordinates (x^i, y^i), $i = 1, \ldots, RES^2$:
    - For each point in the column $\{(x^i, y^i, z^k) \, : \, k = 1, \ldots, ZRES\}$:
        - Determine whether the point corresponds to: fat, skin, fibroglandular tissue, or air.
    - Assign a colour to the pixel based on the proportion of fat, skin, fibroglandular tissue and air in that column.

To determine whether a point corresponds to fat, skin, fibroglandular tissue or air, we need to be able to compute which element (if any) contains any given point. One method of determining if a given point is contained in an element is to compute the $\boldsymbol{\xi}$ value for that point in that element. Recall that for $\boldsymbol{\xi} \in [0,1]^3$ the value of $\mathbf{x}$ corresponding to $\boldsymbol{\xi}$ is given by

$$\mathbf{x}(\boldsymbol{\xi}) = \sum_{i=1}^{8} N_i(\boldsymbol{\xi}) \mathbf{x}_i,$$

where N_i is the i-th trilinear basis function on the cube (4.3), and $\mathbf{x}_i$ are the nodes of the element. This relationship can potentially be inverted, which provides a test for whether a point $\mathbf{x}^*$ is contained in the element: we need to compute $\boldsymbol{\xi}(\mathbf{x}^*)$ and see if it is contained in the unit cube.

However this method involves inverting a nonlinear relationship between $\mathbf{x}$ and $\boldsymbol{\xi}$ (because the 'linear' basis functions for the cube are trilinear not linear), and therefore requires a nonlinear method such as Newton's method. To avoid this, and therefore avoiding issues such as handling possible divergence and choosing suitable stopping criteria, we instead convert the mesh into a tetrahedral mesh, by converting each hexahedral element into 6 tetrahedrons. In linear tetrahedral elements

$$\mathbf{x}(\boldsymbol{\xi}) = \sum_{i=1}^{4} N_i^{\mathrm{TET}}(\boldsymbol{\xi}) \, \mathbf{x}_i^{\mathrm{TET}},$$

where the four tetrahedral basis functions are $N_1^{\mathrm{TET}}(\boldsymbol{\xi}) = N_1^{\mathrm{TET}}(\xi, \eta, \nu) = 1 - \xi - \eta - \nu$, $N_2^{\mathrm{TET}}(\boldsymbol{\xi}) = \xi$, $N_3^{\mathrm{TET}}(\boldsymbol{\xi}) = \eta$ and $N_4^{\mathrm{TET}}(\boldsymbol{\xi}) = \nu$. This is a linear function of $\boldsymbol{\xi}$, so can be to inverted by inverting a single 3 by 3 matrix. An element can be tested to see if it contains a given point $\mathbf{x}^*$ by testing whether the computed value of $\boldsymbol{\xi}(\mathbf{x}^*)$ lies in the canonical tetrahedron. Only one matrix inversion is required with a tetrahedral mesh, at the cost of having 6 times more elements to search through.

This method has been implemented to simulate a CC mammogram and an MLO mammogram. We use a pixel resolution of $RES = 100$ and a sampling resolution $ZRES = 200$. Since this means that the element containing a given point has to be identified for $(ZRES)(RES)^2 = 200000$ different points, computation can take several hours, even with several obvious tricks[10] implemented, which is, in fact, longer than the time taken to compute the deformation. With finer meshes computation would take much longer. Identification of a more efficient algorithm to determine the element containing a given point is a computer science problem which would have a useful application here.

Figure 7.10(a) is the simulated CC mammogram obtained from the CC compressed breast, and Figure 7.10(b) is a simulated MLO mammogram. Since the view is aligned with the mesh in the CC mammogram, the edges of the fibroglandular elements are clearly visible, although blurred by the effect of projecting the data onto the image. (Recall though that accuracy in these simulations of mammograms is not one of our objectives). In the MLO mammogram the view is at approximately 45 degrees to the element edges, so they are not as visible.

We can use the simulated mammograms and the two deformations to solve the mammographic matching problem described in Section 2.3.3. The aim is to determine the curve in the MLO image which is equivalent to a given pixel in the CC image. To do so, we have to sample a number of points in the column in the CC breast corresponding to that pixel, determining for each point the element K containing it and the corresponding value of $\boldsymbol{\xi}$. The position of each point in the MLO compressed breast can then be trivially computed using $(K, \boldsymbol{\xi})$ and the MLO compressed breast nodal values, and projected onto the MLO mammogram. Figure 7.10(a) shows the position of the artificial tumour in the CC mammogram, and Figure 7.10(b) shows the curve in the MLO mammogram which maps to that point. For comparison, the curve in the MLO image which is the same distance from the nipple as the tumour is in the CC image is also plotted. This method is sometimes used by radiologists when analysing mammograms from different views; we see that it is very different to the true equivalent curve. Figure 7.11 displays some similar results with a number of different pixels in the CC image.

Let us return to the curve in Figure 7.10(b) which corresponds to the tumour pixel. A better idea of the distortion of this line can be found in Figure 7.12. Figure 7.12(a) displays the CC-compressed breast and the column (which includes the tumour) which is projected onto the tumour pixel. Using the $(K, \boldsymbol{\xi})$ data for each point in the column, we can compute the equivalent curve in the reference state, i.e. the uncompressed column. This can be seen in Figure 7.12(b), which shows that, although the deformation of the whole breast is very large, the deformation of the column is not very different from simple stretching, especially in the centre of the column. Figure 7.12(c) shows the curve in the MLO-compressed breast, where it has been rotated and re-compressed. Again, the compression does not change the shape of the curve a great deal. This is the reason why the final projected curve

---

[10] For example: if a pixel is found to correspond to a column of air and set to be black, all pixels in the image below that pixel can immediately be set to black as well; a good first guess for the element containing a point is the element which contained the previous point; other good guesses are the last element's neighbours (which have to be computed in a pre-computation stage).

in Figure 7.10(b) is quite close to a straight line.

The curves in Figures 7.12 and 7.11 suffer from our choice of linear basis functions for interpolating position, having a jagged form (because a straight line in the CC mesh will map to a jagged line in the MLO mesh). Since there are more elements nearer the boundary of the mesh, the jaggedness increases towards the edges of these curves. Higher order basis functions and/or finer meshes would result in smoother curves. Still, curves such as these could give very useful information when attempting to analyse mammograms from different views.

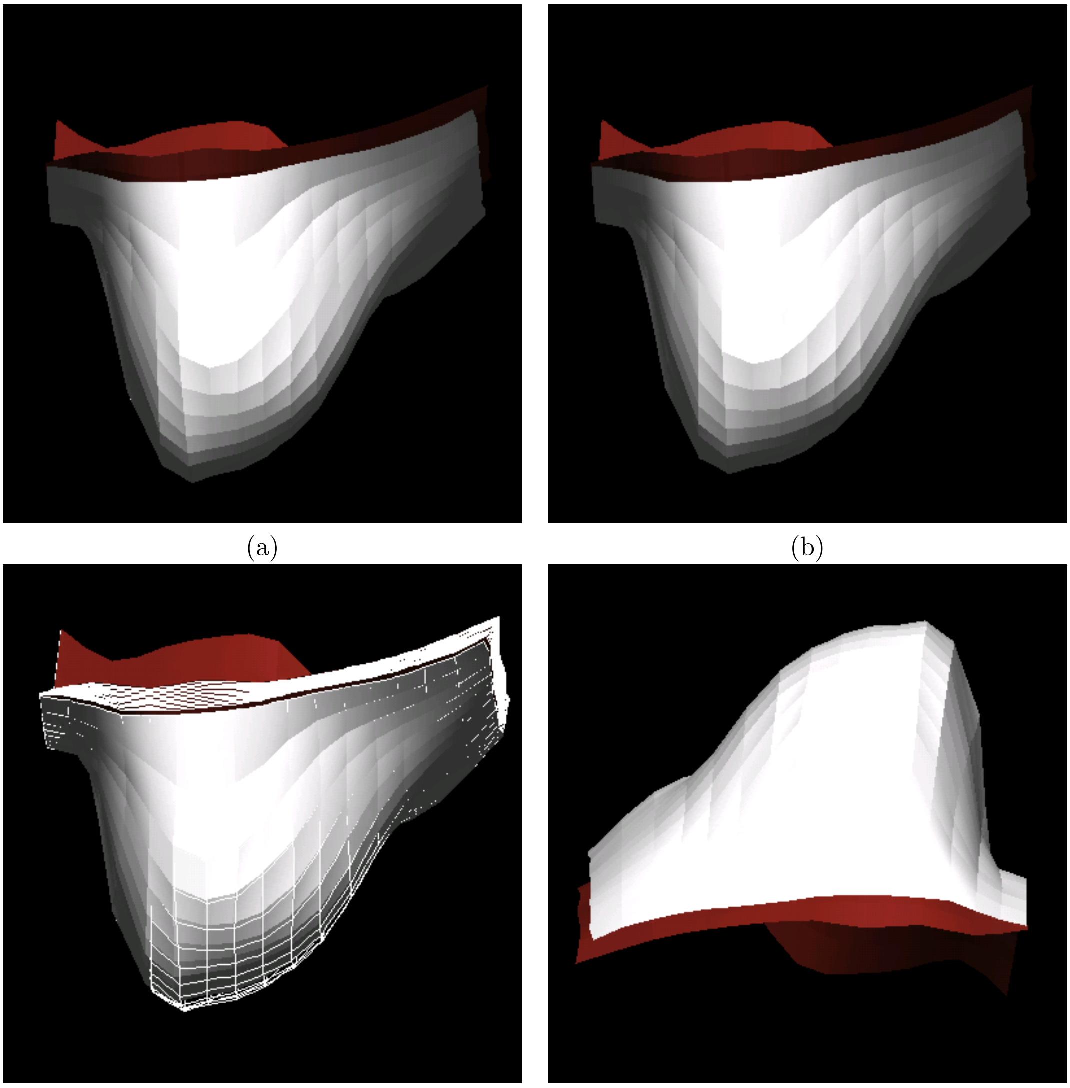

(a)

(b)

(c)

(d)

Figure 7.8: Solving the Prone-to-Supine problem: (a) initial mesh in prone MRI gravity-loaded state, (b) computed reference state, (c) prone (wire-frame) and reference (surface) states overlaid, (d) supine state.

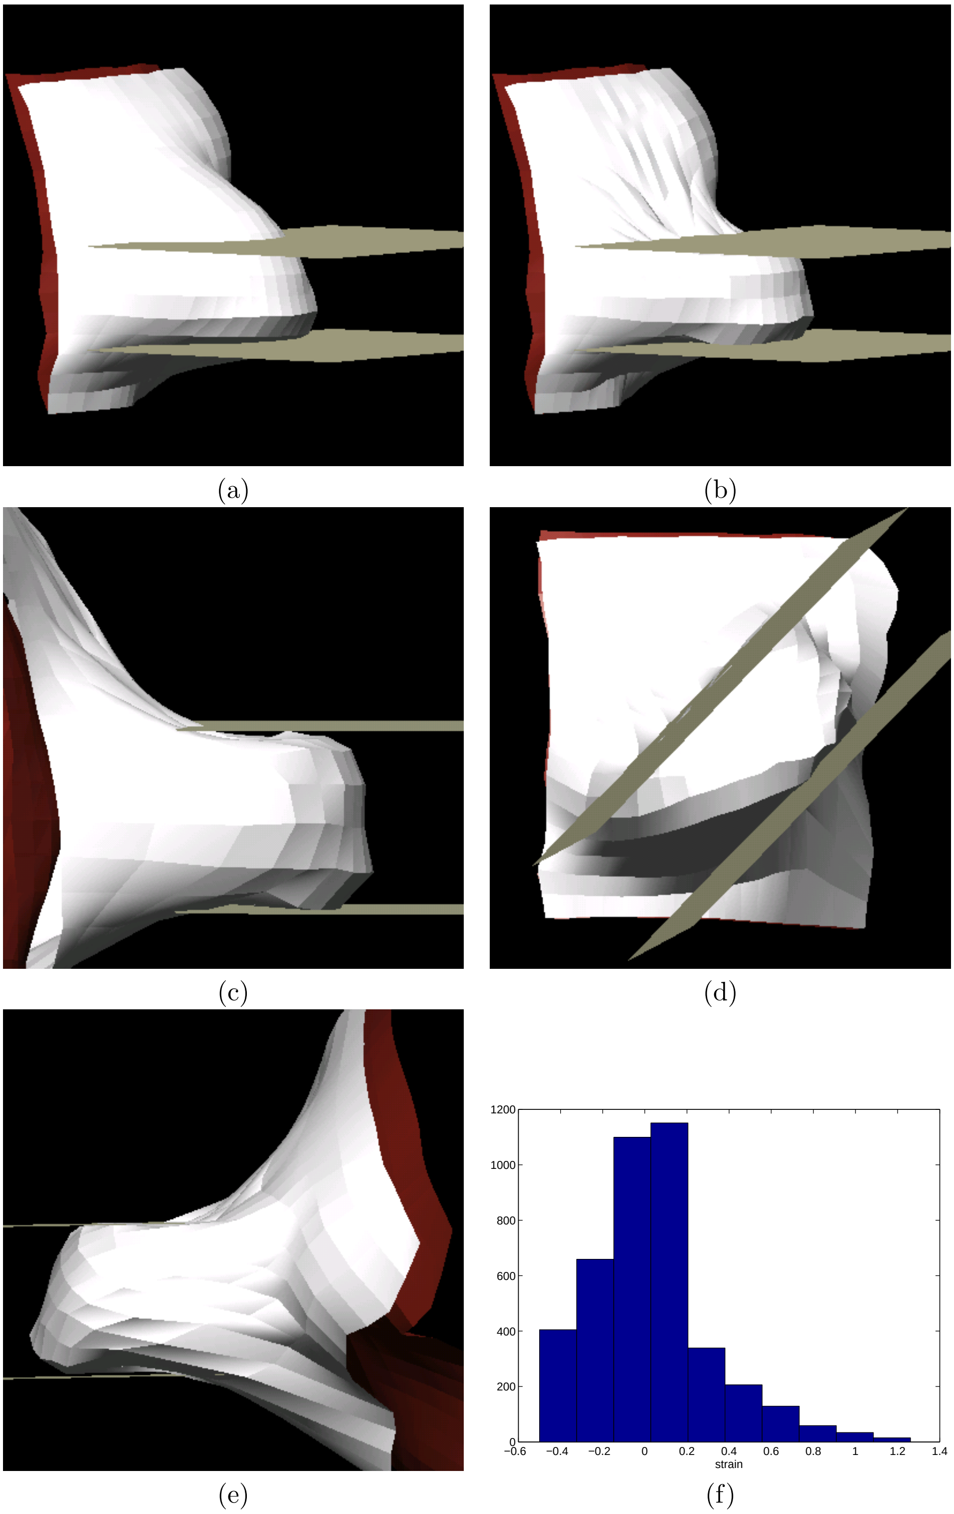

(a) (b) (c) (d) (e) (f)

Figure 7.9: Simulations of mammographic compression: (a) initial mesh before CC compression, (b) CC compressed breast, (c) alternative view of the CC compressed breast, (d) MLO compressed breast, (e) alternative view of the MLO compressed breast, (f) histogram of the strains in the CC compressed breast.

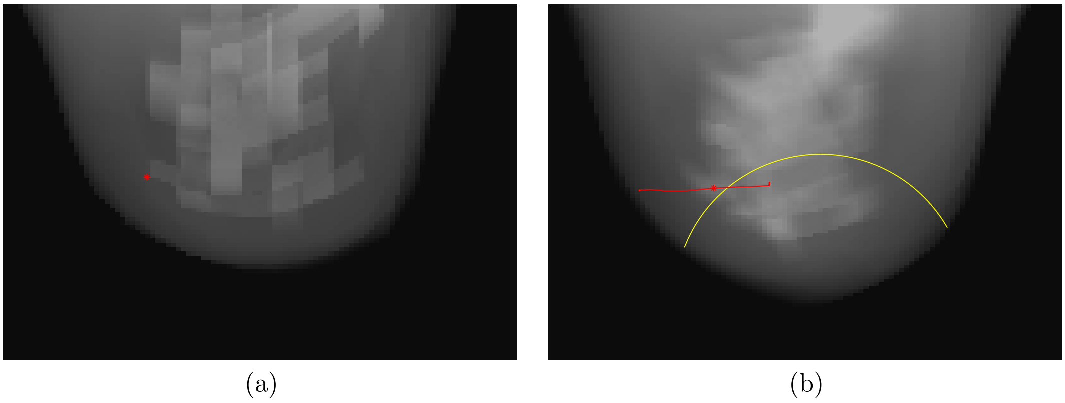
(a) (b)

Figure 7.10: Simulated mammograms: (a) CC mammogram (with position of tumour shown as a red star), (b) MLO mammogram (computed line corresponding the CC tumour point shown as a red line, actual position of tumour shown as a red star, and a guess according to distance from nipple shown as a yellow line). Note that the MLO image does not go far back enough to see the pectoral muscle.

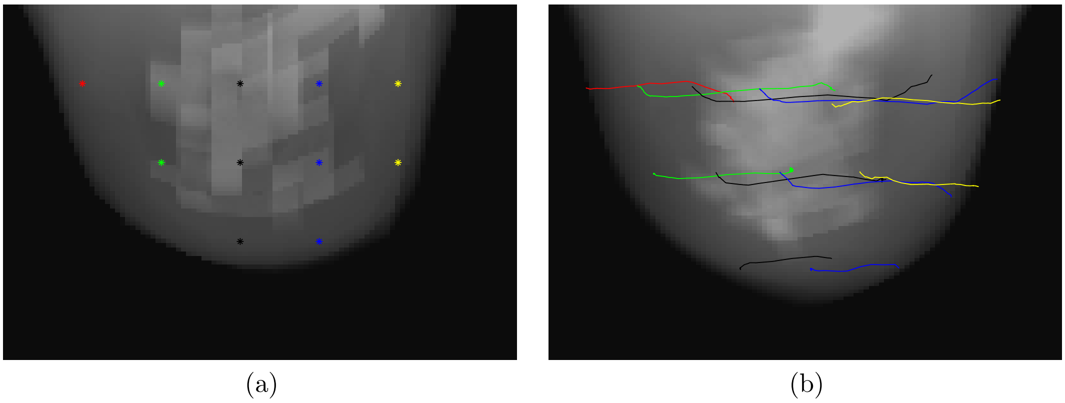
(a) (b)

Figure 7.11: Simulated mammograms: (a) CC mammogram with a selection of pixels, (b) MLO mammogram with equivalent curves to the pixels. Note that the MLO image does not go far back enough to see the pectoral muscle.

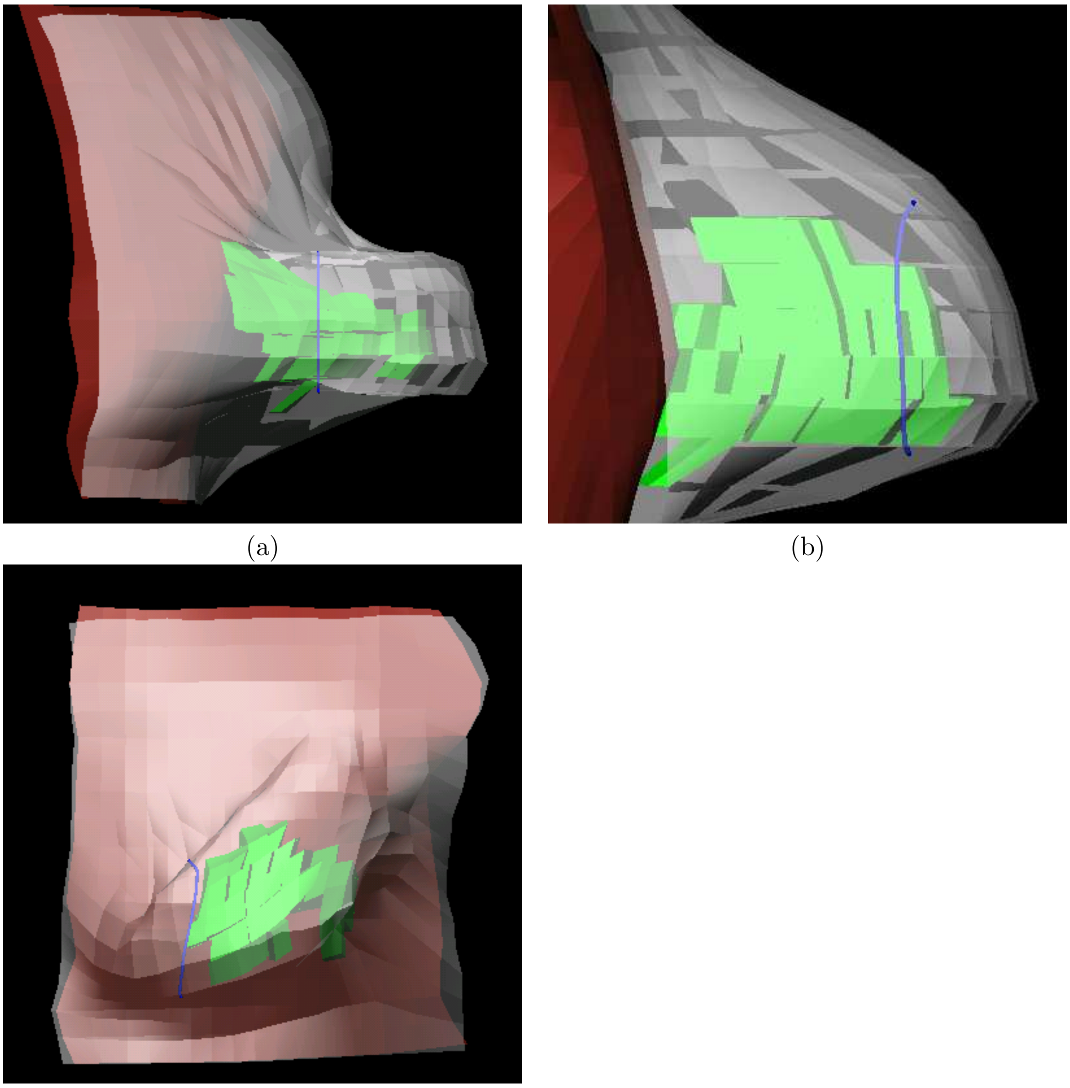

(a)

(b)

(c)

Figure 7.12: CC to MLO matching. In all figures the muscle is the red surface, the fibroglandular region is the green elements, and the skin is the translucent white surface. The blue line represents: (a) the line which is projected onto the tumour pixel in the CC breast, (b) the equivalent uncompressed line in the reference state, (c) the equivalent line in the MLO compressed breast.

7.4 Numerical Experiments

We now use the model to carry out some numerical experiments. We aim to determine the sensitivity of the model to a number of factors, which allows us to come to a more informed decision on whether it is necessary to model a given factor, or provides an idea of acceptable errors in experimental parameters.

Since the prone-to-supine simulations always converge to a solution and are more reliable than the compression simulations, we use this type of simulation in the experiments. The displacements are of course significantly smaller in prone-to-supine deformations than CC or MLO deformations, but the model's sensitivity to parameters can still be established by considering the relative and absolute changes in landmark displacement. We have chosen the (artificial) tumour and the nipple as our two landmarks, and measure the total displacement of each landmark from the prone (MRI) state to the supine state, which involves solving a backward problem followed by a forward problem for each simulation.

7.4.1 Sensitivity of the Model to Material Laws and Parameters

Firstly, we consider the sensitivity of the model to material parameters in the interior of the breast. We continue to use 2nd degree polynomial materials laws (7.2) for fatty and fibroglandular tissue, and study the effect of varying material parameters. Rather than consider changes to the 5 parameters independently, we scale each parameter by constant factor, which means we only need to study a one-dimensional parameter space rather than a five-dimensional one. We want to investigate both small and large changes to the parameters so that either high sensitivity or low sensitivity can be established, and therefore we consider ±10% changes to parameters, together with two-fold and ten-fold decreases and increases.

The results of simulations in which the fatty tissue material parameters are rescaled are given in Table 7.3. For these, fibroglandular and skin material parameters are held constant. Tabulated is the magnitude of the prone-to-supine displacement at the landmarks for each choice of material parameter. Since the norm of the displacement at a landmark being close in two simulations does not guarantee the deformations were similar[11], we also measure the norm of the difference in landmark location between the simulation and the *'base' simulation*, which is the simulation with the 'correct' experimental material parameters (i.e. the 100% row in Table 7.3). Table 7.4 displays the equivalent results of simulations where fibroglandular material parameters are varied and fatty tissue parameters are kept constant.

---

[11] i.e. the displacement at the tumour could be 1mm in two simulations, but this does not mean the final position of the tumour was the same in both simulations. The difference between the final positions could be anything from 0 to 2mm.

It would be expected that fat plays a much greater role in breast-deformation than fibroglandular tissue, especially for this breast, since the total volume of fat is much higher. The results verify and quantify this. They show it is much more important to have little error in fatty tissue material parameters than fibroglandular material parameters. 10% changes to the total stiffness of fat cause approximately 10% changes to the displacement, and larger changes to stiffness cause similar large changes to displacement, but for fibroglandular tissue, a 10% change to stiffness leads to only a 1% change in displacement at both tumour and nipple. For fibroglandular tissue, only when the stiffness is very small or very large is there a significant change in the displacement. The results suggest that much more effort should be expended in determining the mechanical properties of breast fatty tissue *in vivo* than fibroglandular tissue.

Next, we change the fat material law. We replace the 2nd degree polynomial material (7.2) with the one-parameter Neo-Hookean law

$$W = c_1(I_1 - 3) - \frac{p}{2}(I_3 - 1).$$

For the experiment, we use a sample of values of c_1, based around 5 KPa. This value is chosen because it is an approximate average of a number of estimates of c_1 in a survey of experimental results made in [Rui03]. The fibroglandular law remained the 2nd degree polynomial law. The results are given in Table 7.5. They show that the displacements with a Neo-Hookean law are very different to the base simulation, with the displacements much smaller when $c_1 \approx 5$ KPa. This suggests that the size of the deformation in the base simulation is not due to material parameters being too stiff, as was suggested earlier, and are a consequence of the geometry, because the displacements are still small for all Neo-Hookean simulations except when c_1 is very small. The fact that the results in Table 7.5 differ significantly from the base simulation show that the choice of material law and material parameters for fat will have a large impact on the resulting deformation.

Finally, we consider the effect of modelling fat as transversely isotropic, in order to compensate for the reinforcement in the muscle-to-skin direction due to Cooper's Ligaments. We choose the fibre direction to be this direction in the undeformed state (which means making some *a priori* assumptions about the unknown state in the backward problem, as was discussed in Section 5.2.1). We use the 2nd degree polynomial material law for fat (7.2), but augment it with a transverse isotropic term (3.21), with β' in (3.21) set to be 1 and α' varied. The fibroglandular material law was not changed. Results of simulations for a range of values of α', from 0 (no fibre-reinforcement, i.e. the base simulation) to 1000 Pa, are given in Table 7.6. The displacement decreases as the fibre strength increases, as is expected, with the displacement at the nipple changing approximately proportionally with α'. Without experimental data on the true displacement we cannot say if any of these results are more accurate.

Stiffness	Tumour disp.	Tum. rel. base	Nipple disp.	Nip. rel. base	Total mean disp.
10%	8.76	5.91	21.04	12.69	3.55
50%	4.72	1.72	12.73	4.00	1.78
90%	3.26	0.23	9.33	0.55	1.23
100%	3.03	0	8.78	0	1.14
110%	2.83	0.21	8.26	0.52	1.06
200%	1.75	1.29	5.35	3.44	0.66
1000%	0.40	2.63	1.29	7.50	0.15

Table 7.3: Prone to supine displacements (mm) in simulations with **varying fat material parameters**. Tabulated are: magnitude of the displacement at the tumour and nipple, together with the magnitude of these displacements relative to the base deformation (the 100% row), and the total average displacement.

