

THE UNIVERSITY OF SHEFFIELD

PHYSICALLY BASED MODELLING
OF
HUMAN LIMBS

by

Lee Cooper

A THESIS SUBMITTED TO THE FACULTY OF ENGINEERING
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE

September 1998

© Lee Cooper 1998

Contents

Abstract	8
Declaration	9
Acknowledgements	10
1 Introduction	11
1.1 Modelling Limbs for Computer Animation	12
1.1.1 Geometric models	13
1.1.2 Anatomically based models	13
1.1.3 Physically based models	14
1.1.4 A new physically based limb model	15
1.2 Overview of the Thesis	15
1.3 Contributions	16
2 The Anatomy of Human Limbs	18
2.1 General Anatomy	19
2.1.1 Terminology relating to joint motion	20
2.2 The Skeleton	24
2.3 Generating Motion of the Skeleton: Muscles and Tendons	26
2.4 Forces in Musculoskeletal Structures	29
2.5 Muscles Acting in Combination	33
2.6 The Physiology of the Leg	34
2.7 Types of Deformation Occurring in Human Limbs	41
2.7.1 The shape of a limb in the anatomical position	41
2.7.2 Shape variation caused by limb state	42
2.8 Material Properties of Human Tissue	51
2.8.1 Properties of skin	53
2.8.2 Properties of muscle	55
2.8.3 General non-muscular tissue	58
2.8.4 Tendons	58
2.8.5 Properties of bone	59
2.9 Conclusions	59

3	Previous Research	61
3.1	Character Animation	63
3.2	Physically Based Modelling in Animation	66
3.2.1	The animation process	67
3.3	Animation vs. Simulation	68
3.4	Layered Models	69
3.5	Facial Modelling	70
3.6	Achieving Accuracy	71
3.6.1	Anatomically based modelling	71
3.6.2	Anthropometry based modelling	73
3.6.3	Models for surgical applications	73
3.6.4	Biomechanically based limb models	74
3.7	Discussion of Existing Work	75
3.7.1	Processing requirements	78
3.7.2	Skeleton control	78
3.7.3	Control over model deformation	79
4	Modelling the Structure of Limbs	81
4.1	The Modelling and Simulation Process	81
4.2	Case Study Model	82
4.3	Constructing the Model	84
4.4	Preparing the Model for Simulation	87
4.4.1	Mesh generation	87
4.4.2	Meshing in two and three dimensions	88
4.4.3	Existing meshing techniques	90
4.5	Meshing the Leg Model	95
4.5.1	Decomposing the cubes into tetrahedrons	97
4.6	Summary	99
5	Modelling the Physical Properties of Limbs	102
5.1	Modelling Deformable Material	103
5.1.1	Finite difference and finite element models	103
5.1.2	Mass-spring-damper models	104
5.1.3	Other material modelling approaches	106
5.1.4	Summary	106
5.2	Modelling Soft Tissue using Mass-Spring-Damper Models	106
5.2.1	Simulating uni-axial stress-strain tests	106
5.2.2	The issue of collapse in general mass-spring-damper models	113
5.2.3	Eliminating the collapse problem	113
5.2.4	Equalising internal element forces	115
5.2.5	Modelling limited compressibility	119
5.2.6	Modelling non-muscular tissue	120
5.2.7	Modelling muscle tissue	125
5.3	Modelling Tendons	127
5.4	Modelling the Skeleton	129

5.5	Embedded Surfaces	130
5.6	Modelling Collisions	131
5.6.1	Collision detection	132
5.6.2	Collision response	134
5.7	Control Strategies	135
5.8	Simulation Issues	137
5.9	Summary	137
6	Results	139
6.1	Gravity Effects	139
6.2	Tendon Pre-tensioning	140
6.3	Bending of the Knee Caused by an External Object	142
6.4	Activating the “Hamstrings”	147
6.5	Activating the “Quadriceps”	153
6.6	Isometric Contraction of Opposing Muscle Regions	153
6.7	Isometric Contraction of a Single Muscle Region	159
6.8	Summary	159
7	Verification and Discussion	164
7.1	Introduction	164
7.2	Possible Approaches to Verification	165
7.2.1	Observation of simulation results	165
7.2.2	Comparison with real images and animations	165
7.2.3	Improving image based analysis by adding visual cues	166
7.2.4	Direct 3D shape comparisons	166
7.2.5	Comparison of volume data	166
7.2.6	Assessment of physical quantities	167
7.3	Assessing the Leg Model	167
7.4	Discussion of the Verification Study	170
7.5	Discussion of the Structural Model	173
7.6	Discussion of the Material Model	175
7.6.1	Simulation issues	176
8	Conclusions	178
	Bibliography	183
A	Guide to the CDROM	191
B	Preventing Collapse Within Mass-Spring-Damper Models of Deformable Objects	195

List of Figures

2.1	Human body in the anatomical position showing the three main planes of motion.	21
2.2	Motion of joints about the main body axes (adapted from an illustration in [56]).	22
2.3	Motion of joints about transformed axes (adapted from an illustration in [56]).	23
2.4	The soleus in the lower leg.	27
2.5	The biceps in the upper arm.	28
2.6	Static and dynamic states.	29
2.7	Varying angles of attachment in three different knee positions.	30
2.8	The quadriceps acting to stabilise the knee joint.	31
2.9	Forces acting on a lever system.	32
2.10	Bones in the upper and lower leg.	35
2.11	Muscles lying in, or generating motion of the leg.	36
2.12	Muscles of the gluteal region.	37
2.13	Muscles of the thigh.	38
2.14	Muscles of the thigh and knee.	39
2.15	Muscles of the lower leg.	40
2.16	Leg in the rest position (front).	42
2.17	Leg in the rest position (back).	43
2.18	Deformation due to bending at the elbow (1).	45
2.19	Deformation due to bending at the elbow (2).	45
2.20	Effects of muscle activation.	46
2.21	Varying joint state with fixed muscle activation (1).	47
2.22	Varying joint state with fixed muscle activation (2).	47
2.23	Varying joint state with fixed muscle activation (3).	48
2.24	Varying muscle activation with fixed joint state (1).	49
2.25	Varying muscle activation with fixed joint state (2).	49
2.26	Varying muscle activation with fixed joint state (3).	50
2.27	Varying muscle activation with fixed joint state (4).	50
2.28	Tendon effects.	51
2.29	Testing a material sample.	52
2.30	Stress-strain relationships for soft tissue types.	54
2.31	Typical force-length relationship for muscle tissue at a particular activation level.	56

2.32	Typical force-velocity relationship at various muscle activations.	57
2.33	Stress-strain properties of tendon.	58
3.1	Limb model (adapted from an illustration in [35]).	67
3.2	Tension net skin model (adapted from an illustration in [78]).	71
3.3	Cross section of the tri-layer tissue model (adapted from an illustration in [58]).	72
4.1	Case study leg model.	83
4.2	Skin surface mesh.	85
4.3	Internal tissue regions (bones, muscles and tendons).	86
4.4	Simple three dimensional elements.	88
4.5	Two dimensional mesh generation.	89
4.6	Examples of structured meshes.	94
4.7	Generating the cubic approximation.	96
4.8	Alternative approaches to tetrahedralisation of a cube.	98
4.9	Combining cubic elements.	99
4.10	Cubic approximations to the original surfaces of the model.	100
5.1	Forces acting on mass P.	105
5.2	Spring-damper arrangement in a cube composed of five tetrahedrons. . .	107
5.3	Stress-strain test.	108
5.4	Uni-axial compression of the basic material model.	109
5.5	Stress-strain graph for uni-axial compression.	110
5.6	Illustration of element collapse in 2D.	111
5.7	Illustration of tetrahedral element collapse.	112
5.8	Forces in a tetrahedral element.	114
5.9	Forces from linear spring (dashed line) and non-linear spring (solid line) springs.	115
5.10	Uni-axial compression test of the material model with collapse prevention. .	117
5.11	Model featuring collapse prevention.	118
5.12	Volume preserving force.	119
5.13	Images from a uni-axial compression test of the material model with volume preservation.	121
5.14	Graph from a uni-axial compression test of the material model with volume preservation.	122
5.15	Model featuring limited compressibility.	123
5.16	Hill's muscle force model.	125
5.17	The new muscle model.	126
5.18	Activation of an isolated muscle sample.	128
5.19	Tendon model.	129
5.20	Embedded point.	131
5.21	Triangle-vertex intersection test.	132
5.22	Intersection test for closed surfaces (polygons).	133
6.1	Effect of gravity on model deformation.	140

6.2	Effects of tendon pre-tensioning.	141
6.3	Leg bending due to an external object (physical surfaces).	143
6.4	Leg bending due to an external object (geometric surfaces).	144
6.5	Close up view of creasing.	145
6.6	Skin surface around the bent joint.	145
6.7	Flesh bulging around the joint (1).	146
6.8	Flesh bulging around the joint (2).	146
6.9	Muscle deformation around the bent joint.	147
6.10	Physical skin surface during back muscle activation.	148
6.11	Geometric skin surface during back muscle activation.	149
6.12	Internal physical surfaces during back muscle activation.	150
6.13	Internal geometric surfaces during back muscle activation.	151
6.14	Close up of surfaces during back muscle activation.	152
6.15	Skin surface deformation during front muscle activation.	154
6.16	Internal surface deformation during front muscle activation.	155
6.17	Close up of the front muscle region.	156
6.18	Skin at the front of the thigh during maximum muscle activation.	156
6.19	Skin surface before dual muscle contraction.	157
6.20	Skin surface after dual muscle contraction.	157
6.21	Internal surfaces before dual muscle contraction.	158
6.22	Internal surfaces after dual muscle contraction.	158
6.23	Skin surface before muscle contraction.	160
6.24	Skin surface after muscle contraction.	160
6.25	Close up of skin surface before muscle contraction.	161
6.26	Close up of skin surface after muscle contraction.	161
6.27	Internal surfaces before muscle contraction.	162
6.28	Internal surfaces after muscle contraction.	162
6.29	Inside of joint as viewed from the back of the leg.	163
7.1	Real and simulated leg in initial position.	168
7.2	Legs flexing (1).	169
7.3	Legs flexing (2).	170
7.4	Legs flexing (3).	171
7.5	Legs extending (1).	171
7.6	Legs extending (2).	172
7.7	Legs extending (3).	172

Abstract

Existing research into modelling the human body has typically created either oversimplified models for representing characters in computer animation, or accurate models of specific tissue regions for medical/biomechanics applications. The work presented in this thesis attempts to bridge the gap between these two fields by creating a moderately accurate physically based model of an entire limb, combining a range of new and existing techniques.

In order to consider the effectiveness of existing approaches, and make informed decisions during the design of the model, a study is made of the relevant biomechanical issues. This includes consideration of anatomy, the material properties of soft tissue, and the deformation effects that occur at the skin surface. A complete, albeit simplified, physically based model of a leg is presented, which features a skeleton, two musculo-tendon structures (to flex and extend the knee), and a layer of flesh. The model is complete in that it operates in the same way as real limbs: muscle activation exerts forces on the skeleton via tendons, these are combined with other forces (for example due to gravity), and the resulting motion of the skeleton leads to deformation of the soft tissue regions, including the muscles and tendons themselves. Soft tissue behaviour is represented using an enhanced Mass-Spring-Damper approach which models the material properties of real soft tissue including its uni-axial stress-strain characteristics and limited compressibility.

Results are presented from a number of tests which specifically recreate situations in which many existing models behave incorrectly. The leg model is shown to behave in a physically and visually realistic way, exhibiting skin deformation effects which include, creasing around the knee joint when bent, bulging of contracted muscles, wobbling of flesh, and deformation due to contact with external objects. A simple verification study is presented based on video footage of a real leg.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institution of learning.

Acknowledgements

I would like to thank my Mum and Dad and the rest of my family and friends for putting up with me over the past few years. Particular thanks to my brother Brett, whose arms and legs make an appearance in Chapter Two. A special thanks to my supervisor Steve Maddock for his continued support and tireless enthusiasm. Finally I'd like to thank the rest of the Computer Graphics Research Group and members of the Computer Science Department at the University of Sheffield, particularly the support teams who provided much needed help on many occasions.

This research was funded by the EPSRC.

Chapter 1

Introduction

For many years, research into computer based modelling of human limbs and human motion in general has concentrated primarily on *visual* aspects. The models created have focussed on recreating the shape of the skin surface, and then applying techniques to deform this surface in order to attempt to reproduce the same types of effects which occur in real limbs as we perform everyday tasks. For example, creasing of the skin surface around joints, bulging of the flesh around these creases, bulging around tense muscle regions, sagging due to gravity, the effect of tendons at the surface of the skin, wrinkling and so on. Typically, the deformation techniques used are empirical, and consist of a set of rules which describe how to generate a specific effect that has been observed in real life. However, models of this kind have a number of limitations. Firstly, they are non-general, and are only capable of recreating the specific effects that have been designed into them. Secondly, it is often difficult to devise a set of rules which are capable of displaying the effect of several factors acting on a part of the skin surface simultaneously. Also, the results will only ever be as good as the initial observations. As a result, although such models are adequate in certain situations, for example to create simple deformation effects in cartoon style characters [21], they are not capable of generating realistic images or animations of humans. In order to try and overcome these problems, an alternative approach has recently been adopted [101][82], which considers not just the skin surface, but also the underlying anatomy in order to gain an understanding of why these changes in surface shape take place. The hope is that this insight will allow better models to be created, analogous to the way in which an artist studies anatomy in order to be able to paint figures more realistically. Hence, a shift of focus has taken place from the visual aspects to the *functional* aspects of the body. However, this thesis argues that in order to produce models of the human body which are truly realistic, then a further shift of focus is required from the functional to the *physical*. This involves creating models which represent not only the geometric form of the body, but also the relevant physical and mechanical processes which occur within it. The result is less of a computer based “mannequin”, and more of an “artificial human”.

In support of this argument, a physically based model of a limb (the leg) is presented. This consists of four fundamental components, the skeleton, muscle regions, tendons, and general non-muscular tissue. The muscles and tendons are arranged, as in the body,

into *musculotendon* structures which are functional units that generate and transmit force to the skeleton. This force then typically results in motion of the skeleton which in turn causes deformation of the surrounding non-muscular tissue. The leg model presented in this thesis contains two musculotendon structures, one to flex the knee, and one to extend it. Each of the different parts of the model has its own distinct physical and material properties which are based on those that occur in reality. The skeleton is modelled using a rigid body approach in which forces exerted on the bones are combined and result in motion and rotation of the bones within the skeleton structure. These forces can come from a variety of sources including, the joints, the tendons and non-muscular tissue regions which are attached to the surface of the bones, and the weight of the bones themselves. Muscular and non-muscular tissue is represented using an homogeneous, viscoelastic material model with limited compressibility, and in the case of muscle, the ability to generate internal contractile forces is also present. Finally, tendons are modelled as elastic string-like objects which carry only tensile forces. When these tissue regions are combined, the result is a complete physically based model which operates in the same way as real limbs. Namely, that muscles are activated and generate contractile forces. These are then carried via the tendons to attachment points on the skeleton, and combined with the other forces described above. This then typically results in motion of the skeleton, which in turn leads to deformation of the soft tissue regions, including the muscles and tendons themselves. A number of results are presented from tests which show the leg model exhibiting all the main types of skin surface deformation that occur in real limbs. Many of the tests are specifically designed to simulate real situations in which most existing models behave incorrectly or display inaccurate results. Although the potential applications for such models are vast, this thesis concentrates on two main areas : models for use in computer animation applications, and models for medical and biomechanical applications.

1.1 Modelling Limbs for Computer Animation

When generating computer animation of most non-human objects, simple techniques have proven satisfactory to allow computer generated imagery (CGI) to be created which is sufficiently accurate that most people are unable to distinguish it from the real thing. Recent examples of this are the sequences in “Titanic” [28] in which the ill-fated ship is shown docked shortly before setting sail, and also the realistic animated sequences in “Jumanji” [29], which feature herds of elephants, rhinoceros and other wild animals running rampage through streets and inside houses. However, in the case of animation featuring close-ups of humans, although the realism attained has improved significantly over the past decade, it is still not sufficient to fool most people into thinking that it is real. This is due to a number of reasons. Firstly, the human body is an extremely complex object both in terms of its static appearance, and the way in which it moves and deforms, and hence modelling it is difficult. And secondly, we as humans are particularly adept at noticing very slight inaccuracies or peculiarities in human appearance and motion.

1.1.1 Geometric models

Research carried out by the animation community into modelling the human body typically adopts an empirical approach resulting in the creation of *geometric models*. This involves observing the shape or *form* of the body and the way in which it deforms as we carry out everyday tasks, and then attempting to devise rules which recreate this deformation. A typical example of a geometric model is the work done by Thalmann and Magnenat-Thalmann [63] [90] [62]. Their model consists of an articulated skeleton which represents the basic structure of the character, surrounded by a continuous polygon mesh which represents the skin surface. The aim is that the motion of the skeleton is determined by the animator using standard techniques (e.g. keyframing) and then the model is used to deform the skin surface by moving the vertices. In order to achieve this, “Joint-dependent Local Deformation” (JLD) operators are used. These define the way in which the space around a joint deforms as the state of the joint changes, and examples have been implemented which maintain a smooth skin surface around a joint as it moves, and also to create a bulging effect of the flesh close to the joint. Models of this type are particularly attractive for computer animation applications since they are relatively simple and require little computation. As a result they can be used in interactive animation environments, in which the animator can change the pose of the character, and then see the effect that this has on the deformation of the skin surface in real time.

However, geometric models of this type have limitations which typically mean that they only model a limited amount of the “functionality” of real limbs, which leads to unrealistic results which have an artificial appearance. For example, the assumption that skin deformation is caused solely by the state of the skeleton is a gross simplification of real limbs. In actual fact, it is influenced by a complex interaction of a number of factors including muscle activation, the effect of external forces (for example due to contact/collision with an external object), and the dynamic state of the limb, as well as the state of the skeleton (this is discussed in more detail in Chapter Two). A second limitation is that although it can be relatively straightforward to model individual effects e.g. creasing or muscle bulging, attempting to add more effects without interfering with the existing ones becomes increasingly difficult. This is especially true when a number of effects act on the same area of skin simultaneously. Another fundamental reason for the limited realism of geometric models is that since the approach is empirical, then the effects generated will only ever be as good as the observations made and the rules used to implement them. Furthermore, the model will only be able to reproduce the intended deformation within the “envelope” of situations that were initially considered. Hence, if the limb adopts a new state, then the results will be undefined, and often inappropriate.

1.1.2 Anatomically based models

An interesting development to geometric models that attempts to overcome these problems is referred to as *anatomically based modelling* [82][101]. Here the emphasis is now on the way in which the underlying musculature of the limb (or the particular part of the body being considered) moves and changes shape. The shape of the skin surface

is then determined based on this, for example by dynamically “moulding” an elastic skin around the geometry that represents this musculature. In a model produced by Scheepers *et al.* [82], another important addition enables the shape of the muscles to be changed independently of the skeleton state. This then allows subtle effects like “tension” to be displayed, in which a muscle becomes tense (and changes shape slightly) before an “explosive” action is carried out.

The results generated by this kind of approach show a marked improvement over standard geometric models, particularly if the underlying musculature is modelled accurately. However, the approach is still fundamentally a geometric one, and hence still suffers from the problems described above. For example, effects due to external forces and dynamic effects are still not considered.

1.1.3 Physically based models

In order to try and overcome the limitations of geometric models, a theoretical approach has been adopted by some researchers which uses *physically based models* of human limbs. These operate by simulating physical behaviour in order generate realistic deformation. Gourret *et al.* [39] present a physically based model of a hand which consists of a skeleton covered by a single layer of deformable tissue. The tissue is modelled using the Finite Element Method (FEM) which is an approach commonly used in engineering to model the deformation of mechanical components during use. Results are presented which show realistic images of the hand grasping a deformable ball (which is also modelled using FEM). The advantage of this approach is that the deformation is not prescribed in advance by the model designer (as in the case of geometric models), but *emerges* from the simulation based on the material properties of the ball and the hand. As a result the approach is general, in that it can cope equally well with any shaped ball, or any position of the hand (as long as the situation doesn’t represent a physically unrealistic one). In this case, a significant simplification is that the soft tissue in the hand is modelled by a single homogeneous region of tissue, rather than the many layers of tissue in a real hand. However, since the hand is primarily composed of bone surrounded by a thin layer of flesh, then this is thought to be a reasonable approximation, and the results support this.

Consideration of Gourret’s approach suggests that physically based models have significant advantages over geometric ones. Surprisingly however, when it comes to modelling limbs using this approach, or indeed any other physically based approach, very little work exists. That which does has typically been produced by the medical or biomechanical community and attempts to model a very specific aspect e.g. a single tissue region. One of the few less specific models that has been produced is a model by Chen and Zelter [22] of the gastrocnemius muscle in the leg. Here the approach used combines a FEM model to represent the structure of the muscle, and an accurate muscle force model (by Zajac [107]). Muscle force models have their origins in biomechanics, and model the dependency of the force generated by a muscle on its overall length, the rate of change of the length, and its current level of activation. The resulting model accurately represents both the structure and the material properties of the muscle, and as a result it behaves in a realistic manner when simulated.

1.1.4 A new physically based limb model

The work presented in this thesis attempts to build on the work by Chen and Zelter to produce a complete physically based model of a human limb which contains muscle, tendons and other tissue regions. The advantages of such a model over the geometric approaches described earlier would be numerous. For example, the model would be “fully functional” in that it should be possible to activate the muscle regions in order to perform actions in the same way as in real limbs. It would also be capable of **automatically** reproducing realistic motion and deformation of the skin surface, and even the underlying tissue regions. This includes all the main effects that are seen in real limbs, and in particular, subtle effects, for example wobbling of flesh, which are particularly difficult to recreate in other types of models, but are important to give an impression of realism. Another fundamental issue is that the process of constructing such a model only requires that the structure of the relevant tissue regions, and their material properties be specified. Other issues, for example how the different parts of the model function together are not required, since the operation of the model also emerges automatically from the simulation as a consequence of the behaviour and interconnections of the various parts. Hence, this approach represents another shift in emphasis, from the *functional* to the *physical*. This thesis argues that this shift in approach is necessary in order to produce models of human limbs which are capable of exhibiting the same effects seen in real limbs, under the same conditions. In order to test this idea, a complete physically based model of a leg is produced.

1.2 Overview of the Thesis

This thesis is organised into eight main chapters. Chapter Two begins the main body of the thesis by performing a study of real human limbs, and the particular aspects which are applicable to this research. The main reason for doing this is that it is necessary to have a good understanding of the object being modelled, due to the theoretical (as opposed to empirical) approach being adopted. However, the knowledge gained is also useful when assessing the effectiveness of existing models which are described in the subsequent chapter.

In Chapter Three a number of approaches to modelling limbs and other parts of the body are presented and discussed. This work comes from a number of sources including the computer animation community and the medical and biomechanics community. After discussing the various pieces of work, the chapter then argues that a physically based approach would represent an effective way of realistically modelling human limbs. Finally, the chapter then considers some of the implications that would accompany such a model when used in real applications.

Chapters Four and Five present new work that has been carried out to implement the physically based limb model. Chapter Four describes the general structural modelling procedure that was used. This involves: creating a model which contains geometric descriptions of the required tissue regions, pre-processing the model in order to prepare it for simulation, simulating the model to obtain the desired results, and then rendering images from these results. One aspect of the deformable material

model that is used to represent muscle and general non-muscular tissue, is that it is based on small tetrahedral elements. Hence, an important part of the pre-processing stage is to subdivide the soft tissue regions into meshes of tetrahedrons. A number of existing methods for doing this are discussed. However these are rejected for a number of reasons, and an alternative method is presented based on a structured cubic grid.

Chapter Five describes the techniques that were used to model and simulate the physical aspects of the various parts of the model. This includes a new deformable material model used to represent soft tissue. In order to aid in the development of this, a novel approach was used which involved performing simulated uni-axial stress-strain tests of the type commonly used on engineering materials. This is also used on the final model to ensure that it is sufficiently robust under the range of strains commonly encountered in limbs, and also to ensure that the properties match those of the required tissue types. Methods are also presented for modelling the skeleton, tendons, and the contractile behaviour of muscles.

Chapter Six presents results from a number of tests that were carried out on the leg model. Many of these tests specifically recreate situations in which existing approaches to limb modelling behave poorly.

Chapter Seven contains two sections. The first describes a verification study that was performed on the limb model. This compares images from the simulations to images grabbed from video footage of a real leg. The second section then discusses the results as a whole, and the techniques used, and suggests future work. Finally, Chapter Eight draws a number of conclusions relating to various aspects of the work.

The CDROM that accompanies this thesis contains many image and animations from the research. All of the images of the model that feature in Chapters Four, Five and Six are present on the CDROM and in some cases they allows effects to be seen that are lost in the paper reproduction. The animations also show the dynamic behaviour of the model which is an important aspect of this work. For a description of the contents of the CDROM see Appendix A.

1.3 Contributions

A number of contributions are made concerning the realistic modelling of human limbs:

A new approach to limb modelling is proposed which focuses primarily on recreating the physical aspects of real limbs as opposed to the form and function which are the basis of most existing approaches. More specifically, this involves recreating the tissue structures within a limb, and also modelling the material properties of these tissue structures. As the model is simulated, it then automatically recreates the way in which a real limb operates (the function), which in turn leads to realistic deformation of the skin surface (the form).

As a result of the application of the above approach, a new model of a leg is produced which contains two musculoskeletal structures representing those that exist at the front and back of the thigh, and lead to flexing and extending of the knee joint. Tests carried

out on the model produce realistic results which support the new approach. However, in reality, the process of developing the leg model reveals that significant account has to be taken of the function of real limbs, in order to be able to simplify the model by representing only the required tissue regions.

A study is made of real limbs. This presents a list of the main factors which cause limbs to deform as they perform everyday actions. A list is also presented of the main types of deformation that are exhibited. These lists serve as “benchmarks” when discussing both existing models, and the new model presented here.

The issues involved in creating the structural aspects of limb models are discussed, and a meshing procedure is presented which firstly approximates the geometry of the relevant tissue regions using a cubic grid, and then subdivides these regions into tetrahedrons.

A new deformable material model is presented based on a mass-spring-damper approach. Several enhancements are described which make it suitable for modelling soft tissue. This includes the addition of constraints to prevent collapse, and to limit compressibility. The robustness of the model is tested by performing simulated uni-axial stress-strain tests. This model is based on published work (a copy of the relevant paper is presented in Appendix B for completeness).

Finally, a verification study is carried out on the leg model which involves comparing images rendered from simulation data, to images of a real limb grabbed from video footage.

Chapter 2

The Anatomy of Human Limbs

In trying to create computer based models of the motion and deformation of the human body, or in fact any deforming object, there are two main approaches which can be used. The first is to create the model based purely on observations of how the object appears under a range of different circumstances. This is an empirical approach. The actual details of *why* a particular part of the body looks as it does are unimportant, and the object is essentially regarded as a “black box”. Here, the experimental procedure may consist of subjecting the real object being modelled to a range of stimuli, and then measuring how its shape changes as a result. The resulting empirical model of the object would then attempt to map inputs representing the stimuli to an output shape based on the sample data obtained. When using such a model to predict deformation, the accuracy obtained will typically vary based on the input conditions. For example, for input conditions which are close to the sample data, we would expect highly accurate results. For input conditions which were somewhere in between a number of sample points, then we would also expect good results, with the actual accuracy being dependent on the particular interpolation method used and the nature of the object being modelled (i.e. whether it exhibited large changes in shape from small changes in input conditions). However, for input conditions which are a long way from surrounding sample points, or which are outside the envelope of conditions covered by the samples, then significant divergence may occur from the real behaviour.

The second approach, referred to as theoretical, attempts to look inside the black box which is the object being modelled, in order to find out why it behaves as it does. This involves examining the nature of the processes which are taking place so that they can be understood in terms of known physical laws and established theories, and then creating a model based on this understanding. In this case the resulting model is typically a highly simplified version of the real thing, since it will only be possible to analyse and describe a small proportion of its real complexity. For this reason the output produced by the model may often not match a particular experimental observation in the same way that an empirical model would, since the simplifications and assumptions in the model will lead to inaccuracy. However, we would expect the model to represent the general behaviour of the object for all the conditions under which the assumptions hold, and the inaccuracies introduced by the simplifications are small. In the modelling of limbs, both these approaches have been used and each has

been shown to have certain advantages and disadvantages. A discussion of this work is presented in Chapter Three.

This thesis concentrates on a theoretical approach, and for this reason this chapter has been included with the aim of giving an introduction to the known theory of how human limbs function. As well as being essential for the approach taken, this knowledge is also useful when we look at the other, possibly empirical, approaches in order to analyse their effectiveness, and suggest why they are or are not successful.

One of the main challenges when constructing a theoretical model is deciding exactly which aspects of the object need to be modelled in order to achieve suitable accuracy with minimum model complexity. In the case of modelling human limbs, there is a vast amount of knowledge and understanding which could be considered. For example, we could consider the skeleton and general tissue structures (anatomy), the structure within tissue regions, the way in which muscle activation leads to motion of the limbs (biomechanics) and so on. The decision regarding which areas to consider was primarily based on which processes are required to be understood in order to be able to model the motion of the skeleton as a whole, and the various deformable tissue regions that surround it. The information provided in this chapter is aimed at enabling us to do this. The chapter consists of three main sections. The first describes the basic anatomy of limbs, and the mechanics of their motion. Notably, how muscle contraction is converted into bending of joints, and how the muscles within limbs are arranged in order to achieve the required types of motion. The human leg is used as a case study. The second section steps outside the body, and analyses the types of deformation that occur at the surface of limbs, and considers how the internal structures lead to, or influence this deformation. Finally, the third section looks at the material properties of the various tissue types, and considers how these properties influence the observed deformation.

2.1 General Anatomy

The human body is composed of complex structures of many different types of tissue including bone, cartilage, tendons, ligaments, muscle, nerve fibres, blood vessels, fatty tissue, skin, hair etc. The structures that exist in different parts of the body are primarily determined by the particular functions that that part of the body performs. In the case of the skull, the shape is based around the need to protect the brain, and also to provide structure for the face. For the torso, the skeletal structures found there provide support and protection for the internal organs, whilst still allowing sufficient flexibility of movement by having a large number of articulations.

In the case of limbs, the primary function is the generation and control of forces. For example, the arms and hands allow us to lift, manipulate, and throw objects. Whilst in the case of the legs, the forces generated are large enough to enable us to stand upright, and move the body as a whole. Since limbs have this mainly “mechanical” role, their structure is usually more straightforward than in other parts of the body. For this reason, we consider the limbs as an ideal starting point for computer based modelling of the human body. It is then hoped that the techniques developed in this area will be applicable to modelling other more complex parts of the body.

The study of the mechanical aspects of the human body is known as biomechanics, and as its name suggests, it draws on knowledge from a wide range of disciplines including medicine, physics, material science, physiology, biology and so on. This chapter provides an introduction to this field. It begins with a description of the basic terminology that is used to describe joint motion, and the various types of joint that exist. It then moves onto describing the various tissue and skeleton structures that exist in limb themselves, and in particular, how muscles and tendons are arranged with respect to joints in order to generate and control motion of limbs and of the body as a whole. An important discussion is provided which lists the main causes and effects of limb motion and deformation. Finally, it looks at each of main tissue types that are considered relevant when modelling the motion and deformation of limbs, and describes their material properties.

2.1.1 Terminology relating to joint motion

In order to describe the position of anatomical features and the motion that occurs at joints, it is necessary to introduce a naming convention that is commonly used in medicine. The method is based on a human in the *anatomical position* as shown in figure 2.1, with the arms and legs lying straight, the palms facing forwards, and all the other joints fully extended. Motion of the body is always described in terms of the joint rather than the body segment or segments that are moving. For example, we may describe flexing of the hip but not flexing of the thigh. To simplify matters, we break down motion of joints that may have several degrees of freedom into a series of motions which occur in a single degree of freedom. If we consider motion of a joint in a single degree of freedom, then the type of motion that is occurring can be described by its *axis of rotation* and its *plane of motion*. If we consider the body in the anatomical position, then there are three main axes that run through it, each with a corresponding plane of motion. These are shown in figure 2.1.

The first axis runs in a side to side direction relative to the body and is called a *mediolateral axis* (sometimes known as a frontal or transverse axis). If we consider motion of a joint about this type of axis, for example the rotation that occurs at the hip joint when one of the legs is swung backwards and forwards, then any point on the thigh will move around the axis of rotation along a plane which is oriented vertically through the body such that it divides it into left and right sections. This plane is known as a *sagittal plane*. The second axis passes through the body from front to back and is known as an *anteroposterior axis*. The corresponding plane is known as a *frontal plane* since it divides the body into front and back parts. Finally, the third type of axis passes vertically through the body from top to bottom, is known as a *longitudinal axis*, and has a corresponding plane of motion, known as a *transverse plane*.

These expressions are used to describe the motion of joints from the anatomical position. A number of examples are shown in figure 2.2. In figure 2.2a, we see examples of motion about a mediolateral axis (along a sagittal plane). These are rotation of the hip joint (i), the knee joint (ii) and the elbow joint (iii). In figure 2.2b, we see motion about an anteroposterior axis (along a frontal plane). These are rotation of the hip joint (i) and the shoulder joint (ii). Finally, figure 2.2c shows motion about a longitudinal

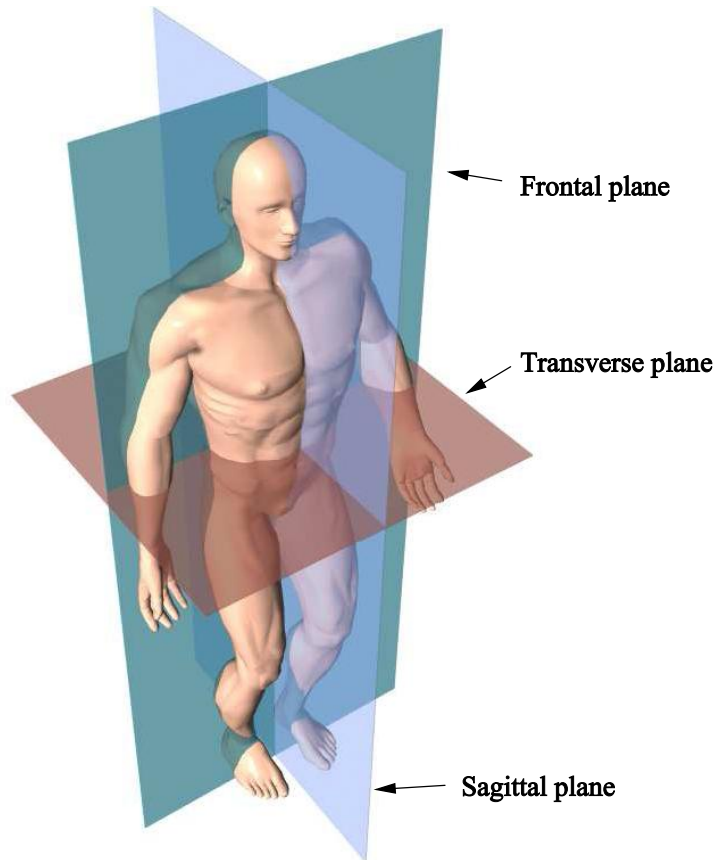


Figure 2.1: Human body in the anatomical position showing the three main planes of motion.

axis (along a transverse plane). These are rotation of the shoulder and hip joints (i) (note that the elbow and knee joints are bent to help illustrate the types of rotation occurring at the shoulder and knee joints) and rotation of the forearm (ii).