Stiffness	Tumour disp.	Tum. rel. base	Nipple disp.	Nip. rel. base	Total mean disp.
10%	3.21	0.30	9.63	0.92	1.23
50%	3.12	0.15	9.17	0.42	1.18
90%	3.05	0.03	8.85	0.07	1.15
100%	3.03	0	8.78	0	1.14
110%	3.01	0.03	8.69	0.10	1.13
200%	2.86	0.23	8.16	0.66	1.08
1000%	2.24	1.02	6.09	2.84	0.88

Table 7.4: Displacements (mm) in simulations with **varying fibroglandular material parameters**. Tabulated are: magnitude of the displacement at the tumour and nipple, together with the magnitude of these displacements relative to the base deformation (the 100% row), and the total average displacement.

Stiffness c_1 (%, KPa)		Tum dis.	Tum. rel. base	Nip dis.	Nip. rel. base	Tot. mean dis.
10%	0.5	6.60	3.82	12.61	5.96	2.32
50%	2.5	2.11	1.10	3.85	5.40	0.69
90%	4.5	1.24	1.84	2.25	6.73	0.40
100%	5	1.12	1.95	2.04	6.92	0.37
110%	5.5	1.03	2.04	1.87	7.07	0.33
200%	10	0.58	2.46	1.06	7.80	0.19
1000%	50	0.12	2.91	0.25	8.55	0.04

Table 7.5: Displacements (mm) in simulations with **varying Neo-Hookean fat material parameters**. Tabulated are: magnitude of the displacement at the tumour and nipple, together with the magnitude of these displacements relative to the base deformation (2nd degree polynomial material laws), and the total average displacement.

Stiffness, α' (Pa)	Tum. disp.	Tum. rel. base	Nip. disp.	Nip. rel. base	Tot. mean disp.
0	3.03	0	8.78	0	1.14
10	3.01	0.02	8.72	0.06	1.13
50	2.95	0.09	8.49	0.30	1.11
100	2.88	0.16	8.26	0.53	1.08
500	2.45	0.62	6.76	2.05	0.92
1000	2.09	0.99	5.58	3.23	0.78

Table 7.6: Displacements (mm) in simulations with **varying transversely isotropic fat material parameters**. Tabulated are: magnitude of the displacement at the tumour and nipple, together with the magnitude of these displacements relative to the base deformation (2nd degree polynomial material laws), and the total average displacement.

7.4.2 Effect of Skin on the Model

Now we look at the effect of skin on the model. We begin by considering variations in material parameters in the skin material law (7.3). Since only a and c in (7.3) have the dimensions of stiffness, we scale just these two parameters, and hold b constant.

Table 7.7 displays the results of varying the skin material parameters. It is clear the stiffness in the skin has very little effect on the deformation, with the only significant changes in the tumour's deformed location coming when the skin stiffness is increased or decreased by a factor of 10. This suggests that the skin, in the form it is currently modelled, is not a major factor in governing the deformation. The insensitivity of the results to small changes in skin material parameters suggest high accuracy in the measurements of these parameters is not necessary, although the fact that the tumour position does change as the skin stiffness decreases to 10% means we cannot necessarily neglect skin.

However, these conclusions can only be made about the current form of the model, and the results may instead provide an insight into a failing of the model. It is known that human skin is under tension, which we have thus far neglected, since we assume there are no stresses in the reference state. It might be expected that skin plays an important role in the deformation, so the fact that the current model is so insensitive to the skin stiffness could be an indication that this assumption should not have been made. Tension in the skin will be discussed in Chapter 9.

Next, we investigate the impact of skin on computation time. We discussed in Section 7.1.2 how the method used to model skin introduces long thin elements, which are badly-shaped and expected to be detrimental to computation. Table 7.8, where the computation time and the number of Newton iterations required is given for a number of simulations, quantifies this. For these simulations we have used coarser meshes than the full mesh. This is because the full mesh has a large number of badly-shaped elements, which include both the thin skin elements and badly-shaped elements due

Stiffness	Tumour disp.	Tum. rel. base	Nipple disp.	Nip. rel. base	Total mean disp.
10%	3.02	0.14	8.96	0.31	1.15
50%	3.05	0.03	8.89	0.12	1.15
90%	3.03	0.01	8.81	0.03	1.14
100%	3.03	0	8.78	0	1.14
110%	3.02	0.01	8.76	0.02	1.14
200%	2.99	0.04	8.60	0.17	1.12
1000%	2.79	0.27	7.76	1.03	1.02

Table 7.7: Displacements (mm) in simulations with **varying skin material parameters**. Tabulated are: magnitude of the displacement at the tumour and nipple, together with the magnitude of these displacements relative to the base deformation (the 100% row), and the total average displacement.

to the method used for mesh generation (which was essentially morphing a cuboid onto the breast). Here, we would like to study the effect of the thin skin elements separately from that of the other badly-shaped elements, and these coarser meshes have less badly-shaped elements, so are better suited for these experiments. We have used 4 meshes. One has 288 elements, which include thin skin elements. The second mesh also has 288 elements, but no thin skin elements. The remaining two meshes contain 512 elements, one with thin skin, one without. The 288-element meshes are considerably better shaped than the 512-element meshes. In these simulations the skin material law is set to be the same as the fatty tissue material law so that the solutions are identical.

Table 7.8 clearly displays the impact of thin skin elements. For the coarser mesh, the computation takes 5 times longer on the mesh with skin, and more than three times as many Newton iterations are required. The results are not as pronounced for the finer mesh, but there is still a significant discrepancy.

For clinical software, interior badly-shaped elements could be removed by improving the mesh generation procedure. However, thin skin elements are integral to the method we have used to model skin and would still be required, potentially slowing down all computations. In Chapter 8 we will consider how to model skin differently to remove the need for thin skin elements. The results suggest this would bring about small gains in the current full mesh, but nevertheless it is vitally important, as it would allow completely well-shaped meshes to be used in the future. It would also enable tetrahedral meshes (together with quadratic basis functions) to be used, something that is not feasible with thin skin elements.

Mesh	Computation Time (s)	Number Newton Iterations
288 elements, thin skin elements	358	19
288 elements, no skin elements	71	6
512 elements, thin skin elements	434	13
512 elements, no skin elements	334	10

Table 7.8: Comparison of the computation time and the number of Newton steps taken, for meshes with thin skin elements/no thin skin elements.

7.4.3 Effect of a Hard Tumour on the Model

Finally, we look at the effect of including a large hard tumour in the finite element model. Recall that the MR image set did not contain a tumour, so that all elements in the mesh were assigned to be either skin, fat or fibroglandular, and we stated in Section 7.3 that a small stiff tumour is unlikely to affect the gross deformation. We investigate this by artificially choosing one element to be a tumour element (as shown in Figure 7.13), assigning it the same material law as the fibroglandular elements, and varying its stiffness.

The tumour element is obviously larger than a normal tumour, but this means that if the model is insensitive to the large artificial tumour we can say it is likely to be insensitive to tumours in general. Simulations have been carried out with the tumour element set to be 1-20 times stiffer than the neighbouring fibroglandular elements. The results are given in Table 7.9, and show that the model is highly insensitive to tumour stiffness. The tumour has very little impact on the gross deformation (for example, even when the tumour is 20 times stiffer than the fibroglandular tissue, the change in the nipple location is just 0.2%), and very little effect on the local deformation. Only when the tumour is 10 times stiffer than the surrounding tissue is there even a 5% change in the tumour location. Tumours are known to be stiffer than healthy breast tissue, but experimental data (e.g. [WHDK99]) suggest that they are less than 10 times stiffer. Since such a large tumour does not significantly affect the simulations, the results suggest it is not necessary to include small hard tumours in the models, or at least that their exclusion will not be a major source of error.

7.5 Discussion and Further Work

In this chapter we have carried out each of the major stages necessary in modelling clinical deformations of the breast. In this final section of the chapter we review the performance of the methods used, and suggest related areas for further research in this project.

The first task that had to be performed was the segmentation of the MR images, in order to gather the information required to build the model. We saw that edge-detection of the skin surface

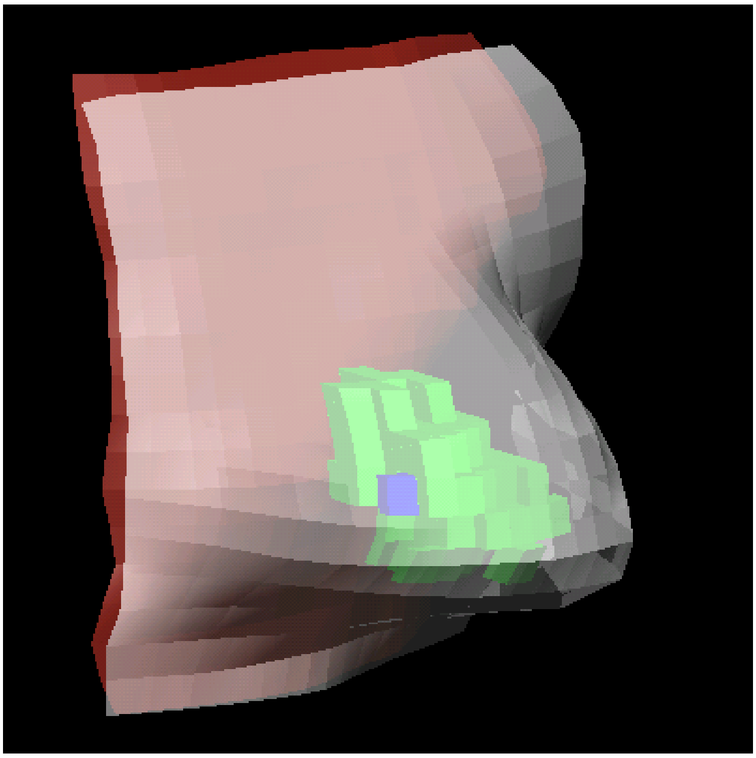

Figure 7.13: The initial mesh (muscle shown red, skin shown translucent white, fibroglandular elements shown green), altered to include a large tumour element (shown blue).

was trivial, but automatically detecting the edge of the pectoral muscle is a much more difficult task, for which current methods in the literature do not appear totally accurate or robust. For this work we had to settle for manually segmenting the pectoral muscle, which would obviously be an unacceptable method in clinical software. An automatic (or semi-automatic), accurate and robust method of pectoral muscle segmentation is a key stage in modelling breast deformation using the methods chosen for this thesis, and this is therefore an important area on which future work could be carried out. Related to this, further investigation into the effect of the boundary conditions on the results will give an insight into whether it is necessary to include an accurate representation of the pectoral muscle on the back surface of the mesh. If the shape of the back surface of the mesh has little effect on the resulting deformations, the required accuracy in the edge-detection algorithm can be lowered.

We also briefly surveyed methods of segmenting the interior of the breast into the two major tissue types. Using the results in Section 7.4.3, we can say that it is likely that it is not necessary to include a tumour in a model as a separate 3D finite-elastic sub-region, and therefore automatic methods for segmenting tumours or regions of interest in MR images are not critical in this particular

Stiffness	Tumour disp.	Tum. rel. base	Nipple disp.	Nip. rel. base	Total mean disp.
100%	3.03	0	8.78	0	1.14
150%	3.02	0.02	8.75	0.04	1.14
200%	3.02	0.03	8.76	0.03	1.14
500%	3.00	0.08	8.76	0.02	1.14
1000%	2.96	0.14	8.74	0.05	1.13
2000%	2.92	0.24	8.77	0.02	1.13

Table 7.9: Displacements (mm) in simulations **with one tumour element, varying the tumour material parameters**. Tabulated are: magnitude of the displacement at the tumour (landmark) and nipple, together with the magnitude of these displacements relative to the base deformation (the 100% row), and the total average displacement.

line of research. A 3D confidence-connected method was used to segment the fatty tissue from the fibroglandular tissue, which gave semi-successful results. However, it is unclear whether this method is more accurate than even a simple thresholding or connected-thresholding approach. A comparison of different segmentation methods, which would require a suitable truth model to be chosen, has not been attempted here, and is left for works dedicated to image analysis. An improved or more accurate method for segmenting interior tissues is unlikely to qualitatively affect the resulting deformations, and is perhaps also unlikely to have a large quantitative effect. Still, an accurate segmentation method will be important when validation studies of the simulated deformations are made.

Relatively unsophisticated methods were also employed for mesh-generation. Since a linear interpolation scheme was used for displacement in the FE method, tetrahedral meshes would have suffered from locking and were not viable, so a hexahedral mesh had to be used. Also, a surface-based method had to be a chosen over a voxel-based method, due to the need for the surface normal to vary smoothly. Each of these choices led to the mesh being extremely badly-shaped: the surface-fitting approach is responsible for the initial creation of the badly-shaped elements, and the mesh being hexahedral meant these elements could not easily be refined into well-shaped elements without introducing hanging nodes. This had a large effect on subsequent simulations, with simulations with the breast mesh taking longer than with cubes with similar numbers of elements. As part of the work for this thesis, attempts were made to produce better quality meshes. This included attempting to refine the mesh in the manner illustrated in Figure 7.4(d), but such a method was quickly found to be unsuitable when an element had to be refined more than once.

Another potential area for further work is the development of new algorithms, or application of current algorithms, for producing well-shaped hexahedral meshes of the breast. Alternatively, quadratic basis functions could be used to interpolate position (together with linear basis functions for pressure), which would enable tetrahedral meshes to be used. In this case mesh-generation would likely become a straightforward procedure (using freely-available or commercial triangulation packages), but computation time would be increased. In any case, it would be important to study the effect of using higher-order interpolation schemes in order to verify the accuracy of the numerical method.

The prone-to-supine deformations simulated in this chapter were seen to be somewhat smaller than might be expected for a typical post-menopausal woman, which may have been due to the image set being of a patient/volunteer who was either pre-menopausal or had relatively small breasts, which would mean the geometry of this particular breast was the cause of the small deformations. The most important area for further research in this project is validation of the results using clinical data, which would allow us to verify hypotheses such as this. Experiments need to be devised where the prone, supine and compressed shapes of a volunteer's breasts are imaged, using several volunteers. Ideally, MR images would be taken of the breast in each of the different positions, so that the predicted displacements in the interior can be validated, but photographic images would still allow the surface displacements to be validated, which is also important, especially for the contact problem. Validation of the simulation of the reference state is potentially feasible as well, for example, by taking MR images of a volunteer whose breast is in a waterbath, to remove the effect of gravity.

Some simple phantom experiments were attempted for this thesis. A small number of heterogeneous cubic porcine skin-based phantoms were produced, and imaged (both externally using photography and internally using an ultrasound probe) under various levels of compression. However, the material laws for the phantom were not known, and can only be determined through careful experimental testing and the solution of inverse parameter-fitting problems (as described in Section 7.2.2), which was well beyond the scope of our experiments. Because of this, the results of these experiments did not provide any useful quantitative data and have not been included in this thesis. For some validation of breast deformations using phantom experiments see [RNN04].

In this chapter we performed a number of numerical experiments to study the effect of various factors on the deformations. We observed that the simulations were relatively sensitive to material parameters for fatty tissue, but relatively insensitive to fibroglandular and skin material parameters. Further numerical tests could be undertaken, but these experiments were hampered by the fact that we did not know the true deformation. Clinical validation would therefore open up a huge range of potential areas of further research. They would allow the experimental material laws and parameters that we have used in this chapter to be studied and verified, and far more powerful parameter searching experiments to be made. Also, any of the assumptions described in Section 7.2.1 could be investigated, and in particular simplifications such as zero-displacement on the pectoral muscle and local incompressibility could be studied. If the simulated deformations do not agree with the experimental results, extensions to the model such as non-zero displacement boundary conditions, a method of modelling blood loss/gain during mammography, or frictional contact could be implemented. Clinical experiments would also be vital to validating the CC-to-MLO predictions made by the model, which suggested that the transformation from a point in a CC mammogram to a curve in an MLO mammogram is dominated by the rotation, rather than the two compressions.

Finally, and related to this, further work could be made into improving the method used to solve the contact problem. There is not a great deal of scope for improving the finite element method on the forward and backward problems, but improved methods for solving contact problems with smooth but essentially discontinuous plates would be important for simulating mammography accurately and quickly.

Chapter 8

Modelling Skin as an Elastic Membrane

8.1 Introduction

We saw in Chapter 7 that the inclusion of skin as a layer of thin elements has a highly detrimental effect on computation time, which is as expected, since badly-shaped elements are well-known to negatively affect computation time and accuracy. However, it would be expected that the skin plays an important role in the deformation of the breast, and so in this chapter we develop a more efficient method to model the skin, which does not require long thin elements.

The key observation is that since the thickness of the skin is very small compared to the typical length-scale, quantities such as strain and stress would be expected to be approximately constant over the thin direction. By performing an asymptotic analysis of the equations with the skin thickness as a small parameter, the three-dimensional integrals over the skin volume, which had previously been required to include skin in the model, can be replaced by two-dimensional integral equations over the skin surface. In essence, the asymptotic analysis replaces the numerical integration over the skin thickness with an analytic, mathematical integration, albeit an approximate one.

As we are now concentrating on one particular facet of the model, these final chapters (excluding Chapter 10) are somewhat removed in content to the majority of this thesis. Also, these chapters involve more mathematics than the previous chapters, and as such are different in style. Therefore, they should perhaps be considered, to some extent, as a separate sub-project to the main body of work.

There has been a great deal of research (see, for example, [Nio85]) in mechanical deformations of *shells*, which is a general term for a thin-walled structure. Shell theory is predominantly used in structural mechanics and therefore concerned with shells which have non-zero bending stiffness, i.e. can support out-of-plane shear stresses. Membranes are examples of shells with no bending stiffness, which can only support membrane (in-plane) forces. There are many different types of shell model, arising through different assumptions made on the physical properties of the shell. Similarly, various assumptions or simplifications have led to different membrane models being used in biomechanical applications (for a review, see [Hum03]). In this chapter we wish to rigorously determine exactly how a 3D strain energy can be converted to a strain energy for a membrane, and it will be useful to begin from first principles rather than shell theory.

Rather than begin with the asymptotics of the full three-dimensional thin solid, we first consider the problem in two dimensions, in order to introduce the major stages in a simplified scenario, and also because the 2D case will be useful for further analysis later on. We will also derive the equations governing a one-dimensional nonlinear membrane (a string) in two-dimensional space, so that the results of the asymptotic analysis can be compared to the string. As well as deriving the equations governing a membrane or a string containing an elastic interior, we also consider the problem of the membrane/string containing an (incompressible, gravity-loaded) fluid. There are a number of reasons for this. The fluid interior is much simpler to model than an elastic interior, and allows us to derive membrane equations on the skin surface without having to deal with the added complexity of internal equations and the coupling between the membrane and the interior. The fluid interior will be important in Chapter 9, when we describe a backward problem for skin. Also, it might be hypothesised that the skin not only plays a major role in, but governs, the deformation of the breast. In this case it would be important to consider a model of the breast consisting simply of an elastic membrane containing an incompressible fluid, to compare with the full model.

Throughout this chapter we will assume that we know the undeformed stress-free configuration of the elastic body, and only attempt to remodel the forward problem. A new backward problem is considered in Chapter 9.

8.2 Membranes and Thin Elastic Solids in 2D

8.2.1 The Equations of a 1D String

Let us begin by considering a 1D elastic membrane in 2D space, or in other words, an elastic string, which we assume undergoes nonlinear deformations. Suppose initially that the string is unstretched and lies on the X-axis between $X = 0$ and $X = L_0$. Setting $s = X$ as a parameterisation of the string, let the deformed position be given by $\{\mathbf{x}(s) = (x(s), y(s)) \,:\, s \in [0, L_0]\,\}$. We use s to make it explicit that this is a parameterisation along a curve. We assume the endpoints are fixed on the x-axis, $\mathbf{x}(0) = (0, 0)$, $\mathbf{x}(L_0) = (L, 0)$, for some L. This is illustrated in Figure 8.1.

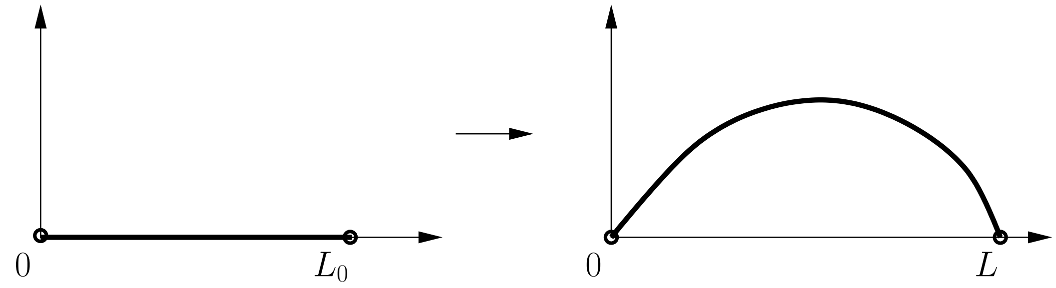

Figure 8.1: The undeformed and deformed states for the string.

The important quantity fundamental to categorising the behaviour of the string is the tangent vector

$$\mathbf{t} = \frac{\mathrm{d}\mathbf{x}}{\mathrm{d}s} = \begin{bmatrix} x_s \\ y_s \end{bmatrix}.$$

$\mathbf{t}$ is not normalised unless $s = \tau$, the arclength along the deformed curve. The magnitude of $\mathbf{t}$ measures the local stretching of the string, since s is arclength in the undeformed string. We can also define a corresponding non-unit (outward) normal

$$\mathbf{n} = \begin{bmatrix} -y_s \\ x_s \end{bmatrix}$$

and unit tangents and normals

$$\hat{\mathbf{t}} = \begin{bmatrix} x_\tau \\ y_\tau \end{bmatrix} = \frac{\mathbf{t}}{\|\mathbf{t}\|} \qquad \text{and} \qquad \hat{\mathbf{n}} = \begin{bmatrix} -y_\tau \\ x_\tau \end{bmatrix} = \frac{\mathbf{n}}{\|\mathbf{n}\|} = \frac{\mathbf{n}}{\|\mathbf{t}\|}.$$

Now, suppose there exists a strain energy function W for the string. We saw in Chapter 3 that strain energy can be considered as a function of the local stretching, so we assume W is a function of magnitude of $\mathbf{t}$

$$W \equiv W(\|\mathbf{t}\|).$$

Suppose we have a system containing the string for which the equations of motion are obtained by minimising an energy of the form $\int_0^{L_0} W \, \mathrm{d}s + \mathcal{K}(\mathbf{x}, \mathbf{t})$, where $\mathcal{K}$ is the remaining energy.

CLAIM: If all other terms in the total energy only contribute to forces in the normal direction, then $\frac{\partial W}{\partial \|\mathbf{t}\|}$ is constant. Let $T = \frac{\partial W}{\partial \|\mathbf{t}\|}$, which can be interpreted as the tension in the string. It follows that, since $T = \frac{\partial W}{\partial \|\mathbf{t}\|}$ is a function of $\|\mathbf{t}\|$ and is constant, the stretch $\|\mathbf{t}\|$ is also constant.

PROOF: Consider the force obtained from the strain energy, which is given by the *Euler*(W), where *Euler* is the Euler-Lagrange operator. Using the fact that $\frac{\partial W}{\partial \mathbf{x}} = 0$ as $W \equiv W(\|\mathbf{t}\|)$ only,

$$Euler(W) \equiv \frac{\mathrm{d}}{\mathrm{d}s}\left(\frac{\partial W}{\partial \mathbf{t}}\right) - \frac{\partial W}{\partial \mathbf{x}} = \frac{\mathrm{d}}{\mathrm{d}s}\left(\frac{\partial W}{\partial \|\mathbf{t}\|}\frac{\partial \|\mathbf{t}\|}{\partial \mathbf{t}}\right) = \frac{\mathrm{d}}{\mathrm{d}s}\left(T\hat{\mathbf{t}}\right).$$

Since $\frac{\mathrm{d}}{\mathrm{d}s} = \|\mathbf{t}\|\frac{\mathrm{d}}{\mathrm{d}\tau}$, and using

$$\frac{\mathrm{d}\hat{\mathbf{t}}}{\mathrm{d}\tau} = -\kappa\hat{\mathbf{n}},$$

where κ[1] is the curvature and $\hat{\mathbf{n}}$ the unit *outward* normal, this becomes

$$Euler(W) \equiv \|\mathbf{t}\|\left(\frac{\mathrm{d}T}{\mathrm{d}\tau}\hat{\mathbf{t}} - T\kappa\hat{\mathbf{n}}\right).$$

Hence, if there are no other tangential forces, we must have $\frac{\mathrm{d}T}{\mathrm{d}\tau} = 0$. □

Note that the strain energy of a material-linear string should satisfy tension T being proportional to strain $\|\mathbf{t}\| - 1$. Hence, a material-linear string would have a strain energy of

$$W = \frac{1}{2}c(\|\mathbf{t}\| - 1)^2,$$

where c is a material constant representing the stiffness of the string.

The Nonlinear Membrane Enclosing a Fluid

Now consider the nonlinear deformation of a 1D membrane enclosing an incompressible fluid. This introduces the constraint that the area under the string has to be constant (A, say), but no extra forces. Since the area is given by

$$\int_{x=0}^{x=x(L_0)} y\,\mathrm{d}x = \int_{s=0}^{s=L_0} x_s y\,\mathrm{d}s,$$

the full minimisation problem is

$$\min \int_0^{L_0} W(\|\mathbf{t}\|) + \lambda x_s y\,\mathrm{d}s,$$

where λ is a Lagrange multiplier corresponding to the area constraint, subject to $\mathbf{x}(0) = 0, \mathbf{x}(L_0) = (L, 0)^{\mathrm{T}}$, and the integral constraint $\int x_s y\,\mathrm{d}s = A$.

---

[1] Since $\hat{\mathbf{t}} \cdot \hat{\mathbf{t}} = 1 \Rightarrow \frac{\mathrm{d}\hat{\mathbf{t}}}{\mathrm{d}\tau} \cdot \hat{\mathbf{t}} = 0$, the rate of change of the unit tangent is proportional to the unit normal. κ is defined to be this constant of proportionality, and measures the curvature of the curve at any point. Explicitly, it is equal to $(y_{ss}x_s - x_{ss}y_s)/(x_s^2 + y_s^2)^{3/2}$.

Letting $\mathcal{L} = W + \lambda x_s y$, the Euler-Lagrange equations for this problem are

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}s}\left(\frac{\partial \mathcal{L}}{\partial \mathbf{t}}\right) - \frac{\partial \mathcal{L}}{\partial \mathbf{x}} = 0 \qquad &\Rightarrow \qquad \frac{\mathrm{d}}{\mathrm{d}s}\left(T\hat{\mathbf{t}}\right) + \frac{\mathrm{d}}{\mathrm{d}s}\left(\left[\begin{array}{c} \lambda y \\ 0 \end{array}\right]\right) - \left[\begin{array}{c} 0 \\ \lambda x_s \end{array}\right] \\ &\Rightarrow \qquad \frac{\mathrm{d}}{\mathrm{d}s}\left(T\hat{\mathbf{t}}\right) - \lambda \|\mathbf{t}\| \hat{\mathbf{n}} = 0 \\ &\Rightarrow \qquad \frac{\mathrm{d}}{\mathrm{d}\tau}\left(T\hat{\mathbf{t}}\right) - \lambda \hat{\mathbf{n}} = 0 \\ &\Rightarrow \qquad \frac{\mathrm{d}T}{\mathrm{d}\tau}\hat{\mathbf{t}} + (-T\kappa - \lambda)\hat{\mathbf{n}} = 0. \end{aligned}$$

As in the claim above, T must be constant. Analogously to the hydrostatic pressure in Chapter 3, let us define a fluid pressure

$$P = -\lambda,$$

and so the equilibrium equation is that the product of tension and curvature is equal to the pressure

$$T\kappa = P.$$

We know that the tension (and the stretch $\|\mathbf{t}\|$) are constant, so the curvature is also constant, and we see that the solution for a nonlinear string covering an incompressible fluid is an arc of an circle.

The Nonlinear Membrane with Gravity

Next we consider the nonlinear membrane enclosing an incompressible fluid which is acted upon by gravity. We have to determine the gravitational potential energy (GPE) of the fluid[2].

$$\mathrm{GPE} = -\iint_{\Omega^f} \rho \mathbf{g} \cdot \mathbf{z} \, \mathrm{d}V,$$

where ρ is the fluid density and $\Omega^f \equiv \Omega^f(\mathbf{x})$ is the volume enclosed by the string, as shown in Figure 8.2.

Now, since $\boldsymbol{\nabla} \cdot \left(\frac{1}{n+1}(\mathbf{g}\cdot\mathbf{z})\,\mathbf{z}\right) = \mathbf{g}\cdot\mathbf{z}$, where n is the dimension of the space, we can use the divergence theorem to say that

$$\mathrm{GPE} = -\iint_{\Omega^f} \rho \mathbf{g} \cdot \mathbf{z} \, \mathrm{d}V = -\frac{1}{3}\oint \rho\,(\mathbf{g}\cdot\mathbf{z})\,(\mathbf{z}\cdot\hat{\mathbf{n}})\,d\tau.$$

On $y = 0$, $\mathbf{z} = (z_1, 0), \hat{\mathbf{n}} = (0, -1)$, so $\mathbf{z}\cdot\hat{\mathbf{n}} = 0$. (In fact, even if this wasn't true, we can neglect this term from the energy, as it is an integral over a fixed curve and therefore constant and does not influence the minimisation problem). Elsewhere,

$$\mathbf{z} = \mathbf{x}, \quad \hat{\mathbf{n}} = \left[\begin{array}{c} -y_s \\ x_s \end{array}\right] \frac{1}{(x_s^2 + y_s^2)^{1/2}} \qquad \text{and} \qquad \mathrm{d}\tau = \left(x_s^2 + y_s^2\right)^{1/2} \mathrm{d}s.$$

---

[2] Note the minus sign in this equation, which means the GPE increases as the body moves away from the source of gravity, as would be expected.

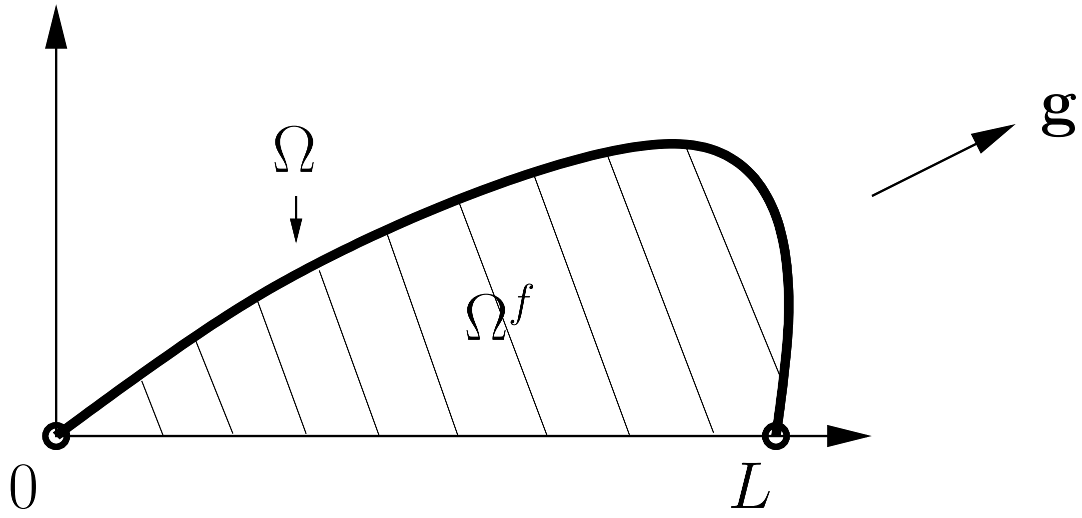

Figure 8.2: The deformed string Ω enclosing a gravity-loaded fluid Ω^f.

The gravitational potential energy is therefore

$$\mathrm{GPE} = -\int_0^{L_0} \frac{1}{3}\rho(\mathbf{g}\cdot\mathbf{x})(\mathbf{x}\cdot\mathbf{n})\,\mathrm{d}s,$$

and the minimisation problem is

$$\min \int_0^{L_0} W(\|\mathbf{t}\|) + \lambda x_s y - \frac{1}{3}\rho(\mathbf{g}\cdot\mathbf{x})(\mathbf{x}\cdot\mathbf{n})\,\mathrm{d}s,$$

subject to $\mathbf{x}(0) = 0, \mathbf{x}(L_0) = (L,0)^{\mathrm{T}}$, and $\int x_s y\,\mathrm{d}s = A$. It can be shown after some (surprisingly non-trivial) algebra that

$$\mathit{Euler}\left(-\frac{1}{3}\rho(\mathbf{g}\cdot\mathbf{x})(\mathbf{x}\cdot\mathbf{n})\right) = \rho(\mathbf{g}\cdot\mathbf{x})\mathbf{n},$$

so that gravity only contributes to forces in the normal direction on the string. The full Euler equations are therefore $T =$ constant, and

$$T\kappa = P + \rho\mathbf{g}\cdot\mathbf{x},$$

which means that the product of tension and curvature is equal to the effective pressure caused by the gravity-loaded fluid.

8.2.2 The Thin Elastic Solid in Two Dimensions

Now that we have derived the equations of a one-dimensional string in 2D space, we are ready to consider the deformation of a thin elastic solid in two dimensions (as a simplification of the three-dimensional thin skin problem which we wish to investigate). We suppose the thin solid is hyperelastic, and therefore the equations governing its deformation are those derived in Chapter 3. We will attempt to reduce the two-dimensional problem to a one-dimensional one, by performing an asymptotic analysis of the equations over the thin direction, the aim being to convert 2D integrals to 1D surface integrals and reduce the equations for the skin to something of the same form as the string equations.

We initially assume the undeformed solid is rectangular, with length L_0 and thickness $h_0 \ll L_0$, and given by $(X, Y) \in [0, L_0] \times [-h_0/2, h_0/2]$, as illustrated in Figure 8.3. We define $\epsilon = h_0/L_0$ as a dimensionless small parameter, and rescale the thin Y direction using the new variable $\hat{Y}$ satisfying $\epsilon\hat{Y} = Y$.

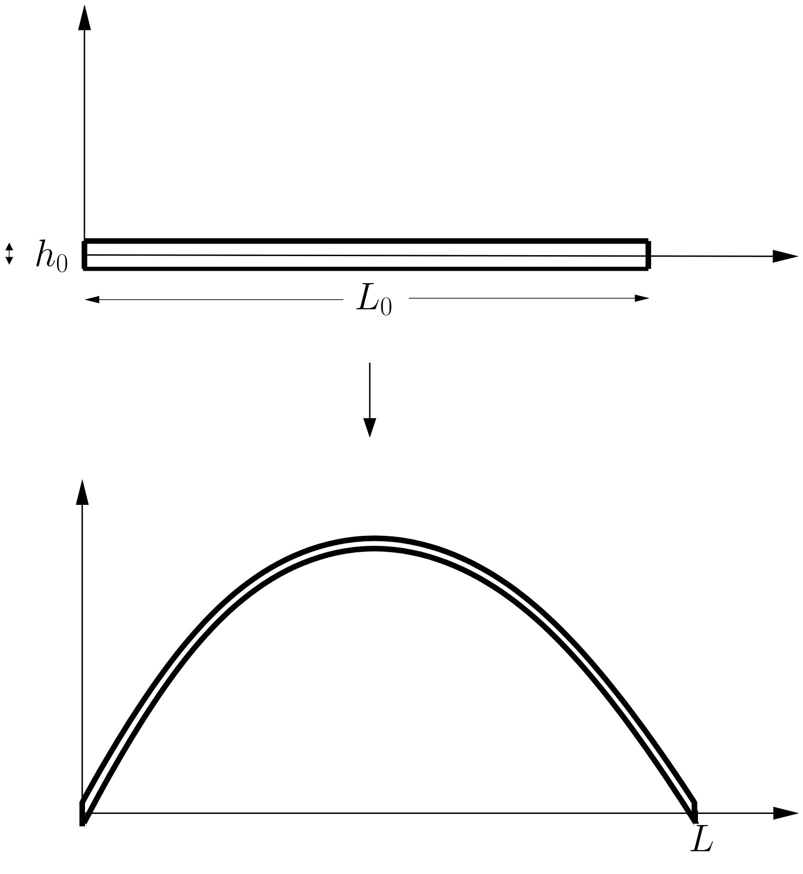

Figure 8.3: The undeformed and deformed states of a 2D thin elastic solid.

The asymptotic expansion of the deformed position $\mathbf{x}(\mathbf{X})$ is

$$\mathbf{x}(X, \hat{Y}) = \mathbf{x}^{(0)}(X, \hat{Y}) + \epsilon\mathbf{x}^{(1)}(X, \hat{Y}) + \epsilon^2\mathbf{x}^{(2)}(X, \hat{Y}) + \ldots$$

We will make the fundamental assumption that the leading order term of the deformation is independent of the thin direction variable, except possibly near the fixed boundaries. This means the deformed membrane is also thin, which is only ever not the case in extraordinary situations, none of which we need to consider.

Assumption: $\mathbf{x}^{(0)}$ satisfies $\mathbf{x}^{(0)}(X,\hat{Y}) \equiv \mathbf{x}^{(0)}(X)$ everywhere except near the fixed ends of the solid.

Now, the deformation gradient $F^i_M = \frac{\partial x^i}{\partial X^M}$ has a leading order component given by

$$F^{(0)} = \begin{bmatrix} x^{(0)}_X & x^{(1)}_{\hat{Y}} \\ y^{(0)}_X & y^{(1)}_{\hat{Y}} \end{bmatrix}.$$

We define the following vectors to simplify the notation.

$$\mathbf{t} = \frac{\partial \mathbf{x}^{(0)}}{\partial X} \qquad \text{and} \qquad \mathbf{h} = \frac{\partial \mathbf{x}^{(1)}}{\partial \hat{Y}}.$$

$\mathbf{t}$ is the tangent vector of the image of the middle line, as shown in Figure 8.4. If we define $\tau \equiv \tau(X)$ to be arclength along this curve then we can write the tangent as $\mathbf{t} = \tau' \hat{\mathbf{t}}$, where $\hat{\mathbf{t}}$ is the normalised tangent vector. $\tau' = \|\mathbf{t}\|$ represents the local stretching of the middle line. $\mathbf{h}$ is the (highest order component of the) image of a vertical line ($X = const$) in the undeformed state. We currently cannot say whether it is normal to $\mathbf{t}$ or not.

Using $\mathbf{t}$ and $\mathbf{h}$, $F^{(0)}$ has the form

$$F^{(0)} = \begin{bmatrix} \mathbf{t}(X) & \mathbf{h}(X,\hat{Y}) \end{bmatrix}.$$

We see that the right Green-Lagrange tensor is, to leading order

$$C^{(0)} = \begin{bmatrix} \|\mathbf{t}\|^2 & \mathbf{t}\cdot\mathbf{h} \\ \mathbf{t}\cdot\mathbf{h} & \|\mathbf{h}\|^2 \end{bmatrix}.$$

Let us initially suppose the thin material is compressible (we will consider the changes required for incompressibility shortly). The two principal invariants are, to leading order

$$\begin{aligned} I^{(0)}_1 &= \operatorname{tr}(C^{(0)}) = \|\mathbf{t}\|^2 + \|\mathbf{h}\|^2 \\ I^{(0)}_2 &= \det C^{(0)} = \|\mathbf{t}\|^2\|\mathbf{h}\|^2 - (\mathbf{t}\cdot\mathbf{h})^2 \end{aligned}$$

There are a number of methods with which we could now proceed. The standard approach is a direct asymptotic analysis of the equation of state (in this case: Cauchy's equation (3.10)). This involves computing the form of the Cauchy stress in terms of the quantities $\mathbf{t}$ and $\mathbf{h}$, expanding each component of the Cauchy stress in ϵ, defining a curvilinear coordinate system on the deformed body

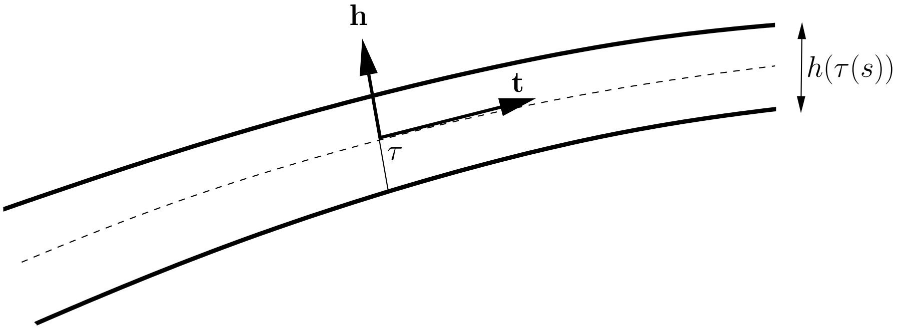

Figure 8.4: The tangent vector **t** of the mid-surface, together with **h**, which represents the image of a vertical line in the undeformed state. Initially we must assume **h** could be non-orthogonal to **t**, but our analysis will show **t** and **h** are orthogonal.

and writing Cauchy's equation in the new coordinate system. Taking the highest order terms for each of the two components of the equation results in two ODEs which can be solved analytically, with careful use of the boundary conditions, for the leading order stresses. From these stresses equations relating **t** and **h** can be inferred.

We have used this method to derive the asymptotic equations for the thin solid, but we omit the details, and instead show how the same equations can be derived in a far more straightforward manner by considering the energy minimisation problem.

There are a number of energy minimisation problems we could consider. In general they take the form[3]

$$\begin{aligned} &\underset{\mathbf{x}}{\text{minimise}} \quad \iint_{\Omega_0} W(I_1, I_2)\,\mathrm{d}S_0 + \mathcal{K}\left(\mathbf{x}, \frac{\partial \mathbf{x}}{\partial X^M}\right) \\ \Rightarrow \quad &\underset{\mathbf{x}}{\text{minimise}} \quad \int_{X=0}^{X=L_0} \int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2} W(I_1, I_2)\,\mathrm{d}\hat{Y}\epsilon\,\mathrm{d}X + \mathcal{K}\left(\mathbf{x}, \frac{\partial \mathbf{x}}{\partial X^M}\right) \end{aligned}$$

where, for example, the remaining energy $\mathcal{K}$ could arise from an elastic interior, or a fluid incompressibility constraint.

In the asymptotic analysis, we take the leading order components of each term, and split the minimisation as a minimisation over $\mathbf{x}^{(0)}$ and $\mathbf{x}^{(1)}$ independently.

$$\underset{\mathbf{x}^{(0)},\,\mathbf{x}^{(1)}}{\text{minimise}} \int_{X=0}^{X=L_0} \int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2} W(I_1^{(0)}, I_2^{(0)})\,\mathrm{d}\hat{Y}\epsilon\,\mathrm{d}X + \mathcal{K}^{(0)}\left(\mathbf{x}^{(0)}, \frac{\partial \mathbf{x}^{(0)}}{\partial X^M}, \mathbf{x}^{(1)}, \dots\right)$$

First we prove that the variation of $\mathbf{x}^{(1)}$ results in the equations determining **h**.

CLAIM: Assuming $\mathcal{K}^{(0)}$ is independent of $\mathbf{x}^{(1)}$, the variation of $\mathbf{x}^{(1)}$ gives $\frac{\partial W}{\partial \mathbf{h}} = 0$ everywhere.

---

[3] Note that in these equations $\hat{Y} = \pm L_0/2$ corresponds to $Y = \pm h_0/2$, i.e. respectively the top and bottom surfaces of the 2D membrane.

PROOF: Assuming the only term in the leading order energy which contains $\mathbf{x}^{(1)}$ is the strain energy, and considering a variation of $\mathbf{x}^{(1)}$ given by $\delta\mathbf{x}^{(1)}$

$$\begin{aligned}\delta E &= \int_{X=0}^{X=L_0}\int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2} W\left(\mathbf{h}+\frac{\partial(\delta\mathbf{x}^{(1)})}{\partial\hat{Y}}\right) - W(\mathbf{h})\ \mathrm{d}\hat{Y}\epsilon\mathrm{d}X \\ &= \int_{X=0}^{X=L_0}\int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2} \frac{\partial W}{\partial\mathbf{h}}\cdot\frac{\partial(\delta\mathbf{x}^{(1)})}{\partial\hat{Y}}\,\mathrm{d}\hat{Y}\epsilon\mathrm{d}X \\ &= \int_{X=0}^{X=L_0}\left[\frac{\partial W}{\partial\mathbf{h}}\cdot\delta\mathbf{x}^{(1)}\right]_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2}\epsilon\,\mathrm{d}X - \int_{X=0}^{X=L_0}\int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2}\left(\frac{\partial W}{\partial\mathbf{h}}\right)_{,\hat{Y}}\cdot\delta\mathbf{x}^{(1)}\,\mathrm{d}\hat{Y}\epsilon\mathrm{d}X\end{aligned}$$

Since $\delta\mathbf{x}^{(1)}$ is completely arbitrary for $X \neq 0, L_0$ (there are no boundary conditions imposed on $\mathbf{x}^{(1)}$, so no boundary conditions on $\delta\mathbf{x}^{(1)}$), it follows that

$$\begin{aligned}\frac{\partial}{\partial\hat{Y}}\left(\frac{\partial W}{\partial\mathbf{h}}\right) &= 0 \qquad \forall X\in(0,L_0),\quad \forall\hat{Y} \\ \frac{\partial W}{\partial\mathbf{h}} &= 0 \qquad \forall X\in(0,L_0),\quad \hat{Y}=L_0/2 \\ \frac{\partial W}{\partial\mathbf{h}} &= 0 \qquad \forall X\in(0,L_0),\quad \hat{Y}=-L_0/2\end{aligned}$$

The solution of this trivial differential equation is $\frac{\partial W}{\partial\mathbf{h}}=0$ everywhere. □

This result leads to a lot of information. Firstly, it tells us that the shear and thin components of the 2nd Piola-Kirchoff stress are zero:

CLAIM: $T^{12}=T^{22}=0$.

PROOF: $0=\frac{\partial W}{\partial h_i}=\frac{\partial W}{\partial F_2^i}=S^{2i}=T^{2N}F_N^i=T^{21}t_i+T^{22}h_i$. But $\{\mathbf{t},\mathbf{h}\}$ is a basis for the space (they must be linearly independent, as otherwise the Jacobian $J=0$), so T^{12} and T^{22} must both vanish. □

We can explicitly compute these components of stress using

$$\begin{aligned}0 &= \frac{\partial W}{\partial\mathbf{h}} \\ &= 2\frac{\partial W}{\partial I_1}\mathbf{h}+2\frac{\partial W}{\partial I_2}(\|\mathbf{t}\|^2\mathbf{h}-(\mathbf{t}\cdot\mathbf{h})\mathbf{t}) \\ &= \left(2\frac{\partial W}{\partial I_1}+2\frac{\partial W}{\partial I_2}\|\mathbf{t}\|^2\right)\mathbf{h}+\left(-2\frac{\partial W}{\partial I_2}(\mathbf{t}\cdot\mathbf{h})\right)\mathbf{t},\end{aligned}$$

These equations completely determine $\mathbf{h}$ in terms $\mathbf{t}$.

$$T^{12} = -2\frac{\partial W}{\partial I_2}(\mathbf{t}\cdot\mathbf{h}) = 0 \qquad \Rightarrow \qquad \mathbf{t}\cdot\mathbf{h} = 0 \tag{8.1}$$

$$T^{22} = 2\frac{\partial W}{\partial I_1} + 2\frac{\partial W}{\partial I_2}\|\mathbf{t}\|^2 = 0 \qquad \Rightarrow \qquad \frac{\partial W}{\partial I_1} + \frac{\partial W}{\partial I_2}\|\mathbf{t}\|^2 = 0 \tag{8.2}$$

Importantly, we have shown that $\mathbf{t}$ is orthogonal to $\mathbf{h}$, and found a relationship (8.2) between $\|\mathbf{t}\|$ and $\|\mathbf{h}\|$. (Recall that $W = W(\|\mathbf{t}\|, \|\mathbf{h}\|)$ so there is some $\|\mathbf{h}\|$ dependence in (8.2)). We have shown that $\mathbf{h}$ is completely determined by $\mathbf{t}$.

Note that we have derived equations for $\mathbf{h}$ without having to consider the bulk deformation. Also, since $\mathbf{h}$ is orthogonal to $\mathbf{t}$ and has norm determined by $\|\mathbf{t}\|$, we can say $\mathbf{h} \equiv \mathbf{h}(X)$ only. It follows that

$$\iint_{\Omega_0} W \,\mathrm{d}V_0 = \int_{X=0}^{X=L_0}\int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2} W(I_1^{(0)}, I_2^{(0)})\,\mathrm{d}\hat{Y}\epsilon\,\mathrm{d}X = \epsilon L_0\int_{X=0}^{X=L_0} W(I_1^{(0)}, I_2^{(0)})\,\mathrm{d}X,$$

so we have analytically integrated over the thin direction and converted the 2D integral to a 1D integral. Also, since $\|\mathbf{h}\| \equiv \|\mathbf{h}\|(\|\mathbf{t}\|)$, we can say that $W \equiv W(\mathbf{t})$, and in fact $W \equiv W(\|\mathbf{t}\|)$. Therefore

$$\iint_{\Omega_0} W(I_1^{(0)}, I_2^{(0)})\,\mathrm{d}V_0 = \epsilon L_0\int_{X=0}^{X=L_0} W(\|\mathbf{t}\|)\,\mathrm{d}X,$$

which is of the form of the strain energy of the string, and we have accomplished our aim: we have converted the double integral of the hyperelastic strain energy to a single integral, and written the strain energy in terms of the norm of the tangent only. We now know how to convert a hyperelastic strain energy to a string strain energy: we must use (8.2) to eliminate $\|\mathbf{h}\|$ from W.

Now we can consider some specific energy minimisations:

Constant area underneath the solid (incompressible fluid)

Here we have an elastic solid containing an incompressible fluid, which again introduces a constraint that the area contained under the thin elastic solid is constant. Using a Lagrange multiplier λ, the full minimisation problem is

$$\min \int_{X=0}^{X=L_0}\int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2} W(I_1, I_2)\,\mathrm{d}\hat{Y}\epsilon\,\mathrm{d}X + \lambda\int_{X=0}^{X=L_0}\frac{\partial x(X, -L_0/2)}{\partial X} y(X, -L_0/2)\,\mathrm{d}X,$$

the second term being an integral over the lower surface of the thin solid.

To leading order this is

$$\min \epsilon L_0\int_{X=0}^{X=L_0} W(\|\mathbf{t}\|)\,\mathrm{d}X + \lambda\int_{X=0}^{X=L_0} x_X^{(0)} y^{(0)}\,\mathrm{d}X.$$

This is formally the same problem as with a nonlinear string, just with a different strain energy. The Euler-Lagrange equations are $\frac{\partial(\epsilon L_0 W)}{\partial\|\mathbf{t}\|} = \text{const}$, which means that the tension

$$T := \frac{\partial(\epsilon L_0 W)}{\partial\|\mathbf{t}\|} = 2\epsilon L_0 \frac{\partial W}{\partial I_1}\|\mathbf{t}\| + 2\epsilon L_0 \frac{\partial W}{\partial I_2}\|\mathbf{h}\|^2\|\mathbf{t}\| = \text{const},$$

and

$$T\kappa = P,$$

where $P = -\lambda$. Note that, if we assume that the pressure is $O(1)$, then since κ is $O(1)$, T must be $O(1)$, so any material constants in W must be $O(\epsilon^{-1})$ for consistency.

Gravity acting on the solid

If gravity is acting on the thin elastic solid, we need to add the following term to the energy

$$\text{GPE}^{(s)} = -\iint_{\Omega_0} \rho_0^{(\mathrm{s})}\mathbf{g}\cdot\mathbf{x}\,\mathrm{d}S_0.$$

Here $\rho_0^{(\mathrm{s})}$ is the density of the thin solid. To leading order this is

$$\int_{X=0}^{X=L_0}\int_{\hat{Y}=-L_0/2}^{\hat{Y}=L_0/2} \rho^{(\mathrm{s})}\mathbf{g}\cdot\mathbf{x}^{(0)}\,\mathrm{d}\hat{Y}\epsilon\,\mathrm{d}X = \epsilon L_0 \int_{X=0}^{X=L_0} \rho^{(\mathrm{s})}\mathbf{g}\cdot\mathbf{x}^{(0)}\,\mathrm{d}X,$$

since $\mathbf{x}^{(0)}$ is independent of $\hat{Y}$. But this is $O(\epsilon)$, which is higher order than the strain energy integral, and therefore we can say that, as would be expected, gravity acting on the solid is a higher-order effect and does not affect the leading order solution.

Gravity acting on the fluid

If the thin elastic solid contains a fluid which is gravity-loaded, we need to add to the energy the gravitational potential energy of the fluid, which is an integral over the volume below the thin solid:

$$\text{GPE}^{(f)} = -\iint_{\Omega^f(\mathbf{x})} \rho^{(\mathrm{f})}\mathbf{g}\cdot\mathbf{z}\,\mathrm{d}V,$$

where Ω^f is the volume enclosed by the thin elastic solid. Exactly as in Section 8.2.1, the solution is

$$T := \frac{\partial(\epsilon L_0 W)}{\partial\|\mathbf{t}\|} = \text{const} \quad\Rightarrow\quad T := 2\epsilon L_0 \frac{\partial W}{\partial I_1}\|\mathbf{t}\| + 2\epsilon L_0 \frac{\partial W}{\partial I_2}\|\mathbf{h}\|^2\|\mathbf{t}\| = \text{const}$$

and

$$T\kappa = P + \rho^{(\mathrm{f})}\mathbf{g}\cdot\mathbf{x}^{(\mathbf{0})}.$$

Incompressibility

If the thin elastic solid is incompressible there is a new constraint $I_2 = 1$ and a new unknown p, the internal pressure in the thin solid. The new strain energy is $W(I_1) - \frac{p}{2}(I_2 - 1)$. Note that this means $p = O(\frac{1}{\epsilon})$, so we write $p = \frac{1}{\epsilon}p^{(-1)} + \ldots$. It can easily be shown using the same argument as before that we must still have $\mathbf{t} \cdot \mathbf{h} = 0$. The previous two equations determining $\|\mathbf{t}\|$ and $\|\mathbf{h}\|$, together with the constraint $I_2 = \|\mathbf{t}\|^2\|\mathbf{h}\|^2 = 1$, now become three equations for three unknowns, the two components of $\mathbf{h}$ and $p^{(-1)}$:

$$\begin{aligned} \mathbf{t} \cdot \mathbf{h} &= 0 \\ 2\frac{\partial W}{\partial I_1}\|\mathbf{h}\|^2 - p^{(-1)} &= 0 \\ \|\mathbf{t}\|\|\mathbf{h}\| &= 1 \end{aligned}$$

These equations hold for any deformation.