The examples of joint motion described above, are all about a single axis, and are all being initiated from the anatomical position. The situation becomes slightly more complex if we consider compound motion i.e. motion of a joint about one axis, followed by motion of the same or a different joint about another axis. In these cases the same convention is used, but the axes themselves may be rotated by an earlier rotation at a joint. Figure 2.3 shows motion of joints about axes which have been transformed by earlier rotations. Figure 2.3a shows the arm after the shoulder has been rotated by ninety degrees about a mediolateral axis such that the arm is extended out to the front with the hand at shoulder height. If we now rotate the shoulder about a longitudinal

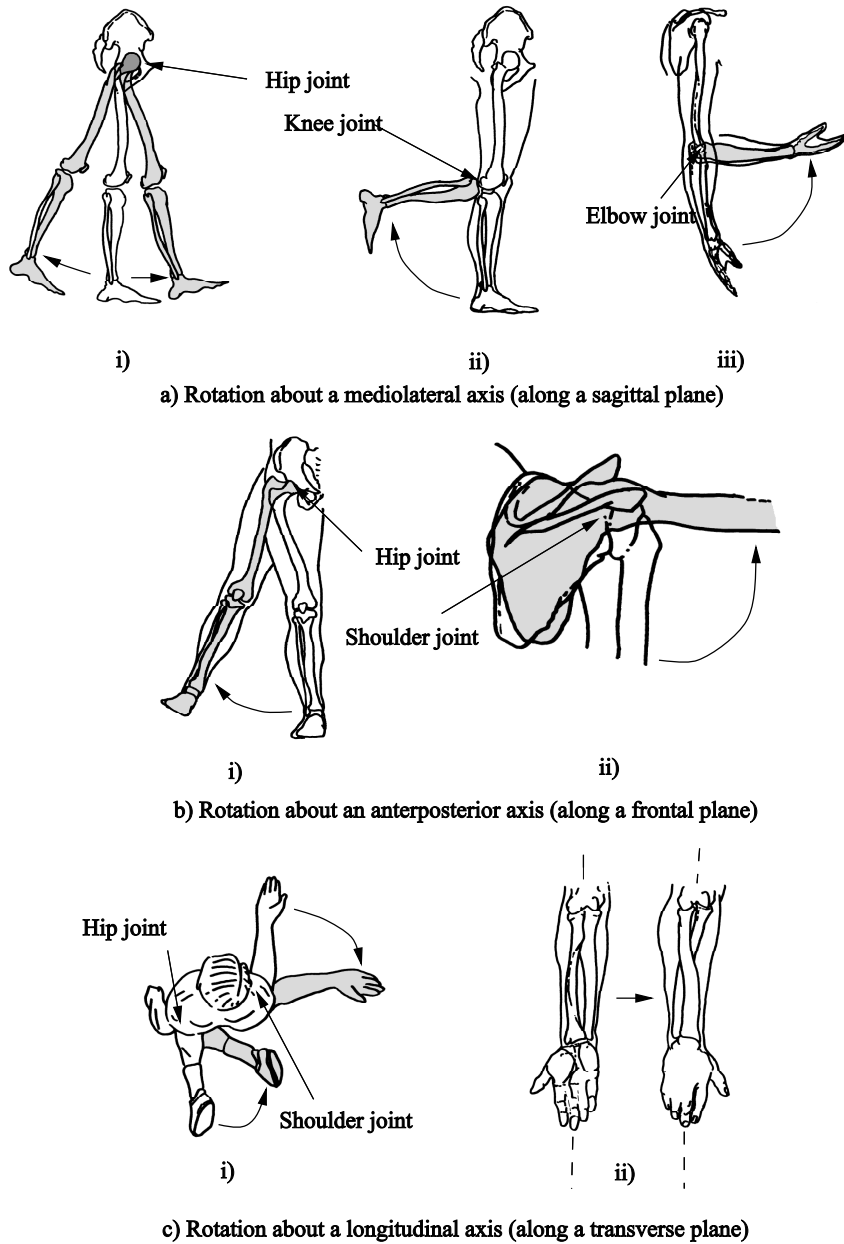





Figure 2.2: Motion of joints about the main body axes (adapted from an illustration in [56]).



a) Rotation of the shoulder about a transformed anteroposterior axis.



b) Rotation of the shoulder about a transformed mediolateral axis.



c) Rotation of the elbow about a transformed mediolateral axis.

Figure 2.3: Motion of joints about transformed axes (adapted from an illustration in [56]).

axis of *the body* such that it is moving in a *transverse* body plane, then this motion is referred to as rotation of the shoulder about an anteroposterior axis since the original anteroposterior axis rotates with the upper arm as it performs the first rotation. In figure 2.3b, we now see an initial rotation about the shoulder of ninety degrees about a mediolateral axis such that the hand moves to the side, again until it is level with the shoulders. This is then followed by a rotation about a longitudinal axis of *the body* such that it moves in a *transverse* body plane. Here, the second motion is described as rotation of the shoulder about a mediolateral axis since this axis has again been rotated along with the first rotation. Finally, figure 2.3c considers rotation of one joint followed by rotation of a *different* joint. Again the same rules apply, and we see rotation of the elbow about a mediolateral axis (now aligned with the longitudinal body axis).

As well as the formal expressions introduced above, a number of shorter expressions are used to describe the most common types of joint motion. The main terms are *flexion* and *extension* which are used to describe motion of a joint about a mediolateral axis. In most cases, the anatomical position represents the limit of motion of a joint in one direction. For example if we consider the knee, then it can only bend such that the lower leg moves backwards. In this situation, flexion describes motion of a joint away from the anatomical position, and extension towards it. Examples of flexion and extension are rotation about a mediolateral plane of the knee (fig 2.2a.ii), elbow (fig 2.3a.iii), shoulder, and wrist. In other joints, motion can occur in either direction away from the anatomical position, for example rotation of the joints between vertebrae, and here rotation in either direction is referred to as flexion. For certain joints special terms are used. For example rotation of the ankle joint about a mediolateral axis such that the top of the foot moves upwards is referred to as *dorsiflexion*, and rotation in the opposite sense is known as *plantar flexion*.

Two other common expressions are *abduction* and *adduction*. These usually describe rotation of a joint about a local anteroposterior axis. Abduction describes motion away from the body and adduction towards. Examples are rotation about an anteroposterior axis of the hip (fig 2.2b.i) and the shoulder (fig 2.2b.ii). As well as terminology to describe joint motion, other expressions are commonly used to describe direction or position. For example: *lateral* (away from the centre), *anterior* (in front), *posterior* (behind), *superior* (above), *inferior* (below), *deep* (away from the surface), and *superficial* (towards the surface).

Having introduced the terminology used to describe most common types of joint motion, we will now move on to discuss the skeleton, and in particular the various types of joint that are found in the body.

2.2 The Skeleton

The skeleton has many important functions [56] including:

- Protecting vital organs in the body.
- Supporting soft tissues.
- Generating red blood cells.

- Acting as a reservoir for minerals e.g. calcium and phosphate.
- Providing attachment points for skeletal muscles.
- Permitting controlled motion about joints.

Although all of these functions are vital for correct functioning of the body, we are primarily concerned with the last two and their relevance in the functioning of limbs. Between many bones, joints exist which allow controlled motion to occur. Joints can be classified into three main categories: *fibrous*, *cartilaginous* and *synovial* [40]. Fibrous and cartilaginous joints are junctures between bones which although not specifically designed to allow large scale relative motion, do allow some movement due to the soft nature of the tissue (ligaments, membrane links or cartilage) that join the two bones together. These joints also have the ability to absorb shock. Examples are the joints between the bones in the skull. The third type, synovial joints (meaning joints with joint cavities), are by far the most common and facilitate motion of the limbs. They are generally classified according to the number of degrees of freedom (DOF) of motion that they allow as described below.

Mon-axial (one degree of freedom) joints.

Mon-axial joints allow rotation in only one plane. There are two main types, *hinge joints* and *pivot joints*.

Hinge joints. Here one of the bones of the joint has a convex surface which fits into a concave surface of another. These allow motion in one plane similar to a door hinge. Examples are the knee joint, elbow joint, and the interphalangeal joints (joints in the fingers and toes).

Pivot joints. A rounded or conical surface of one bone articulates with a depression in another bone. The main type of motion is rotation about a longitudinal axis. An example are the radioulnar joints in the forearm which allow axial rotation along its length.

Bi-axial (two degree of freedom) joints.

These allow rotation in two planes (side to side and back and forth). There are two main types of bi-axial joint, *condyloid joints* and *ellipsoidal joints*.

Condyloid joints. These are joints like those at the heads of the metacarpals (hand bones), an example being the knuckles.

Ellipsoidal joints. Here, the end of one bone has an oval shape which fits into an elliptical cavity in the other bone. Typically motion in one plane is greater than in the other corresponding to the major and minor principal axes of the ellipse. An example is the radiocarpal (wrist) joint.

Tri-axial (three degree of freedom) joints.

These allow rotation about all three axes (side to side, back and forth, and axial rotation). There are three types of tri-axial joint.

Saddle joints. Here, the ends of both bones are saddle-shaped (convex in one direction and concave in the other). This offers more freedom of motion than the ellipsoidal joint. An example is the joint between the trapezium and the metacarpal bones of the thumb (carpometacarpal joint of the thumb).

Ball and socket joints. A ball-like surface of one bone fits into a depression of another bone. These joints permit flexion-extension, abduction-adduction, and rotation. Examples are the hip and shoulder joints.

Gliding (or plane) joints. Here, the two bones have flat or nearly flat articulating surfaces. Although they allow motion in almost any plane, the motion is usually slight. An example is the joint between the sternum (breast bone) and the clavicle (collar bone).

The shape of the articulating surfaces in a joint predominantly determines the number of degrees of freedom that it will have. However, the joints are also influenced significantly by the soft tissue regions surrounding them. The most important of these are ligaments. These are strong fibres which are attached to the bones at either side of it. Their main function is to bind the joint together to prevent dislocation, however they can also inhibit motion in a particular direction.

As well as being categorised by degrees of freedom, each joint also has a range of motion associated with it. This is the total range through which the joint can move before being stopped by bony, ligamentous or muscular structures. If we consider the elbow then the range of motion is limited by bony structures when the joint is extending, and by compression of soft tissue on this inside of the joint when it is flexing.

Having discussed the types of joint that are found in the skeleton, we will now consider how muscle and other tissue types are arranged around the joints of the body.

2.3 Generating Motion of the Skeleton: Muscles and Tendons

If we exclude the effects of external forces, then all the motion that occurs in the human body originates in the contractile forces that are generated by muscles. Typically, these forces act on the bones on either side of a joint generating a torque which leads to rotation of one bone relative to the other. However, the actual mechanics of muscle-bone structures shows a large amount of variation. For example, in many situations the muscle that is responsible for creating a certain type of motion about a joint is too bulky to be situated close to it (as in the case of the muscles which generate motion of the fingers). Here, tendons are used to carry the contractile forces from the end of the muscles to the points of attachment on the bones, in a similar way that rope carries a tensile force from its source to the object which requires moving.

Tendons are tough fibres which take the form of cords or straps, and are made primarily of collagen. Along most of their length they are connected loosely to the tissue that they slide past via lengths of soft tissue which carry nerves, blood vessels and other essential tissue structures. However, in areas where the tendon requires total freedom of movement (for example where it slides around a rigid structure in order to change direction), synovial membranes exist which are disconnected from the tendon itself, and provide a low friction surface past which the tendon can slide. Tendons are extremely strong, having a tensile strength of half that of steel [25]. This means that a tendon with a cross sectional area of 1 cm^2 would be able to support a mass of 600-1000 Kg (roughly 12 people). They are also slightly elastic, although they are only capable of extending by small amounts. In some circumstances this elasticity allows the tendon to store energy, thereby increasing the efficiency of certain types of cyclic motion e.g. running.

The arrangements of muscle and tendons varies in different parts of the body. In some circumstances muscles are attached directly to bones with no need for tendons. In others, tendons will exist at one or both ends of a muscle. Whatever the arrangement, the *musculotendon structure* is responsible for generating and carrying force to the attachment points on the skeleton. The combination of a musculotendon structure and the bones upon which it acts, is known as a *musculoskeletal structure*.

A relatively simple musculoskeletal structure is that based around the soleus muscle in the lower leg. This is shown in figure 2.4a with the surrounding muscles removed for clarity. The soleus is responsible for plantar flexion of the ankle joint. At its upper

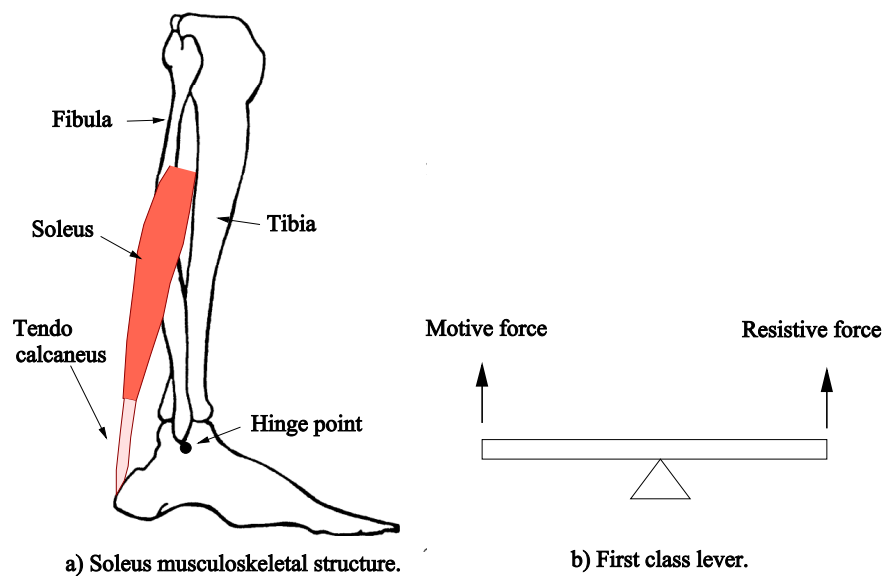


Figure 2.4: The soleus in the lower leg.




Figure 2.5: The biceps in the upper arm.

end the soleus has no tendon, and connects directly to the the tibia and fibula. At its lower end the muscle is attached to a tendon known as the *tendo calcaneus* which carries the force to a point of attachment on the calcaneus (heel). A situation where this muscle would be used is when pressing the accelerator pedal of a car. Here the pedal is exerting a *resistive force* on the end of the foot which is opposing the motion of pressing it down. To oppose this the soleus contracts causing a *motive force* to be exerted on the calcaneus via the *tendo calcaneus*. In this particular musculoskeletal structure the hinge point (located in the ankle) is between the resistive and motive forces, and this is known mechanically as a first-class lever system, as shown in figure 2.4b.

A different type of musculoskeletal structure exists in the upper arm as shown in figure 2.5. This is based around the *biceps brachii* which lies along the humerus. At the upper end, the biceps split into two parts (known as the long head and the short head) each of which is attached via tendons to the shoulder. At the lower end the muscle is attached to a flat thick tendon which connects to the radius bone approximately 2cm from the hinge position in the elbow. This tendon can be felt protruding from the inside of the elbow when it is flexed by ninety degrees and the biceps are tense. If we consider the situation shown in figure 2.5a such that the arm is lifting an object held in the hand, then the force generated by contraction of the biceps brachii (motive force) will be acting to flex the elbow and oppose the weight of the object (resistive force). In this case, the forces are both on the same side of the pivot, with the motive force being closer to the pivot than the resistive force. This system is known as a third-class lever as shown in figure 2.5b.

The final type of lever, a second-class lever, is such that the resistive force is between the pivot and the motive force. However, this type of level is not common in the human body since the separation of pivot and muscle attachment would imply bulky musculoskeletal structures.

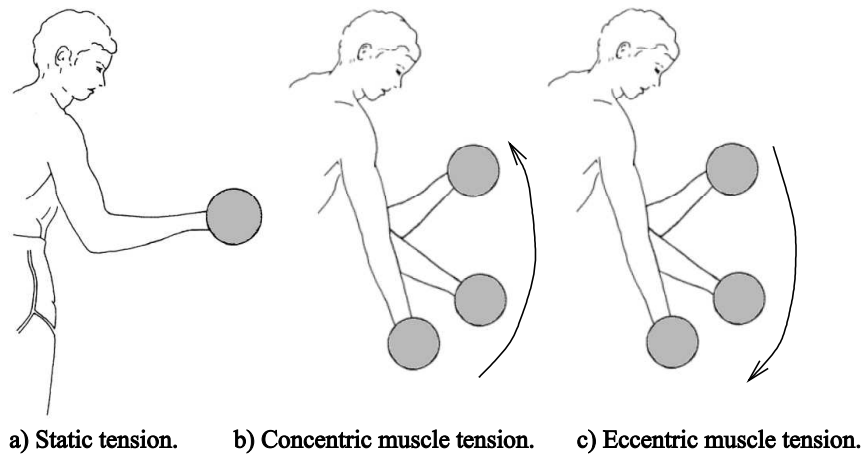


Figure 2.6: Static and dynamic states.

2.4 Forces in Musculoskeletal Structures

If we consider any lever system, then it may be in a state of static equilibrium or a dynamic state (in motion) depending on the forces acting on it. As well as forces from muscles, other forces will act on the bones which are part of a joint, for example, the weight of the bone itself and the surrounding limb segment, or an external force from an object that is in contact with the skin surface. If a joint is in static equilibrium then the sum of the torques on the bones which comprise the joint will be zero. Alternatively, if the forces are such that there is a net torque on one of the bones, then the bone will acquire an angular acceleration in the same direction as the torque.

Figure 2.6 shows three examples of a level system in various static and dynamic states. In fig 2.6a, the muscle is contracting such that the force produced generates a torque which is equal but opposite to the torque caused by the weight of the forearm and the barbell. Hence the system is in equilibrium and this is known as *isometric* muscle tension (tension without motion). In figure 2.6b, the muscle is providing the motive force (it is intending to cause motion), whilst the weight of the forearm and barbell are resistive forces. Here the muscle force creates a torque which exceeds the resistive torques and as a result the forearm rotates upwards. This type of situation in which the motion is in the same direction as the muscle force is called *concentric* muscle tension. Finally in figure 2.6c, the arm is accelerating downwards since the muscle force (resistive) is generating a smaller torque than the other (motive) forces. This is known as *eccentric* muscle tension.

The above has illustrated how the force generated by a muscle is used to perform different types of action. However, the actual effectiveness of the force generated can be dependent on the state of the joint. For example, as the position of a joint changes then the geometry of the musculoskeletal system may also change causing the angle at

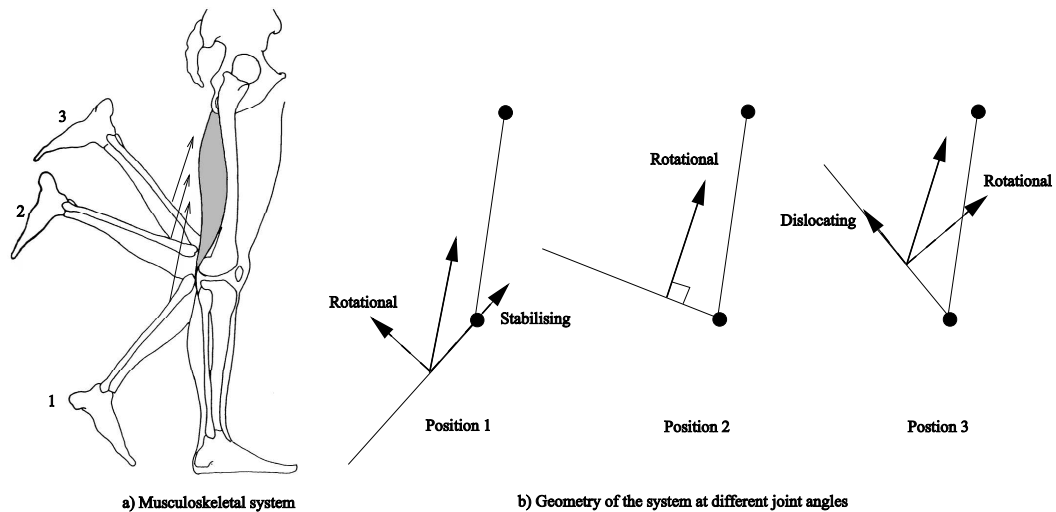


Figure 2.7: Varying angles of attachment in three different knee positions.

which the tendons meet the bones and hence the direction of the forces to change (since the force will always lie along the direction of the tendon at the attachment point). If we consider the force that is being transmitted onto the skeleton from a tendon, then this can be decomposed into two orthogonal forces, one which is parallel to the bone and one which is perpendicular to it. The perpendicular force generates the torque on the bone and hence it is this that performs the useful work. The parallel force does not act to rotate the bone, but instead either pushes the joint together if it acts towards it, or pulls it apart if acts away from it. The former is known as a stabilising force and the latter a dislocating one.

Figure 2.7a shows the lower leg being lifted by the group of muscles known as the hamstrings which lie at the back of the thigh. The situation is such that the tension generated by the muscle is the same in all three of the positions shown. Figure 2.7b shows the geometry of the lever system in the three different positions, and the components of the force at the attachment point in the parallel and perpendicular directions. In position 1 the attachment angle is such that the component of the force which is acting to stabilise the joint is greater than the component acting to rotate the joint. Hence at this position, not all of the force that the muscle is generating is being converted into torque, and the system is not very effective. At position 2 the angle of attachment is now 90 degrees, and this gives the maximum rotary force component with zero stabilising or dislocation component. At this position, all the force is being transferred into torque on the bone, and the musculoskeletal structure is at its most effective. For joint angles beyond 2, a dislocating component occurs and the radial component diminishes (as shown in position 3), and hence the effectiveness again drops.

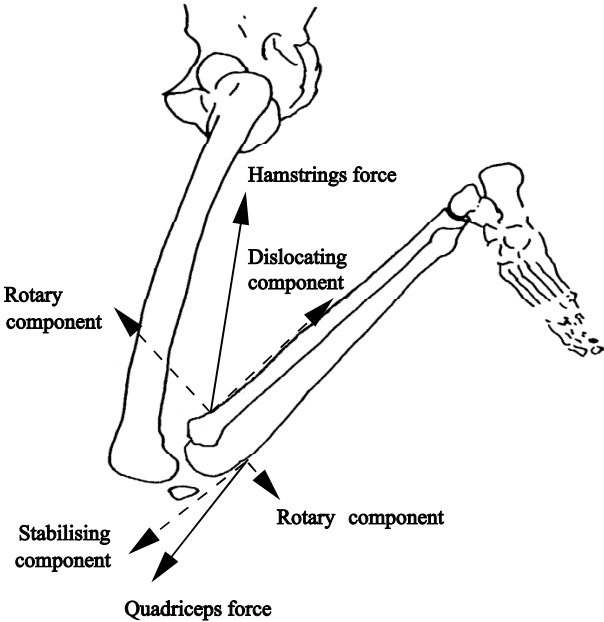


Figure 2.8: The quadriceps acting to stabilise the knee joint.




Figure 2.9: Forces acting on a lever system.

In order to avoid excessive dislocating forces, other muscles around a joint may act along with the muscle generating the primary motive force in order to stabilise the joint. Figure 2.8 shows such a situation, in which slight activation of the muscles at the front of the high (the quadriceps) generates a force which is carried around the front of the knee via tendons and the patella (the knee cap). The shallow angle of attachment of the tendon from the quadriceps at this joint angle generates a force on the bone which acts almost entirely to stabilise the joint, and this cancels out the dislocating component of the force from the hamstrings.

So far we have considered two issues which affect how the force generated by a muscle is converted by a musculoskeletal structure into an applied torque. Namely, the arrangement of level systems, and the angle of attachment of tendons to the skeleton. In order to try and appreciate the scale of forces that are generated within a typical system, we now consider the mechanics of a particular situation. Figure 2.9 shows an arm supporting a barbell with a mass of 5 Kg in static equilibrium (the values shown in the diagram e.g. the dimensions and mass of the forearm, are estimates for an average adult male). In order to compute the tension required by the muscle to hold the system in static equilibrium, we need to consider the torques generated by the motive force (the muscle tension) and the resistive forces (the weight of the arm and the barbell). Since the system is in equilibrium then the sum of the torques will be zero. Hence (considering anti-clockwise to be positive),

$$\begin{aligned}
\sum T' &= 0 \\
T_{\text{muscle}} + T_{\text{weight of arm}} + T_{\text{weight of barbell}} &= 0 \\
F_{\text{muscle}} \times 0.02 - 13.3 \times 0.15 - 49 \times 0.45 &= 0 \\
F_{\text{muscle}} &= 1220N
\end{aligned}$$

Hence the muscle must generate a tension of 1220N to maintain the arm in this position. Hence it can be seen that the muscular force required to lift (or support) an object is often much greater than the weight of the object itself.

When discussing levers, a quantity known as the *mechanical advantage* (MA) is defined as :

$$MA = \frac{\text{resistive force}}{\text{motive force}}$$

In the situation described above, the mechanical advantage has a value of 0.04 since the motive force is much greater than the resistive force. In general, throughout the human body, the point of attachment of the muscle is closer to the pivot than the point of action of the resistive force and hence the muscles are at a mechanical “disadvantage”. However, such lever arrangements do have the advantage that they are very compact. Another quantity relating to levers is known as the range of motion (ROM) advantage. This is equal to the distance moved by the resistive force during an action divided by the distance moved by the motive force. In the above case, the shorter distance between the muscle attachment point and the pivot compared to the distance to the forces on the hand, means that the end of the muscle only has to move a small distance in order to move the hand by a relatively larger distance. The same is also true of the relative speeds of the points of application of the forces, in that a small speed at the muscle attachment point will produce a much greater speed at the hand. These ROM and speed advantages are particularly significant since a muscle will typically only have a small range of lengths over which it can exert its maximum force, and also the tensile force generated by a muscle tends to decrease as the rate of muscle shortening increases. Hence our musculoskeletal structures have clearly evolved towards an optimum configuration which balances the issue of mechanical disadvantage against compactness, ROM advantage and speed advantages.

2.5 Muscles Acting in Combination

So far we have considered the role of single muscles in musculoskeletal structures. However, since muscle tissue is only capable of generating contractile forces, then muscles have to act in groups which can oppose each other’s motion. Typically, when a particular motion is being performed, there will be a single muscle responsible for generating the majority of the motive force, known as the *prime mover*, and a group of smaller muscles known as *antagonists*, which act to counter the action and stabilise the joint. In order to further examine the cooperative role that opposing muscle groups play, we shall consider the muscle activity that occurs to generate flexing of the elbow such that the lower arm is moved quickly upwards from the anatomical position. In order to start

the arm moving the muscles at the front of the upper arm (i.e. the biceps brachii and the brachialis) are activated. This generates a net torque on the lower arm which gives it an angular acceleration upwards, and the arm begins to flex. As the arm approaches the intended rest position a force is needed in the opposite direction to decelerate the arm. This is provided by the muscles at the back of the upper arm (the triceps). The exact magnitude of the forces that are exerted by the muscle groups will depend on the amount of acceleration and deceleration required at the beginning and end of the motion. Performing this motion oneself, it is possible to see that it can be carried out with little effort such that the arm moves to the new intended position very quickly and precisely. This is a simple illustration of how muscles typically have to work together to initiate, control and terminate motion about joints. In most situations there are many muscles which have to be controlled in unison. The fact that we are able to do this subconsciously with extremely fine control over the spatial positioning and velocity profiles of the motion indicates how complex and refined our control of muscle activation is.

Although we have only concentrated on muscles which operate on a single joint so far, it's not uncommon for a muscle to act over two or more joints. The hamstrings (the collection of muscles at the back of the thigh) span the hip and the knee, and act to extend the hip and flex the knee. However, it is possible to isolate the action of the muscles to only one of the joints by using it in combination with the muscles at the front of the thigh (quadriceps femoris). For example, by activating the hamstrings as well as the muscles at the front of the thigh which span the hip joint, then we effectively lock the hip joint giving flexing at the knee only. Alternatively, if the muscles at the front of the thigh which span the knee are activated, then the knee will be fixed, and only extension of the thigh will occur. Of course these are the extreme cases, and it's possible to vary the activation of the muscles involved to obtain a range of motion at the hip and knee. The advantage of this type of system is that only three muscles are required to obtain independent flexing and extending of two neighbouring joints.

Another way in which muscles can act together is to preserve the length of a muscle during a motion in order to maintain optimum muscle strength. An example is when we attempt to produce forceful flexion of the knee. This primarily involves activation of the hamstrings. However as the resulting motion proceeds, then the hamstrings get shorter causing the force generated to decrease due to the nature of the contracting muscle fibres. In order to reduce this effect the muscles at the front of the thigh which span the hip can be contracted causing simultaneous flexing of the hip. This has the effect of increasing the distance that the musculotendon structures at the back of the leg have to travel and hence reduces the degree of muscle shortening.

Having discussed some of the musculoskeletal systems that exist throughout the body, the next section concentrates on a single limb and examines the structures present in more detail.

2.6 The Physiology of the Leg

This section takes the human leg as a case study, and identifies the musculoskeletal structures which exist there. The aim of this is to give a more complete description

of how muscles act together, and also to give an impression of the complexity of the muscle structures that exist in a single limb.

There are five main bones in the upper and lower leg (ignoring the bones in the ankle and foot) as shown in figure 2.10. These are the pelvis, femur, patella (knee-cap), tibia, and fibula. The hip joint lies between the pelvis and the femur. This is a three degree




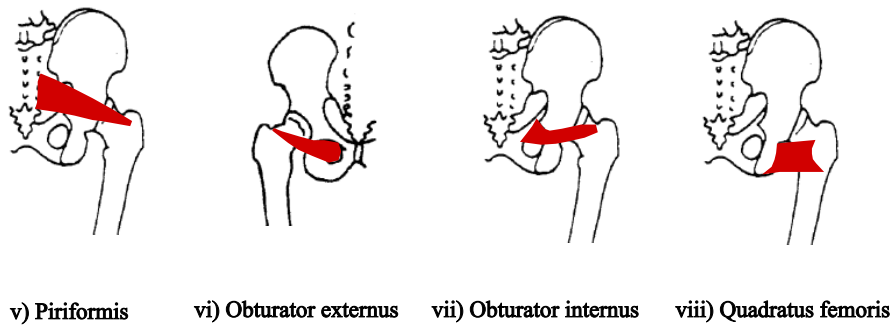

Figure 2.10: Bones in the upper and lower leg.

of freedom ball and socket joint which allows flexing, extending, adduction, abduction, lateral rotation, and medial rotation. The knee is the joint between the femur and the tibia and fibula. This is essentially a single DOF hinge joint allowing extension and flexion, although it also allows limited motion in the other axes. In order to generate motion about these joints there are a total of 26 muscles which lie on or around the leg. As well as these, there are also 12 muscles which lie within the lower leg, and are responsible for motion in the foot. The main muscle regions are the muscles in the *gluteal region* (the buttocks), the collection of muscles at the back of the thigh known as the *hamstrings*, the collection of muscles at the front of the thigh known as the *quadriceps femoris*, and the muscles in the lower leg. The muscles in each of these regions are listed and illustrated in figures 2.11 to 2.15, and have been grouped according to the motion that they generate (this is adapted from [102]).

These images show the position of the muscle within the leg and their points of attachment with the bones. The lighter regions represent tendons, and the darker regions muscle. They illustrate the variation that exists not only in the size and shape of the muscles, but also in the way that they span different joints, and attach themselves to the skeleton. It is clear, based on this study, that the musculoskeletal structures found

GROUP	PRIMARY FUNCTION	NAME	FUNCTION
Muscles of the gluteal region	Extensors, abductors and lateral rotators of the thigh	Gluteus maximus Gluteus medius Gluteus minimus Tensor fasciae latae Piriformis Obturator externus Obturator internus Superior gemellus Inferior gemellus Quadratus femoris	Extends and laterally rotates thigh. Abducts thigh. Abducts thigh. Extends knee. Abducts and medially rotates hip joint. Laterally rotates thigh. Laterally rotates thigh. Laterally rotates thigh. Laterally rotates thigh. Laterally rotates thigh. Laterally rotates thigh
	Main flexors of the thigh	Psoas Iliacus	Flexes thigh. Medially rotates thigh. Flexes thigh.
Muscles of the thigh	Adductors of the thigh	Adductor longus Pectineus Adductor magnus Gracilis Adductor brevis	Adducts thigh. Laterally rotates thigh. Flexes extended thigh. Adducts thigh. Adducts thigh. Extends thigh (posterior fibres) Adducts thigh. Flexes knee. Medially rotates leg. Adducts thigh.
	Extensors of the knee, flexors of the hip	Quadriceps femoris Rectus femoris Vastus lateralis Vastus medialis Vastus intermedius Sartorius	Extends leg. Flexes thigh. Extends leg. Extends leg. Extends leg. Flexes thigh. Flexes leg. Laterally rotates thigh.
	Flexors of the knee, extensors of the hip	Hamstrings Biceps femoris Semitendinosus Semimembranosus	Flexes knee. Extends hip. Flexes knee. Extends hip. Medially rotates flexed knee. Flexes knee. Extends hip. Medially rotates flexed knee.
Muscles of the knee		Popliteus	Laterally rotates femur. Flexes knee.
Muscles of the lower leg	Anterior group	Tibialis anterior Extensor digitorum longus Extensor hallucis	Dorsiflexes ankle. Inverts foot. Extends joints in toes. Dorsiflexes foot. Extends big toe. Dorsiflexes ankle.
	Peroneal group	Peroneus longus Peroneus brevis Peroneus tertius	Everts foot. Plantar flexes ankle. Everts foot. Plantar flexes ankle. Everts foot. Dorsiflexes ankle.
	Posterior group	Superficial Gastrocnemius Soleus Plantaris Deep Tibialis posterior Flexor digitorum longus Flexor hallucis longus	Plantar flexes ankle. Flexes knee. Plantar flexes ankle. Plantar flexes ankle. Plantar flexes ankle. Inverts foot. Flexes phalanges of toes. Flexes phalanges of big toe.

Figure 2.11: Muscles lying in, or generating motion of the leg.



a) Extensors, Abductors and Lateral Rotators of the Thigh.



i) Psoas and Iliacus

b) Flexors of the Thigh.

Figure 2.12: Muscles of the gluteal region.

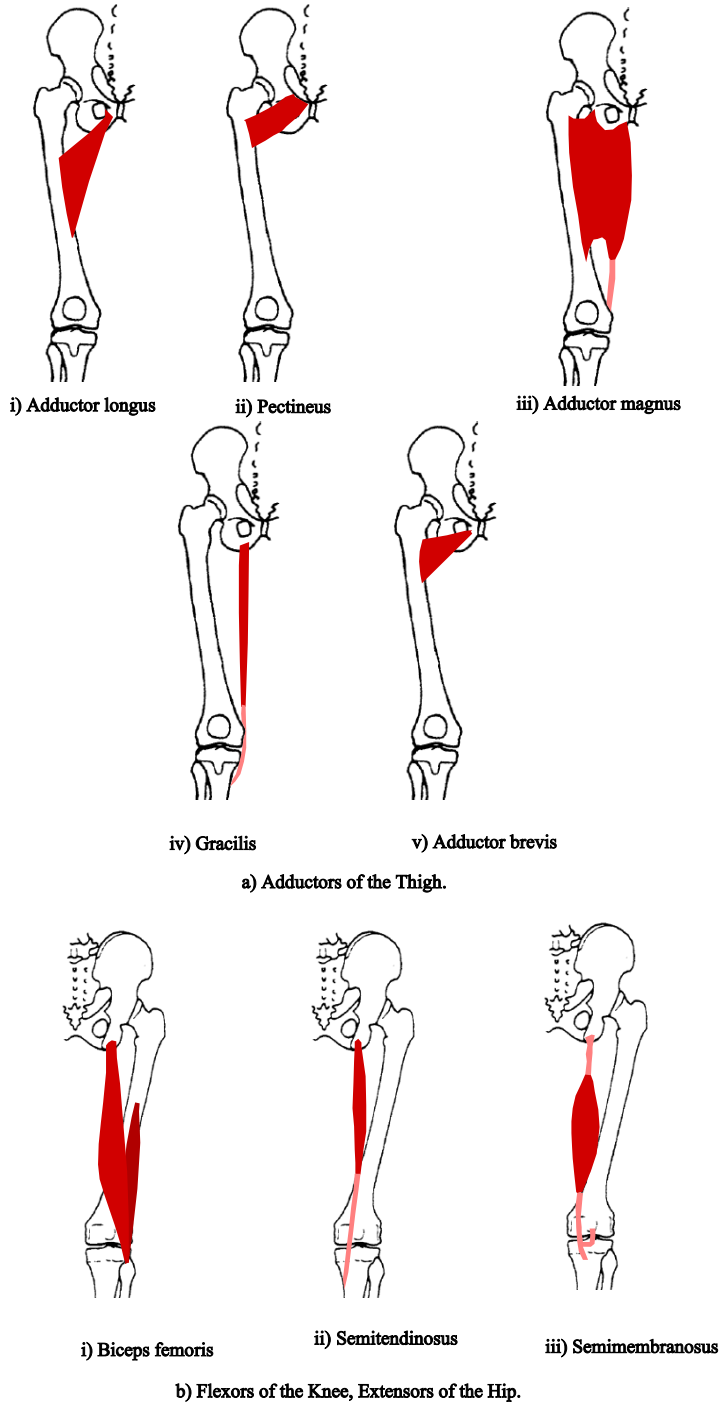


Figure 2.13: Muscles of the thigh.

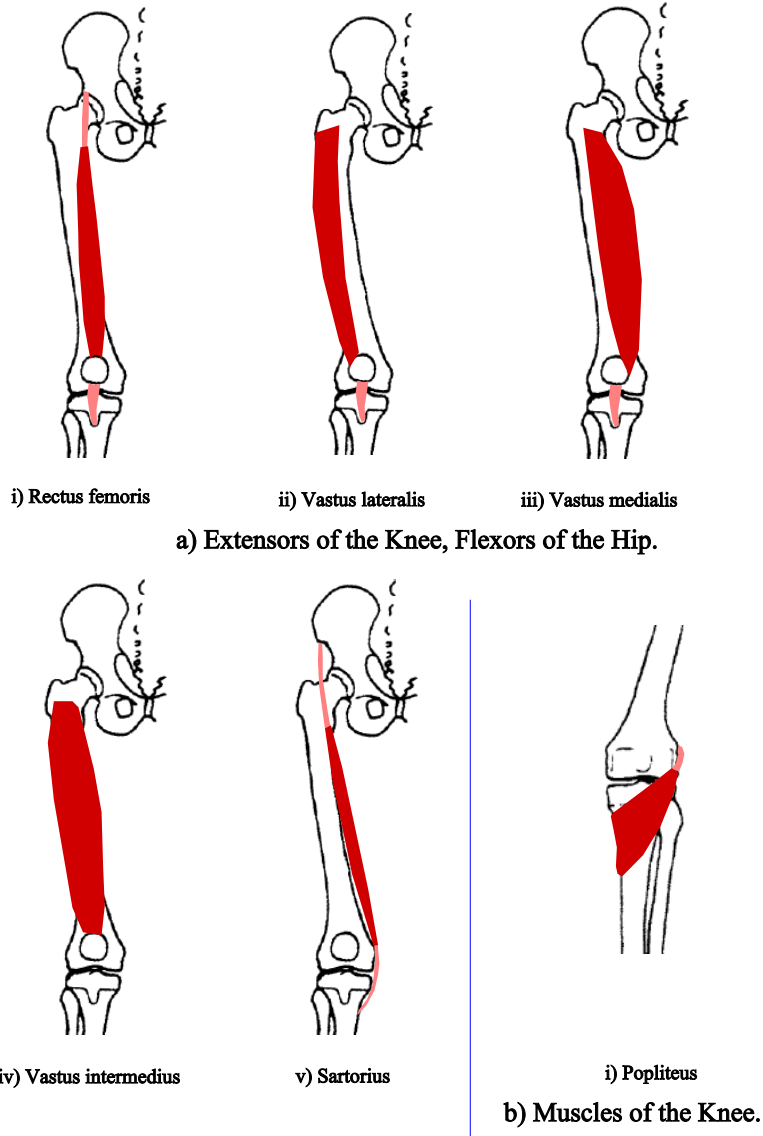


Figure 2.14: Muscles of the thigh and knee.

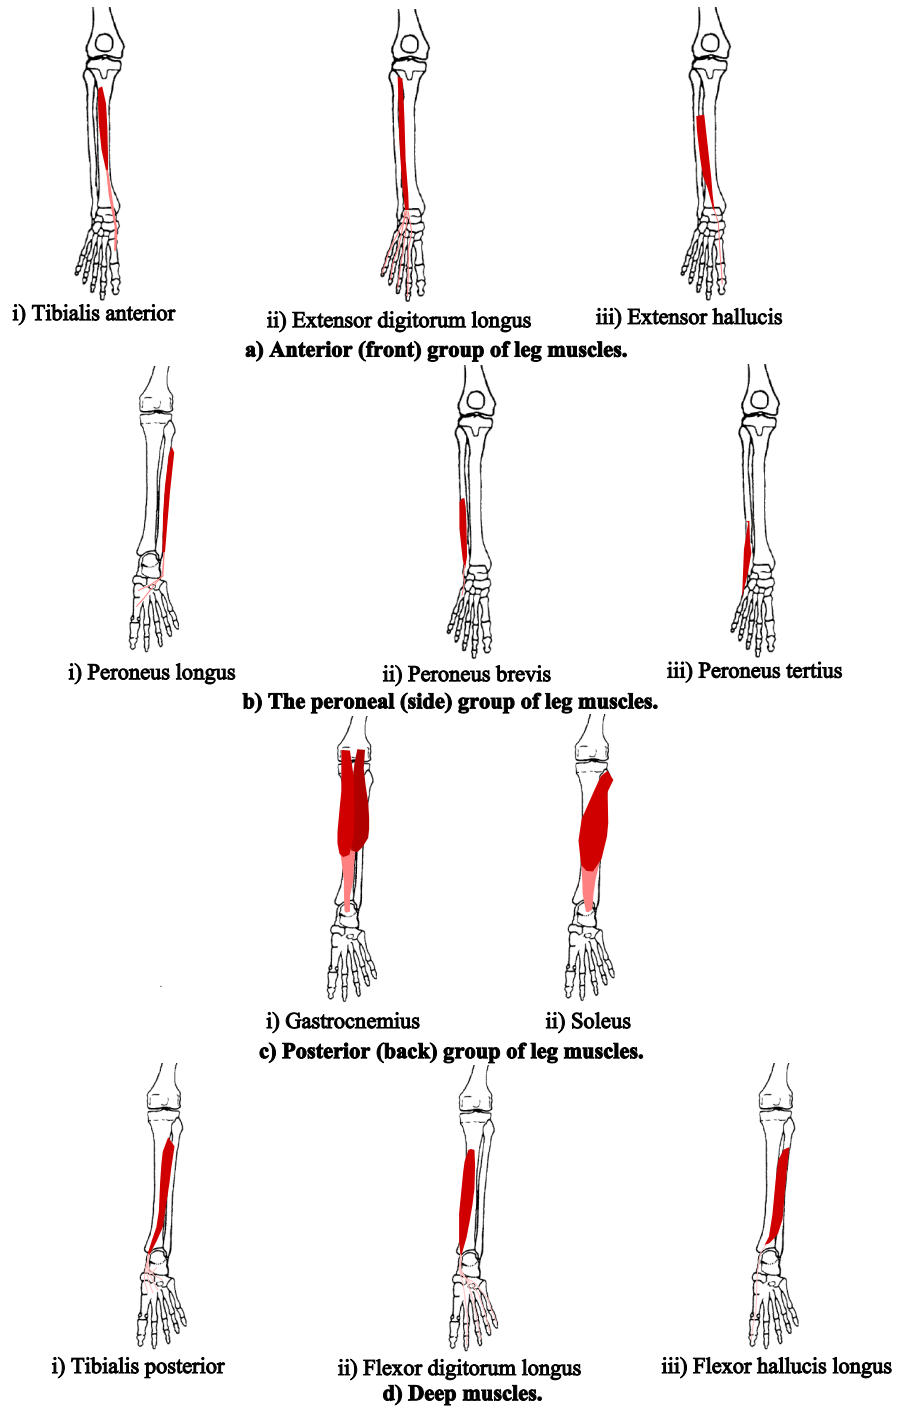


Figure 2.15: Muscles of the lower leg.

in just a single limb are extremely complex. As a result of this, a significant reduction in complexity was required in order to make this research feasible. The details of this are given in Chapter Four. The next section now moves away from anatomy, and considers instead the range of deformation effects that occurs in limbs, and the causes of these effects.

2.7 Types of Deformation Occurring in Human Limbs

This section attempts to provide an analysis of the main factors which influence limb shape, and the range of surface deformation effects that occur as a result. This is considered important since it describes the behaviour that our model is required to reproduce. It is also useful when discussing other research since many existing models make simplifications which prevent them from modelling the full range of influences and deformation effects that occur in reality. It is important at this stage to differentiate between changes in the shape of a limb and changes in the texture or colour of the skin surface. Changes in surface shape are often accompanied by changes in the surface texture due to the skin being stretched or squashed. However, this can be regarded as a separate area of research and will not be dealt with here.

2.7.1 The shape of a limb in the anatomical position

There are two main factors that determine the shape of a limb at a particular instance in time. The first is the structure of the various tissue regions that make up the limb itself, and the second is the deformation of those tissue regions caused by the current state of the limb (for example whether the joints are bent etc). In order to discuss how the tissue structures determine shape, it is useful to consider the limb in a “rest position” with minimal external influences. As an example we shall consider the leg in the anatomical position as shown in figures 2.16 and 2.17, although in actual fact it is not totally free of external influences since it is being influenced by gravity which causes slight sagging of the soft tissue regions.

The rest shape of a limb is determined by three main underlying tissue regions (see [82] for a more detailed discussion). These are the skeleton (consisting of bones and joints), the musculature (consisting of muscles and tendons), and the *panniculus adiposus* which is a layer of fatty tissue located just below the skin surface. The skeleton primarily determines surface form due to the fact that it supports all the other tissue regions. In some cases where the skeleton is close to the skin surface then the shape of the bones can be seen to have a direct influence. This can be seen in figure 2.16 around the knee (due to the influence of the knee cap) and also at the front of the shin (due to the shin bone). Musculature also has a great influence on surface shape since it accounts for most of the volume within the limb. As with bones, in areas where muscles and tendons lie close to the skin surface, their shape becomes more apparent. In figures 2.16 and 2.17 the shape at the front and back of the thigh, and the shape at the back of the calf is determined by the underlying muscles. Also, the skin surface behind the knee shows the structure of the muscles, tendons and ligaments that exist there. Finally, the layer of fat below the skin surface has the effect of smoothing it out,



Figure 2.16: Leg in the rest position (front).

and reducing the prominence of the underlying tissue structures. The thickness of this layer varies at different points in the body, and also tends to be thicker in women than in men.

2.7.2 Shape variation caused by limb state

In the rest state, a limb will adopt a certain shape based on the influences described above. As we perform actions the limb then becomes deformed away from this rest shape. There are four main influences which deform the shape of a limb away from its rest shape. Between them they represent the “state” of a particular limb. These are:

- Skeletal position
- Muscle activation
- Velocity state
- External forces

In real limbs, changes in state are based on complex interactions between each of these four factors. However, many existing models of limbs tend to grossly simplify the situation, which leads to inaccurate results. One common simplification is to determine shape changes based solely on the skeleton state (as seen in [62]). This is typically used



Figure 2.17: Leg in the rest position (back).

in models aimed at animation, with the justification that simplicity is required to enable the model to be used in real time, and high levels of accuracy are not essential. However, this approach rules out many significant types of deformation that occur in real limbs, and are discussed in the rest of this section. Another simplification used is to ignore the velocity state of the limb, and only consider the static deformation of the skin surface (as seen in [39]). This has the effect of neglecting dynamic effects, for example transient wobbling of the flesh after motion has taken place. Although it may be possible to generate realistic static deformation with this kind of model, the results can lack an element of realism. Other models ignore external forces, which prevents effects due to gravity and collisions with external objects from being modelled.

The real way in which these factors interact is complex. However, one of the most important relationships which should be understood is the way in which skeleton state and muscle activation are connected. For example, although muscle activation is the mechanism by which we create voluntary motion of the skeleton, muscle forces are only one of many forces which act on the skeleton. Hence, it is possible for skeleton motion to occur without muscle contraction. A simple example is if the arm is allowed to swing under gravity or if it is lifted by grabbing hold of it. Here, external forces are causing motion of the skeleton. Another consideration is that muscle activation doesn't necessarily lead to joint motion but still has a significant effect on surface shape. For example, if a limb is being held in a fixed position by an external object, then activation of the relevant muscles will not cause motion since the force generated by the muscles will be opposed by the reaction forces from the object. Reproducing the way in which these factors influence one another is important in order to model limbs accurately. We will now consider each of the four influences in more detail and discuss the main types of deformation that each causes.

Skeletal position

Skeletal position has a large influence on deformation. Bending at joints generally causes creasing and bulging of flesh on the inside of the joint and thinning of the flesh outside. Figures 2.18 and 2.19 shows examples of these effects around a bent elbow joint (note that the muscles in the arm are relaxed and hence the deformation show is due entirely to joint motion).

Creasing involves folding of the skin surface along a particular line. The location of such folds is often determined by the underlying anatomy, or alternatively occurs along specific lines of weakness on the skin surface (as seen in the hand and fingers). As well as creasing of the skin, we may also see wrinkling. This is essentially creasing on a smaller scale which involves only the skin itself, and not the underlying tissue regions. Temporary wrinkling of the skin (i.e. not wrinkles which are always present due to ageing etc) occurs when the skin surface is compressed. Bulging of non-muscular tissue around a joint occurs due to the fact that soft tissue is only slightly compressible and hence flesh which is displaced from the inside of a joint is forced to bulge out around the sides of it.



Figure 2.18: Deformation due to bending at the elbow (1).



Figure 2.19: Deformation due to bending at the elbow (2).



Figure 2.20: Effects of muscle activation.

Muscle activation

Muscle activation also has a large effect on surface shape since muscles themselves change shape as they contract. Figure 2.20 shows an arm with the biceps activated. Comparison with figure 2.18 shows a distinct bulge can be seen along the upper arm, and also the shape of the underlying muscles can be seen more clearly at the skin surface.

Although joint state and muscle activation generally lead to the effects described above, the actual dependency of skin shape on these influences is complex as discussed earlier. An example of this dependency can be seen if we consider how the shape of a limb changes as the state of one of the joints is changed, whilst the muscle activation is held at a constant level. Figures 2.21 to 2.23 show the elbow in three different states whilst the biceps activation is kept constant (as perceived by the person). In figure 2.21, the arm is extended and very little bulging is present at the skin surface due to the fact that the muscle is in a stretched state. In figure 2.22 the joint is slightly bent and the muscle is in a partially contracted state leading to some bulging. Finally, in figure 2.23 we see the elbow fully bent and more pronounced muscle bulging. Note, that in this state the muscle force is being opposed by reaction forces from the compressed tissue on the inside of the joint.

Another illustration of the dependency of limb shape on joint and muscle states can be seen if we maintain a fixed joint position whilst varying the activation of the muscles which act upon it. Figures 2.24 to 2.27 show the activation of the biceps and triceps varying whilst the elbow joint is fixed. In figure 2.24, we see the arm with the biceps slightly activated in order to support the weight of the forearm. In figure 2.25 the biceps are highly activated, but motion is being prevented by the other arm. In figure 2.26 the triceps at the back of the arm are now active, but again motion is being prevented by the other arm. Finally, in figure 2.27 both the biceps and triceps are



Figure 2.21: Varying joint state with fixed muscle activation (1).



Figure 2.22: Varying joint state with fixed muscle activation (2).



Figure 2.23: Varying joint state with fixed muscle activation (3).

activated such that they are canceling one another out. In each case we see variation in the surface deformation due to the varying muscle states.

As well as the effects described above, another effect can often be seen due to the action of tendons below the skin surface. This occurs when a muscle is activated, and a tendon is carrying the tension around the inside of a bent joint. Figure 2.28 shows the tendons which connect the biceps to the radius and ulnar bones in the forearm, protruding on the inside of the bent elbow joint. Other places where tendons are prominent at the skin surface are behind the knee, and around the underside of the wrist when the fist is clenched.

Velocity state

The mass and elasticity of flesh give rise to the static deformation effects described earlier, and also to transient dynamic effects for example wobbling of flesh following changes in motion of the skeleton. Although many such effects occurs while limbs are moving, they can be difficult to see unless the motion is slowed down. More easily observed effects occur when motion of a limb is stopped abruptly, for example if the foot is stamped of the floor, and the flesh continues to wobble for a short amount of time before coming to rest. The modes of vibration and their amplitude and frequency are determined by the structure and material properties of the tissue regions in the area of interest, which is also influenced by whether the muscles are tense or relaxed.

External forces

External forces comprise distributed forces which act evenly on all the tissue within a limb (of which gravity is the only practical example), and forces due to contact with external objects. Gravity generally causes sagging of the soft tissue regions. Contact forces cause deformation based on a number of factors including the shape of the object,



Figure 2.24: Varying muscle activation with fixed joint state (1).



Figure 2.25: Varying muscle activation with fixed joint state (2).



Figure 2.26: Varying muscle activation with fixed joint state (3).



Figure 2.27: Varying muscle activation with fixed joint state (4).



Figure 2.28: Tendon effects.

the force applied, and the properties of the tissue in and around the contact area. Again, the properties of the tissue and hence the deformation will vary depending on whether the muscles are tense or relaxed.

This section has highlighted the most prominent deformation effects that occur, for example bulging and creasing. However, it is important to note that these are the extremes in a continuum of deformation. In reality, the effects will range from only slight variations from the rest shape to more obvious surface effects of the type described above. For a model to be effective it should be capable of accurately reproducing this range. The final section in this chapter now looks at the material properties of the main tissue types found in limbs, and in particular considers which properties need to be modelled in order to recreate the deformation effects described above.

2.8 Material Properties of Human Tissue

Analysis of the material properties of human tissue also falls under the field of biomechanics. To do this quantitatively requires application of the established engineering principles of material science [33] [10]. In general, the properties of a material are determined by the number of substances that it is composed of, the way in which these are combined, and their individual properties. In the case of human tissue, complexity in each of these areas typically results in very complex properties. Another complication lies in the fact that we are dealing with living tissue. As a result, factors like whether the tissue is *in vivo* (still in place in a living human), or *in vitro* (removed from the body), and in the case of the latter, the length of time since it was removed from the body, have an effect when trying to determine the material properties. As a result, researchers who attempt to make such measurements have to take great care when setting up the experiments, and also present detailed information relating to the procedures used.

When examining the literature, it is surprisingly difficult to find precise data describing the properties of different tissue types. Where numerical values are presented they are usually embedded in lengthy discussion about the exact part of the body that was being considered. The reason for this is that the properties of, say muscle, can vary greatly from muscle to muscle, and also at different locations within the same muscle. For this reason, the data presented in this section is an amalgamation of results from different sources (including [33] [94] [52]).

The approach taken by many researchers has been to study soft tissue samples from humans and animals using tests which are commonly used in material science. These typically involve testing the deformation of uniform samples under applied loads [80] [54] [41] [105] [52]. One of the simplest tests that can be carried out is a uni-axial stress-strain test, which involves measuring the reaction force exerted by the sample as it is stretched and compressed in one direction.

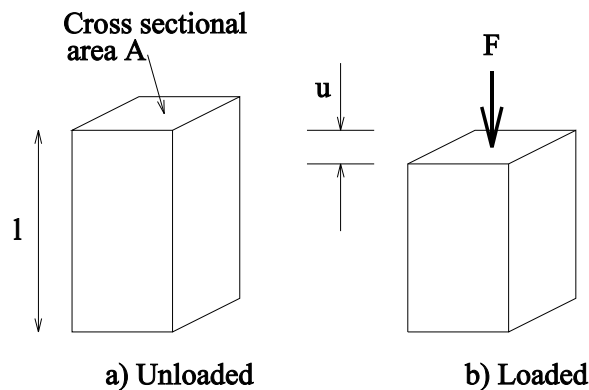


Figure 2.29: Testing a material sample.

Figure 2.29 shows a sample of material at two stages within such a test. As the sample is loaded under tension or compression by applying a force F to the end of the sample, its length changes by an amount u as seen in figure 2.29b. If we want to analyse how the particular material reacts to compressive and tensile forces, then ideally this should be done in a way which is independent of the dimensions of the sample. Clearly, if the cross sectional area of the sample is doubled then the force applied will act on twice as much material and will have only half the affect. Hence, rather than use the absolute force applied, we use the force per unit area which is referred to as the *nominal stress*, σ :

$$\sigma = F/A$$

Similarly, if we double the initial length of the specimen, then the changes in length for a particular force will be twice as large, so we use the change in length per initial unit length known as *nominal strain*:

$$\epsilon_n = u/l$$

For many materials, particularly engineering ones (e.g. metals), stress is found to be proportional to strain for very small values of strain. This behaviour is described as Hookean since it obeys Hooke's Law which says that $\sigma = E\epsilon$. The constant E is called Young's modulus. However, in the case of most real materials the behaviour soon becomes non-linear as the strain is increased. The simplest class of materials are known as elastic materials. Elastic materials deform when an external force is applied, but return back to their original shape when the force is removed. As with Hookean materials, most materials are elastic at low strains. However, if we go beyond the elastic range of a material, then it becomes deformed permanently and will not return to its original length when the applied forces are removed. This is known as plastic behaviour.

Although uni-axial stress-strain tests are the most common, other tests can be carried out to determine the properties of a material more fully. For example, it is possible to study the effect of applying strain to an object simultaneously in different orthogonal directions. Also, there are a number of other tests which can be carried out to determine the response of the material to other modes of stress e.g. shear stress (see [8] for more details).

General properties of soft tissue

The results of tests on human soft tissue samples of the kind described above, have shown that it has non-linear elastic properties. It is also viscoelastic, with the result that it has properties of both an elastic solid and a viscous liquid. The reason for this behaviour is largely due to the fact that it contains a high concentration of water. This viscoelasticity leads to the following types of behaviour:

- Hysteresis : changes in response under cyclic loading and unloading.
- Stress relaxation : reduction in the force opposing a deformation which is held constant over time.
- Creep : the increase in strain over time in a material held under constant load.
- Pre-conditioning : changes in deformation responses to a load when it is applied repeatedly.
- Another important outcome is that soft tissue is only slightly compressible.

The following sections now describe, in detail, the material properties of each of the main tissue types found in limbs, and in particular those involved in musculoskeletal systems.

2.8.1 Properties of skin

Skin is composed of three distinct layers known as the *epidermis*, *dermis*, and *hypodermis*, which have a typical combined thickness of several millimetres. The mechanical properties of skin are largely the result of the interaction of its two main component substances, elastin and collagen. Elastin is very similar to an ideal elastic material with




Figure 2.30: Stress-strain relationships for soft tissue types.

a linear stress-strain response over a wide range of deformations. Collagen has a much stronger stress response to applied loads, and also has a more limited range of elastic deformation. In skin, the proportions of each material type are approximately 72% collagen, 4% elastin, and 20% ground substance (based on the dry weight of dermal tissue [92]). This gives it a typical average density of $1.1g/cm^3$ [30]. Material tests on skin show that the stress response is approximately biphasic with a Young's modulus of $0.5MPa$ for strains below 0.75 and $200MPa$ for strains above. This is shown in figure 2.30 (along with similar graphs for muscle and general non-muscular tissue which are discussed later). This behaviour is thought to be due to the fact that the collagen is arranged in a lattice structure which is surrounded by elastin. At low strains, the collagen fibres in the lattice are not aligned with the direction of stress, and hence the behaviour is dominated by the elastin. However as strain increases, the fibres become aligned and the properties of the collagen begin to dominate [37]. Very often the fibres of the lattice are not symmetrical and form contour lines with fibres becoming aligned along a particular orientation. These lines correspond to lines of anisotropic deformation of the skin called Langer's lines, and here the skin surface is predisposed to creasing and wrinkling. If skin were composed only of collagen and elastic then it

would be very weak under compression since the collagen lattice would collapse. However, the presence of the water bearing ground substances e.g. fat, mean that it is only slightly compressible. This helps to explain the wrinkling behaviour of skin, since as compression forces are applied, its inability to compress results in buckling and hence wrinkling.

2.8.2 Properties of muscle

Skeletal muscle (the muscle used to control body motion) is the most complex of the major tissue types considered here. This is due to its complex structure and the fact that it has active as well as passive properties.

Muscle structure

Muscle is composed of fibres which come in a range of lengths, sometimes stretching the whole length of the muscle, and which are usually 10 to 100 microns in diameter. These muscle fibres are composed of still smaller elements called *myofibrils* that run the whole length of the fibre. Each myofibril is about 1 to 2 microns thick. Hence, a single muscle fibre contains hundreds to several thousands of myofibrils. Along the longitudinal axis of each myofibril is a repeating pattern of short filaments (typically 1 to 2 microns long) called *sarcomeres*. These are the functional unit in the muscle, and contract developing tension along their longitudinal axis. It is this shortening in series of many sarcomeres that creates the shortening of a single muscle fibre which in turn, along with many other muscle fibres acting in parallel, causes muscle contraction (see [33] for more details). A typical average density of muscle tissue is $1.04g/cm^3$ [30], although there is great variation amongst different muscles.

Passive properties

Although muscle is not composed of elastin or collagen, it does show passive material properties similar to those observed in skin and other soft tissues. Figure 2.30 shows a typical stress-strain plot for muscle. As with skin, the plot shows biphasic properties. This is again due to the way in which fibres are stretched into alignment as strain increases. Experiments have also shown that muscle is roughly volume preserving both when acting passively and when contracting [65].

Active properties

Unsurprisingly, the properties of active muscle vary considerably from those of passive muscle. Experiments which measure the force generated by real muscles, have shown that it is dependent on muscle length, velocity of shortening, and activation from the nervous control system. In addition, the force is affected by other phenomena such as fatigue and pre-stretching [76]. One of the classic experiments carried out on muscle tissue is to measure the way in which the contractile force generated at a constant level of muscle activation varies based on the length of the muscle. Figure 2.31 shows a typical plot of *normalized muscle force* against *normalized muscle length* for a sample.

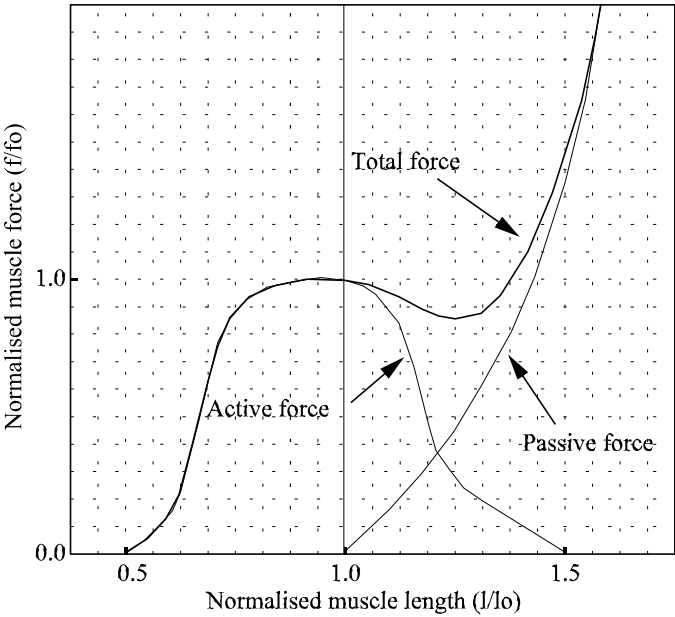


Figure 2.31: Typical force-length relationship for muscle tissue at a particular activation level.

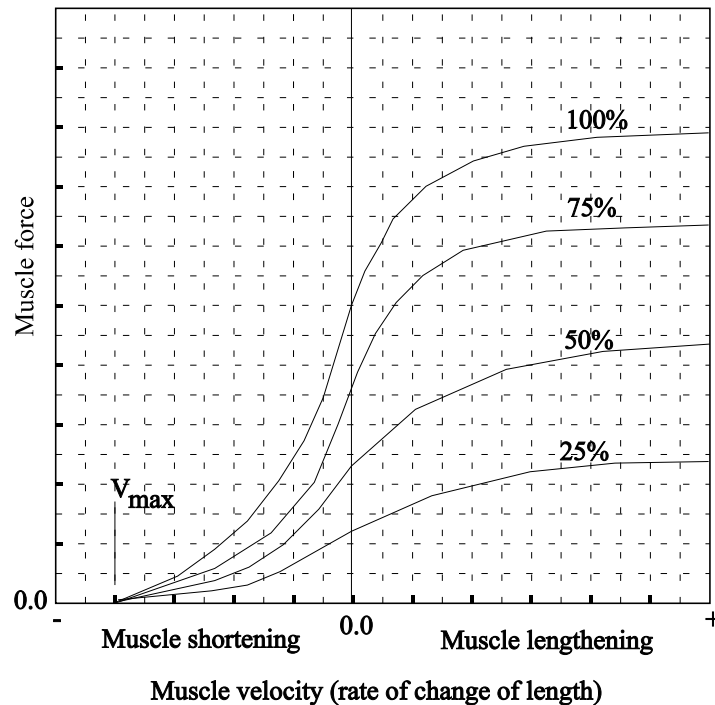


Figure 2.32: Typical force-velocity relationship at various muscle activations.

Normalized muscle force is the force generated at a particular length divided by the force generated at the initial rest length. Normalised muscle length is the current length divided by the initial rest length, hence a value of one represents the muscle at its original length (as indicated by the solid grid line). The graph is essentially the passive plot shown in figure 2.30 combined with the force generated by the active contraction of the muscle. As shown, for a given muscle activation the active force is largest around the rest length of the muscle and then drops off as the muscle length increases or decreases. Another classic experiment examines how the active force varies based on the rate of muscle shortening or lengthening [44]. Figure 2.32 shows a typical set of curves for a muscle sample at four different activation levels. For a given activation level, it can be seen that the force drops as the rate of muscle shortening increases, and increases as the rate of muscle lengthening increases. V_{max} shows the maximum rate of shortening at which the force generated by the muscle drops to zero. At this point the rate of contraction is such that the contracting muscle fibres cannot keep up and the muscle becomes slack. The graphs presented in this section illustrate the significant effects that muscle length and rate of change of length have on the force generated. This backs up the discussion earlier on in this chapter (section 2.4) which described how the geometry of musculoskeletal structures is such that it attempts to minimise

these effects.

2.8.3 General non-muscular tissue

With regard to the mechanical operation of limbs, the remaining tissue regions for example blood vessels, nerve fibres, fat, and so on, can be regarded as a single material. Figure 2.30 shows a typical stress-strain graph for this “agglomerate” of non-muscular tissue (obtained from [52]). Again, the graph is biphasic and shows values lower than passive muscle at low strains and higher at higher strains. The material is again only slightly compressible due to the high concentration of water, and this accounts for the general bulging effects that are seen around areas where flesh is being compressed in a particular direction e.g. on the inside of bent joints. A typical value for the density of this tissue is $0.92g/cm^3$ (the value for adipose fat obtained from [30]). Note that this is less than water (which has a density of approximately $1.0g/cm^3$), whereas the other tissue types considered here all have densities above that of water.

2.8.4 Tendons




Figure 2.33: Stress-strain properties of tendon.

Tendons are chiefly composed of collagen fibres arranged in a lattice. A typical

stress-strain curve for tendon is shown in figure 2.33 (obtained from [94]). This shows approximately biphasic properties with the lower part having a Young's modulus of 36.4 MPa and the upper part 1110 MPa. As with skin, the first phase corresponds to the collagen fibres being misaligned, and the second after they have been stretched into alignment. The graph also shows the stress-strain properties of pure collagen for comparison, which can be seen to match the second phase of tendon. The Young's modulus of tendon can be seen to be much higher than for the other soft tissue types that have been considered. This is necessary since tendons are required to carry large loads, and excessive extension would lead to problems in limb function. A typical average density for tendon is $1.1g/cm^3$ [30].

2.8.5 Properties of bone

Bone tissue shows deformation when subjected to a variety of external loads. Tests have shown that it also has biphasic and viscoelastic properties, behaving elastically at low strain and having plastic characteristics at high strain to the point of fracture [73]. Values for the elastic modulus of bone range between 14 and 22 GPa. Analysis of bone structure reveals that there is a built in safety factor [19] such that under normal conditions the stresses encountered remain well below the fracture point. As a result of this, bones undergo very little deformation during everyday situations, and can be considered as non-deformable rigid body objects. With this in mind, it is only necessary to know a bone's mass and moment of inertia about its main axes in order to be able to model it. These can be obtained for real bones by experiment, or can be approximated from its density (dry bone density is approximately $2.0 g/cm^3$ [66]) and its mass distribution.

2.9 Conclusions

This chapter has considered a number of issues relating to the form and function of real human limbs. The first sections discussed anatomy in general and how tissue structures are arranged in limbs to create musculoskeletal systems. In order to emphasise the complexity of the musculoskeletal structures that exist in a limb, we then looked in detail at the human leg. Having formed a general understanding of limb structure, we then turned our attention to considering the factors that determine the initial shape of a limb, and how this shape changes as the limb moves. To summarise, we defined the main types of surface deformation as:

- Creasing and wrinkling
- Bulging of muscular and non-muscular tissue
- Dynamic effects e.g. wobbling of flesh
- Tendon effects
- Deformation due to gravity

- Deformation due to surface contact

We also determined the main factors that give rise to these effects as:

- Joint state
- Muscle activation
- Velocity state
- External forces

Correctly modelling the way in which these factors interact is important if we are to produce a model which behaves in an accurate and realistic way. However, as discussed, most existing models incorporate only crude approximations to this behaviour. Specific examples of this are given in the next chapter. The final section considered the main material properties of the tissue types involved in musculoskeletal structures, and described how these properties lead to some of the deformation effects discussed in section 2.7.