In the case of constant pressure boundary conditions, the remaining equations are

$$\begin{aligned} \left(2\epsilon L_0 \frac{\partial W}{\partial I_1}\right)\|\mathbf{t}\| - \epsilon L_0 p^{(-1)}\|\mathbf{h}\| &= T \\ T\kappa &= P \end{aligned}$$

8.3 The Thin Elastic Solid in Three Dimensions

Now we are ready to consider a three-dimensional thin elastic solid. We aim to reproduce or generalise the results derived for two dimensions, in particular showing that the 3D integral of the strain energy over the thin solid can be replaced by a 2D surface integral over the mid-surface, and that the strain energy can be considered to be a function of just the tangent vectors on the mid-surface. On the whole, the results can be proved in the same manner to the 2D results, with only minor increases in complexity; however, setting up the equations is more complicated in three dimensions. We need to be able to deal with arbitrary parameterisations of the undeformed and deformed bodies. Such parameterisations lead to the introduction of tangent vectors and metrics, and the language of tensors and differential geometry is required to carry out the analysis. We begin by setting up this framework.

8.3.1 Curvilinear Coordinate Systems

Let $\Omega_0 \subset \mathbb{R}^3$ be the reference state for the thin solid in three-dimensional space, and let $\tilde{\Theta} \subset \mathbb{R}^3$ be a parameter space for Ω_0. Unlike Section 8.2.2, we do not just consider a flat reference state, so we cannot choose $\tilde{\Theta}$ so that tangent vectors on Ω_0 are orthonormal, but we do assume the reference state is of constant thickness h_0, in which case we can choose $\tilde{\Theta}$ such $\tilde{\Theta} = \Theta \times [-h_0/2, h_0/2]$, where $\Theta \subset \mathbb{R}^2$. Let $\boldsymbol{\theta} = (\theta^1, \theta^2, \theta^3)$ be a member of $\tilde{\Theta}$. Let

$$\epsilon = \frac{h_0}{L_0},$$

be a small dimensionless parameter (where L_0 is a typical (long) length-scale), and define $\hat{\theta}^3$ to be the rescaled parameter in the thin direction, given by $\epsilon\hat{\theta}^3 = \theta^3$. Let $\mathbf{X} \equiv \mathbf{X}(\boldsymbol{\theta}) \in \Omega_0$. The mid-surface Γ_0 of the undeformed solid is given by

$$\Gamma_0 = \{\mathbf{X}(\theta^1, \theta^2, 0) \, : \, (\theta^1, \theta^2) \in \Theta\}.$$

Finally, given a deformed configuration of the thin solid Ω, let the deformed coordinates be $\mathbf{x} \equiv \mathbf{x}(\mathbf{X}(\boldsymbol{\theta})) \equiv \mathbf{x}(\boldsymbol{\theta}) \in \Omega$, and let the mid-surface of Ω be Γ. Let $h(\theta^1, \theta^2)$ be the new deformed thickness of the membrane.

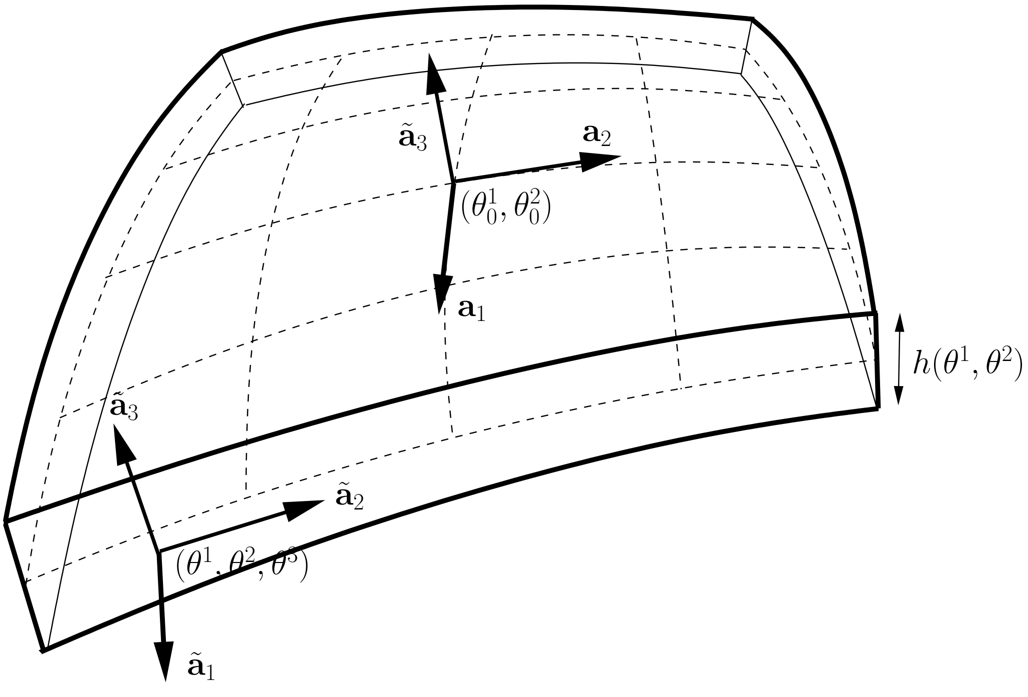

Figure 8.5: The deformed state of a thin 3D elastic object. The middle surface is shown with dotted lines, with each particular dotted line representing a curve of $\theta^1 = const$ or $\theta^2 = const$. At a general point in the object $(\theta^1, \theta^2, \theta^3)$ the three volume basis vectors are shown, and at a point (θ^1, θ^2) on the mid-surface, the two surface basis vectors (tangents) are shown.

It would be natural at this point to use the $\boldsymbol{\theta}$-coordinate system to define three tangent vectors and a three-dimensional metric tensor on each of the spaces Ω_0 and Ω. However, although we will initially wish to work on the 3D spaces and require the 3D metrics and related quantities, at a later point we will be working only on the 2D surfaces, and using the (θ^1, θ^2)-coordinate system and the related 2D metrics and derived quantities, and a little care is needed to avoid confusion or ambiguity later on. Let us introduce the following convention:

- A tensor which has a tilde above it, (e.g. $\tilde{Y}^\alpha$) is assumed to be a tensor over the whole three-dimensional volume (either Ω_0 or Ω), and any indices take values of $1, 2$ or 3. A quantity with a tilde that is derived from a metric tensor is assumed to be derived from the three-dimensional volume metric tensor.
- A tensor without a tilde is assumed to be a tensor over a surface (either Γ_0 or Γ), and any indices only take values of 1 or 2. Implicit summation convention is still used, but indices will be summed from 1 to 2 only. A quantity without a tilde that is derived from a metric tensor is assumed to be derived from the two-dimensional surface metric tensor.

We can now define tangent vectors and metric tensors on both the initial and deformed solid volumes[4]. The tangents vectors $\tilde{\mathbf{A}}^\alpha$ and $\tilde{\mathbf{a}}^\alpha$, for the undeformed and deformed spaces respectively, are derivatives of position, and the metric tensors $\tilde{A}_{\alpha\beta}$ and $\tilde{a}_{\alpha\beta}$ for the two spaces are given by inner products of tangents:

$$\begin{aligned} \tilde{\mathbf{A}}_\alpha &= \frac{\partial \mathbf{X}}{\partial \theta^\alpha} & \qquad \tilde{\mathbf{a}}_\alpha &= \frac{\partial \mathbf{x}}{\partial \theta^\alpha} \\ \tilde{A}_{\alpha\beta} &= \tilde{\mathbf{A}}_\alpha \cdot \tilde{\mathbf{A}}_\beta & \qquad \tilde{a}_{\alpha\beta} &= \tilde{\mathbf{a}}_\alpha \cdot \tilde{\mathbf{a}}_\beta \\ \tilde{A} &= \det(\tilde{A}_{\alpha\beta}) & \qquad \tilde{a} &= \det(\tilde{a}_{\alpha\beta}) \end{aligned}$$

We assume we have chosen the parameter space so that cross-product of the deformed tangents, $\mathbf{a}_1 \wedge \mathbf{a}_2$, is always in the direction of the outward normal of the surface (also making the necessary assumption that the outward normal is a well-defined and continuously-varying quantity).

We can also define corresponding contravariant tangent vectors $\tilde{\mathbf{A}}^\alpha$ and $\tilde{\mathbf{a}}^\alpha$ (members of the dual to the tangent space), and the metric inverses $\tilde{A}^{\alpha\beta}$ and $\tilde{a}^{\alpha\beta}$ by

$$\begin{aligned} \tilde{\mathbf{A}}^\alpha \cdot \tilde{\mathbf{A}}_\beta &= \delta^\alpha_\beta & \qquad \tilde{\mathbf{a}}^\alpha \cdot \tilde{\mathbf{a}}_\beta &= \delta^\alpha_\beta \\ \tilde{A}^{\alpha\beta} \tilde{A}_{\beta\gamma} &= \delta^\alpha_\gamma & \qquad \tilde{a}^{\alpha\beta} \tilde{a}_{\beta\gamma} &= \delta^\alpha_\gamma \end{aligned}$$

The (θ^1, θ^2)-coordinate system also defines surface versions of all these quantities over the two mid-surfaces

$$\begin{aligned} \mathbf{A}_\alpha &= \frac{\partial \mathbf{X}}{\partial \theta^\alpha} & \qquad \mathbf{a}_\alpha &= \frac{\partial \mathbf{x}}{\partial \theta^\alpha} \\ A_{\alpha\beta} &= \mathbf{A}_\alpha \cdot \mathbf{A}_\beta & \qquad a_{\alpha\beta} &= \mathbf{a}_\alpha \cdot \mathbf{a}_\beta \\ A &= \det(A_{\alpha\beta}) & \qquad a &= \det(a_{\alpha\beta}) \end{aligned}$$

---

[4] Note that we use Greek characters for subscripts and superscripts here, which represent an arbitrary curvilinear coordinate system. In Chapter 3 we used lower case Roman letters for coordinates in the deformed space and upper case Roman letters for coordinates in the undeformed space.

the latter pair being 2D determinants. Contravariant surface tangent vectors, $\mathbf{A}^\alpha$ and $\mathbf{a}^\alpha$, and surface metric inverses, $A^{\alpha\beta}$ and $a^{\alpha\beta}$, can be defined similarly. It is important to see that in general, although the tangents $\tilde{\mathbf{a}}_\alpha$ and $\mathbf{a}_\alpha$ are of course the same (if $\alpha = 1$ or 2), and the surface metric tensor $a_{\alpha\beta}$ coincides with a 2D block of $\tilde{a}_{\alpha\beta}$, the components of the inverse metrics are not necessarily equal, ($\tilde{a}^{\alpha\beta} \neq a^{\alpha\beta}$ for $\alpha, \beta < 3$), and neither are the determinants ($\tilde{a} \neq a$).

However, we will assume that the initial shape Ω_0 has constant width and is such that $\tilde{\mathbf{A}}_3$ is normal to the mid-surface and therefore normal to the tangents $\tilde{\mathbf{A}}_1$ and $\tilde{\mathbf{A}}_2$. $\tilde{\mathbf{A}}_3$ is also unit, since we have chosen the parameter space to have the same width as the thin solid. This means that $\tilde{\mathbf{A}}^3 = \tilde{\mathbf{A}}_3$, and that the 3D volume metric is given in terms of the 2D surface metric by

$$\tilde{A}_{\gamma\delta} = \begin{bmatrix} \multicolumn{2}{c}{\multirow{2}{*}{$A_{\alpha\beta}$}} & 0 \\ & & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

It follows that $\tilde{A} = A$, which we note is equal to $\|\mathbf{A}_1 \wedge \mathbf{A}_2\|^2$, and that also $A^{\alpha\beta} = \tilde{A}^{\alpha\beta}$ if $\alpha, \beta < 3$. Note that $\sqrt{A} = \sqrt{\tilde{A}}$ is the factor by which volumes and surfaces transform: $\mathrm{d}V_0 = \sqrt{A}\mathrm{d}\tilde{\Theta} = \sqrt{A}\mathrm{d}\theta^1\mathrm{d}\theta^2\mathrm{d}\theta^3$, and $\mathrm{d}S_0 = \sqrt{A}\mathrm{d}\Theta = \sqrt{A}\mathrm{d}\theta^1\mathrm{d}\theta^2$. We cannot choose the parameter space so that $\tilde{\mathbf{a}}_3$ is also unit or orthogonal to $\mathbf{a}_\alpha$ however, and the purpose of the asymptotic analysis will be to characterise $\tilde{\mathbf{a}}_3$.

Partial derivatives of tensors defined on surfaces expressed by curvilinear coordinate systems are not necessarily tensor quantities. Components of a tensor are basis-dependent, so to define a tensor derivative of a tensor the derivatives of basis vectors need to be computed. This gives rise to the following (non-tensor) quantities, known as *Christoffel Symbols* (of the 2nd kind), defined on the deformed surface[5] by

$$\Gamma^\alpha_{\beta\gamma} = \mathbf{a}^\alpha \cdot \frac{\partial \mathbf{a}_\beta}{\partial \theta^\gamma} \equiv \frac{1}{2} a^{\alpha\delta} \left(\frac{\partial a_{\gamma\delta}}{\partial \theta^\beta} + \frac{\partial a_{\beta\delta}}{\partial \theta^\gamma} - \frac{\partial a_{\beta\gamma}}{\partial \theta^\delta} \right). \tag{8.3}$$

Christoffel symbols are used to define the tensor *covariant derivative* (denoted by a semicolon), which, unlike partial derivatives, do satisfy the tensor transformation laws. For a (2D) contravariant[6] vector field Y^α defined on the mid-surface Γ, the covariant derivative is defined to be

$$Y^\alpha_{\ ;\beta} = Y^\alpha_{\ ,\beta} + \Gamma^\alpha_{\beta\gamma} Y^\gamma$$

and for a contravariant tensor $Y^{\alpha\beta}$

$$Y^{\alpha\beta}_{\ \ ;\gamma} = Y^{\alpha\beta}_{\ \ ,\gamma} + \Gamma^\alpha_{\delta\gamma} Y^{\beta\delta} + \Gamma^\beta_{\delta\gamma} Y^{\alpha\delta}.$$

---

[5] We could define analogous 3D Christoffel symbols over Ω_0 or Ω, or the 2D Christoffel symbols over Γ_0, but the Christoffel symbols over Γ are the only ones we will need to use.

[6] Note that contravariant and covariant vector fields, Y^α and Y_α, represent the coordinates of same quantity with respect to different bases, i.e. $\mathbf{Y} = Y^\alpha \mathbf{a}_\alpha = Y_\alpha \mathbf{a}^\alpha$. The metric tensor (or inverse metric) can be used to raise or lower indices: $Y^\alpha = a^{\alpha\beta} Y_\beta$ and $Y_\alpha = a_{\alpha\beta} Y^\beta$, and similarly with higher-order tensors.

Covariant derivatives of covariant tensors are defined similarly, using $-\Gamma^{\gamma}_{\alpha\beta}$ instead of $\Gamma^{\gamma}_{\alpha\beta}$. Covariant derivatives of scalar functions are simply the partial derivative of the function.

Related to the Christoffel symbols, and also defined on a surface, is the *curvature tensor*[7]

$$b_{\alpha\beta} = \frac{\partial \mathbf{a}_\alpha}{\partial \theta^\beta} \cdot \hat{\mathbf{n}}^{\text{inward}} \tag{8.4}$$

where $\hat{\mathbf{n}}^{\text{inward}} = -\frac{\mathbf{a}_1 \wedge \mathbf{a}_2}{\|\mathbf{a}_1 \wedge \mathbf{a}_2\|}$ is the unit inward normal. The *principal curvatures*, κ_1 and κ_2, are physical, coordinate-system-independent measures of curvature defined to be the eigenvalues of $b^{\alpha}_{\beta} = a^{\alpha\gamma} b_{\gamma\beta}$. Other standard measures of curvature are the mean curvature, $H = \frac{1}{2}\text{tr}(b^{\alpha}_{\beta}) = \frac{1}{2}(\kappa_1 + \kappa_2)$, and the Gaussian curvature, $K = \det(b^{\alpha}_{\beta}) = \kappa_1 \kappa_2$.

8.3.2 Stress Analysis via Energy Minimisation

Having set up our coordinate systems, we are ready to repeat the asymptotic analysis. Let us first consider stress and strain. The metric tensors $\tilde{a}_{\alpha\beta}$ and $\tilde{A}_{\alpha\beta}$ measure length in the spaces, and hence the tensor generalisation of strain is

$$\tilde{E}_{\alpha\beta} = \frac{1}{2}\left(\tilde{a}_{\alpha\beta} - \tilde{A}_{\alpha\beta}\right).$$

If the parameter space was identical to the undeformed space, $\tilde{A}$ would be the identity and this definition of strain would reduce to (3.4). The tensor generalisation of the 2nd Piola-Kirchoff stress is

$$\tilde{T}^{\alpha\beta} = \frac{\partial W}{\partial \tilde{E}_{\alpha\beta}} = 2\frac{\partial W}{\partial \tilde{a}_{\alpha\beta}}.$$

If $\boldsymbol{\theta} = \mathbf{X}$ this is equal to the 2nd Piola-Kirchoff stress, and if $\boldsymbol{\theta} = \mathbf{x}$ it is the Cauchy stress.

Note that since $\tilde{\mathbf{A}}_1$ and $\tilde{\mathbf{A}}_2$ are not necessarily orthonormal, $\tilde{T}^{\alpha\beta}$ and $\tilde{E}_{\alpha\beta}$ do not measure true physical stress or strain. For arbitrary curvilinear coordinate systems the physical components of stress and strain are not tensor quantities, but can be computed by considering the quantity given normalised basis vectors for the component in question. The physical Cauchy stresses can be computed to be [Nas98]

$$\sigma^{(\alpha\beta)} = \tilde{T}^{\alpha\beta}\sqrt{\frac{\tilde{a}_{\beta\beta}}{\tilde{a}^{\alpha\alpha}}}. \tag{8.5}$$

Here, the brackets indicate the quantity is not a tensor, and summation convention is not used.

Since all strain energy functions are assumed to be functions[8] of I_1, I_2 and I_3, we need to be

---

[7] Note that the metric for a surface is sometimes known as the *first fundamental form*. The curvature is the *second fundamental form*.

[8] We only consider the isotropic case for skin in this chapter.

able to compute these invariants. It can be shown that they are given by[9]

$$\begin{aligned} I_1 &= \tilde{a}_{\alpha\beta}\tilde{A}^{\alpha\beta} \\ I_2 &= \frac{1}{2}\tilde{a}_{\alpha\beta}\tilde{a}_{\gamma\delta}\left(\tilde{A}^{\alpha\beta}\tilde{A}^{\gamma\delta} - \tilde{A}^{\alpha\gamma}\tilde{A}^{\beta\delta}\right) \\ I_3 &= \frac{\tilde{a}}{\tilde{A}} \end{aligned} \tag{8.6}$$

It follows that

$$\tilde{T}^{\alpha\beta} = 2\frac{\partial W}{\partial \tilde{a}_{\alpha\beta}} = 2\frac{\partial W}{\partial I_1}\tilde{A}^{\alpha\beta} + 2\frac{\partial W}{\partial I_2}\tilde{a}_{\gamma\delta}\left(\tilde{A}^{\alpha\beta}\tilde{A}^{\gamma\delta} - \tilde{A}^{\alpha\gamma}\tilde{A}^{\beta\delta}\right) + 2\frac{\partial W}{\partial I_3}\tilde{a}\tilde{A}^{-1}\tilde{a}^{\alpha\beta} \tag{8.7}$$

or, in the incompressible case,

$$\tilde{T}^{\alpha\beta} = 2\frac{\partial W}{\partial \tilde{a}_{\alpha\beta}} = 2\frac{\partial W}{\partial I_1}\tilde{A}^{\alpha\beta} + 2\frac{\partial W}{\partial I_2}\tilde{a}_{\gamma\delta}\left(\tilde{A}^{\alpha\beta}\tilde{A}^{\gamma\delta} - \tilde{A}^{\alpha\gamma}\tilde{A}^{\beta\delta}\right) - p\tilde{a}^{\alpha\beta}, \qquad \tilde{a}\tilde{A}^{-1} = 1 \tag{8.8}$$

We assume as in Section 8.2.2 that $\mathbf{x}$ has the form

$$\mathbf{x} = \mathbf{x}^{(0)}(\theta^1, \theta^2) + \epsilon\mathbf{x}^{(1)}(\theta^1, \theta^2, \hat{\theta}^3) + \dots,$$

This means that $\tilde{\mathbf{a}}_1^{(0)} = \frac{\partial \mathbf{x}^{(0)}}{\partial \theta^1}$, $\tilde{\mathbf{a}}_2^{(0)} = \frac{\partial \mathbf{x}^{(0)}}{\partial \theta^2}$ and $\tilde{\mathbf{a}}_3^{(0)} = \frac{\partial \mathbf{x}^{(1)}}{\partial \hat{\theta}^3}$. For convenience we will drop the superscripts (0) on the tangents $\tilde{\mathbf{a}}_\alpha$ and the metric $\tilde{a}_{\alpha\beta}$. The first step is to consider a variation of $\mathbf{x}^{(1)}$.

Consider a general energy minimisation problem of the form

$$\underset{\mathbf{x}}{\text{minimise}} \iiint_{\Omega_0} W(I_1, I_2, I_3)\, \mathrm{d}V_0 + \mathcal{K}\left(\mathbf{x}, \frac{\partial \mathbf{x}}{\partial \theta^\alpha}\right)$$

where $\mathcal{K}$ is the remaining energy and depends on the specific problem being modelled. For the asymptotic analysis, as in 2D, we will take the leading order components of each term, and split the minimisation as a minimisation over $\mathbf{x}^{(0)}$ and $\mathbf{x}^{(1)}$ independently. We assume that to highest order $\mathcal{K}$ is independent of $\mathbf{x}^{(1)}$.

$$\begin{aligned} & \underset{\mathbf{x}^{(0)},\, \mathbf{x}^{(1)}}{\text{minimise}} \iiint_{\Omega_0} W(I_1^{(0)}, I_2^{(0)}, I_3^{(0)})\, \mathrm{d}V_0 \; + \mathcal{K}^{(0)}\left(\mathbf{x}^{(0)}, \frac{\partial \mathbf{x}^{(0)}}{\partial \theta^\alpha}\right) \\ \Rightarrow \quad & \underset{\mathbf{x}^{(0)},\, \mathbf{x}^{(1)}}{\text{minimise}} \iint_{\Theta}\left(\int_{\hat{\theta}^3=-L_0/2}^{\hat{\theta}^3=L_0/2} W(\tilde{\mathbf{a}}_1, \tilde{\mathbf{a}}_2, \tilde{\mathbf{a}}_3)\sqrt{A}\epsilon \mathrm{d}\hat{\theta}^3\right)\mathrm{d}\theta^1\, \mathrm{d}\theta^2 + \mathcal{K}^{(0)}\left(\mathbf{x}^{(0)}, \mathbf{a}_\alpha\right) \end{aligned}$$

CLAIM: The minimisation of $\mathbf{x}^{(1)}$ implies that $\tilde{T}^{3\alpha} = 0$ for $\alpha = 1, 2, 3$ (to leading order).

---

[9] For example, $I_1 = C_{MM} = \frac{\partial x^i}{\partial X^M}\frac{\partial x^i}{\partial X^M} = \frac{\partial x^i}{\partial \theta^\alpha}\frac{\partial \theta^\alpha}{\partial X^M}\frac{\partial \theta^\beta}{\partial X^M}\frac{\partial x^i}{\partial \theta_\beta} = (\tilde{\mathbf{a}}_\alpha \cdot \tilde{\mathbf{a}}_\beta)(\tilde{\mathbf{A}}^\alpha \cdot \tilde{\mathbf{A}}^\beta) = \tilde{a}_{\alpha\beta}\tilde{A}^{\alpha\beta}$. The calculations for I_2 and I_3 are similar.

PROOF: The proof is identical to 2D. Using an identical argument to that used in Section 8.2.2, it can be shown that

$$\frac{\partial W}{\partial \tilde{\mathbf{a}}_3} = 0 \qquad \text{everywhere.}$$

But

$$0 = \frac{\partial W}{\partial \tilde{\mathbf{a}}_3} = \tilde{T}^{3\alpha}\mathbf{a}_\alpha,$$

and $\{\tilde{\mathbf{a}}_\alpha\}$ form a basis for $\mathbb{R}^3$, so each $\tilde{T}^{3\alpha} = 0$. $\square$

The three equations $\tilde{T}^{3\alpha} = 0$ are enough to determine $\tilde{\mathbf{a}}_3$ in terms of $\tilde{\mathbf{a}}_1$ and $\tilde{\mathbf{a}}_2$. The equation $\tilde{T}^{33} = 0$ relates $\|\tilde{\mathbf{a}}_3\|$ to the norms of the surface tangents, and the equations $\tilde{T}^{31} = 0$ and $\tilde{T}^{32} = 0$ imply $\tilde{\mathbf{a}}_3$ is orthogonal to the tangents:

CLAIM: $\tilde{\mathbf{a}}_\alpha \cdot \tilde{\mathbf{a}}_3 = 0$, $\alpha = 1, 2$ (to leading order).

PROOF:

We are attempting to prove something that is independent of the parameterisation (if $\tilde{\mathbf{a}}_3$ is orthogonal to one set of tangent vectors corresponding to one parameterisation of the surface, it will be orthogonal to tangents arising from any other parameterisation). Also, at any particular point we can always pick Θ such that $\{\tilde{\mathbf{A}}_\alpha\}$ are orthonormal at that point. Therefore, it is enough to prove the claim with $\tilde{A}_{\alpha\beta} = \delta_{\alpha\beta}$, which simplifies the problem enormously. In these coordinates $0 = \tilde{T}^{3\alpha} = -2\frac{\partial W}{\partial I_2}\tilde{a}_{3\alpha} + 2\frac{\partial W}{\partial I_3}a\tilde{a}^{3\alpha}$ (or $0 = \tilde{T}^{3\alpha} = -\frac{\partial W}{\partial I_2}\tilde{a}_{3\alpha} - p\tilde{a}^{3\alpha}$) at the point. Dividing through we can reduce this to

$$\tilde{a}_{3\alpha} = \lambda\,\tilde{a}^{3\alpha} \tag{8.9}$$

at the point, where $\lambda = a\frac{\partial W}{\partial I_3}/\frac{\partial W}{\partial I_2}$ or $-p/(2\frac{\partial W}{\partial I_2})$.