The impetus behind this chapter was the need to have a good understanding of real limbs in order to be able to effectively criticise existing models and construct a new model of our own. One of the main insights obtained here is of the complexity that exists in real limbs. This applies to the range of different tissue types involved, the structures that they form, and also their material properties. Clearly, in order to be able to model a limb, a great amount of simplification is required to make the work feasible in terms of its implementation and the amount of processing needed during simulation. In Chapters Four and Five we draw on the knowledge gained here to justify the simplifications that are made in our model.

Chapter 3

Previous Research

As mentioned in the introduction, most of the existing research into modelling parts of the human body has been carried out by either the computer graphics or the medical and biomechanics communities. Models aimed at computer animation applications are typically empirical, and attempt to reproduce the general motion and deformation of a character's skin using simple real time techniques. Real time operation is desirable since computer animation (as with all types of animation) is a highly interactive process in which the animator refines the motion/deformation of a character over many iterations, and hence models which display results quickly aid this process. An example of a large research project centred around animation applications is the European HUMANOID project [9]. This aimed to develop models which could be used by "games designers, multimedia designers and film producers" to produce simulations and realistic animation. The project considered a wide range of problems including general motion and deformation, and also modelling of hair and clothes to give added realism. They also considered the possibility of autonomous models which are capable of understanding high level commands and interacting with one another.

The other main research area is in the development of realistic modelling techniques for use in engineering and medical applications. Here, the requirements usually differ from those for animation. For example, it is often necessary to model internal features as well as the skin surface, and also a much greater degree of accuracy is usually required. In many cases, the latter requirement is not merely a desire but a necessity, since the application may be used in a safety critical scenario. An example of an engineering application is the use of computer models in the design of human environments. For example, in the car industry, a human model may be used to test the ergonomics of, say, a new dashboard layout which is being designed using a CAD system. This offers vast potential improvements in efficiency over the traditional approach of building full size physical mockups. An example of a research project in this area that developed into a commercial product is "JACK" [12]. This is a full model of a human which accurately represents a range of physical and physiological aspects. This includes accurate modelling of the mechanical properties of limb segments and joints (although deformation of the skin surface is not modelled) as well as psychological information relating to how comfortable a particular body pose is. As well as ergonomic design applications, JACK is also intended for applications which require

the modelling of “intelligent agents acting on instructions” which could be used to test, say, the behaviour of a crowd of people entering or leaving a sports arena for safety purposes, “factory or assembly workers” to test the efficiency of working environments, and “semi-autonomous soldiers” for battlefield simulations.

A third potentially far reaching application area, is in medicine. However, due to the infancy of research in this area, and the stringent requirements involved, very little work has resulted in real practical applications. An example of a potential application is the use of computer models for teaching applications. For example, a piece of software could be used as an anatomy instructor to present students with accurate three dimensional models of body parts. These could be viewed from any orientation, made to highlight a particular organ, animated to show dynamic effects, and generally used to portray information in a highly interactive and effective manner. Another more complex application would be to use a computer based model to simulate the effects of surgical techniques. For example, Koch *et al.* [55] and Pieper *et al.* [75] have developed models of the human face which aim to simulate the deformation that results when a particular plastic surgery technique is used. Such a model could be used by a surgeon prior to an operation to assess a number of approaches and determine which is likely to give the best post-operative results. Finally, an even more potentially useful medical application, is one which allows a student or surgeon to simulate an operation in real time. For example Hollands and Trowbridge [49] describe a “keyhole surgery trainer”. This features an “arthroscope” which is an optical device that can be inserted through small incisions and allows a surgeon to see inside, and perform surgery on, a particular part of the body. However, instead of being presented with a real image, the surgeon sees a computer generated one based on a simulation of the part of the body that is being operated on. As well as using an arthroscope, the surgeon is also able to use the standard range of surgical instruments that are used in real operations of this type. This is achieved by attaching them to various sensors and actuators in order to detect the actions carried out by the surgeon, and recreate the forces that would be felt through the instruments. Applications of this type are particularly demanding due to the fact that need to be particularly accurate, and also need to run in real time.

Due to the difference in requirements of animation and medical applications, models created so far in each field have been fairly distinct. However, the gradual progression of medical research into modelling more complete areas of the human body, and the need of the animation community for greater accuracy and hence realism, means that the gap between the two fields is becoming narrower. This thesis aims to continue this process by considering techniques that have typically only been used in medical/biomechanics models, and developing them for the modelling of entire limbs. In the case of animation applications it is hoped that such models would offer new levels of realism that are not possible with existing techniques. In medicine, the immediate potential applications of such a model are not as forthcoming since the wide scope of the model will inevitably necessitate simplifications which reduce the accuracy. However, use in low accuracy applications is one potential area that is considered here.

This chapter now looks in detail at a range of existing models related to this research. The following sections look at models aimed at animation and also considers general issues, for example techniques for representing surfaces, and the control of skeleton

models. The final part of the chapter then considers more accurate models from the medical and biomechanics communities.

3.1 Character Animation

Early models of humans concentrated more on general motion than surface deformation. Withrow [103] and Thalmann *et al.* [88] used simple stick figures to represent characters. As work progressed, improving the appearance of the models became an issue, and three dimensional primitives were used to “fill out” the bodies of the characters. Potter and Willmert [79] used cylinders, however these gave poor results around joints, and hidden surface removal for these primitives is time consuming. Badler [13] Balder used spheres to create his “Bubbleman” whilst Herbison-Evans [42] [43] and Ginsberg and Maxwell [64] used ellipsoids. The attraction of these primitives at the time (before hardware accelerated rendering) was that they could be processed and drawn efficiently to include hidden-face removal and correct intersections. This allowed researchers to compute and display the motion of the figures in real time.

As graphics hardware developed, researchers began using polygon meshes to represent the skin surface. In the simplest cases, each limb segment is represented by a polyhedron which is rigidly attached to the underlying skeleton. However, this tends to result in undesirable effects due to the discontinuities around joints. Nevertheless, this approach is still used in applications where computational requirements, or the number of polygons, has to be kept to a minimum for example in computer games.

In research which was aiming for realism, a natural progression was to use a single continuous polygon mesh for the whole body. The main challenge posed by this approach is how to deform the skin surface to show different character poses. One method is to specify the surface shape for a number of “key” poses, and then to interpolate between them to generate new surface shapes. However, specifying the surface shape of even a single pose is time consuming, and a large number of these are typically required to get the necessary range of poses. A more efficient approach is to use a basic skeleton (defined by line segments connected by multiple degree of freedom joints) to define the pose of the character, and then deform the skin surface around this skeleton. The simplest way of achieving this is to assign vertices on the skin surface to skeleton segments based on proximity and then to transform the points rigidly with the segment. The resulting models are essentially the same as the segmented models described above, but with polygons across the joints. However, this approach also tends to give poor results for large joint angles.

More complex approaches attempt to define how the entire space around a joint deforms as a function of the joint angles, and then embed the skin surface within this space such that the mesh vertices deform along with it. Thalmann and Magnenat-Thalmann [63] [90] [62] use a technique based on *Joint-dependent Local Deformation* (JLD) operators which are individually created for different types of joint. Each vertex in the skin surface mesh is associated with one or more of the segments in the underlying skeleton and JLD’s are used to map vertices onto new transformed positions based on the state of the skeleton. JLD’s have been implemented which maintain a smooth skin surface around different types of joint for a wide range of angles, as well as recreating

common effects for example bulging and creases. They have also been used to create models of entire human bodies. Images of a hand presented in the work show good visual.

Badler *et al.* [14] use a similar approach, but rather than using a general function to represent the deformation around joints, a “free form deformation” (FFD) [83] approach is used. FFD’s are a technique for deforming or warping a region in space (defined by groups of parallelepipeds) by altering the position of a set of control points. Objects to be deformed are placed within this space, and are made to deform with it. One of the main advantages of this method is that it is independent of the particular approach being used to model the surface geometry, and hence it can be used with a wide variety of primitives. In Badler’s work, the skin surface is represented by a polygon mesh embedded within an FFD. The control points are attached to the bones either side of the joint causing the FFD, and hence the skin, to deform as the joint angles change.

Models of the type described above are known as *geometric models* since they utilise simple deformations of the model geometry in order to get the required effects. They are typically aimed at animation applications since they are fast to simulate, however they have a number of limitations. Firstly, the deformation of the skin surface is based solely on the state of the skeleton. For effects like muscle bulging, this usually means that a deformation function is devised which shows only a small degree of bulging when the limb is outstretched, and then as the limb is progressively bent, the bulging increases. Clearly this is a gross simplification of the real situation, and the complex dependency of surface shape on muscle and skeleton state that was discussed in Chapter Two (section 2.7) is not considered. This simplification also means that other effects, for example, dynamic effects and deformation due to external forces are also not present. Another fundamental limitation is that geometric models of this kind are only capable of reproducing the behaviour that has been specifically designed into them. As a result, they will not respond realistically to joint states which were not anticipated or were simply left out of the model. Finally, in the case of models based on deformation operators, another problem is that operators cannot typically be reused on joints with different degrees of freedom, and even where the joints do have the same degrees of freedom, substantial tuning is often required to produce realistic results.

In general, it is difficult to create geometric models which are capable of exhibiting a wide range of effects, based on a large number of input parameters. For example, if we consider trying to model the change in shape of a deformable surface due to another object colliding with it. Here, there are a large number of parameters to consider, for example, the momentum of the object, the angle at which it hits the surface, its shape, the position of the collision on the surface, the properties of the deformable material etc. Trying to prescribe a function which can map such a wide parameter space onto a resulting surface shape poses fundamental problems (although it is being attempted [32]).

An alternative type of model makes use of the computational “prowess” of computers in order to reproduce deformation effects through simulations of known physical laws. These are known as “physically based models”. Gouret *et al.* [39] present a physically based model of a hand consisting of a skeleton (driven by a “classical” control method i.e. parametric, kinematic or dynamic), a layer of deformable flesh, and a

geometric skin surface. The method uses the Finite Element Method (FEM) to model the deformation occurring in the hand as it grasps a deformable ball. FEM is a technique commonly used in engineering when modelling the deformation of components or structures. Gourret uses an approach in which the skeleton, the surrounding flesh, and the ball are all part of the same simulation and hence can influence one another. The result is that the method automatically produces realistic deformation of the hand based on the state of the joints, the shape of the underlying bones, and the external reaction forces from the ball. One disadvantage with the particular method used is that it models only static deformation, and hence dynamic and transient effects are ignored (although this is not necessarily a consequence of using the FEM approach). The approach is also computationally expensive, and so may not be ideal for animation applications, particularly if it is necessary to model an entire body or limb. One of the main simplifications used in the model, is that it only contains a single layer of homogeneous flesh, and discards muscle and other tissue types. However, since the hand is primarily composed of bone surrounded by a thin layer of flesh, this is a reasonable approximation.

Gourret uses a purely physically based approach to determine flesh deformation. However, many models aimed at animation, have combined physically based and geometric techniques in order to try and get the best of both worlds. Chadwick *et al.* [21] present a method which uses a geometric model to generate the basic skin surface deformation, and then applies a physically based approach on top of this to add dynamic inertia effects. The geometric approach used is similar to that used by Badler (described above), in that a skeleton is surrounded by FFD's which deform an embedded polygon mesh representing the skin surface. However, here the deformation is based on an algorithm which takes the joint state as its input and generates new positions for the control points (in a similar way in which JLD operators are used to determine vertex positions). The algorithms used attempt to reproduce simple effects, for example, bulging and creasing around joints, and muscle bulging. This approach is used by Chadwick to model deformation of the skin on an arm based on bending of the elbow joint. The physically based part of the model features point masses which are attached to the control points of the FFD meshes, and are connected to one another using a network of springs and dampers. Point masses, springs and dampers are mechanical elements which can be connected together to form three dimensional meshes, which when simulated, model the behaviour of regions of deformable material. They are "pure" elements in that a point mass (as its name suggests) represents a mass with no physical extent, springs exert forces based purely on their length, and dampers exert forces based purely on their rate of change of length. When the mass-spring-damper (MSD) mesh devised by Chadwick is combined with the FFD model, the effect is that inertia effects in the MSD mesh are superimposed onto the deformation generated by the geometric model. Although this approach reproduces more effects than the purely geometric models described so far, it still represents a gross approximation of real limbs. Again, external forces and muscle activation are not considered (muscle bulge is directly connected to joint angle).

Physically based models as a whole have been applied by many researchers to the problem of generating realistic behaviour in models intended for animation applications.

The next section considers their use in more detail.

3.2 Physically Based Modelling in Animation

Researchers have made use of physically based models in a wide range of situations and at a variety of levels. The simplest way in which they can be applied is to generate secondary effects based on motion that has already been defined by the animator. An example of this is the modelling of hair [26] [51] [57]. Here the motion of the hair is simulated as a post processing stage and is driven by the motion of the head which has been predetermined using a standard animation technique. The result is that the head affects the hair, but not vice-versa. This is a reasonable approximation, since the mass of the hair is typically small compared to that of the head. Another example is the use of cloth models to represent clothes [20][89] [95]. Again these respond to the predetermined motion of the skin surface but do not influence it.

This type of approach is often useful for adding realism to a piece of animation. However it is only suitable in situations where one way interaction is a reasonable approximation to the real thing. In the work by Chadwick described earlier [21], the physically based model is used to determine the dynamic motion of the skin surface based on the motion of the skeleton. Here, strictly speaking, such one way interaction is not a good approximation since in reality flesh has a significant effect on skeleton motion. In other research, rather than being used to generate secondary effects, physically based techniques are used as an integral part of the system. The work described above by Gourret [39] has the potential to be of this type if the motion of the skeleton is determined using a dynamic (force based) approach rather than a scripted one. In this case, the forces that are being exerted on the bones from the flesh would be combined with other motive forces that are acting on the skeleton, and hence the flesh would have a direct influence on skeleton motion.

Another example of this kind of physically based approach is the work by Gascuel *et al.* [35]. Here a limb model is presented which incorporates both a skeleton, and a surrounding deformable layer. Each limb segment consists of a rigid “kernel” which has springs protruding from it in all directions as shown in figure 3.1. The springs “detect” and respond to collisions with external objects, and also determine deformation when the surfaces associated with neighboring links come into contact around joints. The springs are not simulated as with classical mass-spring-damper systems [31][67], but rather their length is determined by an equation which describes the sinusoidal behaviour with exponential decay that is exhibited by a mass connected to a parallel spring and damper arrangement. When the end of a spring “collides” with an object its end point is moved to prevent interpenetration from taking place, and then the deformation is propagated to neighbouring springs. By changing the way in which this is carried out, a variety of effects can be “mimicked”, for example, constant volume and constant surface area deformations. The springs also exert reaction forces on the bones that they are connected to, ensuring that general bending at joints and external collisions also affect motion of the skeleton.

This section has introduced a few ways in which physically based techniques can be incorporated into animation systems. However, their use gives rise to a number of new




Figure 3.1: Limb model (adapted from an illustration in [35]).

issues, for example, whether the skeleton motion is defined entirely by the animator or whether it can be influenced by the physically based modelling technique. In order to explore these issues further, we now consider the animation process more closely, and discuss how the use of physically based models ties in with this process.

3.2.1 The animation process

Animating a computer based character essentially involves instructing it to behave in a desired manner. However, in the case of a character which is required to show general limb motion as well as realistic deformation of the skin surface, the number of variables which need to be specified is potentially huge. As a result, techniques are required to make the animation process more efficient. Geometric and physically based models aim to reduce animator effort by automatically specifying the low level motion and deformation, based on a manageable number of higher level parameters.

In the case of the geometric models describe earlier, the position of the vertices which make up the skin surface is automatically determined from the state of the skeleton. However, a potential problem with this approach is that the animator loses direct control over the animation that is created, and may end up with the motion and deformation that the system produces rather than that which was desired. A solution is to allow the animator to have control at a range of levels, such that low level parameters can be determined automatically for efficiency, but can also be “tweaked” to obtain the exact results required. This allows an intuitive approach to be adopted, in which the animator initially specifies the approximate motion and deformation of the character, and then goes through a series of iterations modifying the model at finer scales, until the desired results are obtained. An analogy is if we imagine an artist shaping a block

of clay to produce a sculpture. The artist will initially shape the clay into the general shape of the finished object, then concentrate on refining local areas, and finally add the surface detail. Unfortunately, attempting to combine physically based simulation techniques with traditional animation approaches in order to give this kind of control, gives rise to problems due to the fact that the techniques are essentially at odds with one another. The next section considers these issues in more depth.

3.3 Animation vs. Simulation

Animation is traditionally concerned with allowing an animator to precisely define the required motion and deformation of a character. Simulation, on the other hand, is based around automatically generating results without the need for user input. Hence, an important issue is to consider how these approaches can be combined to reap the benefits of physically based models (automatic motion specification and realism) whilst still allowing animator control. A simple approach is to manually tweak the results of a simulation. For example, in the work described above by Chadwick *et al.* [21], after the motion of the FFD control points has been determined from the physical simulation of the mass-spring-damper system, the animator could then manually alter the motion of certain control points. However, this is likely to be time consuming, and the animator would have to take care to move the control points in a physically realistic way in order to avoid introducing an artificial appearance back into the animation. An alternative approach is for the animator to “steer” the simulation by imposing external forces or other physical inputs. The results will then remain physically “valid” since the model considers them just like its own internal forces.

An example where this approach has been considered is in the control of physically based skeletons. Devising a model which describes the rigid body physics of an articulated skeleton representing a particular character is relatively straight forward to do [99] [98] [7] [100]. It is also a simple task to run simulations which show the model behaving as an “inanimate” passive object (for example colliding realistically with other objects in the scene, or swinging from a rail). However, in order to make the character perform actions we have to be able to control its motion. The simplest method for doing this is known as *forward dynamics*. This involves explicitly specifying the duration and magnitude of forces and torques that you wish to apply to the figure, in order to make it carry out the required motion. However, this approach proves to be extremely difficult and completely counterintuitive. One reason for this is the lag between applying a force to a limb segment, and it actually moving (since it has to accelerate). An alternative approach is referred to as *inverse dynamics*. Here, the user specifies the required motion of the character, and the system devises a set of forces which when applied over time, will generate this motion. Whilst this doesn’t particularly achieve anything for a single character in isolation (i.e. the motion could have been scripted), it does give a valid way of controlling objects which are interacting with other objects in a physically based simulation.

An alternative approach known as *spacetime constraints* was proposed by Kass and Witkin [104]. Here the animator specifies *what* the character has to do e.g. “jump from A to B, clearing a hurdle in between” and the manner in which it should be done

e.g. “don’t waste energy”, rather than having to specify exactly *how* to do it. The system then generates a valid set of forces and torques which will carry out the action whilst optimising any constraints imposed. Results presented show a simple object (an animated angle-poised lamp) which has to jump over an obstacle, and the authors demonstrate that by applying different criteria the system generates physically realistic motion which also has “character”. This embodiment of subtle effects in the animation e.g. anticipation, squash-and-stretch, and follow-through¹, occurs for the same physical reasons that they occur in real people and animals. For example, anticipation describes the way in which a person might move backward slightly before jumping forward in order to allow more forward momentum to be built up before the motion is carried out.

This type of approach offers an intuitive way of creating realistic motion. Unfortunately however, the computation required quickly becomes impractical for even modest skeletons, and it is sometimes difficult to determine which constraints are needed to get the required results. Nevertheless, a number of researchers are attempting to overcome these problems [23] [72] [61].

3.4 Layered Models

A general approach to the problem of controlling models at different levels which has been used in work containing geometric and physically based techniques is to try and split the model up into sub-parts or layers, such that the animator can “affect” an intermediate result before it is passed onto the next layer. An example of a layered model that has already been introduced is the work by Chadwick [21]. Here the specification of the skeleton, the warping of the space around joints by the FFD’s, and the physically based mass-spring-damper system are all distinct layers. The authors say that this type of approach ...

...provides a methodology where more emphasis is placed on constructing the character so that less emphasis is needed in actually scripting an animation.

Another example of a layered model is the work by Turner [93]. The main feature of this work is a physically-based elastic skin surface which “contracts around” underlying geometric features. The skin surface is modelled using a Finite Difference Method (FDM) approach. Beneath the skin are a number of entities which represent the musculature and tissue layers within the character. The first is a standard articulated skeleton. Attached to this are muscles which are modelled using implicit surfaces in the shape of spheres, cylinders and super-ellipses [17] together with global deformation functions [18]. The skin surface wraps around the whole object due to its internal tension, but is prevented from penetrating the muscles by the use of reaction constraints [77]. The paper shows images of a model of the human torso which features realistic looking skin deformation. However the approach may not be particularly successful at modelling features that occur on limbs, for example creasing, since the skin

¹These were terms introduced in a book describing the techniques used by Disney animators to bring characters to life [91].

tension will tend to pull the skin surface away from the inside of the creases (although the use of spring constraints is suggested to prevent this). Also, no mention is made of animating the underlying muscle primitives based on joint motion in order to model muscle bulging for example.

So far we have presented examples of work aimed mainly at modelling the deformation of limbs and of the torso. However, a large amount of work has also been carried out into modelling the appearance and motion of the face, particularly when talking. Although the anatomy of the face is quite different from that of limbs (a relatively thin layer of flesh attached to a bony substrate (the skull) which contains muscle arranged in sheets and other other complex structures), many of the techniques used are still relevant.

3.5 Facial Modelling

In a landmark paper [96], Waters presents a geometric approach to modelling deformation of the skin surface during facial expressions, based on consideration of anatomical features. In particular, this considers the various forms of muscle that occur in the face, and the way in which they influence surrounding tissue. A number of muscle “primitives” are presented e.g. linear, sphincter, elliptical sphincter, sheet, and are placed at specific positions beneath the skin surface corresponding to the location of real muscles. When activated, these muscles apply displacements to any vertices that lie within their *zone of influence*, based on parameters which describe the form of the displacement function for that particular muscle type, for example, the general “intensity” of the deformation and its fall off characteristics. In the case of vertices which are influenced by more than one muscle, the displacements are applied in succession, however the author describes problems with this approach, for example, muscles moving points out of the zone of influence of subsequent muscles. This approach is similar to the “space warping” techniques described earlier, however rather than being associated with a joint, here the deformation is associated with individual muscles.

Platt [78] uses a physically based approach to deform the skin surface based on entities representing muscles. Here the skin surface is modelled as an elastic sheet in the shape of a face, known as a *tension net*. This is then distorted by simple linear elements (representing muscles), which are attached to the tension net at one end and to the underlying bone surface at the other, and simply exert a force on the skin node in the direction of the bone attachment. This is shown in figure 3.2 which represents a patch of skin above the eye. Note the muscle elements pull upwards on the skin modelling the actual muscles in this part of the face. The main advantages of this approach are that the tension net automatically distributes the deformation caused by a single muscle over the area surrounding the skin attachment point, and also blending of the deformation from different muscles occurs automatically.

Lee, Terzopoulos and Waters [58] present a multi-layer, physically based model of the face. This consists of a network of mass-spring-damper elements which represents the three main deformable tissue regions in the face, notably the subcutaneous muscle layers, and the layers associated with the skin (dermis and epidermis). As shown in figure 3.3, the muscle layer is represented by two layers of hexahedral elements (six

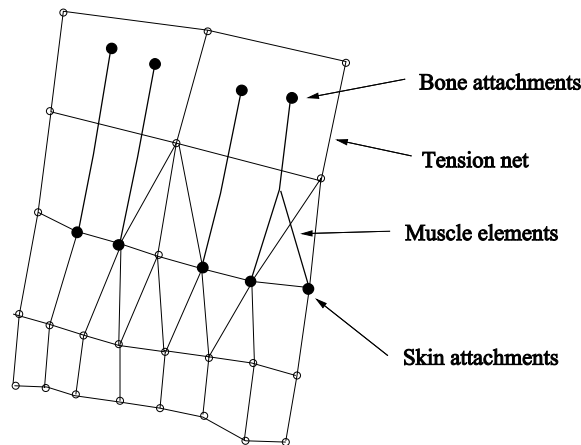


Figure 3.2: Tension net skin model (adapted from an illustration in [78]).

sided deformable elements consisting of eight point masses connected by springs along edges, face diagonals, and body diagonals). The dermal layer is represented by a single layer of pentahedral elements (five point masses connected by springs along edges, and face diagonals). Finally, the epidermal layer is simply represented by the plane of mass-spring-damper elements on the surface. The springs in each of the three layers are given different strengths to represent the differences in elasticity found in the real tissue types, and a volume preservation force is added to the elements to represent the limited compressibility of soft tissue. As shown in figure 3.3 muscle elements similar to those used by Platt (above) are connected between points in the muscle layer and attachment points on the surface of the skull. This work is particularly interesting since it models three tissue regions which are distinct both geometrically and in terms of their material properties.

3.6 Achieving Accuracy

So far we have considered a range of models which are, on the whole, influenced equally by a number of factors, for example, simulation speed and accuracy. We now turn to work in which accuracy is particularly important.

3.6.1 Anatomically based modelling

A relatively new approach involves developing complex geometric models which are based on a thorough understanding of the relevant anatomical features. Scheepers *et al.* [82] describe their work as ...

... an approach to human figure modelling similar to the one taken in artistic anatomy - by analyzing the relationship between exterior form and the

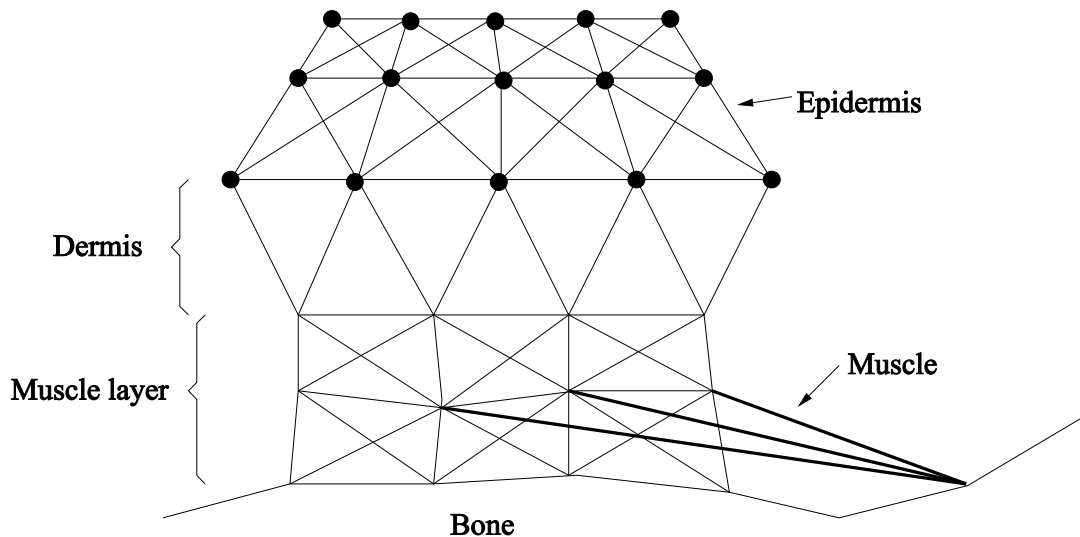


Figure 3.3: Cross section of the tri-layer tissue model (adapted from an illustration in [58]).

underlying structures responsible for creating it, surface form and shape change may be understood and represented best.

The model is based on a skeleton which is modelled to a greater degree of accuracy than in most other work. For example, it features accurate geometric representations of the major bones in the body (although many of the smaller ones are merged into single bones), and the joint models are also more accurate than usual, for example, the shoulder contains representations of both the clavicle and the scapula. For the musculature, a range of primitives (including ellipsoids and bicubic patches) are used to represent the wide variation in muscle shape that exists, and also the geometry of the tendons. A large number of muscles are modelled, in a way that accurately reproduces the musculature in real humans. In order to deform the muscles and tendons, complex scaling, bending, and other deformation techniques are used which incorporate relevant considerations, for example, the need to preserve muscle volume and to maintain tendons at a constant length. This deformation is based on joint angles as with many of the geometric models described earlier, however an important addition is a “muscle tension parameter” which allows control of individual muscle shapes independently of the skeleton position (specifically control over the ratio of muscle width to its length). This allows isometric muscle contraction to be modelled. An important omission for the work is that it contains no skin surface model, however, research is in progress which uses bicubic patches to form a surface which “moulds” itself around the musculature using an implicit surface approach. Images are presented which show very realistic looking static deformation. Wilhelms and Van Gelder [101] present a similar

model with the addition of an elastic skin surface which is attached to the underlying musculature.

Anatomically based modelling is an interesting development on the geometric models that were mentioned earlier. They retain the advantages of geometric models in that they are (relatively) quick to compute and hence could be used in interactive animation environments. The addition of separate control over muscle contraction is an important feature based on a more thorough understanding of real anatomy. However they are still geometric models at heart, and hence suffer from many of the problems associated with the other geometric models that have been discussed. For example, they only exhibit deformation effects that have been specifically designed into the model, and hence other effects like sagging of flesh under gravity, wobbling of flesh, and deformation due to external contact/collision are not modelled. Another problem is that the models have to be built manually from a selection of primitives rather than using general surface representations.

3.6.2 Anthropometry based modelling

Anthropometry is a branch of physical anthropology which studies measurements of the human body. Some researchers have used results from this field to produce models which can be used to represent a range of different body shapes rather than a single “average” shape [11] [40]. As well as considering variation in physical dimensions, accurate anthropometric scaling also requires variation in physical properties to be modelled, for example, variation in the range of motion of joints, mass of body segments, and material properties of different tissue types. Currently this issue is considered in only a few specialist research areas, however, as computer models of humans begin to be used in real applications, then it is likely to become an essential requirement.

3.6.3 Models for surgical applications

Koch *et al.* [55] describe “a prototype system for surgical planning and prediction of human facial shape after craniofacial and maxillofacial surgery for patients with facial deformities”. The basic approach of the system is to first of all extract data describing the shape of the skin, the internal soft tissue regions, and the skull using photogrammetric and computer tomography (CT) scans of the patient. The former approach refers to methods of extracting surface shape based on external measurements, for example, made by laser scanners. The latter refers to a scanning technique generates three dimensional volume data of a patient’s body. From this a model is constructed featuring the skull, a single layer of soft tissue, and the skin surface. The skin surface is represented using a FEM model of a flexible elastic surface, whilst the flesh is represented using a single layer of springs which connect the skin to the skull. The stiffness of the springs is determined based on the CT data such that it represents an average value of the elasticity of the various tissue regions that the particular spring passes through (although how the tissue stiffness is deduced from CT measurements is not revealed). The aim of the system is to predict changes in skin surface shape caused by changes to the underlying bone structure during surgery. Hence, the system generates a final static deformation solution (using an energy minimisation approach) based on

the initial geometry of the face, the soft tissue property information provided by the CT scan, and the changes made to the skull surface geometry. Images are presented which show a simulated operation using the Visible Human Data [74]. A good aspect of this approach is that the model is based on real data from scans (which is of course necessary if it is to represent a real human face). Also, the soft tissue properties vary based on the CT scan data. Simplifications with this approach are that it lumps the muscle and other tissue layers within the face into a single region of deformable material (although in this case the region is not homogeneous). Also the method used to determine the arrangement of the spring mesh which represents the flesh is tailored to the specific topology of the face in that it relies on the skin surface being a similar shape to the front of the skull such that the soft tissue can be regarded as a deformed “slab” of almost uniform thickness. As a result, the method would be difficult to apply to other parts of the body where this was not the case.

Pieper *et al.* [75] describe a model for simulating plastic surgery procedures, including incision (defining a hole), excision (modifying the hole) and suturing (closing the hole). Again, the model consists of the skull, a single layer of soft tissue and the skin surface. The process involves drawing the shape of the incision on the skin surface (in the form of straight line segments) and specifying the tissue which is to be removed, and then the model computes the final deformation after suturing. A three dimensional FEM model of homogeneous linearly elastic material is used to generate the deformation of the soft tissue.

Sagar *et al.* [81] present a model of the human eye and the surrounding features for use in a surgical simulator. The simulator forms part of a teleoperated micro-surgical robotic system being developed for eye surgery. The model consists of a physically-based FEM model of the cornea coupled with a geometric model of the other main features of the eye including the iris, retina, eyelids, eyelashes and surrounding facial skin surface. The FEM model of the cornea is based on large deformation incompressible three dimensional elasticity theory.

3.6.4 Biomechanically based limb models

Chen and Zelter [22] produce a physically based model of muscle action based on a biomechanical force model created by Zajac [107], and a structural FEM model. The aim of this work is to simulate changes in muscle shape under various conditions, as well as accurately modelling the muscle’s force behaviour. Results are presented from simulations of the gastrocnemius muscle in a frog, and the medial gastrocnemius muscle in a human. The physical shape of the muscle regions was obtained in two ways. Firstly, a manual approach was used for the frog muscle which involved dissecting and measuring a real sample. Secondly, the shape of the bones and muscles in a human calf was derived semi-automatically from a sequence of MRI scans in which tissue surfaces were extracted from the raw data, and then “skinned” to create polygonal models.

The technique presented involves constructing a FEM model of the muscle region which models its basic passive properties, and then “active” forces are exerted on the mesh nodes based on the Zajac model of muscle behaviour. As the system is dynamically simulated, muscle activation then leads to muscle deformation. A number of

results are presented. The first are from simulations involving the frog muscle which recreate the classic force-length and force-velocity experiments that were described in Chapter Two (section 2.8.2). Here the authors report a good correspondence with published results. The second set of results are from simulations of the human muscle model. The first of these shows the gastrocnemius muscle being activated. Unfortunately however, the full musculoskeletal structure was not modelled during this simulation, and a simplified situation was used in which both ends of the muscle are attached to the soleus. Nevertheless, the muscle is shown sagging under gravity, and then becomes taut when it is activated. Collision response is also included between muscles and the skeleton using *reaction constraints* [77]. Verification of the shape changes that occur is qualitative based on videos of live frog and human subjects. The authors suggest applications in biomechanical and medical simulation as well as computer animation, and they summarise by saying:

By emphasising the *physical model*, we have taken a step beyond the original computer graphics goal of making a virtual actor, up to goal of making an *artificial person*. By studying the anatomy, the form will be revealed through the function.

Hing [45] considers the use of musculotendon actuators to control the motion of a physically based skeleton. These actuators generate forces which are based on accurate biomechanical models of real muscles, again based on the work by Zajac [107], such that the force is dependent on muscle activation, muscle extension, and rate of change of length. Although the force properties of muscle are accurately simulated, the geometry is not. The geometric form of the musculotendon actuators consists of simple tapered cylinders which deform based on the position of the muscle-skeleton attachment points, the angle of the joint(s) that the muscles span, and the requirement to preserve volume. The authors considers a number of control strategies for dynamically determining muscle activation in order to achieve some required goal (for example, to hold a steady skeleton pose under fixed or dynamic loading). The work is interesting since it creates complete physically-based musculoskeletal models based on accurate research in biomechanics. However, the fact that the musculoskeletal systems are modelled only as geometric entities and are not part of the physical simulation (other than due to their contraction forces) is a gross simplification since, as described by Hing, muscle mass can account for 40-45% of the total body mass [76], and as a result has a significant effect on skeletal motion. Hing attempts to compensate for this by adding extra mass to the skeleton segments. However, this does not account for the dynamic redistribution of muscle mass for different limb states, and also fails to model the fact that muscles are not rigidly attached to the skeleton.

3.7 Discussion of Existing Work

This chapter has described a range of work aimed at modelling different parts of the body. The simplicity of many of the models indicates that much of this research is still in its infancy and further developments are required. Many of the techniques presented for animation purposes are crude geometric models which attempt to recreate the most

obvious deformation effects that are seen in real limbs for example creasing and muscle bulging. Although these may be suitable for modelling stylised deformations in cartoon style characters, they are too simplistic to accurately model the complexity of deformation that occurs in human limbs. Based on this realisation, a number of researchers have turned to alternative approaches. Anatomy based models are an interesting development of geometric models, which effectively use an understanding of the geometry of the musculature beneath the skin surface, and the way in which it deforms and moves as the skeleton state changes, in order to determine the resulting shape of the skin surface. Hence they represent an important shift of attention from modelling based on the visual **form** of the skin surface, to modelling based on the **function** of the limb, and the way in which the musculature interacts with the skeleton to generate motion of the limb. However, even though these models recreate certain types of deformation very realistically, and may even be sufficient for many animation applications, they are still over simplified models of real limbs and fail to encompass a number of important causes and effects which are crucial for realism. For example, they fail to model the dependency that skin surface shape has on muscle activation and skeleton state, and also do not exhibit inertia effects or deformation due to contact with external objects.

The medical and biomechanical models presented represent the flip-side of the coin. Here the main focus of attention is typically on the physical aspects of the object being modelled, and as a result the complexity required to accurately model these physical characteristics leads to a necessary narrowing of the scope in terms of the geometry of the objects being modelled. For example, in the facial plastic surgery applications, the physical properties of the skin and soft tissue regions are typically modelled using accurate FEM models, whilst the geometry of the face is typically represented as a simple “slab” of soft tissue which is deformed to approximately match the shape of the face and the skull. Very little research has considered physically based models which feature more than one distinct tissue region (an exception being the face model developed by Lee, Terzopoulos and Waters [58] which contains distinct muscle, and skin regions).

One of the most advanced models discussed is the work by Chen and Zelter [22]. This contains an accurate muscle force model and a FEM based model of the structure of the particular muscles being represented. The muscle model generates forces which deform the structural model based on the same mechanism of muscle deformation that occurs in real life. Hence the focus is on function, but more specifically the way in which forces are generated and interact with the material that the objects are made up of. The obvious evolution of the work by Chen and Zelter is to use their muscle model within a complete physically based model of a musculoskeletal system.

This type of model represents a new level beyond the geometric and anatomically based limb models discussed so far, since it models not only the structure but also the physical properties of the tissue regions. This is important since it is likely to generate more realistic results, but also since another change in focus is involved, from the functional to the physical. For example, a description of this type of model need only include the geometry of the tissue regions and their material properties, in order to be able to reproduce it. Hence, the emphasis is shifted away from having to understand how the system works, but rather concentrates on assembling the tissue structures in

the same way in which they are assembled in real limbs (e.g. muscles are attached to tendons and tendons are attached to bones), and also modelling the material properties of the tissue regions, including the ability of muscle to generate its own forces. As a result, when the model is simulated, and the muscles are activated, the models automatically recreate the function of real limbs. Hence, the structural and physical properties of the model give rise to the function, which in turn gives rise to the form.

Having developed a full musculoskeletal model, a further development would be to add a layer of flesh and possibly a skin surface. Rather than being a biomechanical model of a musculoskeletal system, we can now consider it to be a fully functioning physical model of a limb. Models of this type have a number of advantages over geometric models. Firstly, the deformation emerges from the simulation rather than having to be prescribed by the model designer. This has a number of implications. Firstly, it is not necessary to design specific skin deformation models for the different joints within a character, or for a new character with different bodily dimensions, since as long as the structure of the new joint and the properties of the surrounding tissue regions are accurately modelled then the resulting deformation should **automatically** be realistic. Another important advantage is that the various factors which contribute to limb deformation, i.e. muscle state, skeleton state, velocity state and external forces, are all handled within a single unified approach. Hence, the models are general in that they can generate deformation from each of these factors acting individually as well as any combination of them acting simultaneously.

Geometric models are not general since they rely on explicit descriptions of the deformation that should be displayed given a particular limb state. The result is that if the limb adopts a state that is not covered by the deformation rules then the results will be undefined. Also, attempting to devise a deformation model which integrates all four of the limb state factors described in Chapter Two is problematic to say the least. Catering for external collisions alone presents difficulties (as mentioned in section 3.1) due to the large number of contributing factors e.g. surface properties, object shape etc. Because of these problems, current geometric models attempt to recreate relatively straightforward effects first, for example, creasing and muscle bulging, and then attempt to add other effects incrementally. However, this process involves specifying new rules whilst trying not to interfere with the existing ones, which inevitably leads to unwieldy over complex models.

This discussion gives rise to the argument that in order to progress to producing accurate models of entire limbs which incorporate the causes and effects of skin deformation that were described in section 2.7, then a physically based model of the type discussed above, which is grounded in an understanding of anatomy, biomechanics, and the material properties of human tissue, is the best approach. This thesis supports this argument by presenting a complete physical based model which realistically incorporates these causes and effects. However, when considering the suitability of such a model in real applications, a number of important issues have to be considered. These are dealt with in the following sections.

3.7.1 Processing requirements

One of the main problems with a physically based approach is that the accuracy comes at a price, and even for modest simulations, large amounts of processing power are required making the models slow to simulate. For example, many of the physically based models discussed in this chapter represent only a single region of deformable tissue, but may take hours to simulate only a few minutes of behaviour. In relation to animation applications, this problem has resulted in most physically-based models being suitable only for off-line post-processing situations rather than being used in real time animation systems. However, although at the time of writing this is still a relevant issue, it is likely to become less significant in the future for a number of reasons. Firstly, the processing capabilities of computers continue to rise at an increasing rate, and hence simulation rates will increase correspondingly. Secondly, physically based models of this type lend themselves to parallel processing, since the computation can be divided into independent computational “packets” corresponding to geometric divisions in the model. Hence, as the cost of multi-processor systems decreases over time, they may become a viable platform on which these models can be simulated in real or near real time. Another reason is that since the research is at an early stage, then the techniques are relatively unrefined, however as research progresses then optimisations are likely to be devised which will further increase simulation speed. And finally, as physically based models become more widespread, then the approach of implementing the most processor hungry computations in hardware may be adopted, in a similar way to the situation with hardware rendering.

3.7.2 Skeleton control

Another important issue, particularly in animation applications, is how the overall motion of the skeleton in such models might be controlled. The discussion in this chapter has shown that in most models, the motion of the skeleton is determined by a standard kinematic approach, for example key framing. However, other approaches have been used and may be preferred by some system, for example force or torque based control. In medical applications, a range of control requirements might exist. For example, in a biomechanical simulation of a human running, the preferred techniques may range from direct kinematic animation of the skeleton, to activation of the muscles. Alternatively, in a surgical simulation system, since the body is typically anaesthetized during an operation, the control requirements may simply involve activation of individual muscles to test the operation of musculoskeletal structures that have been operated on. Clearly, the range of possible control scenarios for such a model is vast. Hence, in order to avoid limiting the model, we propose an approach based on using a complete physically based skeleton which is fully integrated with the rest of the model, but which may be overridden in certain ways depending on the situation. Examples of different control scenarios and the way in which these would be handled are described below.

Direct kinematic control

In order to use the model in a system where kinematic control is preferred, then an approach can be used in which we override the motion of the skeleton which is determined by the rigid body simulation, and take direct control over this ourselves. This downside of this is that the motion of the skeleton is no longer influenced by muscle activation or the action of the non-muscular soft tissue regions which would normally add to the realism of the motion. Another issue is that since the control is being performed directly on the skeleton, rather than via the muscle activations, then the capability of the model to simulate muscle contraction, and the effects that accompany it, is not being utilised. A possible solution to this problem, would be for the system which is controlling the skeleton motion to also control the muscle activations (as in Scheepers' work [82]). For example, in an animation system in which the character motion is scripted, the muscle activations could also be choreographed such that they appeared to be causing the particular motion of the skeleton. This would also allow the animator to activate muscles before a motion is carried out in order to express muscle tension.

Force/torque based control

If the preferred control method is via the application of forces and torque to the skeleton, then this can be accommodated into the model in a number of ways. The first is to ignore all forces acting on the skeleton from the limb model. Hence, the control system takes complete control over the motion of the skeleton. Again, this would restrict the model since we would not be making use of the facility whereby the flesh and musculotendon structures influence the skeleton motion. However, as with direct kinematic control, the muscle activations could be controlled separately to produce convincing results. Another possible approach would be to allow forces from the soft tissue to be transmitted to the skeleton, but not forces from the tendons. This might be used with a control technique in which the forces/torques applied were intended to represent the motive forces generated by the musculoskeletal systems. Again muscle activations could be controlled independently. The advantage of this approach over the previous one is that it would still allow external forces, for example gravity or forces due to surface collisions, to be fed back to the skeleton (and hence to the control system). Finally, another alternative would be to leave the model fully functional, and to simply combine the control forces and torques with the model forces.

Muscle control

The highest level of control (and also the most difficult in terms of the controller), would be to enable the complete model, and to generate the required motion via the muscle activation parameters. This is considered in Chapter Five (section 5.7).

3.7.3 Control over model deformation

One of the more difficult issues to solve relating to the use of the model within animation systems is the ability to control the deformation of the flesh. The simplest scenario

for a complete physically based model of the type being considered is if the animator requires the physically realistic looking motion that is automatically generated through the simulation, since here no control by the animator over the deformation of the model is required. If the animator requires non-realistic motion however, then some means of control is necessary. Since this thesis is primarily concerned with the realistic modelling of limbs, then this issue will only be touched upon here, and we will return to it in the final discussion in Chapter Seven. One option is to allow the physical properties of the model to be changed. This could be used to generate non-realistic behaviour that was still physically based. For example it may be possible to recreate a number of comic effects that have been used in traditional hand drawn cartoons in which objects are made to exhibit behaviour which intentionally goes against that experienced in real life. For example, gravity could be made much higher than in real life, causing flesh to sag and objects to appear very heavy, or the bones could suddenly be made elastic causing them to bend as the character moved.

For cases where the animator requires more general control over the skin surface deformation, then more complex techniques are required. Ideally these should allow control at a range of scales as discussed in section 3.2.1. One possibility might be to apply force fields around the limb which are effective over a certain range of influence and can be used to pull or push the surface to obtain the required shape. The issue of control at different scales could then be addressed by varying the range of influence of the force fields in order to produce local or more widespread control.

Chapter 4

Modelling the Structure of Limbs

The previous chapter examined research that has been carried out into the modelling of limbs and other related areas. It concluded with a discussion which described a shift in emphasis from the form and function of limbs, to the physical aspects of a limb, in order to generate more accurate models which produce more realistic results. This approach, which is being adopted in this thesis, involves focusing on modelling the structure of the various tissue regions within a limb, and the material properties of those tissue regions. If these are modelled sufficiently accurately, then it is hoped that the required function and as a result the changes in form of the limb will emerge from the simulation of the model.

In this chapter and the following one we present a new model of a human leg. This is a complete physically based model which incorporates a skeleton, two musculotendon structures and a layer of non-muscular “flesh”. This chapter firstly considers a number of issues relating to the modelling process as a whole, and the details of the leg model that was developed. It then describes the methods that were used to generate the structure of the model. The following chapter then describes the methods used to model the material properties of the various tissue regions, and also looks at general issues relating to how it was simulated etc.

4.1 The Modelling and Simulation Process

The general process that was used in the development of the limb model involves the following stages:

Specifying surface geometry. In order to generate the model, it is first necessary to define the geometry of the structures within it. This includes the shape of the skin surface, the shape of the skeleton and muscle regions, and the path of the tendons through the non-muscular tissue.

Pre-processing. Before simulating the model, it has to be processed into a form that is suitable for the material modelling methods used. This typically involves making structural simplifications, and discretising the regions of deformable tissue into simple geometric elements.

Simulation. Here the behaviour of the model is simulated in the required environment.

Typically the simulation will take place in an interactive environment which displays the current state of the model as the simulation progresses (usually in a simple form which can be quickly rendered). Various user interactions may take place during simulation, for example, activation of muscle regions or control over external forces or objects.

Post-processing and viewing. This involves analysing the results of the simulation.

For example, we may wish to view the simulation at full speed (if it was not simulated in real time), or use more realistic rendering techniques (for example involving shadows and texture mapping). Visualisation techniques could also be applied, for example, parts of the model might be removed or made semi-transparent in order to show the deformation of the internal tissue regions.

The first two steps of this process are now dealt with in this chapter, and the last two in the following chapter.

4.2 Case Study Model

A model of a leg was used to test out the approaches described in this thesis. A leg was used rather than an arm since it contains a greater proportion of soft tissue regions, and is functionally slightly more straightforward (at least in terms of the operation of the knee joint compared to the elbow). Figure 4.1 shows a two dimensional representation of the structure of the leg model. The model is based on a simple skeleton consisting of two bones which are connected at the knee. The bone above the knee represents the femur. The lower bone is an amalgamation of the bones in the lower leg (tibia and fibula), the knee cap (patella) and the bones in the foot. Attached to this skeleton structure are two musculotendon systems. The first of these is located at the back of the thigh and is connected via tendons to attachment points on the upper and lower bones. This system represents the muscles which comprise the hamstrings, and its function is to flex the leg. The second system is located at the front of the thigh, and is again connected by tendons to the upper and lower bone. This represents the musculotendon systems found in the quadriceps, and its function is to extend the leg. The tissue structures described above are surrounded by a region of soft tissue which represents the various non-muscular tissue regions found in the leg, including fat, skin, blood vessels, nerve fibres etc.

Clearly the model represents a gross simplification of the complexity of the tissue structures found in a real leg (which were discussed in Chapter Two). The main simplification is that only those tissue regions which are primarily involved in the biomechanical functioning of the leg are considered. Since we are concerned with the motion and deformation of the limbs then this is considered reasonable. Another simplification is to model only two musculoskeletal structures. This was done to reduce the complexity of the model as a whole in order to make the work feasible. The intention was that initial efforts should determine the effectiveness of the proposed techniques, and then further work can be carried out to create more complex models. The decision to make




Figure 4.1: Case study leg model.

both of these musculotendon structures act on the knee, but in opposite directions, enables us to examine the ability of the model to simulate the cooperation between similar opposing systems in real limbs. The justification for simplifying the skeleton was based on the fact that, at this stage, since we are only considering musculotendon structures which act on the knee, there is little point modelling other articulations for example the joints in the foot. A final important structural simplification was the decision not to include a distinct skin surface. This was a borderline case since, as shown in Chapter Two, skin has distinct properties from the underlying soft tissue, and hence has a significant affect on the deformation and general behaviour of limbs. However, because the differences in material properties between skin and other soft tissue types are not too extreme, and structurally, the thickness of the skin region is small compared to the thickness of the other soft tissue regions, it was decided that it was reasonable to combine these into a general non-muscular tissue region.

The next section now considers how the leg model was constructed.

4.3 Constructing the Model

As mentioned at the start of the chapter, the first stage in the construction process involves defining the geometry of the model. In this case, closed polygon meshes were used to represent the surfaces of the tissue regions, and polylines were used to represent the paths of the tendons. These representations were chosen for their simplicity, and also due to the fact that they are supported by many software tools and applications. A disadvantage associated with them is that they are both piecewise linear approximations and hence have only C^0 continuity. However, the problems associated with this (i.e. crude surface silhouettes) can be reduced by using a suitably large number of small polygons (and line segments in the case of the tendons).

Models representing limbs can be obtained using a number of techniques. One approach is to create them manually using standard techniques, for example “lofting” (see [97] for details). This involves defining the shape of the object being modelled by specifying a number of cross sections at regular intervals, and then automatically fitting a parametric surface through the cross sections in order to construct a smooth surface. An alternative approach is to laser scan the surface of a real limb, or a model of a limb. In the case of a real leg, this would only provide a skin surface model. To generate models of internal surfaces, anatomical models (of the type commonly used by medical students) could be scanned. A more hi-tech approach is to use magnetic resonance imaging (MRI) scans of living people (or cadavers) to obtain volume data. Three dimensional surfaces can then be extracted from this data using visualisation techniques. Existing data of this type is available from a number of sources including the Visible Human Project [74]. In the case of medical applications, the use of three dimensional scanners opens up the possibility of constructing a model based directly on the patients own internal tissue structures.

For the leg model presented here, a number of sources were used. Figures 4.2 and 4.3 show the skin surface and internal tissue structures of the leg model. The skin surface model was obtained from the Cyberware company internet site at www.cyberware.com. Cyberware manufacture a range of state of the art laser scanning systems. The skin



(a) Skin mesh (1).



(b) Skin mesh (2).




(c) Skin mesh (3).




(d) Skin mesh (4).

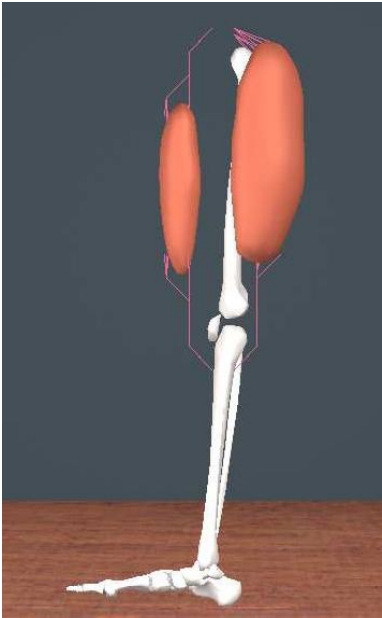
Figure 4.2: Skin surface mesh.



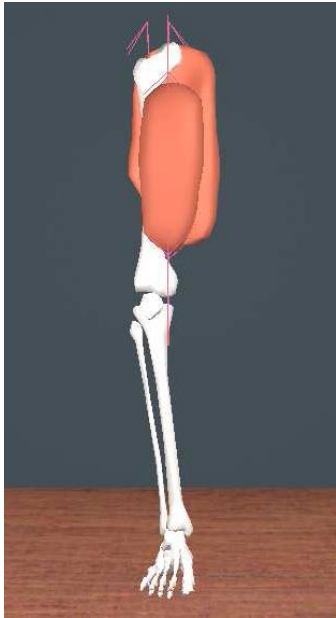
(a) Internal regions (1).



(b) Internal regions (2).



(c) Internal regions (3).



(d) Internal regions (4).

Figure 4.3: Internal tissue regions (bones, muscles and tendons).

surface shown is composed of 15338 evenly spaced triangles with a mesh resolution of approximately 1.2cm throughout. The skeleton models were also obtained from a free model site on the internet and have a combined total of 5208 triangles. The muscle regions were generated manually using the lofting procedure described above. These are loosely based on the quadriceps and hamstring muscle regions found at the front and back of the thigh, but are not an accurate representation. The front muscle is composed of 648 triangles and the back muscle 4660. Finally, the tendons shown were also constructed manually using a 3D model editing package.

4.4 Preparing the Model for Simulation

Having obtained the surface representations, the next stage involves pre-processing the geometry such that it is in a form to which the deformable material modelling techniques can be applied (these are presented in the next chapter). This primarily involves subdividing the tissue regions into simple geometric shapes using a mesh generation process.

4.4.1 Mesh generation

Subdivision into simple geometric elements is a requirement in most methods of modelling deformable material [34][15][22][39][85][87]. One of the most popular approaches is called the Finite Element Method (FEM) [71], and as its name suggests, it approximates the continuous space of the object being modelled using a finite number of simple elements. Generally, situations where subdivision is necessary are ones in which a set of differential equations have been formulated which describe the behaviour of some phenomena within a geometrically complex domain, and a solution to these equations is required at some or all of the points within the domain. In all but the simplest cases, it is not possible to obtain the exact solution analytically, and hence an approximate solution has to be obtained using numerical methods. Problems of this type occur in computational fluid dynamics, heat transfer problems, stress analysis in engineering components, acoustics modelling, and modelling of electric and magnetic fields. Because of the wide range of dependent applications, mesh generation has become a very important and active area of research.

When subdividing a model, a number of three dimensional elements are commonly used. Figure 4.4 shows three examples which are a tetrahedron (four faces), a pentahedron (five faces) and a hexahedron (six faces). Each of the examples shown is irregular, and in general, a more accurate approximation can be obtained by using irregular rather than regular elements. Although it is common to mesh an object using combinations of elements, the problem can be simplified by considering only tetrahedrons. A tetrahedron represents the simplest three dimensional object, and all other elements can be created by combining them together. In order to further simplify the discussion of meshing, and the particular requirements that are needed for this application, it is also useful to consider the problem of two dimensional meshing initially, and then extend the ideas to three dimensions. The two dimensional equivalent of meshing a three dimensional region using tetrahedrons is to subdivide a general two dimensional

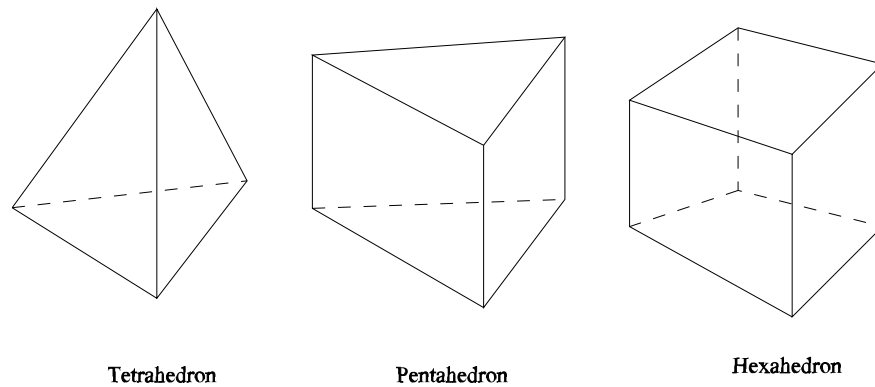


Figure 4.4: Simple three dimensional elements.

region into triangles. Triangles and tetrahedrons are members of a family of geometric objects known as *simplices*, which are the simplest elements that can be used to fill a region in N-dimensional space.

4.4.2 Meshing in two and three dimensions

Ideally, the meshing process should be automatic such that the mesh generator receives input describing the geometry of the object that is being modelled, and generates output containing a list of elements which can be used as the basis for the simulation. Figure 4.5a shows an example of two dimensional input geometry that contains a range of features that a mesh generator should be capable of processing. Figure 4.5b shows a possible output mesh. The features shown are described as follows:

- A** is a closed polygon representing the outer boundary of the domain.
- B** is a hole interface. This is a closed polygon representing a region in which meshing is not required.
- C** is a separator interface. This is a closed polygon representing a boundary between two regions which require meshing.
- D and E** are cuts. These are open polylines which impose edge constraints (i.e. the final mesh should contain these edges).
- F** is an example of a node constraint (i.e. the final mesh should incorporate these nodes).
- G** is a specified region which requires a different density of meshing than the rest of the object.

In the output mesh:




Figure 4.5: Two dimensional mesh generation.

H is a node which has been added as a result of subdividing edges in the input geometry. This occurs when the scale of the input edges is high compared to the required mesh scale.

I is a node which has been added to the internal region as part of the meshing process.

If all the edges in the input geometry are present in the final result without being subdivided, then this is known as a *constrained triangulation*. If some of the edges are represented by at least two edges in the final result then it is known as a *conforming triangulation*. In the latter case, points are added to edges splitting them into two or more edges. These points are known as *Steiner points*. In three dimensions the boundary elements take the form of triangles (simplicial boundary elements), and we may also have edges and points present to act as constraints in the meshing process.

The output mesh shown in figure 4.5b is likely to be one of a set of possible triangulations that could have been produced. The choice as to which is the most desirable is typically governed by the application. One obvious criteria determining whether a particular mesh is “good” is the number of triangles that it is made up of. In general, we would like this to be as small as possible, since a larger number of elements will require more computation. However, since we are representing general curved objects using polygons (in two dimensions) or polyhedrons (in three dimensions), then the accuracy of the approximation will inevitably reduce as the number of elements is reduced. Another important criteria is the shape or *aspect ratio* of the elements. In most applications (e.g. FEM), the aspect ratio of the elements produced by the mesh generator is important since it has an influence on the accuracy of the simulation (badly shaped elements lead to “sensitivity” in the methods for solving partial differential equations which increases the numerical error). In the case of deformable material modelling, an implication of this is that meshes composed of well proportioned elements are generally more robust than meshes containing badly proportioned ones.

In the case of triangular elements, a number of definitions of aspect ratio have been used. These include the ratio of the element height to the length of the base, the ratio of the largest inscribed circle to the smallest circumscribed circle, and the size of the smallest element angle. For tetrahedral elements, measures based on similar geometric considerations have been suggested including, the size of the smallest solid angle, and the ratio of the largest inscribed sphere to the smallest circumscribed sphere. A detailed discussion of tetrahedral shape measures is given in [60]. Whatever the particular measure used, the general requirement is that it should attain a maximum value if the element is regular (all edges are the same length), and approaches zero for poorly shaped elements i.e. triangular elements which are close to the degenerate case (having zero area), and tetrahedral elements which are close to degenerate tetrahedrons (having zero volume).

In order to create three dimensional models of the deformable tissue regions in limbs, many of the input geometry features described above are required. For example, the skin surface is defined using a closed polyhedral boundary. Bones, and other non-deformable regions are represented using hole interfaces (defined by close polyhedra). Deformable regions of tissue which have distinct physical properties from the surrounding tissue (e.g. the muscle regions in the leg model) are represented using separator interfaces (defined by closed polyhedra). For the tendons, the path through the model can be defined by a series of edge constraints. A number of other features may also be used. For example, node constraints could be used to mark points of interest within the limb model, or to steer the meshing routine to create a certain mesh arrangement in a particular region. Another use of node constraints may be to induce the mesh generator to create a high mesh density in regions where large amounts of deformation are expected. This would then enable the deformation in these regions to be resolved more accurately during simulation.

As mentioned above, mesh generation may produce constrained or conforming meshes. However, research has shown that it is not always possible to generate a constrained mesh of a three dimensional region [1], and hence the result will typically be a conforming mesh which minimises the number of elements whilst maintaining the aspect ratios below a reasonable limit.

4.4.3 Existing meshing techniques

Research into the modelling of soft tissue has made use of a range of meshing techniques. Koch *et al.* [55] use a simple three dimensional mesh to represent the layer of soft tissue beneath the skin surface of the face. However, the problem is trivial since the region is topologically simple (essentially a deformed “slab” of material), and corresponding grids exist on the inside and outside surfaces of the tissue region (i.e. the skin surface and the boundary between the soft tissue and the skull) due to the particular construction method used. Piper *et al* [75] are faced with a similar problem in which a layer of tissue representing the soft tissue of the face requires meshing for use in a system to simulate plastic surgery. Here, a two dimensional mesh composed of triangles and quadrilaterals is first of all constructed on the skin surface based on the arrangements of incisions that are to be made during the particular surgical technique being simulated. Meshing of

the three dimensional tissue region is then achieved by “extruding the surface inward” along normals to the skin surface. The result is a mesh composed of a single layer or pentahedral and hexahedral elements.

Unfortunately, simple approaches like the ones described above are inappropriate in our case due to the complex topology and structure of the tissue regions, and also the fact that there are several of them. Hence it is necessary to consider more complex methods used in other applications.

Delaunay based approaches

Research into automatic meshing in two dimensions has been fruitful and a number of useful general purpose algorithms have been developed. Many approaches are based on the two dimensional Delaunay triangulation algorithm which generates a triangular mesh from a set of input points. Chazelle [2][3] describes an $O(n)$ optimal algorithm for triangulating a polygon based on a constrained Delaunay triangulation technique. Ruppert [4] describes Delaunay based bounded aspect ratio approach.

Unfortunately research into meshing in three dimensions has shown the problem to be much more complex than in two dimensions (the general problems involved in tetrahedralisation of non-convex polyhedra is dealt with in more detail in [1]). This is essentially due to the fact that a number of statements which are valid in two dimensions do not hold in three dimensions:

- In two dimensions, if we consider a polygon with n edges, then a constrained triangulation will have $n - 2$ triangles. In three dimensions however, tetrahedralisations of polyhedra do not follow a similar rule.
- In two dimensions all polygons can be triangulated. In three dimensions however, there are polyhedra which cannot be meshed using tetrahedrons unless extra edges are added. In other words all polyhedra have a conforming tetrahedralisation but not a constrained one.
- If we apply the Delaunay triangulation algorithm to a two dimensional set of points the resulting triangulation will be such that it maximises the minimum angle over all possible triangulations. In three dimensions however this does not occur.

The last point is of particular importance since this property of two dimensional Delaunay meshes is the basis of most successful approaches. However, in three dimensions the same guarantee does not apply, and in general, naive application of the Delaunay algorithm will produce badly proportioned meshes.

These problems mean that the techniques used successfully in two dimensions, do not enjoy the same success in three dimensions. Techniques that do manage to generate a solution, generally have to add large numbers of Steiner points, resulting in overly complex meshes containing large number of elements.

Simplification to two dimensions

Because of the problems associated with three dimensional meshing, many applications simplify a three dimensional problem into a problem that can be solved in two dimensions. For example FEM based stress/strain analysis in engineering components which have uniform cross sections or radial symmetry is often carried out along two dimensional planes through the object with the assumption that variations perpendicular to these planes are insignificant. However, this approach is again not suitable in our case due to the relative complexity of the model.

Manual meshing

In two dimensions manual meshing of simple input geometries is feasible albeit time consuming. The quality of such meshes is also often high since it is relatively easy to visually identify poorly shaped elements and then change the local mesh to improve this. In three dimensions however, this process is extremely difficult even for simple surface geometries due to the problem of visualising three dimensional spaces using two dimensional display devices. A number of commercial systems attempt to overcome some of the problems associated with three dimensional mesh generation by using semi-automatic tools which are often part of a suite of applications that accompany FEM software. These are typically based on mesh generation algorithms which are steered by user input. For example they might operate by manually partitioning complex shapes sections which can be meshed automatically, or manually tweaking areas in which the mesh generator does not perform well. Such systems rely on well designed graphical interfaces which aid the operator in visualising the mesh. Unfortunately however, these systems are expensive, and are typically only used for high budget engineering projects.

Chen and Zelter [22] model a single muscle region which has a fairly simple geometry (basically a cylindrical shape with pinched ends) by using “a user-assisted finite element mesh generator” which creates a model composed of just four irregular hexahedrons. However, this approach is not feasible in our case due to the complexity of the model.

Gourret *et al.* [39] construct a model of a hand which contains a single layer of deformable flesh surrounding a rigid skeleton. Although the authors do not state how the deformable regions were meshed, the geometry of the situation is sufficiently simple for the process to have been carried out manually or using a simple automatic technique.

Experimental mesh generators

A number of research packages are available from various institutions. Very often these are experimental projects and are made freely available to the research community for evaluation purposes. Three of these were evaluated for use in our model, and the findings are presented below.

Quickhull [16] is a software package produced at the University of Minnesota. This can be downloaded via anonymous ftp from:

`ftp://geom.umn.edu/pub/software/qhull.tar.Z`

The application takes a set of points in N dimensional space as input, and generates a triangulation using an implementation of the Delaunay method. Unfortunately, as it stands this package would need substantial development to make it into a complete three dimensional mesh generator suitable for our application. For example, the Delaunay process does not create its own internal points, and also the algorithm does not consider edge or face constraints.

Geompack [5] is a software package produced at the University of Alberta. This can be downloaded via anonymous ftp from:

```
ftp://menaik.cs.ualberta.ca/pub/geompack
```

This is a system for generating two and three dimensional meshes, and is capable of handling the range of input features that were described in section 4.4.2. Unfortunately however, tests showed that the system had several restrictions which made it unsuitable for the leg model.

QMG (Quality Mesh Generator) is produced at the Computer Science Department at Cornell University, and can be downloaded from:


```
ftp://ftp.cs.cornell.edu/pub/vavasis
```

This contains a two dimensional mesh generator (based on Delaunay triangulation) and a three dimensional mesh generator based on an octree approach devised by Mitchell and Vavasis [68]. Tests were again performed using this system on a simplified version of the leg geometry (in which the skin mesh was reduced to approximately 600 polygons and the skeleton to approximately 100). Unfortunately, the system generated results which contained many tens of thousands of small elements. Hence, this represented a significantly non-optimal meshing of the model and meant that it was unsuitable since the simulation requirements would have been excessive.


Structured meshes

So far we have only considered the use of unstructured meshes containing irregular elements. An alternative approach is to use a more constrained mesh based on a regular repeating structure. Typically this can then be distorted to obtain a better approximation to the region being modelled. In general, unstructured meshes are preferred since they are capable of modelling a particular region more accurately using fewer elements than a structured mesh. Figure 4.6 shows three examples of a structured mesh based on a square grid which is subdivided into triangles. In 4.6a the grid is simply applied over the object being meshed and cubes which fall wholly outside the object are discarded. The result is a mesh which is a fairly crude approximation to the initial object. In 4.6b a similar approach is used but the grid points close to the object boundary are moved in order to obtain a more accurate approximation. Finally 4.6c shows a structured mesh that has been obtained from a parameterised representation of the ellipse.


Chadwick *et al.* [21] create a three dimensional structured cubic mesh composed of mass-spring-damper elements which is used to model the deformation of the soft tissue around an animated character's limbs. Here only a single region is required to be meshed and the surface which defines the region that requires meshing is defined by a simple



a) Regular triangular mesh based on a cubic grid



b) Adjusting points to improve the approximation



c) Parameterised triangulation

Figure 4.6: Examples of structured meshes.

cuboid aligned with the bones in the limb. Holton and Alexander [50] suggest creating a library of simple three dimensional shapes which have associated meshes, and then combining these to produce more complex objects resembling tissue regions. Typical shapes might include spheres, cones, and cylinders, which can all be meshed easily using three dimensional structured meshes of the type shown in figure 4.6. However, although this is possible in simple situations, it gives rise to the problem of how to connect neighbouring meshes so that there is a smooth transition from one primitive to the next. Also, if we suppose that this method is used to model the muscle regions in the leg model (which are fairly simple), then the passive flesh which is a topologically more complex region i.e. with holes and interfaces etc, would still require meshing, and in order to do this we would probably have to resort to one of the unstructured approaches described above. In general, consideration of this problem leads to the conclusion that it is probably necessary to consider the whole model during the meshing process in order to obtain a “global” solution with continuous meshing across interfaces rather than attempting to combine “local” solutions that have been obtained for each region in isolation.