Our proof is restricted to the case $\lambda \geq 0$, which will always be true for material laws considered for skin in this thesis. Define a (non-physical, non-tensor) vector $\mathbf{V} = (\tilde{a}_{32}, -\tilde{a}_{31}, 0)^{\mathrm{T}}$. Then, $\tilde{a}_{\alpha\beta}V^\beta = \lambda\tilde{a}(\tilde{a}_{11}\tilde{a}_{32} - \tilde{a}_{12}\tilde{a}_{31}, \tilde{a}_{21}\tilde{a}_{32} - \tilde{a}_{22}\tilde{a}_{31}, 0)^{\mathrm{T}}$, which is equal to $\lambda\tilde{a}(-\tilde{a}_{32}, \tilde{a}_{31}, 0)$ using (8.9) and the symmetry of $\tilde{a}_{\alpha\beta}$. Therefore $V^\alpha\tilde{a}_{\alpha\beta}V^\beta = -\lambda\tilde{a}\left((\tilde{a}_{31})^2 + (\tilde{a}_{31})^2\right) \leq 0$. But the metric is positive definite, so also $V^\alpha\tilde{a}_{\alpha\beta}V^\beta \geq 0$, with equality if and only if $\mathbf{V} = 0$. Hence $\mathbf{V} = 0$. $\square$

Since $\tilde{\mathbf{a}}_3$ is orthogonal to $\tilde{\mathbf{a}}_1$ and $\tilde{\mathbf{a}}_2$, the metric $\tilde{a}_{\alpha\beta}$ is of the form

$$\tilde{a}_{\alpha\beta} = \begin{bmatrix} a_{\gamma\delta} & \\ & \tilde{a}_{33} \end{bmatrix}$$

The stress $\tilde{T}$ is therefore a function of $a_{\alpha\beta}$, the 2D surface metric, and $\tilde{a}_{33}$, the normal stretch (squared). The equation $\tilde{T}^{33} = 0$, which, explicitly, is

$$\begin{aligned} 0 &= 2\frac{\partial W}{\partial I_1}A^{33} + 2\frac{\partial W}{\partial I_2}\tilde{a}_{\gamma\delta}\left(\tilde{A}^{\gamma\delta} - \tilde{A}^{3\gamma}\tilde{A}^{3\delta}\right) - 2\frac{\partial W}{\partial I_3}\frac{a}{A}\tilde{a}^{33}, \\ \Rightarrow 0 &= 2\frac{\partial W}{\partial I_1} + 2\frac{\partial W}{\partial I_2}a_{\gamma\delta}A^{\gamma\delta} - 2\frac{\partial W}{\partial I_3}\frac{a}{A\tilde{a}_{33}} \end{aligned} \tag{8.10}$$

can be used to relate $\tilde{a}_{33}$ and $a_{\alpha\beta}$. In the incompressible case we can use the incompressibility constraint $a\tilde{a}_{33} = A$ to relate $\tilde{a}_{33}$ to $a_{\alpha\beta}$, and the equation $T^{33} = 0$, which is

$$0 = 2\frac{\partial W}{\partial I_1} + 2\frac{\partial W}{\partial I_2}a_{\gamma\delta}A^{\gamma\delta} - \frac{p}{\tilde{a}_{33}}, \tag{8.11}$$

can be used to eliminate p. Either way, we can say that $W \equiv W(a_{\alpha\beta}, \tilde{a}_{33}(a_{\alpha\beta})) \equiv W^{\text{mem}}(a_{\alpha\beta})$, where we write $W^{\text{mem}}(a_{\alpha\beta})$ to emphasise that we have derived a new energy which is a function of the surface metric only.

Also, since $a_{\alpha\beta} \equiv a_{\alpha\beta}(\theta^1, \theta^2)$ only (to leading order), we can say that to leading order W^{mem} is independent of θ^3, and so

$$\begin{aligned} \iiint_{\Omega_0} W(I_1^{(0)}, I_2^{(0)}, I_3^{(0)})\,\mathrm{d}V_0 &= \epsilon L_0 \iint_{\Gamma_0} W^{\text{mem}}(a_{\alpha\beta})\,\mathrm{d}S_0 \\ &= \epsilon L_0 \iint_{\Theta} W^{\text{mem}}(a_{\alpha\beta})\sqrt{A}\,\mathrm{d}\theta^1\mathrm{d}\theta^2 \end{aligned}$$

Again, we have analytically integrated over the thin direction, to convert the volume integral to a surface integral, and determined how to convert a 3D hyperelastic strain energy to a 2D surface membrane strain energy (we have to set $\tilde{a}_{13} = \tilde{a}_{23} = 0$ and use (8.10) or (8.11) to eliminate $\tilde{a}_{33}$).

8.3.3 Analysis of Specific Problems

Elastic solid containing an incompressible fluid

To impose the constraint that the volume under the membrane is constant we need to add a constraint term $\lambda \iiint_{\Omega^f} \mathrm{d}V$ to the energy, where Ω^f is the deformed volume of fluid under thin solid. Using the divergence theorem on $\frac{1}{3}\mathbf{x}$, but ignoring integrals over fixed surfaces, this is equal to $\iint \frac{1}{3}\mathbf{x}^{(0)} \cdot \hat{\mathbf{n}}\,\mathrm{d}S$, where the integral is evaluated on the lower surface of the thin solid. Taking leading order contributions, and using $\hat{\mathbf{n}}\mathrm{d}S = \frac{(\mathbf{a}_1\wedge\mathbf{a_2})}{\sqrt{a}}\sqrt{a}\mathrm{d}\theta^1\mathrm{d}\theta^2 = (\mathbf{a}_1\wedge\mathbf{a_2})\mathrm{d}\theta^1\mathrm{d}\theta^2$ (as we assume $\mathbf{a}_1\wedge\mathbf{a_2}$ is outward pointing), the leading order energy is

$$\begin{aligned} &\epsilon L_0 \iint_{\Gamma_0} W(I_1^{(0)}, I_2^{(0)}, I_3^{(0)})\,\mathrm{d}S_0 &+ \frac{1}{3}\lambda \iint_{\Theta} \mathbf{x}\cdot(\mathbf{a}_1 \wedge \mathbf{a_2})\,\mathrm{d}\theta^1\mathrm{d}\theta^2 \\ = &\epsilon L_0 \iint_{\Theta} W^{\text{mem}}(\mathbf{a}_1, \mathbf{a}_2)\sqrt{A}\,\mathrm{d}\theta^1\mathrm{d}\theta^2 &+ \frac{1}{3}\lambda \iint_{\Theta} \mathbf{x}\cdot(\mathbf{a}_1 \wedge \mathbf{a_2})\,\mathrm{d}\theta^1\mathrm{d}\theta^2 \end{aligned}$$

The Euler-Lagrange equations are (if $\mathcal{L}$ is the integrand for the total energy)

$$Euler(\mathcal{L}) \equiv \frac{\partial}{\partial \theta^1}\left(\frac{\partial \mathcal{L}}{\partial \mathbf{a}_1}\right) + \frac{\partial}{\partial \theta^2}\left(\frac{\partial \mathcal{L}}{\partial \mathbf{a}_2}\right) - \frac{\partial \mathcal{L}}{\partial \mathbf{x}^{(0)}} = 0$$

After a little algebra, it can be shown that the Euler-Lagrange operator on the constraint energy is

$$Euler\left(\frac{1}{3}\mathbf{x}^{(0)} \cdot (\mathbf{a}_1 \wedge \mathbf{a_2})\right) = -\mathbf{a}_1 \wedge \mathbf{a_2},$$

and so the Euler-Lagrange equations for this case are

$$\frac{\partial}{\partial \theta^\alpha}\left(\epsilon L_0 \sqrt{A} \frac{\partial W^{\mathrm{mem}}}{\partial \mathbf{a}_\alpha}\right) - \lambda \mathbf{a}_1 \wedge \mathbf{a}_2 = 0,$$

or, setting$P = -\lambda$ as usual

$$\frac{\partial}{\partial \theta^\alpha}\left(h_0 \sqrt{A} \frac{\partial W^{\mathrm{mem}}}{\partial \mathbf{a}_\alpha}\right) + P \mathbf{a}_1 \wedge \mathbf{a}_2 = 0.$$

Elastic solid containing a gravity-loaded incompressible fluid

Here, we need to add to the energy the gravitational potential energy of the fluid, which is an integral over the volume below the thin solid.

$$\mathrm{GPE}^{(f)} = -\iiint_{\Omega^f} \rho^{(\mathrm{f})} \mathbf{g} \cdot \mathbf{z} \, \mathrm{d}V$$

where $\rho^{(f)}$ is the fluid density. Using the divergence theorem, and taking the leading order component,

$$\mathrm{GPE}^{(f)} = -\frac{1}{4}\iint_\Gamma \rho^{(\mathrm{f})}\left(\mathbf{g} \cdot \mathbf{x}^{(0)}\right)\left(\mathbf{x}^{(0)} \cdot \hat{\mathbf{n}}\right) \mathrm{d}S = -\frac{1}{4}\iint_\Theta \rho^{(\mathrm{f})}\left(\mathbf{g} \cdot \mathbf{x}^{(0)}\right)\left(\mathbf{x}^{(0)} \cdot (\mathbf{a}_1 \wedge \mathbf{a}_2)\right) \mathrm{d}\theta^1 \mathrm{d}\theta^2.$$

Again we neglect the term which is an integral over the lower fixed surface, as it is constant and does not affect the variational formulation.

After some more algebra, it can be shown that, if $\mathcal{G} = -\frac{1}{4}\rho^{(f)}(\mathbf{g} \cdot \mathbf{x})(\mathbf{x} \cdot (\mathbf{a}_1 \wedge \mathbf{a}_2))$,

$$Euler(\mathcal{G}) = \rho^{(f)}(\mathbf{g} \cdot \mathbf{x})\mathbf{a}_1 \wedge \mathbf{a}_2,$$

so the new Euler-Lagrange equations are

$$\frac{\partial}{\partial \theta^\alpha}\left(h_0 \sqrt{A} \frac{\partial W^{\mathrm{mem}}}{\partial \mathbf{a}_\alpha}\right) + \left(P + \rho^{(\mathrm{f})} \mathbf{g} \cdot \mathbf{x}^{(0)}\right) \mathbf{a}_1 \wedge \mathbf{a}_2 = 0. \tag{8.12}$$

8.3.4 Analysis of the Governing Equations

Recall that in two dimensions the tension $T = \frac{\partial W}{\partial \|\mathbf{t}\|}$ was constant, and the tension multiplied by the curvature was equal to the effective pressure $T\kappa = P + \rho^{(\mathrm{f})}\mathbf{g} \cdot \mathbf{x}^{(0)}$. We now prove the corresponding three-dimensional surface results. Since the 2D equations were written in terms of tensions (i.e. forces), it is necessary to first derive the in-plane forces. This is the stress integrated over the thin direction in the deformed solid, which is equal to the stress multiplied by the new thickness, h. Using the incompressibility constraint on a small volume, we must have $h_0\sqrt{A} = h\sqrt{a}$, so the new thickness is $h = h_0\sqrt{A}/\sqrt{a}$. The in-plane force tensor, also known as the *stress-residual*, is therefore[10]

$$F^{\alpha\beta} = T^{\alpha\beta} h_0 \frac{\sqrt{A}}{\sqrt{a}} = 2h_0 \frac{\partial W^{\mathrm{mem}}}{\partial a_{\alpha\beta}} \frac{\sqrt{A}}{\sqrt{a}}. \tag{8.13}$$

CLAIM: The stress-residual has zero covariant divergence, and the trace of the product of the stress residual and the curvature tensor is equal to the effective pressure, i.e.

$$\begin{aligned} F^{\alpha\beta}{}_{;\beta} &= 0 && (8.14)\\ F^{\alpha\beta} b_{\alpha\beta} &= P + \rho^{(\mathrm{f})}\mathbf{g} \cdot \mathbf{x}^{(0)} && (8.15) \end{aligned}$$

PROOF:

In terms of stress-residuals, equation (8.12) is

$$\frac{\partial}{\partial \theta^\alpha}\left(F^{\alpha\beta}\mathbf{a}_\beta\sqrt{a}\right) = -(P + \rho^{(\mathrm{f})}\mathbf{g} \cdot \mathbf{x}^{(0)})\mathbf{a}_1 \wedge \mathbf{a}_2,$$

which means that

$$\frac{\partial}{\partial \theta^\alpha}\left(F^{\alpha\beta}\sqrt{a}\right)\mathbf{a}_\beta + F^{\alpha\beta}\sqrt{a}\frac{\partial \mathbf{a}_\beta}{\partial \theta^\alpha} = -(P + \rho^{(\mathrm{f})}\mathbf{g} \cdot \mathbf{x}^{(0)})\mathbf{a}_1 \wedge \mathbf{a}_2. \tag{8.16}$$

We can take the inner product this with $\mathbf{a}^\gamma$ ($\gamma = 1$ or 2) so that the right-hand side vanishes:

$$\begin{aligned} \frac{\partial}{\partial \theta^\alpha}\left(F^{\alpha\beta}\sqrt{a}\right)\mathbf{a}_\beta \cdot \mathbf{a}^\gamma + F^{\alpha\beta}\sqrt{a}\frac{\partial \mathbf{a}_\beta}{\partial \theta^\alpha} \cdot \mathbf{a}^\gamma &= 0\\ \Rightarrow \qquad \frac{1}{\sqrt{a}}\frac{\partial}{\partial \theta^\alpha}\left(F^{\alpha\gamma}\sqrt{a}\right) + F^{\alpha\beta}\Gamma^\gamma_{\alpha\beta} &= 0 \end{aligned}$$

The first equation in the claim now follows by using the equivalence of this quantity with the standard definition of covariant divergence: given a tensor Y and metric $a_{\alpha\beta}$

$$Y^{\alpha\beta}{}_{;\beta} \equiv Y^{\alpha\beta}{}_{,\beta} + \Gamma^\beta_{\delta\beta}Y^{\alpha\delta} + \Gamma^\alpha_{\delta\beta}Y^{\beta\delta} \equiv \frac{1}{\sqrt{a}}\frac{\partial}{\partial \theta^\beta}\left(Y^{\alpha\beta}\sqrt{a}\right) + Y^{\beta\delta}\Gamma^\alpha_{\beta\delta},$$

---

[10] The stress residual $F^{\alpha\beta}$ should not be confused with the deformation gradient F^i_M. F in the remainder of this chapter, and whole of the next, will always refer to the stress residual.

which is a simple relationship to prove.

The second result follows easily by taking the inner product of (8.16) with the deformed unit inward normal so that the term with the partial derivative of $F^{\alpha\beta}$ vanishes, noting that

$$(\mathbf{a}_1 \wedge \mathbf{a}_2) \cdot \hat{\mathbf{n}}^{\text{inward}} = -\|\mathbf{a}_1 \wedge \mathbf{a}_2\| = -\sqrt{a},$$

and using the definition of the second curvature tensor (8.4). $\square$

Chapter 9

Membrane Results and the Backward Problem for Skin

In Chapter 8 we derived the equations governing the deformation of one- and two-dimensional membranes in two and three-dimensional space, and determined how to convert the equations for a thin elastic solid to membrane equations. We now use these results to simulate the deformation of the breast. Firstly, in Section 9.1.1, we describe how the new equations are solved using the finite element method, which amounts to deriving the contributions to the finite element Jacobian and load vector arising from the new terms in the total energy. The results of simulations with the breast using a membrane model of skin are given in Section 9.1.2, which we relate to results obtained in Section 7.4.2.

The second part of this chapter, Section 9.2, is concerned with the backward problem for skin, where we finally drop the assumption that there are no residual stresses in the skin. We will see that fluid-filled membrane simulations lead us naturally to the conclusion that residual stresses are important and should not be neglected, or at least that there is little point in including skin in the model if residual stresses are not also accounted for, as was suggested by the results in Section 7.4.2. In this section we first introduce the complete backward problem for skin, and then attempt to approximate the solution by solving a simplified problem.

9.1 Numerical Results

9.1.1 The Finite Element Implementation

In this section we detail the calculations that have to be made when the finite element method is used to solve the new problems. We have to consider the contributions the new energies make to the finite element global load vector and global stiffness matrix. We will consider the fluid interior contribution and the elastic membrane contribution separately, since the membrane model can be used with either a fluid or an elastic interior.

This section is largely for completeness and for reference, and since the calculations here are technical and algebraic but do not provide insight into the models, the section may be skipped without loss of continuity. (Although the reader should perhaps note that (9.1), when compared to (8.8), shows how the 2D surface stress, containing a term which encompasses the pressure $p^{(-1)}$ in the skin, is computed in the new membrane equations).

Gravity-Loaded Fluid

The contribution to the energy is

$$\mathcal{K} = \frac{\lambda}{3}\iint_{\Theta} \mathbf{x}\cdot(\mathbf{a}_1\wedge\mathbf{a}_2)\,\mathrm{d}\theta^1\mathrm{d}\theta^2 - \lambda A_0 - \frac{\rho^{(f)}}{4}\iint_{\Theta}(\mathbf{g}\cdot\mathbf{x})\mathbf{x}\cdot(\mathbf{a}_1\wedge\mathbf{a}_2)\,\mathrm{d}\theta^1\mathrm{d}\theta^2,$$

where $A_0 = \frac{1}{3}\iint_{\Theta}\mathbf{X}\cdot(\mathbf{A}_1\wedge\mathbf{A}_2)\,\mathrm{d}\theta^1\mathrm{d}\theta^2$ is the initial 'volume' [1]. Let $\lambda' = \frac{\lambda}{3}$ and $\mathbf{g}' = \frac{\rho^{(f)}}{4}\mathbf{g}$, and let $\mathbf{n} = \mathbf{a}_1\wedge\mathbf{a}_2$ be the non-unit normal. The energy contribution is then

$$\mathcal{K} = \iint_{\Theta}(\lambda' - \mathbf{g}'\cdot\mathbf{x})\,\mathbf{x}\cdot\mathbf{n}\,\mathrm{d}\theta^1\mathrm{d}\theta^2 - \lambda A_0.$$

We need to calculate the directional derivative of $\mathcal{K}$.

$$D_{\delta\mathbf{x}}(\mathcal{K}) = \iint_{\Theta} -(\mathbf{g}'\cdot\delta\mathbf{x})(\mathbf{x}\cdot\mathbf{n}) + (\lambda' - \mathbf{g}'\cdot\mathbf{x})(\delta\mathbf{x}\cdot\mathbf{n} + \mathbf{x}\cdot\delta\mathbf{n})\,\mathrm{d}\theta^1\mathrm{d}\theta^2,$$

where

$$\begin{aligned}\delta\mathbf{n} := D_{\delta\mathbf{x}}(\mathbf{n}) &= \mathbf{a}_1\wedge\delta\mathbf{a}_2 + \delta\mathbf{a}_1\wedge\mathbf{a}_2\\ &= \frac{\partial\mathbf{x}}{\partial\theta^1}\wedge\frac{\partial(\delta\mathbf{x})}{\partial\theta^2} + \frac{\partial(\delta\mathbf{x})}{\partial\theta^1}\wedge\frac{\partial\mathbf{x}}{\partial\theta^2}.\end{aligned}$$

Let α_I be the I-th finite element variable, i.e. the I-th component of (4.12). (We use α_I instead of a_I in this chapter to avoid confusion with the metrics and tangent vectors). Then the contribution to the I-th finite element equation is given by

$$\mathrm{LoadVec}(\mathcal{K})_I := D_{\delta\mathbf{x}_I}(\mathcal{K}) = \iint_{\Theta} -(\mathbf{g}'\cdot\delta\mathbf{x}_I)(\mathbf{x}\cdot\mathbf{n}) + (\lambda' - \mathbf{g}'\cdot\mathbf{x})(\delta\mathbf{x}_I\cdot\mathbf{n} + \mathbf{x}\cdot\delta\mathbf{n}_I)\,\mathrm{d}\theta^1\mathrm{d}\theta^2,$$

---

[1] A_0 is not the true initial volume as it does not include a contribution to the integral over the fixed lower surface

where $\delta\mathbf{x}_I$ is the basis function for α_I.

We also need to compute the Jacobian of $\mathcal{K}$, given by

$$\mathrm{Jac}(\mathcal{K})_{IJ} = \frac{\partial \mathrm{LoadVec}(\mathcal{K})_I}{\partial \alpha_J}$$

This satisfies

$$\begin{aligned}\mathrm{Jac}(\mathcal{K})_{IJ} &= \iint_\Theta -(\mathbf{g}'\cdot\delta\mathbf{x}_I)\left(\frac{\partial\mathbf{x}}{\partial\alpha_J}\cdot\mathbf{n}+\mathbf{x}\cdot\frac{\partial\mathbf{n}}{\partial\alpha_J}\right)-\left(\mathbf{g}'\cdot\frac{\partial\mathbf{x}}{\partial\alpha_J}\right)(\delta\mathbf{x}_I\cdot\mathbf{n}+\mathbf{x}\cdot\delta\mathbf{n}_I)\\ &\quad+(\lambda'-\mathbf{g}'\cdot\mathbf{x})\left(\delta\mathbf{x}_I\cdot\frac{\partial\mathbf{n}}{\partial\alpha_J}+\frac{\partial\mathbf{x}}{\partial\alpha_J}\cdot\delta\mathbf{n}_I+\mathbf{x}\cdot\frac{\partial(\delta\mathbf{n}_I)}{\partial\alpha_J}\right)\mathrm{d}\theta^1\mathrm{d}\theta^2\end{aligned}$$

We also need to consider the constraint equation. It arises from the directional derivative with respect to λ

$$0 = D_{\delta\lambda} = \frac{1}{3}\iint_\Theta \mathbf{x}\cdot(\mathbf{a}_1\wedge\mathbf{a}_2)\,\mathrm{d}\theta^1\mathrm{d}\theta^2 - A_0.$$

Let I_λ be the index satisfying $\alpha_{I_\lambda}=\lambda$. This constraint equation is the I_λ-th finite element equation

$$\mathrm{LoadVec}_{I_\lambda} := \frac{1}{3}\iint_\Theta \mathbf{x}\cdot(\mathbf{a}_1\wedge\mathbf{a}_2)\,\mathrm{d}\theta^1\mathrm{d}\theta^2 - A_0 = 0.$$

The contribution this equation makes to the Jacobian is

$$\mathrm{Jac}(\mathcal{K})_{I_\lambda J} = \frac{\partial\mathrm{LoadVec}_{I_\lambda}}{\partial\alpha_J} = \frac{1}{3}\iint_\Theta \frac{\partial\mathbf{x}}{\partial\alpha_J}\cdot\mathbf{n}+\mathbf{x}\cdot\frac{\partial\mathbf{n}}{\partial\alpha_J}\,\mathrm{d}\theta^1\mathrm{d}\theta^2.$$

Finally, we need to include the component of the Jacobian which arises from differentiating the finite element equations with respect to λ. For $I\neq I_\lambda$:

$$\mathrm{Jac}(\mathcal{K})_{II_\lambda} = \frac{\partial\mathrm{LoadVec}_I}{\partial\alpha_J} = \frac{\partial\mathrm{LoadVec}_I}{\partial\lambda} = \frac{1}{3}\iint_\Theta \delta\mathbf{x}_I\cdot\mathbf{n}+\mathbf{x}\cdot\delta\mathbf{n}_I\,\mathrm{d}\theta^1\mathrm{d}\theta^2,$$

and for $I=I_\lambda$

$$\mathrm{Jac}(\mathcal{K})_{I_\lambda I_\lambda} = \frac{\partial\mathrm{LoadVec}_{I_\lambda}}{\partial\lambda} = 0.$$

Integrated Thin Elastic Skin

The contribution to the energy by a thin elastic skin is

$$\begin{aligned}\mathcal{I} &= \iiint_{\Omega_0} W(I_1^{(0)},I_2^{(0)},I_3^{(0)})\,\mathrm{d}V_0\\ &= h_0\iint_\Theta W^{\mathrm{mem}}(a_{\alpha\beta})\sqrt{A}\,\mathrm{d}\theta^1\mathrm{d}\theta^2.\end{aligned}$$

The directional derivative of $\mathcal{I}$ is

$$\begin{aligned}D_{\delta\mathbf{x}}(\mathcal{I}) &= h_0\iint_\Theta \frac{\partial W^{\mathrm{mem}}}{\partial a_{\alpha\beta}}\delta a_{\alpha\beta}\sqrt{A}\ \mathrm{d}\theta^1\mathrm{d}\theta^2\\ &= h_0\iint_\Theta \frac{1}{2}T^{\alpha\beta}\delta a_{\alpha\beta}\sqrt{A}\ \mathrm{d}\theta^1\mathrm{d}\theta^2.\end{aligned}$$

Now, $a_{\alpha\beta} = \mathbf{a}_\alpha \cdot \mathbf{a}_\beta$ so

$$\delta a_{\alpha\beta} = D_{\delta\mathbf{x}}(\mathbf{a}_\alpha \cdot \mathbf{a}_\beta) = \mathbf{a}_\alpha \cdot \delta\mathbf{a}_\beta + \mathbf{a}_\beta \cdot \delta\mathbf{a}_\alpha,$$

where $\delta\mathbf{a}_\alpha = \frac{\partial(\delta\mathbf{x})}{\partial\theta^\alpha}$. Since this quantity is summed with $T^{\alpha\beta}$, which is symmetric, it can be replaced by $2\mathbf{a}_\alpha \cdot \delta\mathbf{a}_\beta$, so

$$D_{\delta\mathbf{x}}(\mathcal{I}) = h_0 \iint_\Theta T^{\alpha\beta}\mathbf{a}_\alpha \cdot \delta\mathbf{a}_\beta \sqrt{A}\, \mathrm{d}\theta^1 \mathrm{d}\theta^2.$$

Now, let us assume incompressibility, which means the full strain energy is of the form $W(I_1, I_2) - \frac{p}{2}(I_3 - 1)$. However we know the strain energy is a function of the metric $a_{\alpha\beta}$ and the normal stretch (squared) $\tilde{a}_{33}$, and we know the relationship between $a_{\alpha\beta}$ and $\tilde{a}_{33}$ from the incompressibility constraint $a\tilde{a}_{33} = A$. So, if we consider $\tilde{a}_{33}$ as a function of the surface metric in W^{mem}, we are manually imposing the constraint, and the pressure term in the strain energy is no longer required. As long as we use $W^{\mathrm{mem}}(a_{\alpha\beta})$ with $a_{\alpha\beta}$ and $\tilde{a}_{33}$ *not* independent, we just need to use the material part of the strain energy $W(I_1, I_2)$.

We need to compute the strain invariants in terms of $a_{\alpha\beta}$ and $A_{\alpha\beta}$ (i.e. not in terms of $\tilde{a}_{\alpha\beta}$ and $\tilde{A}_{\alpha\beta}$):

$$\begin{aligned} I_1 &= \tilde{a}_{\alpha\beta}\tilde{A}^{\alpha\beta} \\ &= a_{\alpha\beta}A^{\alpha\beta} + \tilde{a}_{33} \\ &= a_{\alpha\beta}A^{\alpha\beta} + \frac{A}{a} \end{aligned}$$

(Compare with (8.6)). A similar calculation gives

$$I_2 = \frac{1}{2}a_{\alpha\beta}a_{\gamma\delta}\left(A^{\alpha\beta}A^{\gamma\delta} - A^{\alpha\gamma}A^{\beta\delta}\right) + \frac{A}{a}a_{\alpha\beta}A^{\alpha\beta},$$

from which the surface stress can be computed be (c.f. (8.8))

$$\begin{aligned} T^{\alpha\beta} &= 2\frac{\partial W}{\partial I_1}\frac{\partial I_1}{\partial a_{\alpha\beta}} + 2\frac{\partial W}{\partial I_2}\frac{\partial I_2}{\partial a_{\alpha\beta}} \\ &= 2\frac{\partial W}{\partial I_1}\left(A^{\alpha\beta} - \frac{A}{a}a^{\alpha\beta}\right) + 2\frac{\partial W}{\partial I_2}\left(a_{\gamma\delta}\left(A^{\alpha\beta}A^{\gamma\delta} - A^{\alpha\gamma}A^{\beta\delta}\right) - \frac{A}{a}a^{\alpha\beta}a_{\gamma\delta}A^{\gamma\delta} + \frac{A}{a}A^{\alpha\beta}\right). \end{aligned} \tag{9.1}$$

With this form of $T^{\alpha\beta}$, the contribution to the I-th finite element equation is

$$\mathrm{LoadVec}(\mathcal{I})_I := D_{\delta\mathbf{x}_I}(\mathcal{I}) = h_0 \iint_\Theta T^{\alpha\beta}\mathbf{a}_\alpha \cdot (\delta\mathbf{a}_\beta)_I \sqrt{A}\, \mathrm{d}\theta^1 \mathrm{d}\theta^2,$$

where $\delta\mathbf{x}_I$ is chosen to be the basis function for the I-th finite element variable α_I, and $(\delta\mathbf{a}_\alpha)_I = \frac{\partial(\delta\mathbf{x}_I)}{\partial\theta^\alpha}$.