4.5 Meshing the Leg Model

Because of the problems associated with generating unstructured meshes, an alternative structured mesh approach was adopted based on a cubic grid which is subdivided into tetrahedrons. The method used approximates the regions corresponding to the bone, muscle and soft tissue using the cubes within the grid. However, rather than use the surfaces in this cubic approximation as the tissue surfaces, we maintain the original geometry, and simply deform it based on the deformation of the grid. This is similar to the FFD technique introduced in section 3.1. This essentially gives us two layers to our model, the **physical layer** which is involved in simulation and collision detection etc., and the **geometric layer** which is used to represent the surfaces of the different tissue regions. The two are linked by embedding the geometric surfaces within the cubic mesh such that they deform along with it (the technique used is described in section 5.5 of the following chapter).

As mentioned above, the meshing process involves firstly determining the cubic approximation to the surfaces which make up the leg model and then subdividing the cubes into tetrahedrons. Obtaining the cubic approximation involves four stages as described below (and illustrated in figure 4.7 for a two dimensional model). Initially, a cubic approximation to the skin surface is obtained. This is done by effectively laying a grid over the model and removing those squares which lie wholly outside the skin surface as shown in figure 4.7a. We then perform the same operation for the internal regions as shown in 4.7b. Note, that the dotted line shows the silhouette of the upper bone region which partially overlaps the back muscle region. Stage three involves eroding the bone regions until there is no overlap between a particular bone and the neighbouring bone or muscle regions (see figure 4.7c). *Erosion* describes a process in which the cubes at the surface of a region are switched so that they belong to the surrounding region until a particular requirement is met. Finally, stage four involves eroding the internal soft tissue regions in order to prevent neighbouring regions from being in direct contact with





Figure 4.7: Generating the cubic approximation.

one another (figure 4.7d). When considering neighbouring muscle and bone regions, bone is eroded in preference to the muscle unless the bone is only a single block thick in which case the muscle will be eroded. When considering regions where the bone is close to the skin surface, erosion of the bone will take place even if the region is a single element thick, as seen in the foot. However, when this takes place, the nodes which were part of the bone region are identified as “bone points” (shown in figure 4.7d as filled in circles) and are still considered to be part of it even though the bone has no physical volume at this point. The result of this is that the bone point will still be constrained to move along with the bone during simulation, as if it were part of the same rigid body. This final stage is necessary to main skeletal support in areas which are thin compared to the mesh resolution. The final stage also involves approximating the tendon paths in the original geometry using a polyline based on the grid points (as shown in figure 4.7d).

The general aim of this meshing approach is to produce isolated muscle and bone regions which are surrounded by a layer of passive tissue that is at least one cube thick. However, in certain situations the method will give poor results, for example if the erosion process remove a significant part of one of the regions. In this case, the scale and orientation of the initial mesh is changed and the method applied again until satisfactory results are obtained. Hence this is a semi-automatic approach based on experiment.

4.5.1 Decomposing the cubes into tetrahedrons

After creating the cubic mesh, the next step is subdivide the cubes into tetrahedrons. This can be done in a variety of ways. Figure 4.8 shows three possible tetrahedralisations of a cube. A and B show approaches which involve firstly splitting the cube into two wedge elements, and then decomposing each of these into three tetrahedrons. Mesh B is the same as A, but with one of the wedges turned upside down. This is likely to be more effective since it will help to cancel out the anisotropic effects in each of the wedges. Meshes A and B contain six tetrahedrons. C shows an alternative mesh based on a central tetrahedron with four identical tetrahedrons attached to each of its four faces.

The approach used in our model was C. This was considered to be the most appropriate since it contains fewer tetrahedrons, which are more regular than those in meshes A and B. As with the other approaches, the behaviour of this mesh is likely to show anisotropic effects, due to the way in which the tetrahedrons are arranged. For example, if we consider the corners, then for four of them there are three axis aligned edges coming together, whilst for the other four there are three axis aligned edges and also three face diagonal edges. The effect of this can be reduced by ensuring that when the cubic mesh as a whole is subdivided, the meshing applied to neighbouring cubes is rotated by ninety degrees. This is illustrated in figure 4.9. Figure 4.9a shows cubes combined without rotation, whilst 4.9b shows neighbouring cubes which have been rotated by ninety degrees about the vertical axis (although rotating about the other axes would achieve the same result). This approach also has the benefit that the faces of tetrahedrons in neighbouring cubes match up with one another, which prevents these

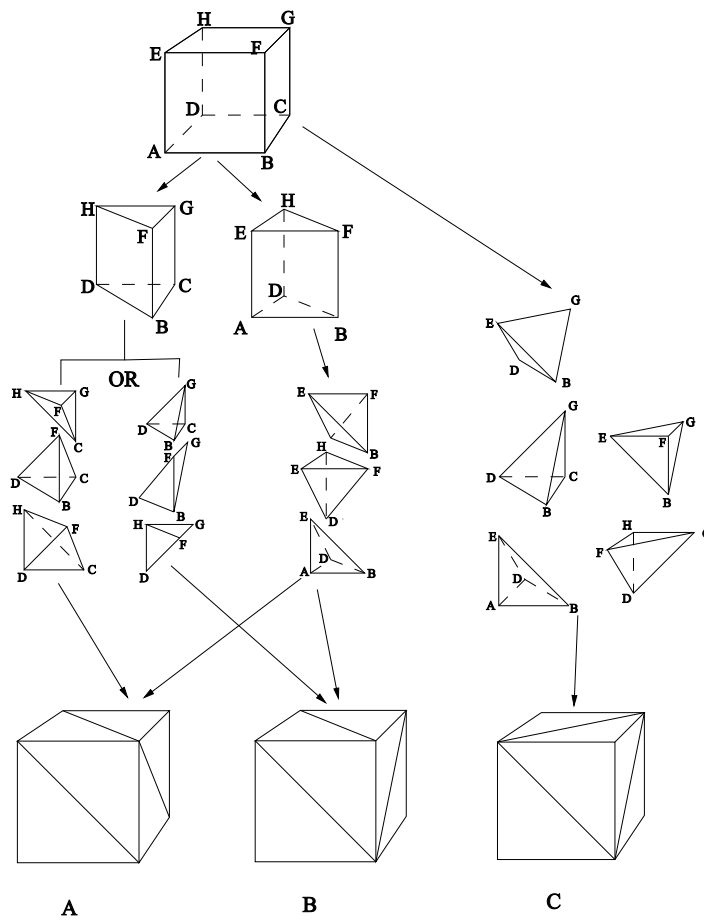


Figure 4.8: Alternative approaches to tetrahedralisation of a cube.

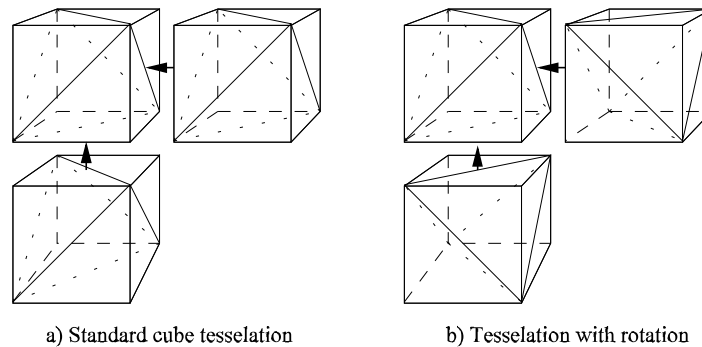


Figure 4.9: Combining cubic elements.


tetrahedrons from intersecting as the mesh is deformed.

Figure 4.10 shows the cubic approximations to the skin and internal surfaces obtained by applying the above approach to the leg model shown at the start of this chapter (figures 4.2 and 4.3). Note that the situation described above in which the bones are reduced below unit thickness occurs around the ankle area in the lower bone (as seen in figure 4.10(c)), and here fixed bone constraints are applied as described above (section 4.5). The final mesh is composed of 959 cubes which decompose into 4795 tetrahedrons. These break down into, 60 tetrahedrons for the front muscle, 260 for the back muscle, 55 for the upper bone, 55 for the lower bone, and 4365 for the passive tissue.


The modelling of the tendons is as follows. The tendon connecting the back muscle to the upper bone is composed of 8 “strands” each made of up of a single line segment. The tendon connecting the back muscle to the lower bone is composed of 6 strands each made up of 5 line segments. For the front muscle, the tendon connecting it to the upper bone is composed of 6 strands made up of 4 line segments, and the tendon connecting it to the lower bone is composed of 6 strands made up of 5 line segments. This approach of composing a tendon up using multiple strands, was used to enable the tendon to split up at its ends such that it could be made to attach to the muscle at multiple points in order to distribute the forces evenly. A total of 11 “bone points” were generated by the meshing process.

4.6 Summary


A method is presented for decomposing the deformable regions within the leg model into tetrahedrons which involves firstly generating a cubic approximation to the structure of the model, and then subdividing the cubes which are part of the deformable regions into tetrahedrons. The original geometry is retained and deforms with the cubic approximation. This results in distinct physical and geometric layers to the model. The physical layer is represented by the cubic approximations to the original surfaces and



(a) Skin approximation (front).



(b) Skin approximation (back).



(c) Internal regions.

Figure 4.10: Cubic approximations to the original surfaces of the model.

the paths of the tendons which are shown in figure 4.10. This is used to simulate the behaviour of the model including the deformation of the soft tissue, and the effects of collisions with external objects. The geometric layer is the “visible” part of the model and is made up of the original high resolution surface geometry.

This approach allows us to mesh topologically complex models consisting of different tissue regions as required. However, it has the obvious disadvantage that the physical surfaces represent a poor approximation to the initial surfaces. One area in which this can be predicted to cause a problem is when part of the skin surface is in contact with an external object, or with another part of the skin surface i.e. when creasing occurs, since there will be an obvious gap between the surfaces in contact corresponding to the discrepancy between the physical and geometric surfaces at this point. In order to reduce this effect at the back of the knee, points at the surface of the physical skin mesh were moved towards the underlying geometric surface, in order to reduce the discrepancy. This is evident in the images of the model shown in figure 4.10(b).

Chapter 5

Modelling the Physical Properties of Limbs

The modelling approach used in this thesis is based on simulating the physical behaviour of a limb. The discussion in section 3.7 concluded that this requires two aspects to be modelled, the structure of the tissue regions, and their physical and material properties. Chapter Four presented various techniques that have been developed to model the structure of a limb, and also showed these techniques being applied to the leg model. This chapter now concentrates on modelling the material properties and hence the physical behaviour of human tissue, and again shows how these techniques were applied to the leg model.

One of the main requirements of this work is a suitable approach for accurately representing the behaviour of the deformable soft tissue regions. The first section in this chapter begins by discussing various deformable material models that have been used by other researchers considering soft tissue modelling and other similar problems. In section 5.2 we present a new approach based on a mass-spring-damper (MSD) technique, and describe how this is applied to representing the muscle and flesh in the leg model. Some of the techniques presented here are based on published work that is provided in Appendix B for completeness. The following sections then consider other aspects of the limb model including, the modelling of tendons, modelling the skeleton, and techniques for responding to contact between the skin surface and external objects. Finally, section 5.9 gives a summary of the chapter.

Simplifying material properties

As discussed in Chapter Two, the material properties of the various tissue types that exist in human limbs shows a great deal of complexity. However, in order to accurately simulate the behaviour of musculoskeletal systems and limbs as a whole under normal conditions (i.e. carrying out behaviour which would not normally lead to injury or generation of excessive forces), much of this complexity does not need to be incorporated into the model. Typically, simplifications can be made in a number of areas. One example is the way in which material properties show variation within a single tissue region. Here, we can hypothesise that it is reasonable to remove this variation and

model the tissue regions as homogeneous regions with material properties based on some kind of average of the values found within the real tissue region (whether this is true or not will be revealed when testing the model). Another important simplification is to model only those material properties which are relevant to the biomechanical functioning and deformation behaviour that we are trying to reproduce. For example, properties which describe subtle types of behaviour e.g. creep (the gradual increase in strain over time in a material held under constant load), do not play a significant role and can be ignored. Finally, having determined that a particular property is important and should be modelled, very often much of the complexity associated with it can be justifiably ignored. For example, since we are modelling normal behaviour (as defined above), then the stress-strain properties of say, bone, associated with behaviour close to the point of fracture can be disregarded. In the sections that follow this process of simplification is considered for each of the component tissue regions within the leg model, in order to reduce the complexity of the leg model as a whole.

5.1 Modelling Deformable Material

As mentioned above, a key requirement of this research is the need for a deformable material model to represent the soft tissue within the model. This section describes the main approaches that have been used by other researchers.

5.1.1 Finite difference and finite element models

The Finite Difference Method (FDM) and Finite Element Method (FEM), are established ways of modelling deformable material. They have been used in engineering for many years to analyse the behaviour of engineering components which deform during use. FEM and FDM approaches are usually based on material models derived from *elasticity theory* [10]. However, although basic elasticity theory is appropriate for models of engineering materials, its application to the modelling of soft tissue is questionable. For example Fung [33] points out that ...

... the customary application of the infinitesimal theory of elasticity, which is commonly used for analysis of engineering materials in which the strain might be typically 0.2%, is inappropriate or will lead to inaccuracies when applied to soft body tissues which will commonly experience maximum strains between 10% and 40%.

The application of elasticity theory produces a set of partial differential equations which describe how the state of the material evolves over time. These typically describe behaviour in continuous space and time domains. However, in order to simulate the behaviour of a sample of material using a numerical method, it is necessary to discretise these domains. FDM and FEM represent two different methods of performing discretisation of the spatial domain. FEM involves partitioning the object into a finite number of small elements e.g. irregular tetrahedrons, hexahedrons etc. These elements are geometrically simple enough to allow a solution to the partial differential equations

to be generated analytically. The global “solution” is then based on local solutions within each element along with constraints which are imposed at the boundaries between neighbouring elements (see [53] [108] for more details). FDM uses an alternative method of discretisation which involves selecting a number of sample points within the material (often arranged on a regular grid) and reformulating the partial differential equations based on the values at these grid points. FEM and FDM have been applied successfully in the modelling of engineering components, for example bridges, turbines and engine components, to determine how they deform during use. However, the scale of deformation which is modelled in these applications is usually very small. One reason for this is that materials exhibiting large deformation would be unlikely candidates in the design of a bridge, or a solid engine component. As a result, when using FEM or FDM, the material model can usually be simplified based on the assumption that extensive deformation of the object does not occur. In the human body however, large scale deformation of the flesh is commonplace, and is in fact necessary to allow correct functioning of the limbs.

Terzopoulos *et al.* [87] [85] have used elasticity theory and FDM to create general models of deformable material aimed at animation applications. These include elastic models with internal damping, viscoelastic materials, plastic materials, and materials featuring fracture. Interaction with external objects and the prevention of surface self-intersection is included in the model, and is achieved using force fields around external objects and the surface of the deformable materials themselves.

Many researchers have applied FEM to modelling various soft tissue regions in the body (these are discussed more fully in Chapter Three). Gourret *et al.* [39] present a physically based model of a hand grasping a deformable ball. The hand is modelled as single layer of flesh surrounding a skeleton. FEM is used to model deformation occurring in the hand and the ball. Koch *et al.* [55] and Pieper *et al.* [75] present FEM based models of the soft tissue in the face for plastic surgery applications. Sagar *et al.* [81] present a FEM based model of the cornea. Chen and Zelter [22] produce a physically based model of a single muscle using a FEM model.

5.1.2 Mass-spring-damper models

Another type of model which is often used to represent deformable material are mass-spring-damper (MSD) models [106][50][21]. These are based on networks of point masses connected together by springs and dampers. The elements are “ideal” meaning, for example, that the spring elements have no mass or damping, and the masses are considered as points with no volume. If the material behaviour required is principally that of an elastic material with damping (as is the case here), then the masses are usually connected by a spring and damper in parallel. However, it is possible to achieve more complex non-elastic behaviour using other spring-damper arrangements, and introducing other components for example slip and fracture elements (see [86]). As an MSD network deforms, then the springs and dampers exert forces on the masses (which are considered along with external forces) and cause the points to move. Figure 5.1 shows a point mass with position vector \underline{x}_p connected to N other point masses with position vectors $\underline{x}_1 \dots \underline{x}_N$ by damped springs (a spring and damper in parallel). The




Figure 5.1: Forces acting on mass P.

force exerted on particle p by the spring connecting it to particle k is given by \underline{F}_k . \underline{F}_e also acts on p and represents the resultant of all the external forces. Applying Newton's second law to the particle gives:

$$\underline{F}_e + \sum_{k=1}^N \underline{F}_k = m_p \ddot{\underline{x}}_p$$

where m_p is the mass of the particle, and \underline{F}_k is given by:

$$\underline{F}_k = - \left[k_s (|\underline{l}| - r) + k_d \frac{\dot{\underline{l}} \cdot \underline{l}}{|\underline{l}|} \right] \frac{\underline{l}}{|\underline{l}|}$$

where $\underline{l} = \underline{x}_p - \underline{x}_k$, $\dot{\underline{l}} = \dot{\underline{x}}_p - \dot{\underline{x}}_k$, r is the rest length of the damped spring, and k_s and k_d are the spring and damping constants respectively. By using a suitable method of numerical integration the motion of the masses can be computed.

In order to model large blocks of material, basic deformable elements are usually created which are then combined to form the required shape. Chadwick [21] uses cubic elements consisting of eight masses (positioned at the corners of the cube) with springs connecting each mass to every other mass. Holton and Alexander [50] use irregular tetrahedrons as their basic element which are combined to form a cube. When deciding which form of element to use, a number of issues should be taken into account:

- The deformable element should tessellate in three dimensions to fill space.
- The design should be as “elementary” as possible to aid analysis of the element behaviour.
- The element should be stable in isolation, to maximise the stability of groups of elements.

- The number of mass, spring and damper components involved should be kept to a minimum, to minimise the computational requirements during simulation.

5.1.3 Other material modelling approaches

Smith [84] describes a novel, albeit complex, approach to modelling deformable material based on elasticity theory. The method extends the domain that the deformable object is modelled in from three to six dimensions. This then allows a linear theory to model a wide range of phenomena, including quasi-rigid motions (motion of very stiff bodies) as well as the motion of highly flexible objects which exhibit large global deformations.

Gascuel *et al.* [35] present an interesting approach for modelling deformable limb segments based on springs which extend outwards from a rigid “kernel” positioned at the center of the limb. The springs detect and respond to collisions with external objects and a method is described for propagating deformation of one spring to other neighbouring springs. The rules which determine the propagation, can be defined to “mimic” effects like limited compressibility of the surface. This is a simple approach which requires little computation, but it is not considered suitable here since it would be difficult to apply to a limb with more than one deformable tissue layer.

5.1.4 Summary

Of the methods presented above, we consider the Mass-Spring-Damper approach to be the most suitable. FDM and FEM models have been widely used in other research, however the applications tend to involve situations in which the magnitude of the deformation that occurs is small in comparison with that seen in limbs, for example modelling of facial tissue [55][75] and modelling of the cornea [81]. This means that this work has avoided the difficulties involved in using FEM models to recreate large deformations. Use of FEM also makes it difficult to integrate other techniques, for example in [87] [85] a complex force-based approach is required to prevent FEM based deformable models from intersecting one another. Other complex approaches for example the work by Smith [84], also suffer from similar problems. MSD models on the other hand represent a good compromise between complexity and effectiveness, and have the scope to model the range of material properties required for human tissue. They are also intuitive, and as a result can be extended easily to incorporate new techniques. The next section now presents a new MSD model for use in representing soft tissue within limb models.

5.2 Modelling Soft Tissue using Mass-Spring-Damper Models

5.2.1 Simulating uni-axial stress-strain tests

In order to examine the properties of current MSD models and develop them further for use in soft tissue modelling, we adopted a novel approach which involved performing simulated uni-axial stress-strain tests on samples of material. The samples were




Figure 5.2: Spring-damper arrangement in a cube composed of five tetrahedrons.

cylindrical with a height of 15cm and a diameter of 10cm. The method involved firstly creating a simple cylindrical surface model using a standard 3D object modelling package. Then meshing techniques developed in Chapter Four were employed to create a physical model based on this surface representation. A mesh resolution of 2.5cm was used, which resulted in a model composed of 120 tetrahedrons.

A simple test MSD deformable element model was used which consisted of four point masses at each of the vertices of the tetrahedron, each of which was connected to the other vertices (along the tetrahedron edges) using a total of six springs and dampers acting in parallel. This resulted in an arrangement of spring and dampers within each cube as shown in figure 5.2. Here each of the lines (solid and dashed) represents a spring-damper combination. The spring constants were set to 10.75, and the damper constants set to 0.2. These values were chosen since they made the material relatively stiff, and also resulted in its motion being damped down quickly. A simulated stress-strain test was then carried out on the material sample. This was achieved by attaching the top and bottom faces to rigid blocks (as shown in figure 5.3), by constraining the point masses lying on these surfaces to move along with the adjacent block. The specific constraint used involved setting the velocity of the points to be the same as the velocity of the blocks themselves.

The stress-strain test involved moving the upper block up and down in order to stretch and compress the sample whilst the lower block remained in a fixed position. The test consisted of two phases involving firstly compression and then extension of the material sample. For the compression phase, the block was slowly moved down until the sample was compressed to approximately half its original length, and then was slowly moved back up to the rest position. Similarly, the extension phase involved moving the block upwards until the sample length was approximately twice its original length and then moving it slowly back to the original position. The complete procedure took about 12 seconds (of simulation time). The nominal stress, σ_n , and nominal strain, ϵ_n ,




Figure 5.3: Stress-strain test.

experienced by the material at a given point during the test is given by:


$$\sigma_n = \frac{F}{A}$$

$$\epsilon_n = \frac{d - l_o}{l_o}$$

Where F is the reaction force acting on the blocks from the MSD mesh, A is the initial cross sectional area of the sample (as shown in figure 5.3), d is the current length of the sample (also shown in figure 5.3), and l_o is the sample's initial length. When determining F it is only necessary to consider the force acting on one of the blocks since the force acting on the other is the same but acts in the opposite direction. The actual force was determined by considering each of the constrained points, and adding together the forces acting on these points from the springs connected to them. Since these points can be considered to be part of the block itself, this then gave us the resultant force acting on the block from the material. The vertical component of this forces was then extracted (small sideways components tended to occur due to anisotropy in the material model) to give F . Figure 5.4 shows the state of the geometric and physical surfaces at various stages through the test, and figure 5.5 shows the stress-strain graph obtained.

The graph shows the values that were obtained during the whole of the stress-train test and hence the curve shown represents a complete cycle of compressing the sample, and then taking it back to its original length.


The images of the model, and the stress-strain graph show that this simple arrangement of MSD elements results in a material model which does not behave in a similar



(a) Physical surface (1).




(b) Geometric surface (1).



(c) Physical surface (2).



(d) Geometric surface (2).



(e) Physical surface (3).






Figure 5.5: Stress-strain graph for uni-axial compression.




Figure 5.6: Illustration of element collapse in 2D.

way to normal elastic materials when subjected to compression. As the sample is compressed the stress initially becomes negative as expected (corresponding to the sample pushing the two blocks apart). However, as the strain approaches -0.4 , the stress then begins to level off. As we decrease the strain further, the stress can be seen to suddenly rise back to zero at a strain of around -0.47 . If we continue to compress the sample, the stress becomes negative, but again rises sharply back to zero at a strain of approximately 0.55 . The reason for this behaviour can be seen by observing the images in figure 5.4. Here figures (c) and (d) show the sample at a strain value of -0.4 at around the point where the stress-strain graph levels off and (e) and (f) show the sample at a strain of -0.47 after the stress has increased back to zero. In (c) and (d) the elements can be seen to be compressed in a concertina-like fashion as we would expect. However, in (e) and (f) the structure of the mesh has “collapsed” causing elements to lose their spatial configuration and overlap one another. The mechanism of this collapse can be best appreciated by considering the same effect in two dimensions as shown in figure 5.6.

Here we see a single triangular element (denoted by ABC) which is part of a general triangular mesh that is being compressed in a vertical direction. Figure 5.6(a) shows the element in an uncompressed state. However, as the compression is applied, then the forces on point B cause it to become closer to the opposite edge (AC). As this compression increases and the element begins to flatten out, then the springs which should be resisting the compression (AB and BC) become less and less aligned with the vertical direction and hence become less and less effective as shown in figure 5.6(b). As a result, the effect becomes steadily worse until eventually point B is pushed passed the opposite edge (AC). As this happens the forces acting on B from the springs along AB and BC now begin to act downwards pushing the point even further below the edge until a configuration is reached where the length of the springs are such that the forces cancel one another out and the system returns to equilibrium. The visual effect of this is that the mesh appears to suddenly “flip” into an alternative configuration. For three dimensional tetrahedral elements, the mechanism is essentially the same (see




Figure 5.7: Illustration of tetrahedral element collapse.

figure 5.7). As an element is compressed such that one of the points (P) approaches the opposite face, the resistance to further compression decreases and eventually the point passes through the plane containing the opposite face and flips into an alternative stable configuration.

Having described the mechanism of collapse we shall now return back to the stress-strain graph shown in figure 5.5. As pointed out before, initially the stress decreases almost linearly with the decreasing strain, but then it begins to level out due to the springs becoming flattened, which reduces the resistance to further compression. As the compression is increased, elements within the mesh eventually collapse and adopt a new stable state which has the effect of relieving the stress that has built up, and even causing it to go positive for a moment. However, rather than individual elements collapsing at different moments, the effect is that a complete layer of elements collapses simultaneously. This is due to the fact that the mesh is uniform, and as a result at any given time all the elements within a particular layer will be at approximately the same stage of compression. Then, as the compression increases, one of the elements which is in a marginally more advanced state than the others suddenly flips, causing the surrounding elements to flip until all the elements in that layer have collapsed. Hence, the initial rise in stress is caused by a number of layers in the mesh collapsing in this way. As the mesh is then compressed even further, the stress begins to build up again (in a negative sense) until another layer of elements collapses at which point the build up of stress is again relieved. The graph also shows the hysteresis effect that occurs as the block is moved back towards the initial position and the strain increases back to zero. Now, we see that the stress becomes positive until it reaches a maximum value and then drops suddenly. This is due to the collapsed elements being pulled back into their correct configurations and hence we see it occurring in two distinct places corresponding to the two distinct collapses that occurred.

5.2.2 The issue of collapse in general mass-spring-damper models

Although the problem of element collapse is typically not acknowledged in other research based on MSD models, we suspect that it occurs in all naive models that are based on a network of linear springs. The situation described above is really a combination of two problems. First of all, the elements collapse such that they adopt an incorrect spatial configuration. And secondly, after doing this they flip into an alternative stable state such that even if the compressive force is removed, then the elements will still not return to their original shape, but rather have to be “pulled” back by applying tensile forces.

5.2.3 Eliminating the collapse problem

In order to prevent collapse occurring in the tetrahedral elements (and in MSD models in general) it is necessary to impose some kind of constraint. There are various ways of doing this. For example, adjustments can be made to the position or the velocity of the point masses or alternatively constraint forces can be imposed upon them. A problem with constraints which act upon position or velocity is that it is difficult to resolve situations in which constraints from different elements affect the same point. For example an element which moves a point to prevent itself from becoming collapsed may result in another element attached to that point becoming collapsed. Methods which rely on forces to impose constraints do not suffer from this problem, since forces from different elements are added to one another to produce a resultant force, and the resulting motion of the mass is a compromise between these constraints.

With this in mind, a force based constraint technique was developed to prevent collapse from occurring. It involves exerting a repulsion force on the point masses from the face which lies opposite them. The force is such that its magnitude tends to infinity as the point tends towards the face. The technique was arrived at by considering what happens in real solid materials as they become compressed. Here, the material typically behaves in a linear manner for very small compressive strains but then non-linear effects arise, and as compression increases, the reaction forces generated by the material become larger and larger until we are unable to compress it any further. Figure 5.8, shows a repulsion force, \mathbf{R}_A , from the face BCD acting on point A. The force can be seen to act in a direction perpendicular to the face it is associated with. Note that although this is called a “face” repulsion force, in actual fact it is necessary for the force to be applied if the point approaches any part of the *plane* containing the face BCD since this would represent an attempt to completely compress the material inside the element. The function which determines the magnitude of the face repulsion force, f , is given by:

$$\begin{aligned} f &= k(x_o - x) + \frac{kx_o^2}{x} - kx_o \\ &= k\left(\frac{x_o^2}{x} - x\right) \end{aligned} \tag{5.1}$$

Where x is the current distance of the point from the opposite face, x_o is the distance to the face when the element is in the rest position, and k is a constant which determines




Figure 5.8: Forces in a tetrahedral element.

the strength of the repulsion force. This expression combines the term for a linear spring with a non-linear term which tends to infinity as x tends to 0. The $-kx_0$ term is also needed to ensure that $f = 0$ at $x = x_0$. A graph of the force exerted by this spring as a function of distance (x), is shown in figure 5.9 (for $k = 1$, $x_0 = 1$). The force for a linear spring is also shown (dashed line).

The actual force acting on point A, \mathbf{R}_A (as shown in figure 5.8), is then given by ¹ :

$$\mathbf{R}_A = k_r \left(\frac{PA_0^2}{PA} - PA \right) \frac{\mathbf{PA}}{PA} \quad (5.2)$$

where P is the projection of point A onto the plane containing the face BCD, k_r is a constant representing the strength of the face repulsion, and PA_0 is the distance

¹Note we are using bold type to represent a vector e.g. \mathbf{PA} , and standard type to represent a length e.g. PA .

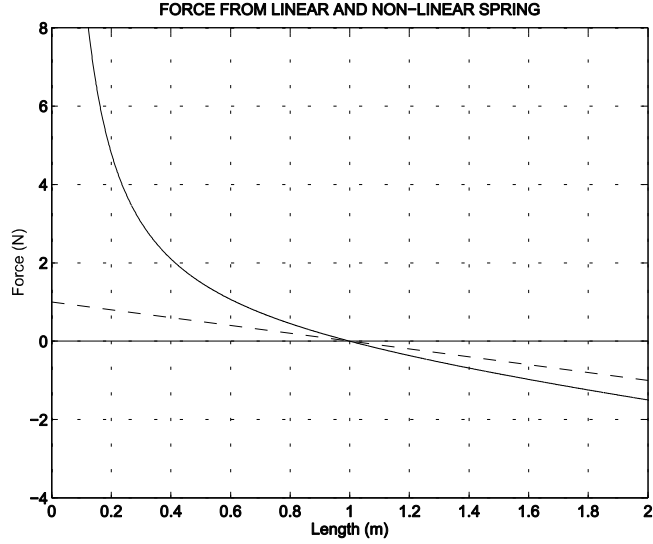


Figure 5.9: Forces from linear spring (dashed line) and non-linear spring (solid line) springs.

between the point and the face when the element is in its rest position i.e. with no forces applied.

5.2.4 Equalising internal element forces

Forces like the face repulsion one described above represent a force generated by the material within the element on a particular point, and hence since these are *internal forces* it is necessary to have a further set of forces which act on the other vertices such that the net resultant force and torque on the element is zero. This is satisfied in the case of spring and damper elements since they exert equal and opposite forces on the points that they are connected to. In order to satisfy this requirement for the face repulsion forces, the following approach is used (as illustrated in figure 5.8). Here, as before, \mathbf{R}_A is the face repulsion force acting on point A. In order to cancel out the effects of this force for the element as a whole, forces F_B , F_C and F_D , are applied to the points B, C, and D, such that each of these forces acts parallel to R_A (perpendicular to the face BCD). The magnitude of these forces is obtained by taking moments about an edge. For example, to compute F_B we take moments about CD,

$$\begin{aligned} \mathbf{F}_B d_{CD}^B + \mathbf{R}_A d_{CD}^P &= 0 \\ \mathbf{F}_B &= -\mathbf{R}_A \frac{d_{CD}^P}{d_{CD}^B} \end{aligned} \quad (5.3)$$

and similarly for \mathbf{F}_C and \mathbf{F}_D . The set of forces \mathbf{R}_A , \mathbf{F}_B , \mathbf{F}_C and \mathbf{F}_D , then give a resultant force and torque of zero. Hence, to summarise, the new element model now

contains forces from the springs and dampers along each edge, face repulsion forces acting on each vertex based on the distance to the opposite face, and reaction forces which are required to balance the overall force and torque generated by the repulsion forces. The total force acting on point A, \mathbf{T}_A , is given by the following:


$$\begin{aligned}
\mathbf{T}_A = & k_s(AB - AB_o) \frac{\mathbf{AB}}{AB} + k_s(AC - AC_o) \frac{\mathbf{AC}}{AC} + k_s(AD - AD_o) \frac{\mathbf{AD}}{AD} \\
& + k_d \left(\dot{\mathbf{A}}\mathbf{B} \cdot \frac{\mathbf{AB}}{AB} \right) + k_d \left(\dot{\mathbf{A}}\mathbf{C} \cdot \frac{\mathbf{AC}}{AC} \right) + k_d \left(\dot{\mathbf{A}}\mathbf{D} \cdot \frac{\mathbf{AD}}{AD} \right) \\
& + \mathbf{R}_A \\
& - \mathbf{R}_B \frac{d_C^B D}{d_{CD}^P} \\
& - \mathbf{R}_C \frac{d_D^C B}{d_{DB}^P} \\
& - \mathbf{R}_D \frac{d_B^D C}{d_{BC}^P}
\end{aligned} \tag{5.4}$$

Where \mathbf{R}_A is the face repulsion force for point A (as given in equation 5.2), \mathbf{R}_B is the face repulsion force for point B (and similarly for \mathbf{R}_C and \mathbf{R}_D), k_s is the element spring constant, and k_d is the element damping constant. Here, the first line contains terms corresponding to the forces from the springs, the second line contains the forces from the dampers, the third line contains the face repulsion force, and the final three lines are forces required to balance the face repulsion forces acting on B , C and D respectively.

In order to test the effectiveness of the face repulsion approach, another stress-strain test was performed. Images of the sample during this test are shown in figure 5.10, and figure 5.11 shows the stress-strain graph. It can be seen that the elements now resist collapse for a range of compressive and tensile strain. As a result, the large hysteresis effect that was visible in the previous test is no longer present, although slight separation of the lines still exists due to the dampers which exert different forces at a particular sample length depending on whether the length of the sample is increasing or decreasing. The stress-strain graph shown in figure 5.11 now also shows another measure that was taken during the simulation, which is the change in volume of the sample. The volume of the sample at a given time was computed by summing the volume of each of the individual tetrahedral elements. This was then converted to a percentage change in volume according to,


$$V_{\text{percent change}} = \frac{V}{V_o} \times 100$$

Where V is the current volume of the sample and V_o is the initial volume at zero strain. The volume-strain graph in figure 5.11 shows that the volume is approximately proportional to the strain. This occurs since, as shown in the images in figure 5.10, the sample acts as if it is being scaled in the vertical direction. This behaviour shows that the material model is compressible and does not attempt to preserve its volume when deformed. However, since muscle and soft tissue are only partially compressible, and




(a) Physical surface (1).

(b) Geometric surface (1).



(c) Physical surface (2).

(d) Geometric surface (2).



(e) Physical surface (3).

(f) Geometric surface (3).

Figure 5.10: Uni-axial compression test of the material model with collapse prevention.




Figure 5.11: Model featuring collapse prevention.




Figure 5.12: Volume preserving force.

this is the cause of many effects which occur in real limbs (especially flesh and muscle bulging), then it was necessary to further enhance the model to incorporate limited compressibility.

5.2.5 Modelling limited compressibility

This is implemented as another force constraint, which acts on each point mass in the tetrahedral elements, and this time attempts to move the point to the required distance from the opposite face such that the initial volume of the element is preserved. Figure 5.12 shows a tetrahedral element with the volume preserving force, \mathbf{V}_A , acting on point A. This is given by:

$$\mathbf{V}_A = k_v \left(\frac{3V_o}{Area_{BCD}} - PA \right) \frac{\mathbf{PA}}{PA} \quad (5.5)$$

Where V_o is the original volume of the element in the rest state, $Area_{BCD}$ is the current area of the face BCD, and k_v is a constant representing the level of volume preservation. Here, the first term in the brackets gives us the distance that the point is required to

be away from the opposite face (or more precisely the plane containing the opposite face) in order to attain the same volume as in the rest position. Hence, this force can be thought of as that generated by a spring whose natural length is constantly altered to that required in order to preserve volume. Again, as with the collapse prevention forces, since the volume preserving force represents an internal one, then it is necessary to remove its effect on the element as a whole by adding element equalisation forces to the other vertices (as described in section 5.2.4).

This approach is similar to one used by Lee, Terzopoulos and Waters [58]. A useful feature of this approach, particularly when considering its use in animation applications, is that the method is controllable and hence different levels of compressibility can be obtained by varying the “strength” of the volume preserving force (via k_v).


Again, a stress-strain test was performed to assess the performance of this technique. This time, both the collapse prevention and the volume preserving forces were present. Figures 5.13 and 5.14 show images of the physical and geometric surfaces during the test, and figure 5.15 shows the stress-strain and volume-strain graphs for this material.

The volume-strain graph now shows that for the particular volume force constant used during the test, the volume of the sample remained within 80% of its original value down to a compressive strain of -0.5 (half the original length of the sample). The effect of the volume force can also be seen in the images resulting in the sample bulging and becoming thinner as it is compressed and stretched.

The material model now displays the main types of behaviour that are required to model the soft tissue within the model. These include, damped elastic behaviour, the ability to robustly model large deformations (due to the inclusion of the collapse prevention technique), and the ability to model limited compressibility (due to the volume preservation technique). The complete set of constants which determine the material properties are now given by, the spring coefficient k_s , the damping coefficient k_d , the face repulsion constant (i.e. the strength of the face repulsion forces) k_r , and the compression constant (strength of the volume preserving forces) k_v .


5.2.6 Modelling non-muscular tissue

In Chapter Two the notion of general non-muscular tissue was presented which represents an agglomerate of many different tissue types e.g. blood vessels, nerve fibres, fat etc. Material properties for this “tissue type” were also given based on an average of those of the component materials. The material model described above incorporates the main properties that were considered with regard to the biomechanical operation and deformation effects required. These are damped elastic behavior and limited compressibility. As a result, this material model was applied to represent the behaviour of the non-muscular tissue in the leg model, and in actual fact, the sample which was tested in section 5.2.5 above is the same material used in the leg model. As well as representing the basic material behaviour, another important feature of the tissue model is its ability to simulate large scale deformations. Soft tissue in real limbs typically experiences a range of nominal strains during normal operation between -0.4 and +0.4 [52]. Hence, this is regarded as the minimum required “working range” for the deformable material in the limb model. The graph presented in figure 5.15 shows that




(a) Physical surface (1).

(b) Physical surface (2).




(c) Physical surface (3).

(d) Physical surface (4).




(e) Physical surface (5).

Figure 5.13: Images from a uni-axial compression test of the material model with volume preservation.




(a) Geometric surface (1).

(b) Geometric surface (2).



(c) Geometric surface (3).

(d) Geometric surface (4).



(e) Geometric surface (5).

Figure 5.14: Graph from a uni-axial compression test of the material model with volume preservation.




Figure 5.15: Model featuring limited compressibility.

the approach developed here is robust within this range, and in fact the test showed that strains well beyond this range were achieved without any indication of divergence from the required properties.

The model also incorporates effective modelling of the limited compressibility of non-muscular tissue due to its high proportion of water. However, achieving the same level of compressibility in the model used, would require an extremely high compression constant. This would typically lead to large forces within the elements which would require small time steps and hence excessive computation in order to maintain stability. As a result, the requirement was relaxed to the more practical one that the volume of the elements should remain within $\pm 20\%$ of their initial volume for the working range of the model (determined in the previous section). The graph in figure 5.15 shows that model successfully achieves this requirement.

Despite the closeness of the model to the real properties described above, there were areas in which approximations were made. Firstly, although real soft tissue shows significant spatial variation in its material properties, our model considers regions of soft tissue to be homogeneous. This is justified by the fact that the deformable tetrahedral elements within the model are relatively large compared to the scale of blood vessels and nerve fibres etc. and hence the differences between the individual tissue regions would tend to become insignificant at this scale.

Another simplification is that the biphasic properties of soft tissue under tension were not modelled. Instead, as seen in the stress-strain graph in figure 5.15, only the initial phase was modelled. This is considered appropriate, since the effects of the second phase only become significant at strains in excess of 0.7, which is beyond the required working range of the model. Hence, as seen in the stress-strain graph, a single phase was modelled with an elasticity modulus corresponding to that of the first phase of real non-muscular tissue.

Another important simplification was the decision not to include skin as a separate soft tissue region in the model. In reality, skin has distinct properties from the underlying soft tissue and as a result has a significant effect on the motion and deformation of a limb. However, in this case it was considered that combining the skin and the underlying non-muscular tissue into a single region was a reasonable approximation. The justification for this is that the properties of skin are sufficiently similar to those of the other soft tissue types for strains within the working range of the model. Also, the thickness of the skin region is small compared to the thickness of the other soft tissue regions that are being considered.

An issue that has been ignored so far is the modelling of the mass density of the material. This was simply achieved by calculating the original volume of the particular tissue region by summing the volume of the tetrahedrons that it is composed of. This value was then multiplied by the density of real tissue (the value of 0.92 g/cm^3 presented in Chapter Two was used) in order to find the total mass of the tissue region. Finally, this mass was divided equally amongst all the point masses that were part of the MSD elements representing the region.




Figure 5.16: Hill's muscle force model.

5.2.7 Modelling muscle tissue

This typically results in a reduction of

Modelling the behaviour of muscle tissue is an active research area. Many models have been generated by biomechanics research, and have concentrated on aspects relating to the overall force generated under certain conditions (as discussed in Chapter Two). Notably, the dependency of force on muscle activation, overall length, and rate of change of length of the muscle. Unfortunately however, this work typically does not consider the shape of the muscle, and the changes in shape that accompany contraction. An example of this kind of research into force generation, is the muscle model developed by Hill [36] which was formulated based on experiments on isolated muscle samples, and is illustrated in figure 5.16. This is composed of the same types of basic mechanical element (i.e. springs and dampers) that are used in the MSD material model, however here they are being used to model the way in which a single force (the contractile muscle force) varies rather than to simulate deformable material. The model consists of a passive elastic element in parallel with a contractile component. The elastic element represents the various tissue membranes found in muscle tissue. The contractile component represents the “active force generator” in a muscle and produces a contractile force corresponding to the forces generated by chemical reactions in the sarcomeres of real muscle. This force is proportional to the muscle activation up to a limit which represents the maximum strength of the muscle. The tendon is treated as an elastic element in series with the muscle as shown.

One piece of research that has considered modeling both the force and the geometric/deformation aspects of a muscle is the work by Chen and Zelter [22]. Here a force

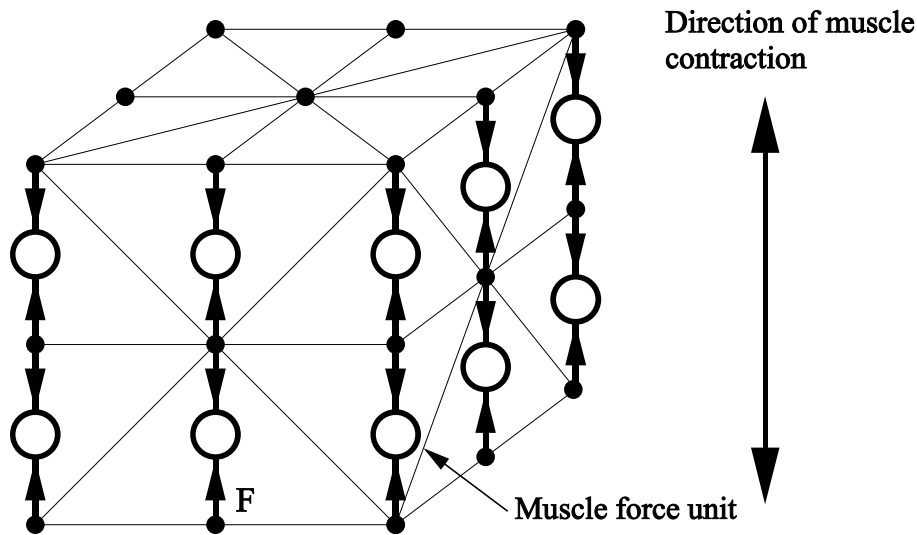


Figure 5.17: The new muscle model.

model is combined with a FEM model representing the structure of the muscle in order to consider the interaction of these two aspects. The integration of these models is achieved by applying force generating units along edges in the FEM model which are aligned with the direction of muscle contraction. The resulting material properties are then determined by the interaction of the passive material properties of the FEM model combined with the effects of the force units. The actual force model used is based on a refined version of Hill's model developed by Zajac [107].

In order to model the muscle regions in our leg model, a new approach was adopted based on a simplified version of Chen and Zelter's model. It involves using the basic soft tissue model that was developed at the start of this chapter, and applying contractile force generating units along edges which are aligned with the required direction of muscle contraction. This is illustrated in figure 5.17. Here we see part of the mesh of tetrahedral elements associated with a particular muscle region. Note, that each unit generates equal but opposite forces on the point masses that it is connected to. Hence, the resultant force and torque generated by the force units is always zero, and extra steps to cancel out net forces and torques (of the type required by the face repulsion forces) are not required. In our case the contractile forces generated by the units are simply proportional to the activation parameter of the muscle region. Initially, a more complex approach using Hill's model as the force generating unit was considered but was not thought to be necessary for a number of reasons. Firstly, the tendon spring in Hill's model is not required here, since the tendon is modelled as a distinct component of the limb model. Also, the spring and damper elements in Hill's model were considered unnecessary since they would have been duplicating the spring and damper that already lie along the edges of the elements as a part of the basic material model. Hence, the

contractile element was the only part of the model which was considered necessary.

In order to observe the behaviour of the muscle model, a simple test was carried out which involved activating an isolated sample of muscle tissue. Figure 5.18 shows the results of this test. The procedure used involved simply increasing the activation of the muscle tissue and observing the resulting contraction and deformation. As seen, the behaviour is generally as we would expect, and shows contraction of the sample in the direction of the internal contractile forces. The slight difference in the shape at the top and bottom of the sample are due to slight anisotropy in the material model due to the tetrahedral subdivision process used.

5.3 Modelling Tendons


Tendons are another important part of the musculoskeletal systems in the leg model. The approach used consists of representing tendons as a length of elastic string-like material which is connected at its ends to two of the point masses in the MSD network (typically one which is at the surface of a muscle region and one on the surface of a bone). The tendon follows a predetermined path through the passive tissue regions which is described by a polyline connecting a series of neighbouring mesh points. This path is chosen during the pre-processing stage to represent the best approximation to the path of the real tendon. Figure 5.19 shows a tendon passing through a region of soft tissue which is connected to a bone at one end and a muscle at the other. At the end points the tendon exerts a force on the node with a magnitude equal to the tension in the tendon, and in the same direction as the tendon segment as shown. For nodes along the length of the tendon, the force acting in the node is determined by finding the sum of the forces acting on the point from the two tendon segments that it is connected to. These component forces have a magnitude equal to the tension in the tendon and act in the same direction as the segments. The result is that force components acting along the length of the tendon cancel one another out, and hence the tendon only exerts a force on the surrounding soft tissue in a transverse direction, which is consistent with the tendon sliding through the material.

The tension in the tendon at a particular moment is given by,

$$T = \begin{cases} k_t(L - L_o) & \text{for } L > L_o \\ 0 & \text{for } L \leq L_o \end{cases}$$

where, L is the current length of the tendon and is obtained by summing the length of the individual line segments that it is composed of. L_o is the initial length of the tendon obtained before the simulation begins.


Hence the tendon only carries tension which is proportional to the extension. The value of k_t was initially chosen to represent a modulus of elasticity equivalent to the first phase of real tendon (approximately 35MPa) as presented in Chapter Two. However, this resulted in extremely large forces requiring small simulation time steps (and hence excessive computation) to maintain stability of the model. Hence, the elasticity was reduced to a significantly smaller value of 0.8MPa. The effect of this on the accuracy of the model is discussed in the next chapter.



(a) Physical surface (1).



(b) Geometric surface (1).



(c) Physical surface (2).



(d) Geometric surface (2).



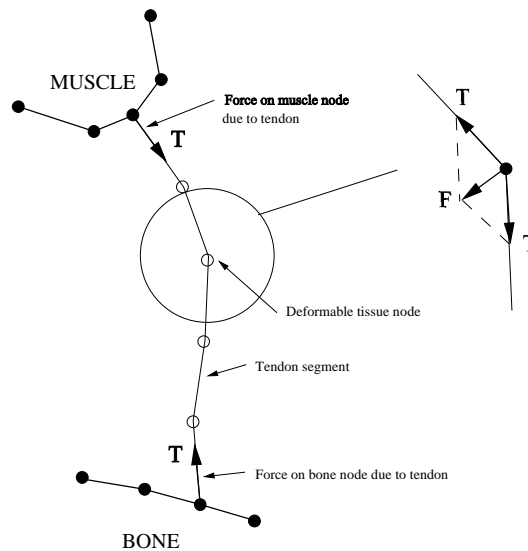


Figure 5.19: Tendon model.

When testing the leg model, one problem that arose initially was that the tendons had a certain amount of “slack” due to the fact that they follow an approximate piecewise linear path between nodes of the tissue mesh which was longer than the real path. As a result when a muscle was activated, the initial muscle contraction was “wasted” (i.e. not transferred to the bone) in taking up this slack in the tendon. In order to reduce this, a tendon pre-tensioning stage was introduced at the very start of each simulation in which the natural length of the tendons was reduced by 20%. This caused the tendon to become taut and adopt a more direct path before any muscle activation occurred. The effect of this is shown in the following results chapter.

5.4 Modelling the Skeleton

As discussed earlier, although bone shows elastic and viscoelastic properties, its stiffness is sufficiently high that no significant deformation takes place during normal use. Hence, the skeleton can be regarded as a collection of rigid body objects. However, it is still necessary to model the mass and moment of inertia properties of the skeleton. One way of doing this is to assume the bones are cylindrical, and to compute these properties based on this simplification. However, since our skeleton may contain arbitrarily shaped bones derived from the initial limb geometry, then this simplification may not be appropriate. Hence, the following approach was adopted. To determine the mass of the bone, we simply sum the mass of each of the cubes that make up its physical representation. The mass of the individual cubes is obtained by multiplying the density of bone tissue (the value of $2.0g/cm^2$ given in Chapter Two was used) by

the volume of the cubes themselves. In the case of the moment of inertia, we compute the value about a particular axis of rotation by again summing the contributions from each of the cubes which make up the bone. For each cube we consider the mass to be concentrated at its center. Hence the total value is given by:

$$I = \sum_{i=1}^n mr_i^2$$

Where, n is the number of cubes which make up the bone, m is the mass of each cube, and r_i is the shortest distance from the center of the i^{th} cube to the line passing through the axis of rotation for the bone.

In order to model the knee joint in the leg model, a simple one degree of freedom joint was used with a fixed center of rotation. This provides a good first order approximation to the behaviour of a real knee joint. The approach used to simulate the skeleton was a standard rigid body approach. This involves computing the forces and torques acting on the bones and then using Newton's second law (and the rotational equivalent) to find the linear and rotational acceleration of the bones about their center of mass. The forces on the bones originate from a number of sources including joint forces, forces from the soft tissue attached to the bone surface, forces from the tendons and the weight of the bone itself.

5.5 Embedded Surfaces

As discussed in the previous chapter, a feature of the meshing approach is that the initial geometric surfaces are retained and embedded in the cubic mesh such that as the physical model deforms, the surface geometry is made to deform with it. This has the advantage that the surface of the model still has an accurate appearance, even though the "actual" physical surface is a fairly crude approximation. The method used to achieve this is similar to the FFD (Free Form Deformation) approach [83] in that the surfaces are embedded within a deforming space (the physical mesh). However, rather than use a series of FFD blocks that allow tri-cubic Bézier deformation, our space is defined by the tetrahedral elements and we assume that the space within each deforms only linearly with zero order continuity between neighbouring elements. This simplification should not present problems since our tetrahedrons are small and have dimensions which are close to the scale of the deformation features that we expect to obtain. The process of deforming the geometric surfaces is relatively simple and consists of two steps. Firstly, during the model pre-processing stage, we consider each of the vertices in the surface meshes and determine which of the tetrahedral elements it is embedded within. This is done by iterating through the list of tetrahedrons and performing a standard inside-outside test. We then find the location of the point inside the particular tetrahedron in terms of a coordinate system which is local to the element and based on its edges as follows. If we consider the point P in the tetrahedron ABCD, as shown in figure 5.20, then the position vector of the point, \mathbf{P} , can be written using the edges of the tetrahedron as basis vectors as shown:

$$\mathbf{P} = \mathbf{A} + u\mathbf{AB} + v\mathbf{AC} + w\mathbf{AD} \tag{5.6}$$




Figure 5.20: Embedded point.

This can then be rearranged into matrix form:

$$\mathbf{P} - \mathbf{A} = \mathbf{M} \begin{pmatrix} u \\ v \\ w \end{pmatrix}$$

where the columns of \mathbf{M} are given by the vectors \mathbf{AB} , \mathbf{AC} and \mathbf{AD} . Therefore:

$$\begin{pmatrix} u \\ v \\ w \end{pmatrix} = \mathbf{M}^{-1} (\mathbf{P} - \mathbf{A})$$

Hence, u , v and w describe the position of the point in the tetrahedron, and this information is stored with the point (along with the index into the tetrahedron list of the particular element). During simulation, when we come to determine the new position of the point, we simply use equation 5.6 along with the precomputed u , v , w and the position vectors of the (potentially deformed) tetrahedron vertices.

5.6 Modelling Collisions

Collision handling is required in two situations for the limb model. Firstly, it is needed to prevent the parts of the skin surface above and below the back of the knee joint from interpenetrating when the knee is bent. Also, it is needed to prevent external objects from intersecting the skin surface. In general the process consists of firstly detecting when a collision or contact situation occurs (collision detection), and then applying constraints to the objects that are colliding to prevent them from intersecting and to cause them to deform realistically (collision response). When discussing these issues in relation to contact between one or more deformable bodies where the surfaces are represented by general non-convex polyhedra, these issues become research topics




Figure 5.21: Triangle-vertex intersection test.

in their own right. As a result, we adopted simple unrefined techniques as described below.

5.6.1 Collision detection

Naive collision detection is a relatively straightforward problem but is computationally expensive and hence much of the research effort has been into considering how it can be optimised [59] [69]. If we consider two surfaces represented by meshes composed of triangles (for simplicity) then the basic intersection test involves determining whether any of the vertices in the two meshes have passed through any of the triangles in the other mesh during the last time step. In actual fact, a truly accurate approach would also test edge-triangle pairs for intersection since it is possible for polyhedral objects to collide edge-on in which case no vertices need intersect. However, it is reasonable to ignore this case and suffer only a small error when it does occur. The basic triangle-vertex test is illustrated in figure 5.21. We shall initially consider the simplified case shown in figure 5.21a in which a vertex passes through a stationary triangle. Here P and P' show the positions of the point at the beginning and end of the time step, and P_0 , P_1 , P_2 are the positions of the triangle vertices. The test for collision is described in the following equation:

$$P + (P' - P)t = P_0 + (P_1 - P_0)u + (P_2 - P_0)v$$

Here the left hand side is a parametric equation for the position along the path between the old and new point positions (where t represents time and goes from 0 to 1), and the right hand side is a parametric equation for the points lying on the surface of the triangle. Hence, if equality exists for $0 \leq t \leq 1$ and values of u and v which lie on the triangle then intersection is taking place during the time step. This is solved by matrix inversion. In the more general case where the triangle vertices are also moving during the time step, as illustrated in figure 5.21b, then the parametric equation becomes more complex:

$$P + (P' - P)t = P_0 + V_0t + ((P_1 - P_0) + (V_1 - V_0)t)u + ((P_2 - P_0) + (V_2 - V_0)t)v$$




Figure 5.22: Intersection test for closed surfaces (polygons).

Where V_i is the velocity of the triangle points. This takes considerably more effort to solve, since it results in a polynomial of order 5 in t . Moore and Wilhelms [69] use a binary search technique to find a solution. In most practical solutions, various techniques can be put in place to reject point-triangle pairs without having to resort to the expense of solving the equation above. For example, Moore and Wilhelms use a simple test to see if the point is on the same side of the triangle at the start and end of the time step and assume that if it is then a collision has not occurred. They also use a bounding box test which determines whether the line segment intersects the bounding box containing the start and end triangles and reject if this is not the case.

The approach above is a general method that can be used for two surfaces. However, in the case where we are considering the intersection of solid bodies which have closed surfaces, then an alternative approach can be used. This involves testing the position of each vertex after the time step to determine whether it is inside the other surface. A simple way of carrying out this inside-outside test is to count the number of intersections that occur between the line segment that joins the point being tested to a point which we know is definitely outside the surface. This is illustrated in figure 5.22 for the equivalent two dimensional situation. Here we are testing the line segment between the point, P , and a point which lies on the x-axis at $+\infty$ (since we know this definitely lies outside the object, and it also simplifies the equations). The line is tested to see how many times it intersects with the surface of the mesh. Here it intersects an odd number of times (three) so we know it must be inside the object. In the case where the line passes directly through a vertex in the polygon (or an edge in the three dimensional situation), care has to be taken to make sure the intersection is not counted more than once. This approach is used in [35] for detecting collisions between two deformable polyhedral bodies. Again, initial tests (e.g. bounding box) can be done to reject trivial cases.

This approach was also adopted in the leg model. However, the situation was optimised considerably by performing collision detection between only those parts of the surfaces which we knew were likely to collide during the simulation. Hence, the actual parts of the surfaces which were tested were the area behind the back of the

knee, for motion in which the leg bends at the knee, and in the case where an external object was likely to make contact with the skin surface, the surface of the object and the skin surface around the likely contact site was tested.

5.6.2 Collision response

Responding to contact or collision requires the application of constraints. A number of methods have been suggested which involve various combinations of position, velocity and force based constraints. Gascuel [35] presents two approaches to handle collision between a deformable and a rigid surface, or two deformable surfaces. In the former case, the intersecting point is simply moved back along its trajectory to a “legal” position on the surface of the body it has intersected. In the latter case, both the point position and the position of the triangle vertices are moved based on an evaluation of the local stiffness of the material around the contact site, again such that the point is moved to a legal position on the surface of the intersected object. This approach is adopted to correctly ensure that deformation of both surfaces occurs as a result of the collision.

In the case of collisions between two rigid bodies, velocity based constraints are typically applied in the form of impulses based on the law of conservation of momentum. However, this is not considered appropriate here since we are concerned with contact that lasts for more than an infinitesimally short time period. Also, the transfer and possible loss of energy that occurs during the collision that is modelled by such velocity constraints is automatically modelled here by the way in which the surfaces of the objects deform and then push one another away as they try to recover their initial shape.

Moore and Wilhelms [69] present a force based approach. This involves introducing temporary springs between the two objects that are intersecting in order to force them apart. One problem associated with this method is that it allows the objects to intersect during the duration of the collision. Also, the stiffness of the spring typically has to be high for it to be effective which requires small time steps to prevent instability. An alternative force based technique introduced by Platt and Barr [77] for responding to collisions between a rigid and a deformable body are known as *reaction constraints*. This involves detecting intersections, and then computing the external force that is required to be added to the intersecting point during the time step in which the intersection occurred, in order to maintain the point in a legal position at the surface of the object. This method is used by Chen and Zelter [22] to prevent their deformable muscle model from penetrating the skeleton. The advantage of this approach is that no intersection occurs, although the approach is difficult to apply to two deformable objects (according to [34]) since the constraints become excessively complex.

The approach adopted for the leg model was to use an inside-outside intersection test, and then apply a position constraint to an intersecting point in order to move it back to the surface of the closest triangle face in the mesh, in a direction along the normal to the face. The velocity of the point was also set to zero. This approach does not intrinsically deform both of the deformable surfaces involved in the collision, however we consider it to be suitable since the high density of the surface meshes involved here

means that if two surfaces are colliding, or are in contact, then interpenetration will be detected for points belonging to both surfaces and each surface will become deformed as a result.

5.7 Control Strategies

In its current form, the deformation and motion of the leg model is determined through the simulation. As a result, the system would not be suitable for animation applications in which the animator wished to manually control low level deformation of the soft tissue regions. In this case a hierarchy of control of the type discussed in Chapter Three would be required. On the other hand, if the requirement is to generate realistic looking animation then a model of this type offers the possibility of greater realism than could be obtained using other techniques.

In order to be able to generate the required skeleton motion, it is necessary to coordinate the activation of the different muscle groups within the model. The issue of control via muscle activation is considered by Hing [45]. Here he discusses the different types of control mechanism that have been observed in animals and humans, and describes approaches that have been researched based on these findings.

Control systems can be classified as being either open-loop or closed-loop. The former type issue control signals to try and produce the required motion but make no attempts to monitor whether the resulting motion actually matches this desired motion. Closed-loop control on the other hand, monitors feedback from sensors whilst the motion is being carried out, and issues further control signals if the motion begins to deviate from that which is required. Observations of biological systems have lead to a number of hypothesized control structures including, motor programs and central pattern generators. These are thought to be open-loop systems which are responsible for rapid types of motion which would not be achievable by closed-loop systems due to the propagation delays involved. Another observation that supports this is the fact that many animals are still able to carry out complex motion even if information from their senses has been cut off (e.g. due to nerve damage etc). Motor programs are predefined sequences of activation for groups of musculotendon systems. They are used to produce quick controlled motion that often would not be possible in a closed-loop system. For example, if a person was given instructions to move their hands towards one another quickly but to stop them before they make contact, then a closed-loop control system would only be able to perform this at a limited rate since it would be relying on feedback information to tell it when to issue the muscle activations to stop the motion of the hands. However, a motor program would attempt to determine the required timing of the activation signals in advance in order to remove the need for feedback.

Central pattern generators are thought to control repetitive coordinated motion, for example during walking and other gaits. These can be thought of as oscillators which send out regular patterns to various muscle groups. It is also thought that the controllers for different limbs are coupled in order to retain the correct phase relationships when the motion of a limb is disrupted, for example due to it colliding with an external object.

Although open-loop systems are thought to be responsible for initiating motion, closed-loop control is also necessary to correct error that occurs during the motion. The simplest kind of closed-loop system is proportional control, in which the current activation signal is changed by an amount proportional to the error between the current position and the desired position of the particular part of the body. However, if we apply this to musculotendon systems it often leads to oscillation about the desired position due to the delay involved in the muscle activation manifesting itself as motion about a particular joint. Because of this more complex closed-loop control strategies are required.

Although we are not restricted to using models observed in real biological systems to control our model, it is likely that a similar level of complexity would be required in an artificial system to produce the types of motion that are observed in real limbs. Because of this, the field of control via muscle activation was considered to be largely outside the scope of this work, and a simple open-loop control system was implemented in which the muscle activations are controlled according to a script. The scripting system contains commands describing how muscle activation should change at certain predetermined times. The following script is from a test that was carried out on the leg model (and is described in the next chapter), which involved activating the back muscle to flex the knee until the lower leg is approximately horizontal, and then deactivating the muscle to allow the leg to fall.

```
% Set initial activation to zero
0.0 BACK_MUSCLE 0.0 0.0

% Increase back muscle activation gradually
0.01 BACK_MUSCLE 0.4 0.99
1.0 BACK_MUSCLE 0.65 1.0
2.0 BACK_MUSCLE 0.85 1.0

% Reduce back to zero
8.0 BACK_MUSCLE 0.0 1.0
```

The format of each line is:

```
<time> <muscle name> <activation value> <time scale>
```

where “time” is the simulation time at which the command should be carried out, “muscle name” is the muscle to apply the command to, “activation value” is the new value that we wish to change to, and “time scale” is the length of time over which we wish the change to take place (change between activation values occurs linearly). The specific values used in the scripts were determined based on experiment. For example, in the case of the above script, the activation levels were altered and the simulation re-run several times in order to make the leg rise at the required speed without overshooting the desired position.

5.8 Simulation Issues

The leg model contains two type of dynamic system, the MSD tissue model and the rigid body skeleton. Simulation of these systems is carried out using a standard Euler based approach, which uses first order approximations to the first and second derivatives (with respect to time) of the point mass positions and rigid body orientations (see [6] for more details regarding this approach).

A number of optimisations were used to reduce the amount of computation required for the leg model, in order to speed up the rate of simulation. The areas which required the most computation are the calculation of the forces generated by the deformable elements, and the collision detection. In the case of the collision detection, the main optimisation used was to only detect for collisions in regions on the skin surface where we knew they were likely to occur, for example behind the knee. In the case of the leg model, this typically reduced the computational requirements for collision testing by approximately 80%. In the case of the element forces, most of the computation is involved in computing the face repulsion forces, the compression forces and the extra forces which are required to ensure that the net forces and torques are zero. Hence, in order to reduce this, a similar approach is used to that in the collision detection in which the collapse prevention, and volume preservation forces are only applied in an element if sufficiently deformed from its rest shape. This is determined by considering the distance of the points from the opposite faces in the element, such that if any of these distances drops to below 70% of its original distance, then these forces are applied. This resulted in large savings in computation since excessive deformation typically only occurs in localised areas around the joint and muscle regions, whilst the rest of the leg remains relatively unaffected. Reductions of approximately 40% in processing requirements of the soft tissue model were achieved, although the actual value varied greatly based on the action being performed by the leg.

5.9 Summary

This chapter has considered how the physical and material properties of the various tissue regions within limbs can be modelled, and these techniques have been applied to the model of a leg. One of the main areas considered was how to model the soft tissue regions in a limb model, and a new enhanced mass-spring-damper (MSD) approach was presented for this purpose. This incorporates a number of important additions to standard MSD models in the form of constraints. The first of these is a technique to ensure that the tissue model remains robust within the range of strains that are likely to be experienced in the leg model. The second one acts to maintain the compression of the tissue within acceptable limits. This is important in order to recreate important effects for example bulging around joints and muscle regions.

In section 5.2.7 a new muscle model was presented. This is based on the soft tissue model with the addition of units which generate contractile forces within the muscle tissue. Sections 5.3 and 5.4 then described the techniques which are used to model the other main elements in the leg model, notably the tendons and the skeleton. Sections 5.5, 5.6 and 5.7 described the techniques used for three other important requirements

of the model. Namely, how the geometric model is embedded within the physically based model, how to model collisions in order to prevent intersection between different parts of the skin surface, and the control strategy used during simulation. Finally, section 5.8 described the details of how the dynamic systems (the soft tissue and the skeleton) were simulated, and also described optimisations that were used to reduce the computational requirements of the model.

Chapter 6


Results

A number of simulations were carried out using the leg model described in the previous chapters. These were designed to test the effectiveness of the tissue modelling techniques and to determine how the leg model behaved as a whole. The simulations were carried out on a PC with a P200MMX processor running Microsoft Windows NT. An interactive simulation system was developed which features a window showing the model being simulated (rendered using OpenGL), and also a graphical user interface which enables most of the variables associated with the model and the simulation to be varied if required. During simulation of a model, the current state of all the geometric and physical surfaces is stored to a file every 25th of a second (of simulation time). In order to view the simulation data, the resulting file is converted into a series of 3D object files (in the 3D Studio mesh format) which store the position of the surfaces at every “frame”. These are then rendered using a commercial rendering package to produce the images shown in this chapter. In most of the images, the majority of the surfaces are invisible so that those being considered can be seen more easily. It is also important to note that the colour and reflectance properties of the surfaces have been chosen to portray surface shape more clearly, and are not meant to represent the real appearance of the tissue types that they represent. For example, in many cases the skin surface is shown with an unrealistically high specular coefficient in order to generate distinct specular highlights on the surface which aid the visualisation process.

The tests which were carried out are now described, and images are presented which were rendered from the simulation data.

6.1 Gravity Effects

One of the initial tests performed simply involved observing the effects of gravity on the model. Figure 6.1(a) shows the leg model in its initial rest position at the start of the simulation. Gravity was then turned on resulting in immediate sagging of the flesh which could be seen to fall noticeably until this was resisted due to its connection with the bones. The flesh could then be seen to oscillate slightly about its new rest position. As well as the motion of the flesh, the entire lower leg swung backwards and again oscillated about a new rest position. The overall motion died down after



(a) Gravity off

(b) Gravity on

Figure 6.1: Effect of gravity on model deformation.

approximately four seconds resulting in the shape shown in figure 6.1(b). The motion of the lower leg is caused by the fact that its centre of gravity lies slightly forwards of the bone (somewhere near the middle of the shin) due to the mass associated with the foot. Hence, a new position is adopted when gravity is applied such that the centre of mass lies directly beneath the knee joint. Overall, the model reacted to gravity in the same way as a real leg. However, the situation shown is slightly “incorrect” since the original skin surface mesh already shows the surface shape with gravity acting, and as a result gravity is effectively being applied twice. One possible solution would be to apply gravity to the physical model before embedding the skin surface mesh within it. However, this was not considered necessary here since the error is only slight.

6.2 Tendon Pre-tensioning

This simulation was carried out to illustrate the effects of the tendon pre-tensioning which was described in Chapter Five (section 5.3). The simulation was started with the leg in its rest position, and the natural length of the tendons set to their original values (corresponding to the lengths of the component line segments). The natural length of the tendons was then reduced over a short period of time to 70% of their initial value. This caused them to contract and adopt a more direct path between the muscle and bone attachment points. This in turn lead to deformation of the surrounding tissue regions as we would expect. Figures 6.2(a) and 6.2(b) show the physical muscle and bone regions before and after pre-tensioning, and figures 6.2(c) and 6.2(d) show the effect on the geometric surfaces.