To calculate the Jacobian we need to differentiate this with respect to the J-th finite element variable α_J

$$\begin{aligned}\mathrm{Jac}(\mathcal{I})_{IJ} &= \frac{\partial(\mathrm{LoadVec}(\mathcal{I})_I)}{\partial\alpha_J} \\ &= h_0 \iint_\Theta \left(T^{\alpha\beta}\frac{\partial\mathbf{a}_\alpha}{\partial\alpha_J}\cdot(\delta\mathbf{a}_\beta)_I + \frac{\partial T^{\alpha\beta}}{\partial a_{\gamma\delta}}\frac{\partial a_{\gamma\delta}}{\partial\alpha_J}\mathbf{a}_\alpha\cdot(\delta\mathbf{a}_\beta)_I \right)\sqrt{A}\;\mathrm{d}\theta^1\mathrm{d}\theta^2\end{aligned}$$

The final calculation which needs to be made is $\frac{\partial T^{\alpha\beta}}{\partial a_{\gamma\delta}}$, for which the following results can be used.

$$\begin{aligned}\frac{\partial^2 I_1}{\partial a_{\alpha\beta}\partial a_{\gamma\delta}} &= \frac{A}{a}\left(a^{\alpha\beta}a^{\gamma\delta} + a^{\alpha\gamma}a^{\beta\delta}\right) \\ \frac{\partial^2 I_2}{\partial a_{\alpha\beta}\partial a_{\gamma\delta}} &= \left(A^{\alpha\beta}A^{\gamma\delta} - A^{\alpha\gamma}A^{\beta\delta}\right) + \frac{A}{a}a_{\xi\nu}A^{\xi\nu}\left(a^{\alpha\beta}a^{\gamma\delta} + a^{\alpha\gamma}a^{\beta\delta}\right) - \frac{A}{a}\left(a^{\alpha\beta}A^{\gamma\delta} + a^{\gamma\delta}A^{\alpha\beta}\right)\end{aligned}$$

9.1.2 Results

In Table 9.1 we repeat the simulations carried out for Table 7.8 using the membrane model of skin instead of a thin 3D elastic skin. For these simulations we use meshes without any thin skin elements, instead including the membrane surface integral contribution to the total energy. Comparing the results for the membrane with the thin-skin results, we see that the new 2D method for modelling skin is a significant improvement over the previous 3D method, both in terms of computation time and the number of Newton iterations required.

Figure 9.1 displays the results of two simulations, one using the thin skin model (Figure 9.1(a)) and one using the membrane model (Figure 9.1(b)), together with the initial undeformed state (Figure 9.1(c)). The two deformed states match closely, partially verifying that the membrane model agrees with the thin solid model. However, we saw in Chapter 7 that the deformations are very insensitive to skin material parameters, so the fact that the simulations match cannot be taken as proof that the models agree.

Next, we can attempt to simulate the deformation of the breast using a fluid-filled membrane model. This involves ignoring the elasticity in the interior of the breast, supposing instead that the interior just consists of a single incompressible fluid, and just modelling the surface deformation. The motivation for modelling the breast as a fluid-filled membrane is as follows. Suppose the deformation of the breast under gravity turned out to be well-approximated by the fluid-filled membrane model. This would firstly give an indication of the extent that skin controls the deformation. It also suggests a new numerical approach for computing breast deformations: rather than solving the full 3D problem, we first compute the much less computationally-intensive 2D fluid-filled membrane problem to obtain an approximation of the deformation of the skin surface, and then use this information as a boundary condition or as part of a more informed first guess in a subsequent 3D computation of the displacement in the interior.

Mesh	Computation Time (s)	Number Newton Iterations
288 elements, thin skin elements	358	19
288 elements, no skin elements	71	6
288 elements, membrane skin	100	6
512 elements, thin skin elements	434	13
512 elements, no skin elements	334	10
512 elements, membrane skin	314	9

Table 9.1: Comparison of the computation time and the number of Newton steps taken, for meshes with thin skin elements/no thin skin elements/membrane skin.

To compute fluid-filled membrane simulations, we create a surface mesh of the skin by taking the front surface of the 3D mesh, assume zero-displacements boundary conditions on the (1D) boundary of the skin, and model the deformation as a membrane enclosing an incompressible gravity-loaded fluid. However, when we try to simulate the deformed state of the membrane with gravity applied in some direction, the finite element method does not converge, even with tiny magnitudes of the gravity vector. The reason for this is that we are assuming the starting state for the membrane is the reference state, which means that there are no stresses in the undeformed starting state, i.e. that the undeformed starting state is of an *unstretched* membrane covering a fluid interior. It is intuitively obvious that trying to compute the deformation when varying the direction of gravity without increasing the internal pressure of such an object may not be a well-posed problem.

This issue does not prevent the numerical method converging when we use a membrane enclosing an elastic interior, because the non-zero stiffness in the interior uniquely determines the displacement. However, this does not mean there is no problem with the modelling. The fact that simulations can be computed in this case hides the fact that the skin is still being modelled as an unstretched membrane lain over an elastic interior, which is physically unrealistic.

In the last part of this thesis we consider how to model skin as a *pre-stressed membrane*. This means we have to determine the stresses in the skin in the initial state. To simplify matters, we decouple the full backward problem from the gravity-loading backward problem, and consider the skin-only backward problem. For this, we assume we know a configuration for the breast which is not loaded by gravity or any other external forces, but for which there are residual stresses in the skin. The questions we seek to answer are how we determine these stresses, what the undeformed shape of the membrane is, if it exists, and how the forward problem can be implemented to account for the residual stresses.

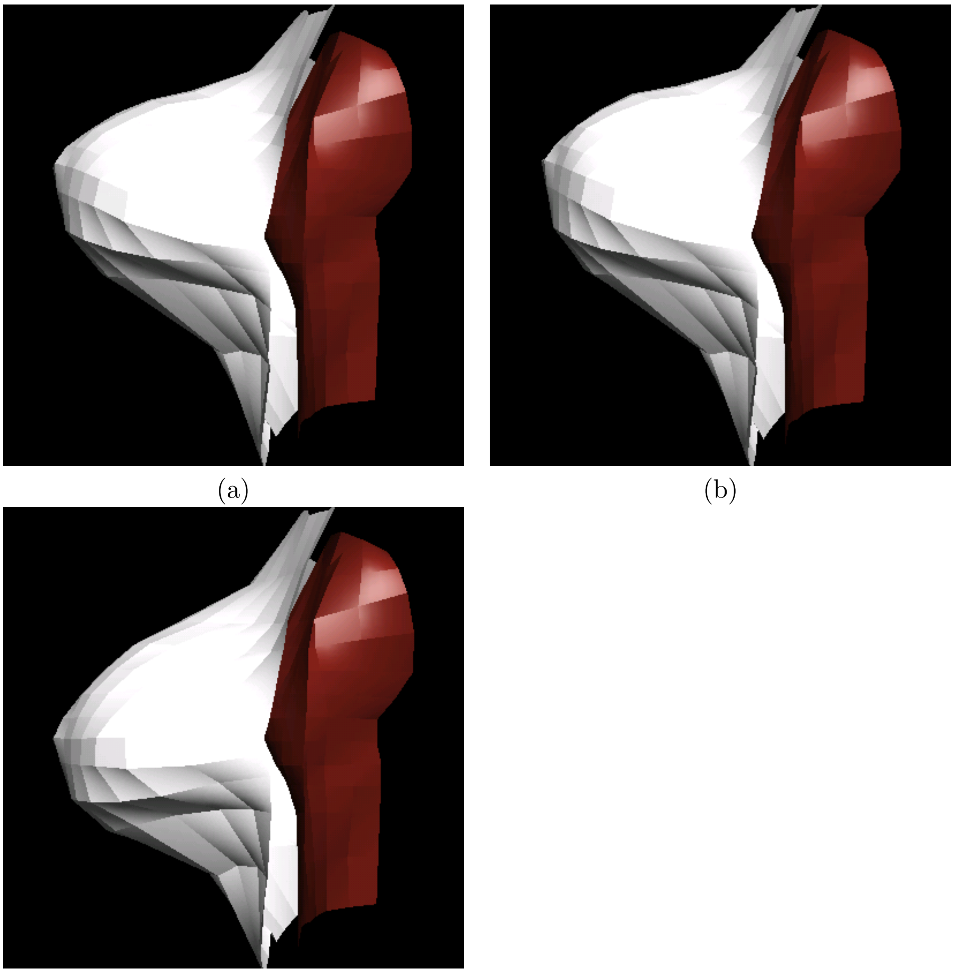

(a) (b)

(c)

Figure 9.1: Comparison of results using the two different methods to model skin: (a) thin skin model, (b) membrane skin model. Both simulations are a forward simulation with gravity applied upwards, from the starting state shown in (c).

9.2 The Backward Problem for Skin

9.2.1 Introduction

Let us begin by clarifying the different possible backward problems for the breast. The breast as imaged in MRI contains internal stresses which arise from at least two and potentially three or more different sources. The first is loading by gravity. In Chapter 5 we assumed this was the only cause of internal stresses and determined how to compute those stresses by explicitly computing a true reference state. The second cause of internal stresses is the pre-stretching of the skin *in vivo*, due to different rates of growth in different tissues. This leads not only to residual stresses in the skin, but to stresses in the interior, which must be present to counter the tensile forces in the skin. Further possible causes of stress include contact with the MR coils, in which case there would be normal surface forces on the skin, and other residual stresses in the interior.

If we allowed for the effects of gravity, tension in the skin and contact with the MR coils (as shown in Figure 9.2(a)), we would have to deal with an extremely complex *backward-contact* problem. This is well beyond the scope of this thesis. The next scenario, which we refer to as the *full backward problem*, is that of the breast loaded by gravity and containing residual stresses in the skin (Figure 9.2(b)). This again is well beyond the scope of this thesis, and for the remainder of this chapter we concentrate on what we call the *backward problem for skin*, where we assume the known shape of the breast is one which has a pre-stressed skin but is not acted upon by gravity, as shown in Figure 9.2(c).

In this case the breast is not under any body or surface forces, but internal stresses are present, so a true reference state cannot exist, and an inverse elasticity calculation analogous to that made in Chapter 5 is not possible. A method of dealing with such problems (see [Ols04] for details) is to assume there is a disconnected stress-free state, or in other words a set of locally stress-free configurations, which do not in general form a globally stress-free state. This method is based on the fact that any point can be made stress-free by relieving it from the main body (i.e. by cutting). Using such a method we can compute strains or deformation gradients from the disconnected stress-free state to the known stressed state, although such strains or deformation gradients will not arise from derivatives of a displacement, i.e. they are *not compatible*, in the sense described in Section 5.1.1.

For the breast we could assume there is a disconnected stress-free state which consists of two separated components: an unstressed skin and an unstressed interior. Since the skin is under tension in the residually-stressed state it will have shrunk in its unstressed case, whereas the interior, being under compression in the stressed state, will expand. This is illustrated in Figure 9.3. To solve the backward problem for skin we would need to determine both of these reference states, whilst also considering how the skin and the interior are connected in the stressed state.

This is obviously another extremely complex problem, and also beyond the scope of this thesis. Instead, we make another simplification, and use a fluid-filled interior instead of an elastic one, as shown in Figure 9.2(d). We assume we know the initial stressed state of the membrane, and the initial fluid pressure $P > 0$, and we wish to determine the unstressed membrane shape (if one exists), which corresponds to reducing the pressure to $P = 0$. The residual stresses in the fluid-filled membrane could then potentially be used to approximate the residual stresses in the membrane-with-elastic-interior model, if we can determine a method of choosing a suitable initial pressure P.

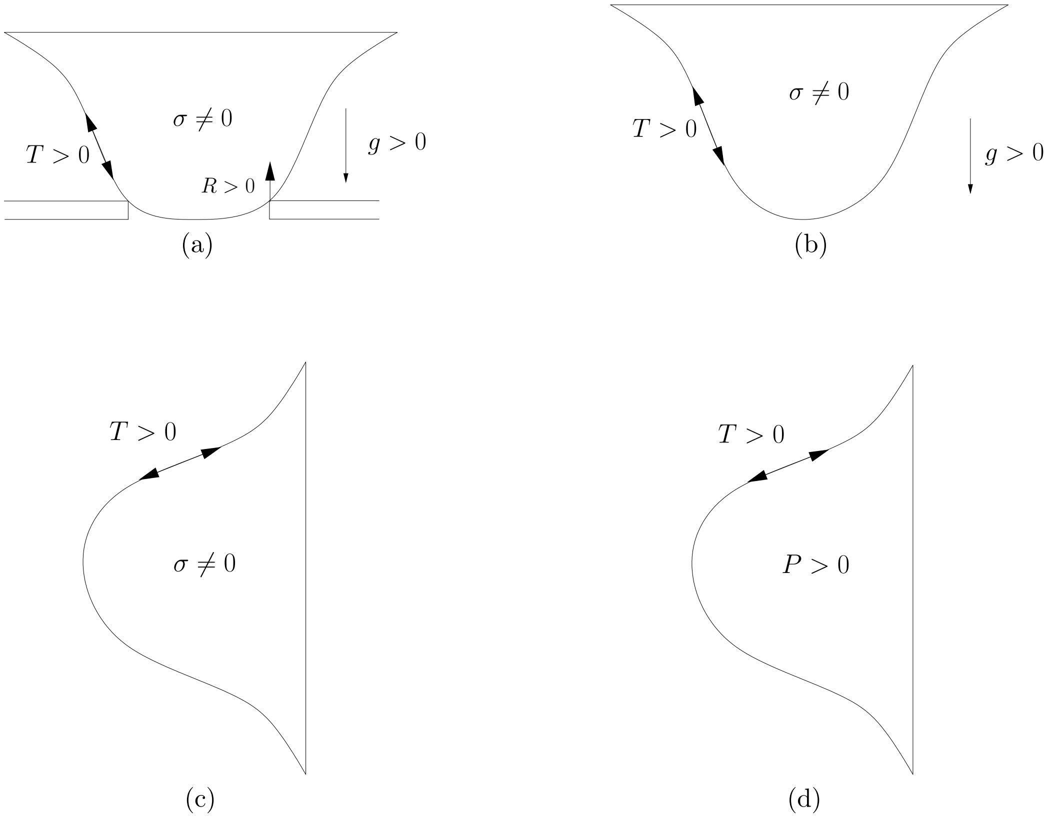

Figure 9.2: The various possible different backward problems. (a) Full backward-contact problem, where the breast is under the effect of gravity g, contains a tension T in the skin, and is in contact with the MR coils (which introduces a surface reaction force R), all of which must be balanced by non-zero internal stresses σ. (b) Full backward problem, with gravity and residual stress in the skin. (c) Skin-only backward problem, where there is no gravity. (d) Skin-only backward problem with an incompressible fluid interior instead of an elastic interior, where the residual stresses are balanced by a single fluid pressure P.

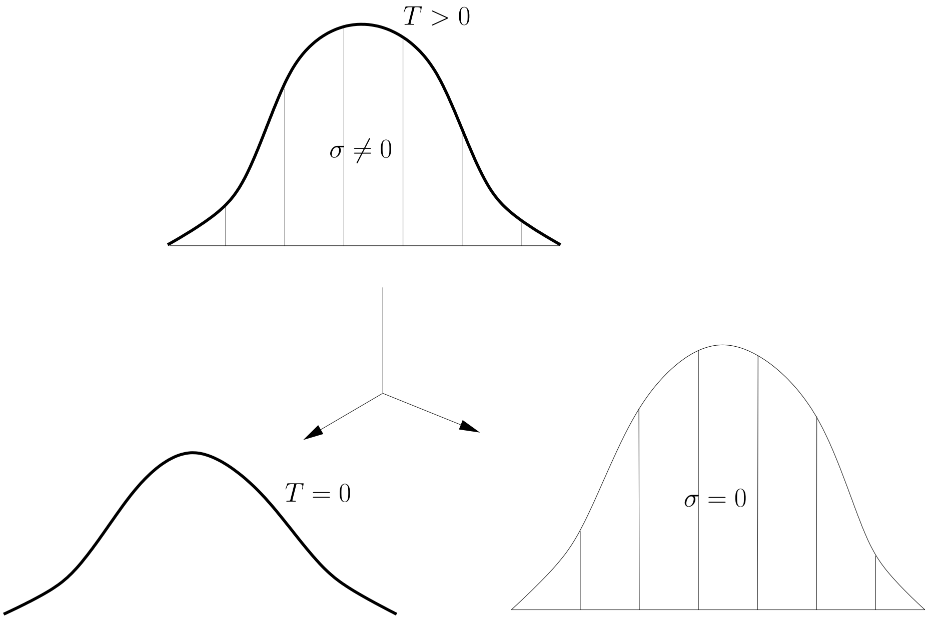

Figure 9.3: A disconnected reference state for the skin-only backward problem: the initial stressed state (top) maps to a shrunken unstressed skin (bottom left) and separate, expanded, unstressed interior (bottom right).

9.2.2 The Backward Problem for Skin using the Fluid Filled Membrane Model

Suppose we have a membrane enclosing an incompressible fluid under a known pressure, with no body or surface forces, in particular no gravity or contact. This is our deformed starting state, denoted by Ω. We wish to compute the stresses in the membrane, determine whether an undeformed state (the shape of the membrane under zero pressure) exists, and compute it if so.

First, let us make a very simple calculation. Consider the 2D case of a string and suppose the deformed state is an arc of a circle with constant curvature κ, and that we know that the internal pressure is P. Using the results and notation from Section 8.2.1, suppose the strain energy is $W = \frac{1}{2}c_1(\|\mathbf{t}\| - 1)^2$. Then the constant tension in the string is $T = c_1(\|\mathbf{t}\| - 1)$, but we know $T\kappa = P$, so therefore the constant stretch is $\|\mathbf{t}\| = 1 + \frac{P}{\kappa c_1}$. Since we know the deformed shape, we know the deformed length (L, say) of the curve, and this stretch can be used to calculate the undeformed ('natural') length of the string, $L_0 = L/\|\mathbf{t}\|$.

Note that in this case the undeformed shape can be any shape with that natural length, since the one-dimensional string equations are independent of the string's starting shape, depending only on the starting length. There are therefore an infinite number of potential solutions for the undeformed shape, although none may exist if position boundary conditions are prescribed. However, also note that one particular zero-stress state is picked out by taking the limit as the pressure $P \to 0$ (with appropriate boundary conditions), i.e. by continuously deflating the membrane.

This analysis required assumptions that the deformed shape took a specific form (an arc of a circle), and that the undeformed state existed. Also, in 3D such a calculation will not be possible, so let us consider how to deal with this problem in general.

It is useful to unify the 2D analysis with the 3D analysis, and write the equations derived in Section 8.2.1 for the string in the same form as the equations for the membrane in 3D. Suppose we have a general parameterisation of $\mathbf{X} \equiv \mathbf{X}(\theta)$ for the string. We need to define metric tensors for the undeformed and deformed strings. These metric tensors are one-dimensional (i.e. scalars), and it is easy to see that they satisfy

$$A_{11} = \left\| \frac{\partial \mathbf{X}}{\partial \theta} \right\|^2 \qquad \sqrt{A} = \left\| \frac{\partial \mathbf{X}}{\partial \theta} \right\| \qquad \text{and} \qquad a_{11} = \left\| \frac{\partial \mathbf{x}}{\partial \theta} \right\|^2 \qquad \sqrt{a} = \left\| \frac{\partial \mathbf{x}}{\partial \theta} \right\|$$

Now, the stretching $\|\mathbf{t}\|$ from the undeformed state to the deformed state generalises to the ratio of the determinants of the metric tensors, so

$$W(\|\mathbf{t}\|) \equiv W\left(\frac{\sqrt{a}}{\sqrt{A}}\right).$$

We can now write the equations determining the string in 2D space or the membrane in 3D space, for any strain energy function (thin solid, string, etc) in the form

$$\min \int_\Theta W(a_{\alpha\beta}, A_{\alpha\beta}, a, A)\sqrt{A}\,\mathrm{d}\Theta + \frac{\lambda}{N}\int_\Theta \mathbf{x}\cdot\mathbf{n}\,\mathrm{d}\Theta$$

where N is the dimension, and the second term is just the area or volume under the membrane. The Euler-Lagrange equations

$$F^{\alpha\beta}{}_{;\beta} = 0 \qquad (9.2)$$

$$F^{\alpha\beta} b_{\alpha\beta} = P \qquad (9.3)$$

must then hold for either dimension, where the indices only range over 1 for the string. If the parameterisation is undeformed arclength in 2D, these equations reduce to tension T is constant and $T\kappa = P$.

Now, consider the backward problem, where we know the deformed state, which means we know $\mathbf{x}$ and $b_{\alpha\beta}$, and we know $\mathbf{x}$ satisfies (9.2) and (9.3). Note that these equations do not need an undeformed state $\mathbf{X}$ to exist in order for a solution $\mathbf{x}$ to exist, they just require $A_{\alpha\beta}$ to be defined. This is a crucial observation: to solve the backward problem for a fluid-filled membrane, it is only necessary to solve for the undeformed metric tensor $A_{\alpha\beta}$, *it is not* necessary to solve for the full undeformed positions $\mathbf{X}$. Also, for the same reasons, subsequent forward problems only require the undeformed metric to be known. So we immediately see that the full computation of the undeformed shape of the membrane is not needed, we only need to compute the undeformed metric. In the context of the simple calculation above, we only need to compute the undeformed length $L/\|\mathbf{t}\|$, we do not need to compute the undeformed position $\mathbf{X}$ (of which there were an infinite number of solutions anyway).

Equations (9.2) and (9.3) arose as two/three (depending on dimension) partial differential equations for the two/three components of the unknown $\mathbf{x}$. For the backward problem the situation is different. Here $\mathbf{x}$ is known and $A_{\alpha\beta}$ is the unknown, and because $F^{\alpha\beta}$ is an algebraic function of $A_{\alpha\beta}$, equations (9.2) and (9.3) become a set of one/two differential equations (9.2), together with a single algebraic equation (9.3), for $A_{\alpha\beta}$.

At this point 2D and 3D become very different, and have to be considered separately. In 2D we have two equations (one differential and one algebraic) for one unknown A_{11}, or equivalently, one unknown $\sqrt{A}$. This is an over-constrained problem and only has a solution if the algebraic equation is consistent with the differential equation. For example, if we choose a sensible parameterisation such as one that is uniform across the deformed curve (i.e. is a multiple of deformed arclength), the covariant derivative reduces to a partial derivative, so the two equations become

$$F^{11} = 2\frac{\partial W}{\partial a_{11}}\frac{\sqrt{A}}{\sqrt{a}} = \text{ const} \qquad\qquad \text{and} \qquad\qquad F^{11}b_{11} = P$$

These are only consistent if b_{11} is constant, i.e. if the deformed state is an arc of a circle, as we would expect. In this case $F^{11}(\sqrt{A})b_{11} = P$ becomes a single algebraic equation for the unknown $\sqrt{A}$. We have therefore shown that the backward problem for a fluid-filled membrane in 2D only has a solution if the deformed shape is an arc of a circle.

In 3D, there are three equations and also 3 unknowns A_{11}, A_{12} and A_{22} (since $A_{\alpha\beta}$ is symmetric). Unlike in 2D, here we do not have an over-constrained problem, and the equations do appear to have to be consistent. In this case, it is better to think of (9.2) and (9.3) as two differential equations and one algebraic equation for the (generalised) in-plane forces $F^{\alpha\beta}$, of which there are again three independent components due to symmetry. Once $F^{\alpha\beta}$ has been calculated, $A_{\alpha\beta}$ can be computed by inverting the nonlinear relationship $F^{\alpha\beta}(A_{\alpha\beta}) = 2h_0\frac{\partial W(a_{\alpha\beta},A_{\alpha\beta})}{\partial a_{\alpha\beta}}\frac{\sqrt{A}}{\sqrt{a}}$.

Assuming the curvature tensor is not zero anywhere we can use (9.3) to eliminate one of the components of $F^{\alpha\beta}$ from the two PDEs, so that (9.2) becomes two differential equations for two unknowns. We see that, unlike in two-dimensions, we need to specify further boundary conditions on $F^{\alpha\beta}$ (or equivalently, on $A_{\alpha\beta}$) in order for the solution to be uniquely determined. We will investigate suitable boundary conditions on the undeformed metric by studying an axisymmetric test case in Section 9.2.5. Before that however, we discuss as an aside the problem of determining the undeformed position $\mathbf{X}$ from the undeformed metric (recall though that this is a problem that would never have to be solved in practical simulations), and then classify the PDEs as elliptic, parabolic or hyperbolic.

9.2.3 Existence of an Undeformed State

The question of the existence of an undeformed state is loosely related to the question of the compatibility of strains. Recall that in Section 5.1.1 we described how a strain is compatible if there is a deformation giving rise to that strain. Equivalently, we can talk about deformation gradients being compatible. In Section 5.1.1 we stated the compatibility condition for linear elasticity. We now have the necessary mathematical framework to be able to state the compatibility condition for finite elasticity. Suppose we have a simply-connected deformed body in 3D, on which we know the strains E_{MN}. Suppose the equivalent Green-Lagrange tensor $C_{MN} = 2(E_{MN} + \delta_{MN})$ is symmetric and positive-definite (if this is not true the strains are certainly not compatible). C_{MN} therefore defines a metric on the space. This metric induces 3D Christoffel symbols, defined analogously to (8.3) using C_{MN} instead of $a_{\alpha\beta}$, on the space. It can be shown (see [Blu89]) that the strain is compatible if and only if the *Riemann-Christoffel curvature tensor*, a fourth-order tensor given by

$$R^M_{NPQ} = \Gamma^M_{NQ,P} - \Gamma^M_{NP,Q} + \Gamma^R_{NQ}\Gamma^M_{RP} - \Gamma^R_{NP}\Gamma^M_{RQ},$$

is zero.

The question of compatibility here was essentially a differential geometry question: given a 3D metric on a 3D space (such a pair is known as a 3D *manifold*), is it possible to define a mapping from that space with that metric to Euclidean space with the identity metric, which is also an isometry (length-preserving)[2]. The question of the existence of an undeformed state for the membrane is also essentially a differential geometry question. Here we have a 2D manifold $(\Theta, A_{\alpha\beta})$. Rather than map this into 2D space, we wish to find an isometric map into 3D space, i.e. isometrically embed the manifold in 3D. This extra freedom is apparent through the fact that here we have the same number of equations as unknowns: three equations (the three independent components of $A_{\alpha\beta}$) for the three unknown components of $\mathbf{X}$, whereas in the previous scenario there were 6 equations (the independent components of C_{MN}) for three unknowns, hence the need for compatibility conditions.

In general a 2D manifold cannot be isometrically embedded in 3D space, although it can be shown that any m-D manifold can be isometrically embedded in $2m$-D space. However, we have the extra information here that the manifold initially arose as a surface in 3D, i.e. there exists an embedding, namely the map $\boldsymbol{\theta} \mapsto \mathbf{x}(\boldsymbol{\theta})$, from $(\Theta, A_{\alpha\beta})$ to $\mathbb{R}^3$. This is not an isometry (as the induced metric, which of course is $a_{\alpha\beta}$, is not equal to $A_{\alpha\beta}$), but it is enough to prove the existence of an isometric embedding, if we use the *Nash-Kuiper embedding theorem* ([Nas55], [Kui55]). This states that, loosely-speaking, if there exists a smooth embedding from a manifold of dimension m to $\mathbb{R}^n$ $(n > m)$, then it is possible to construct a C^1-embedding into $\mathbb{R}^n$ which is (arbitrarily close to) an

---

[2] A smooth mapping from a manifold (M^m, g) (where M^m is a m-dimensional space, g a metric) to $\mathbb{R}^m$ induces a second metric on M^m, given by the Euclidean inner product of first derivatives of the mapping. If the induced metric is equal to g, the mapping is length-preserving and called an isometry

isometry. Therefore, if we assume the deformed surface is smooth, the theorem allows us to say that a continuously-differentiable undeformed state $\mathbf{X}(\boldsymbol{\theta})$ equivalent to the metric $A_{\alpha\beta}$ exists. $\mathbf{X}(\boldsymbol{\theta})$ can then be determined, up to constants of integration, by solving the three PDEs $\frac{\partial \mathbf{X}}{\partial \theta^\alpha} \cdot \frac{\partial \mathbf{X}}{\partial \theta^\beta} = A_{\alpha\beta}$ for the three unknowns.

9.2.4 Classification

Let us recall the theory regarding the classification of a system of first-order PDEs. Suppose $\mathbf{v} \equiv \mathbf{v}(x, y)$ is a two-dimensional vector function of two variables, satisfying a system of PDEs

$$A\mathbf{v}_{,x} + B\mathbf{v}_{,y} + \mathbf{f}(\mathbf{v}, x, y) = 0,$$

where A and B are matrices. Then the PDEs are classified based on the *generalised eigenvalues* of A and B, which are defined to be the values $\lambda \in \mathbb{C}$ satisfying $\det(A - \lambda B) = 0$. The PDEs are elliptic at points in space where the two roots of $\det(A - \lambda B) = 0$ are complex, hyperbolic where the roots are real and distinct, and parabolic where the roots are equal.

CLAIM: Let K be the Gaussian curvature of the deformed state of the membrane, and suppose $b_{\alpha\beta} \neq 0$ everywhere. The equations determining $A_{\alpha\beta}$ are: elliptic if $K > 0$, parabolic of $K = 0$ and hyperbolic is $K < 0$.

PROOF: Since $F^{\alpha\beta}$ is related to $A_{\alpha\beta}$ via an algebraic relationship, we just need to classify the PDEs governing $F^{\alpha\beta}$. Since $b_{\alpha\beta} \neq 0$, one of b_{11}, b_{12} or b_{22} is non-zero. Suppose initially $b_{22} \neq 0$. Then the pressure equation $F^{\alpha\beta} b_{\alpha\beta} = P$ is equivalent to

$$F^{22} = \frac{P}{b_{22}} - \left(\frac{b_{11}}{b_{22}}\right) F^{11} - 2\left(\frac{b_{12}}{b_{22}}\right) F^{12},$$

which can be used to eliminate F^{22} from $F^{\alpha\beta}{}_{;\beta} = 0$, giving

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} F^{11} \\ F^{12} \end{bmatrix}_{,\theta^1} + \begin{bmatrix} 0 & 1 \\ -\left(\frac{b_{11}}{b_{22}}\right) & -2\left(\frac{b_{12}}{b_{22}}\right) \end{bmatrix} \begin{bmatrix} F^{11} \\ F^{12} \end{bmatrix}_{,\theta^2} + \mathbf{f}(F^{11}, F^{12}, \theta^1, \theta^2, P) = 0,$$

where $\mathbf{f}$ is a nonlinear function containing the Christoffel symbols. The PDE is therefore classified based on the roots of

$$\det(A - \lambda B) = \det \begin{bmatrix} 1 & -\lambda \\ \lambda\left(\frac{b_{11}}{b_{22}}\right) & 1 + 2\lambda\left(\frac{b_{12}}{b_{22}}\right) \end{bmatrix} = 1 + 2\left(\frac{b_{12}}{b_{22}}\right)\lambda + \left(\frac{b_{11}}{b_{22}}\right)\lambda^2,$$

which is a quadratic with discriminant $-\frac{4}{(b_{22})^2}\det(b_{\alpha\beta})$. The PDEs are elliptic if both roots of the quadratic are complex, i.e. if the discriminant is negative, i.e. if $\det(b_{\alpha\beta})$ is positive, which is equivalent to K being positive. The PDE is hyperbolic if both roots are distinct and real, i.e. if $\det(b_{\alpha\beta})$ (or equivalently K) is negative, and parabolic if the roots are equal, i.e. $\det(b_{\alpha\beta}) = 0$.

By symmetry, the same result holds if $b_{11} \neq 0$. A similar calculation shows that if $b_{12} \neq 0$, and we eliminate F^{12} instead of F^{22}, the discriminant is $-\frac{4}{(b_{12})^2}\det(b_{\alpha\beta})$, so the result again holds. $\square$

The sign of the Gaussian curvature of a surface indicates whether the inward pointing directions for each of the principal curvatures are aligned. A spheroid has positive Gaussian curvature, whereas a hyperboloid has negative Gaussian curvature. In between these, a cylinder has zero Gaussian curvature, being flat in one direction. The breast surface geometry is that of an inner spheroid-type region, connected to an outer hyperboloid-type region, as shown in Figure 9.4. Therefore the equations for $F^{\alpha\beta}$ for a fluid-filled membrane taking the shape of a realistic breast geometry are mixed hyperbolic-elliptic, being hyperbolic on the outer region (which includes the boundary), and elliptic in the inner region.

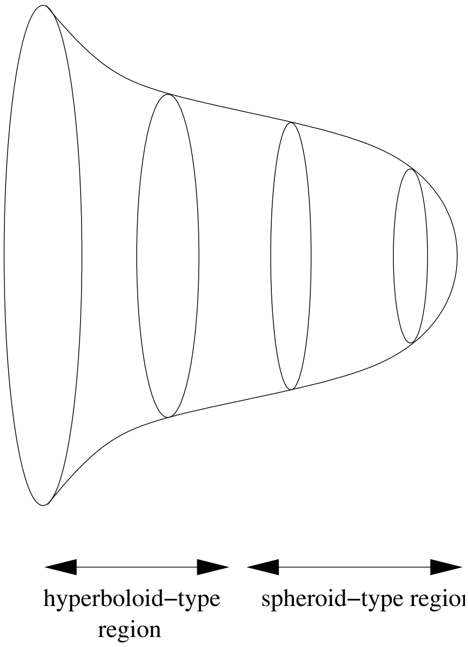

Figure 9.4: The breast surface geometry is of the form of an inner spheroid-type region (positive Gaussian curvature) and an outer hyperboloid-type region (negative Gaussian curvature).

9.2.5 The Axi-symmetric Case

Now, let us consider a membrane which has an axisymmetric deformed state. This is primarily in order to gain an insight into mathematical features of the equations we have derived, but also to show an example of how the equations are formed and can be solved in practise. Let us take the axisymmetric deformed membrane, as shown in Figure 9.5, as the known configuration of the body. Parameterising the surface with the azimuthal angle θ and the axial distance z, a point on the surface in given by

$$\mathbf{x}(\theta, z) = r(z)\mathbf{e}_r + z\mathbf{e}_z,$$

where $\theta \in [0, 2\pi]$, $z \in [0, 1]$ (say). The parameter space is therefore $\Theta = [0, 2\pi] \times [0, 1]$, with $\boldsymbol{\theta} = (\theta, z)$. Here $\mathbf{e}_r$ and $\mathbf{e}_z$ are the standard cylindrical basis vectors: $\mathbf{e}_r = (\cos\theta, \sin\theta, 0)$, $\mathbf{e}_z = (0, 0, 1)$. The functional form of $r(z)$ determines the shape of the surface. For example, for a hemisphere $r(z) = \sqrt{1 - z^2}$, and for a cylinder $r(z) = \text{const}$.

Note that there are two different classes of axisymmetric surface. There are surfaces which have only one bounding curve (for example, the hemisphere), for which $r(1)$ must be zero. In this case, we must also have $r' \to -\infty$ as $z \to 1$ for the surface to be smooth (i.e. not have a sharp point at $z = 1$). The second class of axisymmetric surface are those with two bounding curves at $z = 0$ and $z = 1$ (for example, the cylinder). For these $r(1) > 0$.

Assuming an undeformed state exists and is also axisymmetric, we can use the parameter space to also express the undeformed position:

$$\mathbf{X}(\theta, z) = R(z)\mathbf{e}_r + Z(z)\mathbf{e}_z,$$

for some unknown functions $R(z)$ and $Z(z)$, as shown in Figure 9.5. The parameterisation gives rise to the following tangent vectors and surface metrics

$$\mathbf{a}_\theta = r\mathbf{e}_\theta, \qquad \mathbf{a}_z = r'\mathbf{e}_r + \mathbf{e}_z \qquad a_{\alpha\beta} = \begin{bmatrix} r^2 & 0 \\ 0 & r'^2 + 1 \end{bmatrix}$$

on the deformed state, and

$$\mathbf{A}_\theta = R\mathbf{e}_\theta, \qquad \mathbf{A}_z = R'\mathbf{e}_r + Z'\mathbf{e}_z \qquad A_{\alpha\beta} = \begin{bmatrix} R^2 & 0 \\ 0 & R'^2 + Z'^2 \end{bmatrix}$$

on the undeformed state.

We need to compute the curvature tensor and the Christoffel symbols on the deformed space. The curvature tensor is given by the dot product of $\frac{\partial \mathbf{a}_\alpha}{\partial \theta^\beta}$ and the inward normal $\mathbf{n} = \frac{-\mathbf{e}_r + r'\mathbf{e}_z}{(r'^2+1)^{1/2}}$, it is

$$b_{\alpha\beta} = \begin{bmatrix} \frac{r}{(r'^2+1)^{1/2}} & 0 \\ 0 & \frac{-r''}{(r'^2+1)^{1/2}} \end{bmatrix}.$$

Similar calculations can be made to obtain the Christoffel symbols:

$$\begin{aligned} \Gamma^{\theta}_{\theta\theta} &= 0 & \Gamma^{z}_{\theta\theta} &= \frac{-r'r}{r'^2+1} \\ \Gamma^{\theta}_{\theta z} = \Gamma^{\theta}_{z\theta} &= \frac{r'}{r} & \Gamma^{z}_{z\theta} = \Gamma^{z}_{\theta z} &= 0 \\ \Gamma^{\theta}_{zz} &= 0 & \Gamma^{z}_{zz} &= \frac{r''r'}{r'^2+1} \end{aligned}$$

We can now write down the equilibrium equations for $F^{\alpha\beta}$. Using axisymmetry we assume $F^{\theta z} = 0$, and $F^{\theta\theta}$ and F^{zz} are independent of θ. The θ-component of the covariant divergence of $F^{\alpha\beta}$ vanishes, since $F^{\alpha\beta}$ is independent of θ and the only Christoffel symbols involved are zero ones. $F^{\theta\beta}{}_{;\beta} = 0$ is therefore trivially satisfied. The remaining differential equation is $F^{z\beta}{}_{;\beta} = 0$, which expands to

$$\frac{\partial F^{zz}}{\partial z} + F^{zz}\left(\frac{r'}{r} + \frac{2r''r'}{r'^2+1}\right) - F^{\theta\theta}\frac{r'r}{r'^2+1} = 0 \tag{9.4}$$

The curvature equation $F^{\alpha\beta}b_{\alpha\beta} = P$ is

$$F^{\theta\theta}\frac{r}{(r'^2+1)^{1/2}} - F^{zz}\frac{r''}{(r'^2+1)^{1/2}} = P,$$

which can be used to eliminate $F^{\theta\theta}$ from (9.4), resulting in the uncoupled ODE

$$\frac{\partial F^{zz}}{\partial z} + F^{zz}\left(\frac{r'}{r} + \frac{r''r'}{r'^2+1}\right) = \frac{Pr'}{(r'^2+1)^{1/2}}. \tag{9.5}$$

Using a little analysis of the equations it appears that two boundary conditions are needed to completely determine the undeformed shape, using the following methodology:

1. Specify R at $z = 0$.
2. Compute $A_{\theta\theta}$ at $z = 0$ using $A_{\theta\theta} = R^2$.
3. Compute A_{zz}, $F^{\theta\theta}$ and F^{zz} at $z = 0$ using the three simultaneous algebraic equations $F^{zz} \equiv F^{zz}(A_{\theta\theta}, A_{zz})$, $F^{\theta\theta} \equiv F^{\theta\theta}(A_{\theta\theta}, A_{zz})$, and $F^{\alpha\beta}b_{\alpha\beta} = P$.
4. Solve for F^{zz} for all z by integrating (9.5) using the boundary value computed in step 3.
5. Compute A_{zz}, $A_{\theta\theta}$ and $F^{\theta\theta}$ for all z using the three simultaneous algebraic equations $F^{zz} \equiv F^{zz}(A_{\theta\theta}, A_{zz})$, $F^{\theta\theta} \equiv F^{\theta\theta}(A_{\theta\theta}, A_{zz})$, and $F^{\alpha\beta}b_{\alpha\beta} = P$.
6. Compute $R(z)$ for all z using $A_{\theta\theta} = R^2$.
7. Specify Z at $z = 0$ and solve for $Z(z)$ for all z by integrating $R'^2 + Z'^2 = A_{zz}$.

We see that we need to specify any one of R, $A_{\alpha\beta}$ or $F^{\alpha\beta}$ at $z = 0$ in order to compute the undeformed metric $A_{\alpha\beta}$. To compute the undeformed shape we also have to specify the value of Z at $z = 0$.

Note that this analysis lets us deduce statements such as the cylinder is the only undeformed axisymmetric shape which can deform to a cylinder (because for a cylinder $r' = 0$, which means that (9.5) becomes ${F^{zz}}_{,z} = 0$, so that F^{zz}, and therefore $F^{\theta\theta}, A_{\alpha\beta}$ and R, are all constant in z). Different undeformed cylinders can be inflated to the same deformed cylinder by providing appropriate displacement boundary conditions and an appropriate inflation pressure P. Now, we will shortly study the hemisphere, but first, let us consider the physical components of the two principal forces, which we denote by f^{zz} and $f^{\theta\theta}$. They can be determined from the generalised forces $F^{\alpha\beta}$ by using normalised basis vectors (analogously to (8.5)), and are given by

$$\begin{aligned} f^{\theta\theta} &= r^2 F^{\theta\theta} \\ f^{zz} &= (r'^2 + 1) F^{zz} \end{aligned}$$

Using this transformation in (9.5), the ODE determining f^{zz} can be shown to be

$$\frac{\partial f^{zz}}{\partial z} + f^{zz}\left(\frac{r'}{r} - \frac{r''r'}{r'^2+1}\right) = Pr'(r'^2+1)^{1/2}.$$

Now, consider the hemisphere, for which $r = \sqrt{1-z^2}$. Then it is easy to show that then $r' = -z/r$ and $r'' = -1/r^3$, from which the equations for F^{zz} and f^{zz} can be calculated to be

$$\begin{aligned} \frac{\partial F^{zz}}{\partial z} &= -Pz \\ \frac{\partial f^{zz}}{\partial z} - \frac{2z}{r^2} f^{zz} &= -\frac{Pz}{r^2} \end{aligned}$$

In some situations it can be easier to work with the ODE governing F^{zz}, in others it can be easier to use the ODE for f^{zz}; obviously in this case the former is easier to solve. Suppose we specify $F^{zz} = F_0$ on the boundary. Then the solution for F^{zz} is

$$F^{zz} = F_0 - \frac{P}{2}z^2,$$

from which the solution for f^{zz} can be seen to be

$$f^{zz} = \left(r'^2+1\right)F^{zz} = \left(r'^2+1\right)\left(F_0 - \frac{P}{2}z^2\right) = \left(\frac{z^2}{r^2}+1\right)\left(F_0 - \frac{P}{2}z^2\right) = \frac{F_0 - \frac{P}{2}z^2}{r^2},$$

using $r^2 + z^2 = 1$. However, $r(1) = 0$, so the numerator must also be zero at $z = 1$, as otherwise f^{zz} would blow up at $z = 1$, which is impossible as it is a physical force. We see that we were not in fact free to prescribe one boundary condition at $z = 0$, and the constant of integration must be chosen to enforce regularity instead. In this case, this means that we have to take $F_0 = P/2$, from which we obtain

$$f^{zz} = \frac{\frac{P}{2} - \frac{P}{2}z^2}{r^2} = \left(\frac{P}{2}\right)\left(\frac{1-z^2}{r^2}\right) = \frac{P}{2}.$$

It is straightforward to then show that $f^{zz} = f^{\theta\theta}$, $F^{zz} = \frac{Pr^2}{2}$ and $F^{\theta\theta} = \frac{P}{2r^2}$, which means that $F^{\alpha\beta} = \frac{P}{2}a^{\alpha\beta}$. Now, let us choose a specific strain energy: suppose the membrane is an incompressible thin solid of width h_0 with a Neo-Hookean material law $W = c_1(I_1 - 3)$. The surface stress arising from the equivalent W^{mem} is (using (9.1))

$$T^{\alpha\beta} = 2c_1\left(A^{\alpha\beta} - \frac{A}{a}a^{\alpha\beta}\right),$$

from which we can say that

$$\frac{P}{2}a^{\alpha\beta} = F^{\alpha\beta} = 2c_1h_0\left(A^{\alpha\beta} - \frac{A}{a}a^{\alpha\beta}\right)\frac{\sqrt{A}}{\sqrt{a}},$$

(using (8.13)). The solution to this is $A_{\alpha\beta} = \lambda a_{\alpha\beta}$, for some constant λ determined by

$$\frac{P}{2} = 2c_1h_0\left(\frac{1}{\lambda} - \lambda^2\right)\lambda,$$

i.e.

$$\lambda = \left(1 - \frac{P}{4c_1h_0}\right)^{1/3}.$$

Since $A_{\alpha\beta}$ is a multiple of $a_{\alpha\beta}$, the undeformed state is also a hemisphere, whose radius is given by $R_0 = \sqrt{\lambda} = \left(1 - \frac{P}{4c_1h_0}\right)^{1/6}$. Note that there is no solution for the undeformed radius if $P > 4c_1h_0$, so no solution exists if the specified pressure is too great.

The hemispherical example showed us that, making the assumption of axisymmetry, we cannot always prescribe the initial radius R_0. Instead only one initial radius is possible, any other would lead to a solution that was not regular. In fact, it can be seen that this is true of any surface with one bounding curve (i.e. any surface for which $r(1) = 0$, such as an axisymmetric approximation to the breast shape), by considering the equation $f^{zz} = (r'^2 + 1)F^{zz}$. For smoothness of the surface we must have $r'(1) = -\infty$, so F^{zz} must be zero at $z = 1$ for f^{zz} to be finite.

In this section we have characterised the boundary conditions required when using the simplification of axisymmetry. Ignoring the boundary condition on $Z(0)$ (which has no effect on the shape of the undeformed state, it just specifies where in space the undeformed state is positioned), we have seen that the two distinct types of membrane behave very differently. Membranes with two bounding curves, of which the canonical example is the cylinder, require one boundary condition, which can be on the undeformed radius, R, at either boundary, or on $A_{\theta\theta}$ or A_{zz} at either boundary. For membranes with only one bounding curve, such as the hemisphere, we cannot specify anything (other than $Z(0)$). The undeformed shape is completely determined from the deformed shape and the pressure by regularity. The reason why the cylinder has more freedom is that, in a forward problem, there is an extra surface on which displacements or forces can be prescribed, so a greater range of deformed shapes are possible from any undeformed shape.

9.2.6 Discussion on the Backward Problem using a Fluid-Filled Membrane

Finally, we use the insight gained from the axisymmetric analysis to consider how the non-axisymmetric case can be solved. Suppose we start from a deformed state representing the breast under only residual skin stress. We know the PDEs governing the undeformed state of the skin are elliptic in one 'inner' region, and hyperbolic in another 'outer' region. This is illustrated in parameter space in Figure 9.6. Let the boundary of the two regions be denoted by $\Gamma_{\mathrm{e,h}}$. If we wanted to prescribe boundary conditions on the true boundary of the membrane, we would be on the edge of the hyperbolic region and therefore have to specify both the value of the unknown ($A_{\theta\theta}$ say), and the normal derivative $\left(\frac{\partial A_{\theta\theta}}{\partial n}\right)$ on the boundary. However, if we choose to specify boundary conditions on $\Gamma_{\mathrm{e,h}}$ instead, this problem does not arise, because specifying only the values (and not derivatives) of the unknown on the boundary is enough to be able to solve inwards into the elliptic region. We can then use the computed solution to obtain the derivative of the unknown on $\Gamma_{\mathrm{e,h}}$, and then have two boundary conditions with which to solve outward into the hyperbolic region.

Specifying boundary conditions on the curve of zero Gaussian curvature $\Gamma_{\mathrm{e,h}}$ initially seems unusual and somewhat artificial. However, if we look at images of real breast geometry (for example Figure 9.1(c)), this choice of boundary actually appears quite sensible. There is no obvious boundary of the hyperbolic region, and the curve of zero Gaussian curvature can be seen to be a natural choice for the edge of the breast, with the elliptic region taking up the majority of the breast skin surface[3].

The next question is what the exact form of the boundary conditions should be. We saw that in the axisymmetric case there were two unknowns in the undeformed metric, but regularity meant no boundary conditions could be specified for the hemisphere. In the cylindrical case when one boundary condition was required, it could be on any of the two non-zero components of $A_{\alpha\beta}$. In the non-axisymmetric case we would then expect to have to specify one boundary condition, and would expect that we can specify any one of the three components of $A_{\alpha\beta}$. The natural choice would be to specify the component of the metric along $\Gamma_{\mathrm{e,h}}$ (i.e. specify the equivalent of $A_{\theta\theta}$).

In the latter part of this chapter we have described the need for including residual skin stresses in the model and shown that it is potentially possible to gain approximations to them by considering a backward problem involving a fluid-filled membrane. We leave as open problems the best method for choosing appropriate values of parameters such P, or $A_{\theta\theta}$ on the boundary. Development of

---

[3] Note that this is partly due to the assumption that gravity is not acting on the breast here.

numerical methods for actually solving such problems computationally[4] is left for further work. Even once the residual stresses in the skin have been computed, integrating them into the full model (for which balancing internal stresses need to be computed) is another non-trivial task, one which is also left as further work.

---

[4] There will be a number of numerical issues involved with solving this problem, not least of which is the issue of computing the curvature tensor $b_{\alpha\beta}$ and Christoffel symbols $\Gamma^{\alpha}_{\beta\gamma}$ given only a data set of skin surface points.

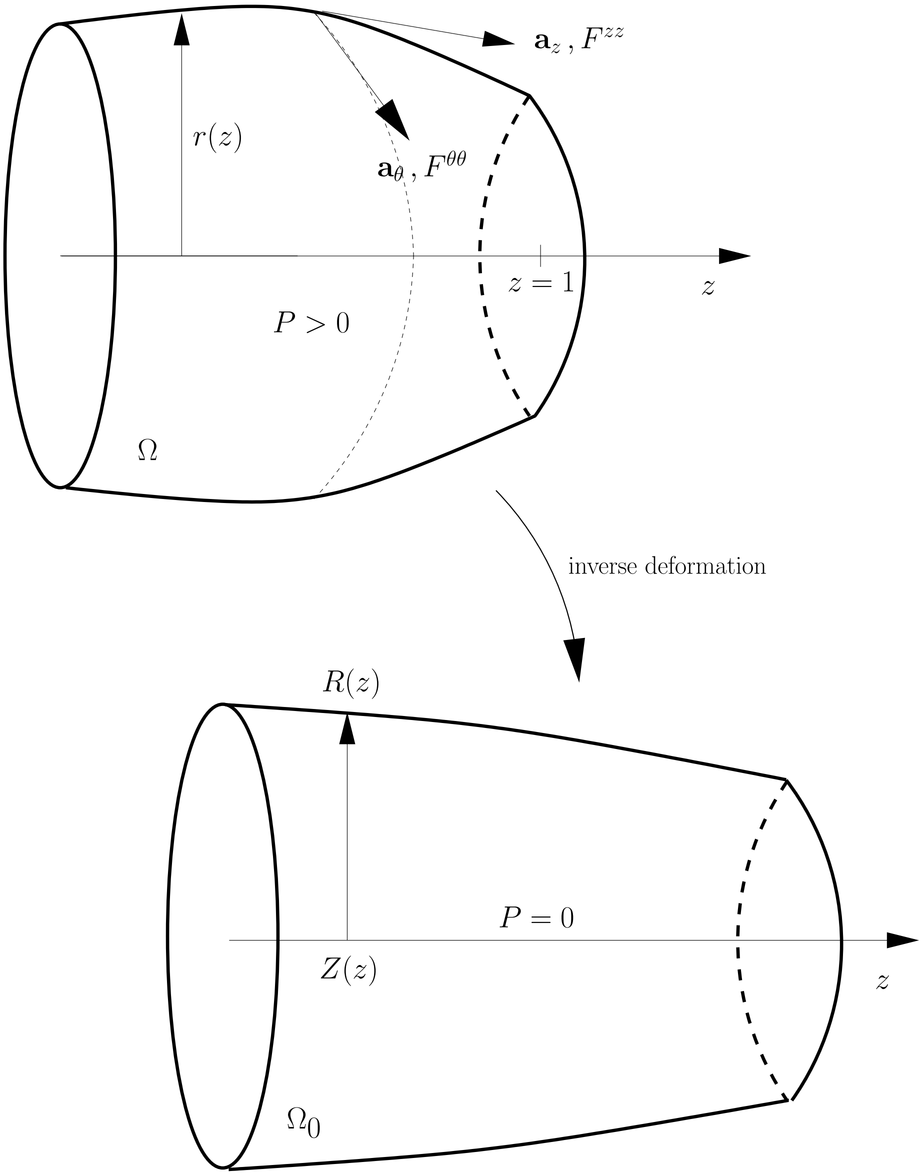

Figure 9.5: General deformed axisymmetric surface Ω and undeformed state Ω_0.

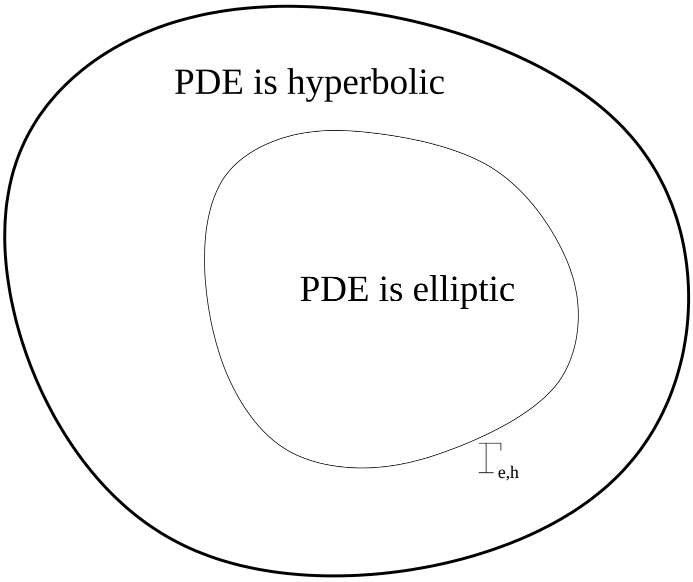

Figure 9.6: The breast surface in parameter space, with inner region in which the problem is hyperbolic, and an outer region in which it is hyperbolic. The internal boundary $\Gamma_{\mathrm{e,h}}$ is the curve of zero Gaussian curvature.

Chapter 10

Conclusions

10.1 Summary

Breast cancer is a major killer of women worldwide, and is a disease for which early diagnosis is the best way to ensure survival. A model of the breast which can predict the displacements of tumours or regions of interest in different clinical scenarios would facilitate diagnosis, and this thesis has been concerned with the problem of developing a patient-specific anatomically-accurate model of the deformation of the breast. Throughout this thesis we have attempted to give an overview of every stage involved in building a clinically-useful model, covering, to various extents, each stage. The most detailed research was focused on the mathematical modelling and numerical analysis aspects of the project.

We began in Chapter 2 by describing the structure of the breast and the modalities used to image the breast, in order to motivate the need for a model. The main applications were seen to be the ability to track the location of a tumour in prone-to-supine and CC-to-MLO deformations, although a number of other applications were identified. We also described previous work in this field, finding them in general to be overly-simplified, and in particular often based on linear elasticity.

In Chapter 3 we derived the equations of finite elasticity, including all of the strong, weak and energy forms, each of which would be required at some point in the thesis. We considered in particular how to set up the governing equations in the cases of hyperelasticity, isotropy and transverse isotropy, and incompressibility. Chapter 4 was then concerned with using the finite element method to numerically approximate the solutions of these equations. Here, we introduced the basics of the finite element method, and described in detail how the method can be used to solve incompressible finite elasticity problems. We described the methods employed for solving nonlinear systems, solving linear systems, and numerical integration, and simulated some simple deformations. We also explained how a simple non-dimensionalisation of the equations can bring about a huge reduction in computation time by acting as an efficient preconditioner of linear systems.

In Chapter 5 we studied an inverse finite elasticity problem, where we assumed we knew the deformed state and wished to compute the undeformed state. We named this problem, and those like it, backward problems. Backward problems are rarely encountered in the literature, and have never been solved for the breast before (as far as we are aware), despite the fact that the backward problem is a necessary first stage when modelling breast deformation. After making some simplifying assumptions such as ignoring residual stresses in the skin, we considered a few methods for solving the backward problem, before explaining how the problem can be solved with a direct reformulation of the finite elasticity equations. Some example simulations verified that the method used was extremely accurate and robust. Initially, however, the solution of the backward problem took considerably longer than equivalent forward problems. This led us to an analysis of the spectra of the matrices involved, and we pinpointed a simple aspect of the method that has a huge impact on the speed of the GMRES method in both forward and backward problems. With a resulting modification, backward problems could be solved in virtually the same time as forward problems.

Next, we studied the contact problem in Chapter 6, in order to be able to model compression of the breast. Restricting ourselves to the frictionless case only, we introduced some of the main concepts through a simple example, discussed how to state the problem, and then derived and evaluated two methods, the penalty and Augmented Lagrangian methods, concluding that the Augmented Lagrangian method was far superior.

Chapter 7 dealt with actual computations of breast deformation. We began by discussing the two major pre-computation steps, segmentation and mesh generation. Edge-detection of the pectoral muscle surface was seen to be the most challenging of the prerequisite segmentation procedures. A relatively unsophisticated method of mesh generation was used in this chapter, and the issue of creating good quality hexahedral meshes was discussed. We then described in some detail the validity of the modelling assumptions made in this thesis, and also surveyed the literature for experimental material laws for breast tissue.

In the second half of Chapter 7 we computed a number of simulations using the model of the breast and the methods developed in the previous chapters. We simulated each of the following deformations: prone to reference state, reference state to supine, reference state to CC compressed, and reference state to MLO compressed. The total prone-to-supine deformation was then utilised in a series of numerical experiments to study the behaviour of the model as material laws and parameters are changed. We also used the simulations of compression to perform CC-to-MLO mammographic matching, visualised on simulated mammograms. The results suggested the position of the MLO curve equivalent to a CC point is dominated by the rotation from CC to MLO rather than the compressions.

The final two chapters in this thesis were mathematical modelling chapters on the specific subproblem of modelling skin. The goal of Chapter 8 was to develop a method of modelling skin as a surface membrane instead of a thin solid (which had been found to be numerically inefficient). By performing a formal asymptotic analysis of the finite elasticity equations, first in two dimensions and then three, we determined the equivalent membrane strain energy to the strain energy for a thin solid, and reduced integrations over a thin solid to integrations over a surface. We also derived the equations governing the shape of a membrane containing an incompressible fluid.

Finally, in Chapter 9 we described how to solve the new equations using the finite element method and performed some simulations on the breast using the membrane model, discovering that, numerically, they were greatly superior. However, it became apparent that it is important to include in the model the effect of residual stresses in the skin, and so we introduced the general backward problem for computing these stresses, before attempting to simplify the problem by using the fluid interior model. This problem was solved and a method was proposed for approximating skin residual stresses using it, but the numerical computation of solutions was left as further work.

10.2 Further Work

Finally, we suggest a number of areas of possible further research. Let us first consider modelling aspects that could be improved or developed. One possibility is investigation into the correct boundary conditions on the pectoral muscle, and methods of implementing a generic non-zero boundary condition to model the fact that the shape of the muscle can change. In addition, the material models could be generalised, for example, visco-elasticity could be assumed instead of hyperelasticity. Also, the assumptions of isotropy could be relaxed, and anisotropy or transverse isotropy used. This, of course, would be limited by the lack of non-isotropic experimental data. A method of modelling blood loss during mammography, where the constraint of incompressibility everywhere would be replaced by a new position-dependent constraint, is another interesting and potentially important modelling task. It may also be necessary to include friction between the skin and compression plates in the model, which would require a more complicated model of contact and extensions to the numerical methods used. Another interesting mathematical modelling problem is the need to (eventually) model breasts containing silicon implants. Here, the domain would break down into an elastic sub-region and an incompressible fluid sub-region. The deformation would have to be computed at the points in the elastic region only, and the surface of the implant would form an internal boundary whose position in the deformed breast is unknown.

There is also plenty of scope for investigating numerical aspects of the model. In this thesis we used only linear interpolation for displacement and piecewise-constant interpolation for pressure;

these need to be compared to results using higher-order basis functions, or with cubic Hermite bases. Also, it would be interesting to compare the accuracy and computation time with a discontinuous Galerkin implementation of the finite element method. This would have the massive advantage of enabling meshes to contain hanging nodes, and mesh-generation would no longer be such a problematic task.

The backward problem, as it was stated in Chapter 5, has less scope for further work, since we devised a method which computed solutions as accurately and as fast as for forward problems. One area for continued work is further analysis of the issue of GMRES being highly sensitive to the structure of the Jacobian (Section 5.3). We were able to partially explain this phenomenon, but a precise mathematical explanation was left as an open problem.

In contrast to the backward problem, the contact problem has a lot of potential for improvement. The Augmented Lagrangian method was shown to be highly superior to the penalty function method, but was not compared to any other methods. Other methods for solving contact problems could be investigated, or alternatively research could focus on determining optimal parameters in the Augmented Lagrangian method. Other possible methods for improving the numerical scheme include replacing the discontinuous constraint with a smooth constraint (which could be carried out by modelling the plate as a 3D object with rounded edges), which might be expected to perform better in simulations, and considering how to deal with the issue of the deformed mesh not being able to align with the constraint (as occurred in Figure 6.11). Adaptive mesh refinement of elements on the surface of the breast near the edge of the plate may lead to increased accuracy here.

Areas for further research arising from the work recounted in Chapter 7 were suggested and discussed in detail in Section 7.5. They included the two most important areas for further work, from the point of view of the goal of clinical software: clinical validation and automatic pectoral muscle detection. Also discussed were segmentation validation, mesh generation and further numerical experiments.

Lastly, the major area for further work from the membrane modelling chapters is mathematical analysis and numerical solution of the various backward problems described in Section 9.2.1. We were able to state and analyse the equations governing the backward problem for a fluid-filled membrane, but did not compute any numerical solutions. The backward problem for skin with an elastic interior, full backward problem (breast affected by residual stress and gravity), and the full backward-contact problem all remain challenging open problems. Each would need to be investigated mathematically to determine governing equations, and numerical methods for computing their solutions proposed.

10.3 Conclusion

The long term aim of this project is to develop software which can accurately predict the position of a tumour in the breast during surgery of mammography. This would involve a number of stages from a range of fields. We have shown that, although such software would still be some years away from completion (requiring a great deal of work in the modelling, validation and bio-mechanical experimentation aspects), the aim is feasible and already computationally tractable. As well as providing an overview of the project, in this thesis we have provided a detailed investigation into the modelling and numerical analysis features involved. The hope is that this can eventually be brought together with further work in the areas identified to produce a powerful and valuable clinical tool.

Bibliography

[AMLdR80] P.G. Agache, C. Monneur, J.L. Leveque, and J. de Rigal. Mechanical properties and Young's modulus of human skin *in vivo*. *Archives of dermatological research*, 269(3):221–32, 1980.

[AMS01] F.S. Azar, D.N. Metaxas, and M.D. Schnall. A deformable finite element model of the breast for predicting mechanical deformations under external perturbations. *Academic Radiology*, 8(10):965–975, 2001.

[Arb] L. Arbach. *Breast MRI: Segmentation and Classification.* `www.engineering.uiowa.edu/~larbach/2003-01/image_understanding/Final_Project_report2.pdf`.

[Aza01] F.S. Azar. *A Deformable Finite Element Model of the Breast for Predicting Mechanical Deformations under External Perturbations.* PhD thesis, University of Pennsylvania, 2001.

[BA02] B. Banerjee and D. O. Adams. Micromechanics-based prediction of effective elastic properties of polymer bonded explosives. In *Proc. 6th Intl. Conf. Electrical Transport and Optical Phenomena in Inhomogeneous Materials (ETOPIM6)*, Snowbird, Utah, 2002.

[BABM02] P.R. Bakic, M. Albert, D. Brzakovic, and A.D.A Maidment. Mammogram synthesis using a 3d simulation. I. breast tissue model and image acquisition simulation. *Medical Physics*, 29(9):2131–2139, 2002.

[BBGS04] T. Barth, P. Bochev, M. Gunzburger, and J. Shadid. A taxonomy of consistently stabilized finite element methods for the Stokes problem. *SIAM Journal of Scientific Computing*, 25(5):1585–1607, 2004.

[BKT02] P. Beal, J. Koko, and R. Touzani. Mesh r-adaptation for unilateral contact problems. *Int. J. Appl. Math. Comput. Sci.*, 12(1):9–16, 2002.

[Blu89] J.A. Blume. Compatibility conditions for a left Cauchy-Green strain field. *Journal of Elasticity*, 21:271–308, 1989.

[BMA$^+$00] C. P. Behrenbruch, K. Marias, P. A. Armitage, M. Yam, N. Moore, R. E. English, and M. Brady. MRI-Mammography 2D/3D data fusion for breast pathology assessment. In *MICCAI*, pages 307–316, 2000.

[BPM$^+$95] S.E. Benzley, E. Perry, K. Merkley, B. Clark, and G.D. Sjaardema. A comparison of all-hexahedral and all-tetrahedral finite element meshes for elastic and elasto-plastic analysis. In *Proceedings, 4th International Meshing Roundtable*, pages 179–191, 1995.

[Che01] Z. Chen. On the Augmented Lagrangian approach to Signorini elastic contact problem. *Numer. Math.*, 88:641–659, 2001.

[Cia78] P. Ciarlet. *The Finite Element Method for Elliptic Problem.* North-Holland, 1978.

[Cou43] R. Courant. Variational methods for the solution of problems of equilibrium and vibrations. *Bull. Amer. Math. Soc.*, 49:1–23, 1943.

[DDP$^+$04] H. Dehghani, M. M. Doyley, B. W. Pogue, S. Jiang, J. Geng, and K. D. Paulsen. Breast deformation modelling for image reconstruction in near infrared optical tomography. *Phys. Med. Biol.*, 49:1131–1145, 2004.

[Dem97] J.W. Demmel. *Applied Numerical Linear Algebra.* Society for Industrial and Applied Mathematics, 1997.

[EdRR$^+$89] C. Escoffier, J. de Rigal, A. Rochefort, R. Vasselet, J.L Leveque, and P.G. Agache. Age-related mechanical properties of human skin: an *in vivo* study. *Journal of Investigative Dermatology*, 93(3):353–357, 1989.

[Emb99] M. Embree. *Convergence of Krylov Subspace Methods for Non-normal Matrices.* PhD thesis, University of Oxford, UK, 1999.

[Eng80] H. Engels. *Numerical Quadrature and Cubature.* Academic Press, 1980.

[Fun68] Y.C. Fung. Biomechanics: Its scope, history, and some problems of continuum mechanics in physiology. *Applied Mechanics Reviews*, 21(1):1–20, 1968.

[Fun94] Y. C. Fung. *Biomechanics: Mechanical Properties of Living Tissues.* Springer, 1994.

[FW04] G.E. Forsythe and W.R. Wasow. *Finite-Difference Methods for Partial Differential Equations.* Dover Publications, 2004.

[GLLG$^+$05] A. Gefen, A. Levchakov, E. Linder-Ganz, S. S. Margulies, and R. Raghupathi. Strain distribution in brain tissue of rats subjected to closed head injury is age-dependent. In *Summer Bioengineering Conference of the ASME Bioengineering Division*, 2005.

[Gou99] P.L. Gould. *Introduction to Linear Elasticity.* Springer, 1999.

[GPS96] A. Greenbaum, V. Ptak, and Z. Strakos. Any convergence curve is possible for GMRES. *SIAM Matrix Anal. Appl*, 17(3):465–470, 1996.

[GW98] G. Golub and A. Wathen. An iteration for indefinite systems and its application to the Navier-Stokes equations. *SIAM Journal of Scientific Computing*, 19(2):530–539, 1998.

[Hay98] P. Hayton. *Analysis of Contrast-Enhanced Breast MRI.* PhD thesis, University of Oxford, UK, 1998.

[HE52] M.R. Hestenes and E.Stiefel. Methods of conjugate gradients for solving linear systems. *J. Res. Natl. Bur. Stand.*, 49:409–436, 1952.

[HFS03] G. Hirota, S. Fisher, and A. State. An improved finite-element contact model for anatomical simulations. *The Visual Computer*, 19(5):291–309, 2003.

[HH04] M. Hyodo and K. Hirahara. GeoFEM kinematic earthquake cycle simulation in southwest Japan. *Pure Appl. Geophys.*, 161:2069–2090, 2004.

[Hir02] G. Hirota. *An Improved Finite Element Contact Model for Anatomical Simulations.* PhD thesis, University of North Carolina, Chapel Hill, USA, 2002.

[HJ91] R.A. Horn and C.R. Johnson. *Topics in Matrix Analysis.* Cambridge University Press, 1991.

[HKS05] Hibbit, Karlson, and Sorenson. *ABAQUS Theory Manual.* `www.abaqus.com`, 2005.

[Hum03] J. D. Humphrey. Continuum biomechanics of soft biological tissues. In *Proc. Roy. Soc. Lond.*, volume 459, pages 3–46, 2003.

[Hun] P. Hunter. *CMISS (Continuum Mechanics, Image analysis, Signal processing and System identification).* `http://www.bioeng.auckland.ac.nz/cmiss/cmiss.php`.

[IJBM06] D. R. Schuler III, J.J.Ou, S.L. Barnes, and M.M.Miga. Automatic surface correspondance methods for a deformed breast. In *Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display: Proc of the SPIE*, 2006.

[ISNC] L. Ibanez, W. Schroeder, L. Ng, and J. Cates. *The ITK Software Guide.* `http://www.itk.org/ItkSoftwareGuide.pdf`.

[JP99] S. Jacob and M. K. Patil. Three-dimensional foot modeling and analysis of stresses in normal and early stage Hansen's disease with muscle paralysis. *J. Rehabil Res Dev*, 36(3):252–263, 1999.

[KB04] G. Ketsetzis and J.M. Brady. Automatic segmentation of T1 parametric maps of breast MR images via a hidden Markov random field. In *Proceedings of MIUA*, 2004.

[KCAR04] S.M. Kwok, R. Chandrasekhar, Y. Attikiouzel, and M.T. Rickard. Automatic pectoral muscle segmentation on mediolateral oblique view mammograms. *Medical Imaging, IEEE Transactions on*, 23(9):1129–1140, 2004.

[KHB01] Y. Kita, R. Highnam, and M. Brady. Correspondence between different view breast X-rays using curved epipolar lines. *Computer Vision and Image Understanding*, 83(1):38–56, July 2001.

[KJL$^+$05] A. Kerdok, P. Jordan, Y. Liu, P. Wellman, S. Socrate, and R. Howe. Identification of nonlinear constitutive law parameters of breast tissue. In *ASME Summer Bioengineering Conference*, 2005.

[Klo02] G. Kloosterman. *Contact Methods in Finite Element Simulations.* PhD thesis, Universiteit Twente, The Netherlands, 2002.

[KMH$^+$04] R. Keiser, M. Mueller, B. Heidelberger, M. Teschner, and M. Gross. Contact handling for deformable point-based objects. In *Proc. Vision, Modeling, Visualization VMV'04*, pages 315–322, 2004.

[Knu01] P.M. Knupp. Algebraic mesh quality metrics. *SIAM Journal of Scientific Computing*, 23(1):193–218, 2001.

[KO88] N. Kikuchi and J.T. Oden. *Contact problems in elasticity: a study of variational inequalities and finite element methods.* SIAM, 1988.

[KS05] J. Kybic and D. Smutek. Computational elastography from standard ultrasound image sequences by global trust region optimization. In *Lecture Notes in Computer Science*, volume 3565, pages 299–310, 2005.

[KTF$^+$98] T.A. Krouskop, T.M.Wheeler, F.Kallel, B.S. Garra, and T.Hall. Elastic moduli of breast and prostate tissues under compression. *Ultrasonic Imaging*, 20(4):260–274, 1998.

[Kui55] N.H. Kuiper. On C^1 Isometric Imbeddings I. *Nederl. Akad. Wet. Proc.*, A-58:545–556, 1955.

[KWT88] M. Kass, A. Witkin, and D. Tezopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[LCSC02] L. Li, Y. Chu, A. F. Salem, and R. A. Clark. Image segmentation and 3D visualization for MRI mammography. In *Proc. SPIE Vol. 4684, Medical Imaging 2002: Image Processing*, pages 1780–1789, May 2002.

[LKC+95] H.D. Li, M. Kallergi, L.P. Clarke, V. Jain, and R.A. Clark. Markov random field for tumor detection in digital mammography. *IEEE Trans. Medical Imaging*, 14(3):565–576, 1995.

[LRM+99] A.J. Lawrence, P.J. Rossman, J.L. Mahowald, A. Manduca, L.C. Hartmann, and R.L. Ehman. Assessment of breast cancer by magnetic resonance elastography. In *Proc. Intl. Soc. Magn. Reson. Med. 7*, page 525, 1999.

[LSBG03] J. Lorenzen, R. Sinkus, M. Biesterfeldt, and A. Gerhard. Menstrual-cycle dependence of breast parenchyma elasticity: Estimation with magnetic resonance elastography of breast tissue during the menstrual cycle. *Investigative Radiology*, 38(4):236–240, 2003.

[MJ95] M. Melloul and L. Joskowicz. Segmentation of microcalcification in X-ray mammograms using entropy thresholding. *IEEE Transactions on Medical Imaging*, 14(3):565–576, 1995.

[MMRT01] J. Martikainen, R.A.E. Mäkinen, T. Rossi, and J. Toivanen. A preconditioner for linear elasticity problems. In *Thirteenth International Conference on Domain Decomposition Methods*, 2001.

[Moo40] M. Mooney. A theory of large elastic deformation. *Journal of Applied Physics*, 11:582–592, 1940.

[Nas55] J. Nash. The imbedding problem for Riemannian manifolds. *Annals of Mathematics*, 63:20–63, 1955.

[Nas98] M. Nash. *Mechanics and Material Properties of the Heart using an Anatomically Accurate Mathematical Model.* PhD thesis, University of Auckland, New Zealand, 1998.

[Nio85] F.I. Niordson. *Shell Theory.* Elsevier Science, 1985.

[NW99] J. Nocedal and S.J. Wright. *Numerical Optimization.* Springer, 1999.

[Ogd72] R. W. Ogden. Large deformation isotropic elasticity I: On the correlation of theory and experiment for incompressible rubberlike solids. *Proc. Roy. Soc. London.*, 326:565–584, 1972.

[Ogd97] R.W. Ogden. *Non-linear Elastic Deformations.* Dover Publications, 1997.

[Ols04] T. Olsson. *Modelling of Residual Stress Fields in Soft Tissues with Application to Human Arteries.* PhD thesis, Linköping University, Sweden, 2004.

[PBSS00] D. Plewes, J. Bishop, A. Samani, and J. Sciaretta. Visualization and quantification of breast cancer biomechanical properties with magnetic resonance elastography. *Physics in Medicine and Biology*, 45:1591–1610, 2000.

[PXP00] D.L. Pham, C. Xu, and J.L. Prince. A survey of current methods in medical image segmentation. *Annual Review of Biomedical Engineering*, 2:315–337, 2000.

[Red93] J. N. Reddy. *An Introduction to the Finite Element Method.* McGraw-Hill, 1993.

[RGC01] M. A. F. Rodrigues, D. F. Gillies, and P. Charters. Realistic deformable models for simulating the tongue during laryngoscopy. In *Medical Imaging and Augmented Reality, 2001. Proceedings. International Workshop on*, pages 125–130, 2001.

[RH73] W.H. Reed and T.R. Hill. Triangular mesh methods for the neutron transport equation. Technical Report LA-UR-73-479, Los Alamos National Laboratory, Los Alamos, New Mexico, 1973.

[RNN04] V. Rajagoal, M.F. Nielsen, and M. Nash. Development of a three-dimensional finite element model of breast mechanics. In *Proc. Ann. Int. Conf. IEEE EMBS*, volume 26(7), pages 5080–5083, 2004.

[ROM$^+$05] D. Raba, A. Oliver, J. Martí, M. Peracaula, and J. Espunya. Breast segmentation with pectoral muscle suppression on digital mammograms. In *IbPRIA (2)*, pages 471–478, 2005.

[Rot54] S. Rothman. *Physiology and Biochemistry of the Skin.* Chicago University Press, 1954.

[Rui03] N. Ruiter. *Registration of X-ray Mammograms and MR-Volumes of the Female Breast based on Simulated Mammographic Deformation.* PhD thesis, Universität Mannheim, 2003.

[RV66] M.D. Ridge and V.Wright. Mechanical properties of skin: A bioengineering study of skin structure. *Journal of Applied Physiology*, 21:1602–1606, 1966.

[SB03] A. Samani and J. Bishop. Measuring the elastic modulus of *ex vivo* small tissue samples. *Physics in Medicine and Biology*, 48:2183–2198, 2003.

[SBRP99] A. Samani, J. Bishop, E. Ramsay, and D. Plewes. A 3-d contact problem finite element model for breast deformation derived from MRI data. In *Proc. ASB, 23rd Annual Meeting (Pittsburg, PA)*, page 72, 1999.

[SBYP01] A. Samani, J. Bishop, M.J. Yaffe, and D. Plewes. Biomechanical 3-d finite element modeling of the human breast using MRI data. *IEEE Transactions on Medical Imaging*, 20(4):271–279, 2001.

[SGMO94] A.P. Sarvazyan, D. Goukassian, E. Maevsky, and G. Oranskaja. Elastic imaging as a new modality of medical imaging for cancer detection. In *Proceedings of the International Workshop on Interaction of Ultrasound with Biological Media*, pages 69–81, 1994.

[She02] J.R. Shewchuck. What is a good linear element? interpolation, conditioning, and quality measures. In *Eleventh International Meshing Roundtable*, 2002.

[SP04] A. Samani and D. Plewes. A method to measure the hyperelastic parameters of *ex vivo* breast tissue samples. *Physics in Medicine and Biology*, 49:4395–4405, 2004.

[Spe04] A. J. M. Spencer. *Continuum Mechanics.* Dover Publications, 2004.

[SS86] Y. Saad and M.H. Schultz. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Statist. Comput.*, 7:856–869, 1986.

[Sul02] E. Suli. *Finite Element Methods for Partial Differential Equations.* `http://web.comlab.ox.ac.uk/oucl/work/endre.suli/`, 2002.

[SZJ$^+$96] R. Skalak, S. Zargaryan, R.K Jain, P. A. Netti, and A.Hoger. Compatibility and the genesis of residual stress by volumetric growth. *Journal of Mathematical Biology*, 34:889–914, 1996.

[TDS$^+$01] C. Tanner, A. Degenhard, J. A. Schnabel, A. D. Castellano Smith, C. Hayes, L. I. Sonoda, M. O. Leach, D. R. Hose, D. L. G. Hill, and D. J. Hawkes. A method for the comparison of biomechanical breast models. In *Proc. IEEE Workshop Mathematical Methods in Biomedical Image Analysis*, pages 11–18, 2001.

[Tre01] L. N. Trefethen. *Spectral Methods in MATLAB.* Society for Industrial and Applied Math, 2001.

[Tur99] E. Turkel. Preconditioning techniques in computational fluid dynamics. *Annual Reviews in Fluid Mechanics*, 31:385–416, 1999.

[vdV03] H.A. van der Vorst. *Iterative Methods for Large Linear Systems.* Cambridge University Press, 2003.

[Ver70] D. R. Veronda. Mechanical characterization of skin-finite deformation. *Journal of Biomechanics*, 3:111–124, 1970.

[VM96] H. Versteeg and W. Malalasekra. *An Introduction to Computational Fluid Dynamics: The Finite Volume Method Approach.* Prentice Hall, 1996.

[VSCS03] P. Vandewalle, F. Schutyser, J. Van Cleynenbreugel, and P. Suetens. Modelling of facial soft tissue growth for maxillofacial surgery planning environments. In *IS4TH*, pages 27–37, 2003.

[WDH$^+$02] J. Weaver, M. Doyley, E. Van Houten, M. Hood, X.Qin, F. Kennedy, S.Poplack, and K. Paulsen. Evidence of the anisotropic nature of the mechanical properties of breast tissue. In *Proceedings of the 44th AAPM Annual Meeting*, 2002.

[WHDK99] P.S. Wellman, R.D Howe, E. Dalton, and K.A. Kern. Breast tissue stiffness in compression is correlated to histological diagnosis. Technical report, Harvard Biorobotics Lab., 1999.

[WM04] C. W. Washington and M.I. Miga. Modality independent elastography (MIE): A new approach to elasticity imagining. *IEEE Transactions on Medical Imaging*, 23(9):1117–1128, 2004.

[WS03] M.A. Wirth and A. Stapinski. Segmentation of the breast region in mammograms using active contours. *Visual Communications and Image Processing*, 5150:1995–2006, 2003.

[ZWB$^+$04] K.H. Zou, S.K. Warfield, A. Bharatha, C.M.C. Tempany, M.R. Kaus, S.J. Haker, W.M. Wells, F.A. Jolesz, and R. Kikinis. Statistical validation of image segmentation quality based on spatial overlap index. *Academic Radiology*, 11(2):178–189, 2004.