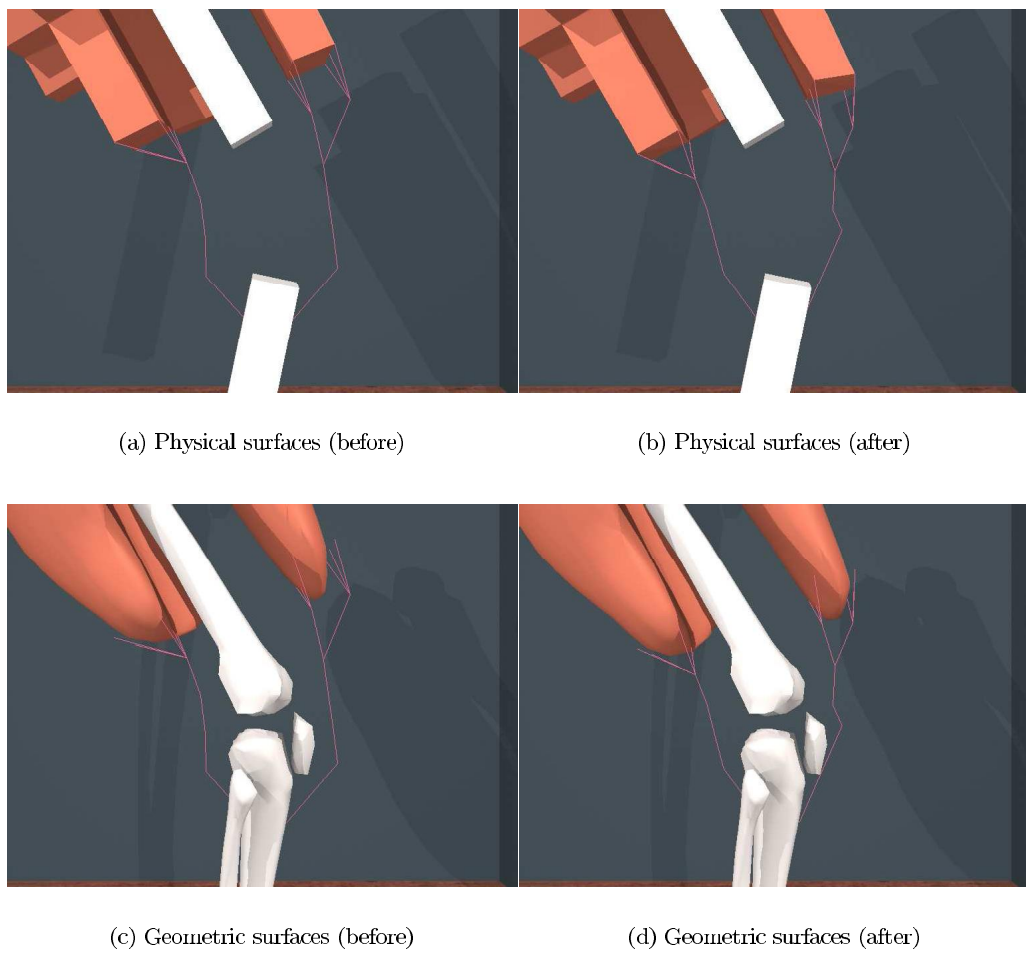



Figure 6.2: Effects of tendon pre-tensioning.

6.3 Bending of the Knee Caused by an External Object

Here the aim was to examine the behaviour of the leg as it bends, and in particular to concentrate on the deformable tissue to determine whether effects such as creasing and bulging occurred around the knee joint. In order to observe this in isolation (without muscle activation), the bending was achieved by pushing the leg using an external object rather than via muscle contraction. The simulation began with the leg at rest and the upper bone fixed at an angle of 30 degrees to the vertical. A cylindrical object was then moved towards the front of the leg just above the ankle. As the object made contact with the skin surface, the flesh around the contact area could be seen to compress and the resulting forces on the bone caused the leg to begin rotating upwards. The object was then used to push the lower leg upwards until it was roughly horizontal (the motion of the external object was scripted). Figures 6.3(a), 6.3(b), 6.3(c) and 6.3(d) show the physical skin surface at four stages during the simulation. The first of these is the initial position before the cylinder makes contact with the skin surface. The second and third show the lower leg at angles of roughly 30 and 45 degrees to the vertical, and the fourth shows the leg in the final position. Figures 6.4(a), 6.4(b), 6.4(c) and 6.4(d) show the state of the geometric skin surface at the same stages during the simulation.

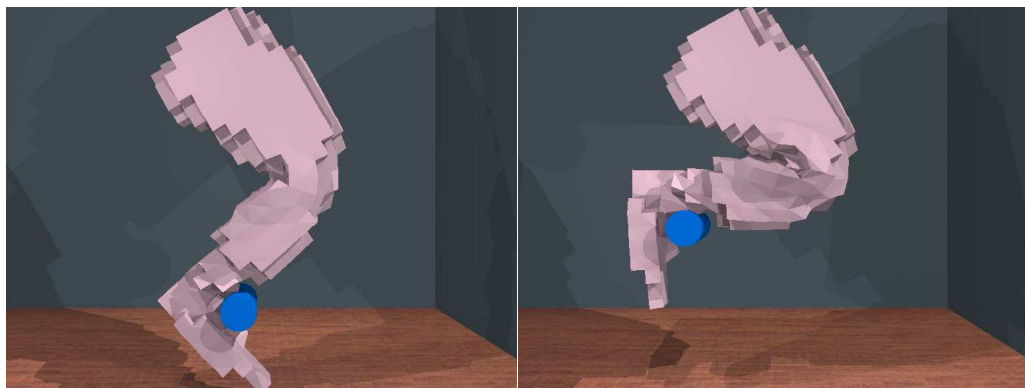
In order to correctly simulate the behaviour during this test, collision detection was present for a region behind the knee in order to cater for the skin-to-skin contact around the crease, and also at the front of the ankle for the object-to-skin contact. The specific collision regions were chosen such that they would be large enough to cover the likely contact area, whilst not being overly large resulting in wasted processing resources. The images show that the model behaved in a realistic manner during this simulation. Creasing around the back of the knee can be observed along with bulging of the flesh around the joint. This indicates that the limited compressibility built into the material model is effective and leads to the same effects that are observed in real limbs.

Figures 6.5, 6.6, 6.7, and 6.8, show regions of interest for the leg in the final position. Figure 6.5 shows the contact that occurs between the physical skin surface above and below the joint, and the bulging that occurs as flesh is pushed out of the contact region. In figure 6.6 the skin surface is shown. This image highlights a minor anomaly in the model, namely that a slight gap exists between the geometric skin surfaces either side of the crease. This is due to the difference between the positions of the physical and geometric surfaces. In figures 6.7 and 6.8 views have been chosen which highlight the bulging around the sides of the joint. The effect can be seen to be very pronounced, almost to the point of been unrealistic. This is again partly due to the fact that the geometric skin surface lies a small distance within the model, in this case at a position where the curvature is higher than at the true surface, and this results in a smaller more curved bulge rather than a larger more rounded one. Finally, figure 6.9 shows the deformation of the muscle regions. When compared with the rest state, it can be seen that the muscle at the front of the leg is stretched slightly by the bending of the leg whilst the muscle at the back experiences slight compression close to the joint.



(a) Initial position.

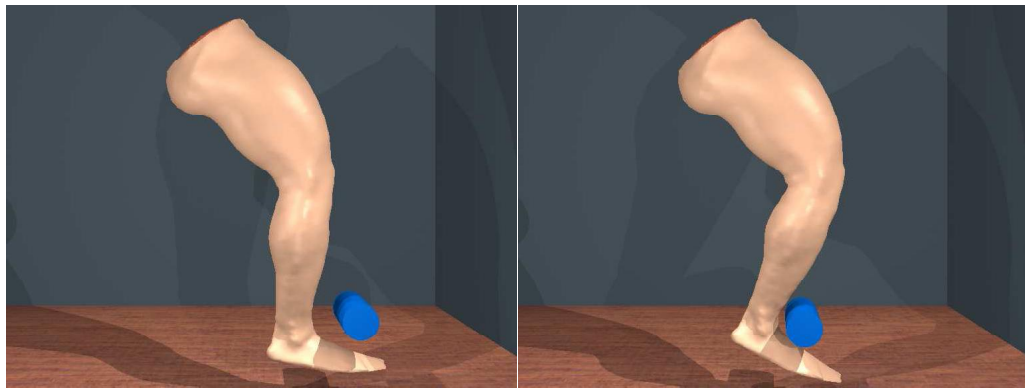
(b) Second position.



(c) Third position.


(d) Final position.

Figure 6.3: Leg bending due to an external object (physical surfaces).



(a) Initial position.

(b) Second position.



(c) Third position.

(d) Final position.

Figure 6.4: Leg bending due to an external object (geometric surfaces).



Figure 6.5: Close up view of creasing.




Figure 6.6: Skin surface around the bent joint.




Figure 6.7: Flesh bulging around the joint (1).



Figure 6.8: Flesh bulging around the joint (2).

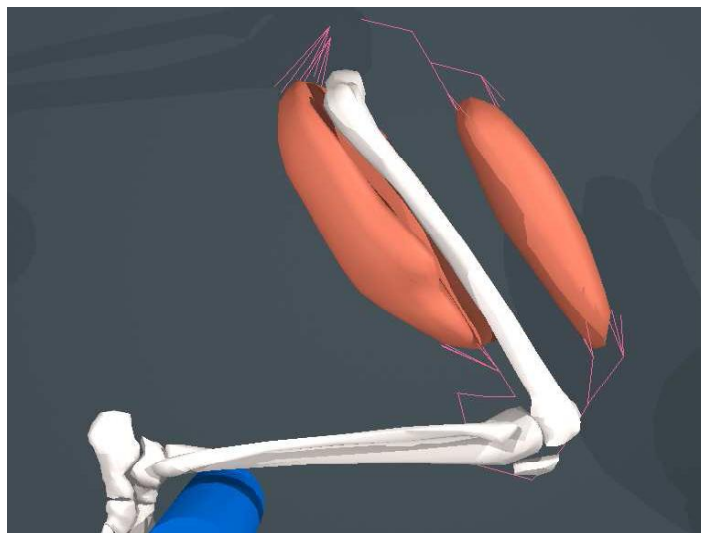
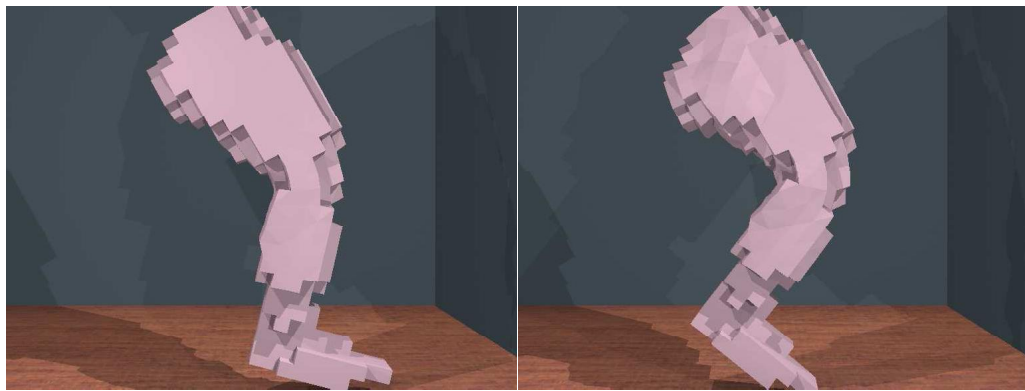


Figure 6.9: Muscle deformation around the bent joint.

6.4 Activating the “Hamstrings”


This simulation involved activating the muscle region at the back of the leg (which represents the hamstrings), and observing the resulting motion of the leg and deformation of the flesh and muscle regions. The parameter controlling the degree of activation of the muscle was increased from zero up to a maximum value over a period of approximately five seconds, was held there for several seconds, and then reduced back to zero again. Figures 6.10(a), 6.10(b), 6.10(c) and 6.10(d) show the physical skin surface at four stages as the leg is lifted due to the muscle contraction. Figures 6.11(a), 6.11(b), 6.11(c) and 6.11(d) show the geometric skin surface. Figures 6.12(a), 6.12(b), 6.12(c) and 6.12(d) show the physical surfaces of the bone and muscle regions, and also the tendons. Finally, figures 6.13(a), 6.13(b), 6.13(c) and 6.13(d) show the geometric surfaces of the bone and muscle regions.

The images show the effects that occur as the model is simulated. As the activation of the muscle increases, it can be seen to contract. Bulging of the muscle is also evident in directions perpendicular to the direction of contraction. As the muscle contracts it pulls on the tendon making it increasingly taut. Evidence of this can be seen by the way in which the tendons attached to it straighten out as the force increases. The tension in the tendon acts on the lower part of the leg, causing bending at the knee. A number of effects are visible at the skin surface. Firstly, the effects around the bent joint that were present in the previous test are also present here as expected. Also, the muscle deformation produces slight bulging at the back of the thigh. Unfortunately this is not particularly distinctive since the shape change is masked slightly by the shape of the buttocks. The effect is more obvious if the images are compared with the images of



(a) Initial position.


(b) Second position.



(c) Third position.

(d) Final position.

Figure 6.10: Physical skin surface during back muscle activation.



(a) Initial position.


(b) Second position.



(c) Third position.


(d) Final position.

Figure 6.11: Geometric skin surface during back muscle activation.



(a) Initial position.


(b) Second position.



(c) Third position.


(d) Final position.

Figure 6.12: Internal physical surfaces during back muscle activation.



(a) Initial position.

(b) Second position.



(c) Third position.


(d) Final position.

Figure 6.13: Internal geometric surfaces during back muscle activation.



(a) Physical skin surface.

(b) Geometric skin surface.



(c) Physical internal surfaces.

(d) Geometric internal surfaces.

Figure 6.14: Close up of surfaces during back muscle activation.

the leg bent without muscle activation (figures 6.4(c) 6.4(d) and 6.6). Finally, images 6.14(a), 6.14(a), 6.14(a) and 6.14(a) show close ups of the fully activated muscle region viewed from behind the leg. The first two figures show the physical and geometric skin surfaces and the last two the physical and geometric muscle surfaces. These images show that significant bulging of the muscle also occurs in a mediolateral direction, and this in turn produces significant bulging at the skin surface. The slight asymmetry that can be seen in the contraction of the muscle region (the left hand side appears to contract more than the right hand side) is due to the fact that there is more muscle bulk on the left than the right, and also fewer tendons are attached to the left hand side of the muscle.

6.5 Activating the “Quadriceps”

The next simulation that was carried out involved activating the muscle region at the front of the leg which represents the muscles in the quadriceps. Again, the simulation involved activating the muscle over a period of time, holding it at maximum activation for several seconds, and then deactivating it. As shown in the following images, the muscle activation chosen was sufficient to almost fully extend the leg (this was arrived at by experimenting with different activation profiles). Figures 6.15(a), 6.15(c) and 6.15(e) show the geometric skin surface at three stages as the leg is extended due to the muscle contraction. Figures 6.15(b), 6.15(d) and 6.15(f) show the physical skin surface. Figures 6.16(a), 6.16(c) and 6.16(e) show the internal geometric surfaces, whilst figures 6.16(b), 6.16(d) and 6.16(f) show the internal physical surfaces.


These images show contraction and slight bulging of the muscle region, straightening of the tendon due to the tension in it, and limited deformation of the skin surface at the front of the thigh. The deformation is less pronounced than in the previous simulation since the muscle region is much smaller, and less contraction occurs since the geometry of the musculoskeletal systems does not change as much. An obvious artifact that appears in figures 6.15(c) and 6.15(e) is the irregular silhouette at the front of the thigh. This reflects corresponding irregularities in the front surface of the muscle region as seen in figure 6.16(f). This effect is thought to be due to anisotropic effects in the muscle and soft tissue models which are not “averaged out” due to the fact that the region is only a single cube thick. Finally, figure 6.17 shows a close up view of the muscle region at the front of the thigh at full contraction, and figure 6.18 shows the geometric skin surface. This shows that significant bulging of the muscle occurs in a lateral direction (across the thigh), and this leads to corresponding bulging of the skin surface.

6.6 Isometric Contraction of Opposing Muscle Regions

This simulation involves contraction of both muscles at the front and back of the thigh such that the forces produced cancel one another out resulting in no motion of the limb. This “state” is typically seen in real limbs when they are tense, for example before an “explosive” motion. Although no motion about the joints occurs, we would still expect to see significant deformation of the tissue regions due to the muscle activation that is taking place. Figures 6.19 and 6.20 show the skin surface before and after muscle contraction. Figures 6.21 and 6.22 show the internal surfaces before and after muscle contraction.


The images show that shape changes occur at the front and back of the thigh and around the paths of the tendons. However, as in a real leg, the deformation is only slight when compared to non-isometric tension experiments in which the joint angles change resulting in greater changes in muscle length. This issue is also noted by Chen and Zelter [22] when assessing their muscle model under similar conditions:

The other result taken from this simulation, interesting from the point of view of computer graphics, is that very little changes in shape are produced




(a) Geometric (1).

(b) Physical (1).



(c) Geometric (2).


(d) Physical (2).



(e) Geometric (3).


(f) Physical (3).

Figure 6.15: Skin surface deformation during front muscle activation.




(a) Geometric (1).

(b) Physical (1).



(c) Geometric (2).

(d) Physical (2).



(e) Geometric (3).

(f) Physical (3).

Figure 6.16: Internal surface deformation during front muscle activation.




Figure 6.17: Close up of the front muscle region.




Figure 6.18: Skin at the front of the thigh during maximum muscle activation.



Figure 6.19: Skin surface before dual muscle contraction.




Figure 6.20: Skin surface after dual muscle contraction.

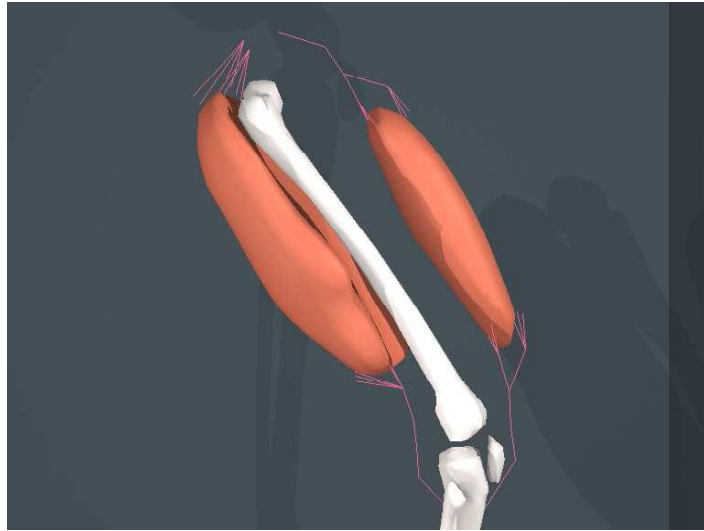


Figure 6.21: Internal surfaces before dual muscle contraction.




Figure 6.22: Internal surfaces after dual muscle contraction.

by a muscle undergoing purely isometric contraction ...hence, to produce larger shape changes for the purposes of computer animation requires either the muscle to lengthen or shorten, or for the contraction to work against shape changes due to external forces such as gravity.

This test also highlighted another subtle effect generated by the model that may not be particularly apparent in the images (the effect is more visible in the animations provided on the CDROM). Namely, as muscles contract (with or without changes in overall length) their form becomes more apparent at the skin surface. More specifically, when the muscle is relaxed the skin surface appears smooth, but with the muscle tense the skin surface displays slight undulations which match those in the surface of the muscle. An example of where this can be seen in a real leg is if the quadriceps are tensed with the leg fully extended. This effect helps to portray the same sense of “tension” that we see in a real leg when muscles are activated.

6.7 Isometric Contraction of a Single Muscle Region

This final simulation was aimed at highlighting the deformation that is caused by the tendons under certain circumstances. The situation involves fully activating the muscle at the back of the thigh, whilst preventing the leg from flexing further than ninety degrees using a fixed cylindrical object which makes contact with the back of the ankle. In the case of a real leg, the tendons at the back of the knee can be seen to protrude when in this state. Figures 6.23 and 6.24 show the skin surface before and after muscle contraction. Figures 6.25 and 6.26 show a close up of the skin surface on the inside of the joint. In figures 6.27 and 6.28 the model has been rendered with the skin surface semi-transparent, and finally figure 6.29 shows the inside of the joint from a viewpoint at the back of the leg.

The images show that significant deformation takes place when the muscle is activated. As the tendon straightens it can be seen to deform the surrounding flesh such that the crease at the back of the joint becomes less pronounced. One difference between the deformation seen in real legs under similar circumstances is that here we see general deformation of the whole region behind the joint rather than more pronounced deformation just around the tendon. This is due to the fact that the tendons in the model lie significantly deeper than in a real leg, and as a result the deformation at the skin surface is more spread out. In a real leg, we also see effects due to two separate tendons which lie behind the knee at either side of the joint.

6.8 Summary

The tests described in this chapter were designed to determine whether the leg model behaved realistically under a range of circumstances. The results show that the model displays realistic behaviour in two ways. Firstly, it correctly models the relationship between the four state factors that were introduced and discussed in Chapter Two (section 2.7). For example, the changes in muscle shape produced are determined by the joint state and the muscle activation states. Also, the shape of the skin surface

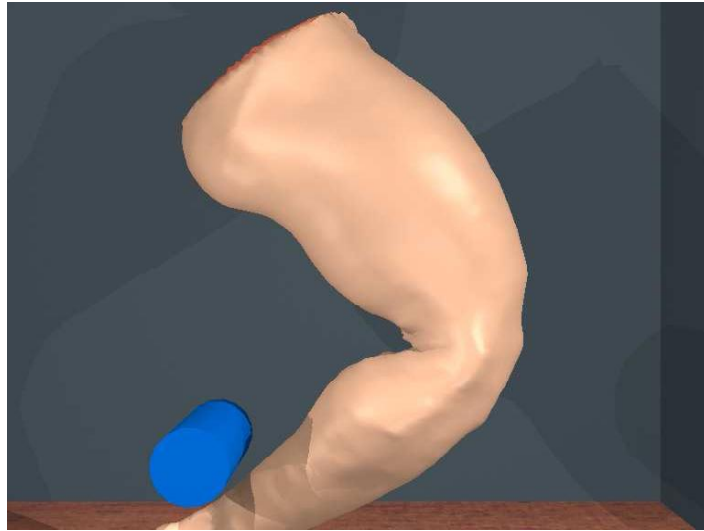


Figure 6.23: Skin surface before muscle contraction.




Figure 6.24: Skin surface after muscle contraction.




Figure 6.25: Close up of skin surface before muscle contraction.




Figure 6.26: Close up of skin surface after muscle contraction.




Figure 6.27: Internal surfaces before muscle contraction.




Figure 6.28: Internal surfaces after muscle contraction.




Figure 6.29: Inside of joint as viewed from the back of the leg.

is influenced by many factors. For example, the skin surface at the back of the knee shows deformation due to the combined influences of knee joint angle, the action of the tendon at the back of the knee, and also gravity. Secondly, the complete range of surface effects that were listed in section 2.7 are exhibited (with the exception of wrinkles).

Chapter 7

Verification and Discussion

This chapter begins by describing a number of approaches that can be used to verify the behaviour of limb models of the type developed here. Results from a simple verification study that was carried out are presented and discussed. The rest of the chapter then discusses the performance of the different parts of the model and the suitability of the techniques used.

7.1 Introduction

In most areas of science and engineering, testing a model to see how well it represents the real thing is an essential part of the modelling process. However, in the field of computer graphics and animation this is rarely done. One reason for this is that much of the research carried out, for example into techniques for animating characters, the aim is not to recreate a specific type of motion that already exists, but rather to provide a tool which allows an animator to create any desired motion. In other areas of research however, modelling real objects and phenomena is the goal, and here verification is necessary.

One area in which this requirement has been recognised is photo-realistic rendering techniques. These model the way in which light interacts with objects, and how this leads to the particular distribution of light within a scene. Much of the early work in this field relied on qualitative comparisons with real situations. For example, an illumination model intended to recreate the appearance of metallic surfaces may be judged by considering whether the images generated look similar to real metallic objects. Although verification approaches of this type can be useful in certain circumstances, they are still subjective, and as a result some researchers have attempted more accurate quantitative assessment of their models. One of the best examples is the work done at Cornell University [38] [24]. Here a physical scene featuring a box with various simple objects placed inside was constructed and used as a bench mark to test radiosity techniques. Accurate measurements of the light coming from different parts of the scene were taken and compared directly to the values generated by the model.

Another area where researchers have attempted to verify their results is in the modelling of human or animal gaits. Hodgins [47] [46] presents animations of a human

model carrying out a range of gaits which are compared to video footage of real people carrying out the same types of motion. Hodgins [48] also presents a method for adapting simulated skeletal motion to different characters, which again features similar comparisons based on video footage.

For models of human limbs of the type presented here, the need for verification will typically depend on the application. For example, models intended for movie special effects will typically not require accurate qualitative comparisons since the application is not a safety critical one. On the other hand, in the case of models for use in medical applications then extensive verification is likely to be essential.

7.2 Possible Approaches to Verification

Comparing the behaviour of a model of even the limited complexity of the one described in this thesis to the real thing is extremely difficult. The main reason for this is that the object being modelling, i.e. a leg, is a three dimensional object containing internal and external surfaces which undergo complex changes in shape over time. Furthermore, some of this dynamic behaviour occurs in a relatively short time scale (significant shape change may occur within a tenth of a second), which presents problems for many techniques. As a result, it is difficult to obtain precise measurements of the motion and deformation of a living human limb which we can compare directly to the output of the simulations, and we typically have to use alternative, less accurate, measurement techniques. Some of the possible approaches are discussed below.

7.2.1 Observation of simulation results

The simplest type of evaluation that can be carried out is to observe the images and animations produced, and consider whether the motion and deformation looks realistic (this was done in the previous chapter). This is probably more effective than it may at first seem, since we have a vast amount of experience observing other people (including their legs), and the human brain is extremely effective at noticing when the appearance of an object differs from what we consider “normal”. It is also possible to perform simple visual tests to determine how accurately the dynamics of a real leg are modelled. For example the frequency of oscillation of the fleshy part of the leg and the speed at which these oscillations die down could be measured and compared directly to a real leg.

7.2.2 Comparison with real images and animations

A more accurate assessment can be carried out by comparing the images from the simulations to photographs of real legs carrying out the same types of motion. Similarly, video footage could be compared with animations to assess the dynamic accuracy. This approach was used in the work described above by Hodgins [47] [46]. Further improvements can be made by ensuring that the conditions under which the images are generated in the real and “virtual” world are as similar as possible e.g. the camera position and settings, lighting conditions etc. This then allow the images to be compared directly, for example difference images could be generated to highlight areas where the

error is high. Images from a range of viewpoints could also be used to compare the deformation of the whole surface.

Taking photographs or video footage of a real limb is a particularly convenient way of indirectly “measuring” surface deformation. However, it suffers from obvious problems in that it only reveals the skin surface shape (as opposed to the shape of the internal surfaces), and also provides a two dimensional image rather than three dimensional data. Nevertheless, it is the most common approach used. Pieper [75] present a model to simulate plastic surgery, and images from the simulation are compared to real post-operative photos. The images were also showed to over a dozen practising plastic surgeons who were able to assess the results based on experience.

7.2.3 Improving image based analysis by adding visual cues

Another way of improving comparisons between images or animations, is to add visual cues in order to make 3D surface shape more apparent. An example might be to increase the shininess of the surface in order to generate specular highlights which aid in the mental reconstruction of the surface shape. This was employed in some of the images that are presented in the previous chapter. An alternative approach is to draw a grid on the real limb, and use texture mapping to apply a similar grid to the skin surface of the simulated limb. These could be made to match up with one another when the limbs are in the rest position, and then the subsequent deformation can be assessed by comparing the distortion of the grid lines.

7.2.4 Direct 3D shape comparisons

Obtaining further improvements over the techniques described above, inevitably means obtaining 3D measurements from real limbs. One way of doing this is to use laser scanning technology. Systems are currently available which are capable of extracting surface geometry to a high degree of accuracy (sub-millimetre). Unfortunately however, these systems often impose restrictions. For example, it is only possible to scan a limited region in space, and hence capturing the movement of an entire leg may not be possible. Also, it typically takes several seconds to perform a scan (depending on the scanning resolution), and hence it may only be possible to capture static deformation, or motion that occurs very slowly. A third restriction is that the method only allows us to extract and compare the deformation of the skin surface. And finally, although such techniques extract the surface shape, they do not extract surface deformation that occurs perpendicularly to the surface normals i.e. we may not be comparing the same points on the surface. This could be overcome by using a facility on many modern scanners which allows not only the position of points on the skin surface to be measured, but also their colour. Hence a grid could be painted on the limb surface enabling this kind of deformation to also be captured in the scan data.

7.2.5 Comparison of volume data

If we wish to compare internal deformation as well as surface deformation, then some kind of volume scanning technology is required. The most commonly used methods are

magnetic resonance imaging (MRI) and computer aided tomography (CAT) scanning. These can be used to generate measurements of tissue density at regularly spaced points throughout a limb. It is then possible, using visualisation techniques, to extract surfaces from this data corresponding to the boundaries between different tissue regions. These surfaces could then be directly compared to simulation results. However, as with scanning techniques, limitations are usually imposed. For example, the apparatus can only operate on the body when it is in the anatomical position due to the limited physical dimensions of the scanning region through which the body passes. Also, a single scan may typically take several minutes, and so capturing data from a moving limb is not possible. Chen and Zelter [22] discuss the possibility of taking a series of MRI scans of a real leg in order to quantitatively assess the deformation of the gastrocnemius muscle model that they have developed. The technique suggested would be to scan the leg under a range of different states of isometric tension whilst also measuring the muscle force generated wherever possible. However, the authors concede that “this would be a big project, and just barely doable with the current MRI technology.”

7.2.6 Assessment of physical quantities

As well as assessing surface shape, another area of importance involves comparing the forces which are occurring in the simulated limb with forces in real limbs. In the case of our limb model a number of physical quantities could be assessed, including the forces occurring at the knee joint, the tension in the tendons at different model states, the material stresses occurring within the various tissue regions, and the forces occurring at the surface of the skeleton due to the skeleton/flesh interaction. Unfortunately, making assessments of this kind is again extremely difficult due to the problems involved in taking measurements from real human limbs. Attachment of strain gauges to tendons and other tissue regions during surgery has been considered, but it is currently not possible to do this in a non-invasive manner for healthy limbs. One example where this general issue has been considered is in the work done by Chen and Zelter [22] in which biomechanically based models of a frog muscle and a human leg muscle are developed. Here the frog muscle is assessed by comparing tension-length and tension-velocity data from the simulations to data measured in real experiments. The authors also consider assessment of the shape changes that occur in the models and state that:

It is much harder to validate the shape changes produced (by the muscle model) than to validate the forces. Videos of live frog and human subjects have been made to serve as checkpoints for the simulations, but these can do so only qualitatively.

7.3 Assessing the Leg Model

Due to the difficulties involved in accurate assessment of models of the kind dealt with here (as discussed above), it was decided to carry out a simple image based comparison of our model with a real leg. The approach used was to film a real leg under controlled conditions using a video camcorder. Images were then “grabbed” from the footage




Figure 7.1: Real and simulated leg in initial position.

using a video capture board. Various steps were taken to ensure that the footage could be compared easily to the research results:

- Since the simulation results only show the right leg, the left leg was covered up during filming.
- The lighting falling on the subject was “softened” (using diffusers) in order to get even illumination of the skin surface, and to prevent shadows from being cast on the surrounding surfaces.
- The camera settings were noted and reproduced when rendering the images from the simulation data.

The following figures show a comparison between images obtained from the video footage and images generated from the simulation results. The motion being carried out is flexion and extension of the leg caused by activation of the muscles at the front and back of the thigh. Figure 7.1 show the legs in their initial position. Here we see that the simulated deformation matches that seen in the real leg quite closely. One obvious difference is the effect of the tendons behind the knee. As mentioned earlier this is due to the depth of the tendon below the skin surface in the model. Figures 7.2, 7.3 and 7.4 show flexing of the leg. In figures 7.3 and 7.4 the main differences are that the effect of the tendons on the deformation at the back of the real knee is not as pronounced as that shown in the model which results in a much more distinct crease



Figure 7.2: Legs flexing (1).

at the back of the real knee. Also, the flesh at the front of the knee is more rounded than that in the real leg. This is due to the fact that skeleton model used is a crude representation of a real skeleton. As a result, the shape determining effect of the bones that is evident in the real images, in particular the knee cap which lies very close to the skin surface, is not evident in the images of the model.

Finally, figures 7.5, 7.6 and 7.7 show extension of the leg. The main difference here is that the muscle bulging shown in the model is greater than that shown in the real leg. This is probably due to a number of reasons. Firstly, the muscle in the model contracts more than in the real leg due to the elasticity in the tendons, and as a result the bulging is greater. Also, the muscle region in the model is concentrated in a region towards the front of the leg rather than being spread out more thinly around the whole of the front of the thigh as in a real leg. Lastly, differences are again apparent due to the irregularities in the surface of the model at the front of the thigh, which were discussed in section 6.5.

The assessment work described above was useful in comparing the main types of deformation. However, its effectiveness was limited due to the low quality of the real leg images used. Hence, higher quality recordings should be used in this kind of work. Also, footage of the leg from a range of positions would have enabled a more thorough comparison of the entire skin surface to be carried out.

The following sections now end the chapter with a general discussion of the techniques used in this research and the results presented.



Figure 7.3: Legs flexing (2).

7.4 Discussion of the Verification Study

The results featured in Chapter Six, and the verification study described above show that in general the model is successful in reproducing the static appearance as well as the motion and deformation of a real human leg. All the main types of skin surface deformation that were discussed in Chapter Two are present. This includes creasing (behind the knee joint when bent), bulging of the flesh (most visible around the sides of the knee joint when bent), bulging of the muscle regions, the effects of tendons (again behind the knee when the back muscle is activated), dynamic effects (the most obvious example seen is wobbling of the flesh when the muscles are deactivated quickly), deformation due to gravity, and deformation due to surface contact.

The dynamic behaviour of the model was also shown to be realistic. For example, the way in which the lower leg moves when either of the muscles is activated. However, the tests performed were fairly limited since each was designed to highlight a specific type of deformation or behaviour. Hence, further work could examine the behaviour of the model under more extreme situations, for example the leg could be made to kick a ball or reproduce the motion that occurs during walking or running. As well as being able to display deformation effects in isolation, the model was also capable of displaying multiple effects simultaneously influencing the same region of the skin surface. For example, in cases where the knee is bent, the skin surface behind the knee displays creasing, deformation due to the action of the tendon, and also slight sagging




Figure 7.4: Legs flexing (3).



Figure 7.5: Legs extending (1).



Figure 7.6: Legs extending (2).



Figure 7.7: Legs extending (3).

due to gravity (although this is barely noticeable).

Many of the effects modelled do not appear in existing geometric models. For example, deformation due to contact with external objects, and wobbling of the flesh. This is due to the fact that they are difficult to define using simple rules. For example, in the case of a crease behind the knee, we know that generally its dimensions will change in a predictable manner as the angle of the knee joint changes, and hence this is relatively easy to incorporate. However, if we consider deformation due to contact with an external object, then the form of the deformation will vary based on the shape of the external object, the position of contact on the skin surface, the current state of the limb (i.e. whether the muscles beneath the contact area are relaxed or tense), and so on. The model presented here is capable of modelling all the variation involved in these effects since the deformation is generated by simulating the behaviour of the tissue rather than having to be explicitly specified. In the case of wobbling of flesh, this is also a particularly difficult effect to incorporate into geometric models for a number of reasons. Firstly, it tends to be more of a global rather than a local effect, in that although it may originate in a particular area (e.g. due to an external collision) it tends to spread out over the surrounding part of the limb until the energy involved becomes dissipated. Wobbling of flesh is one example of a number of subtle, natural effects, that are almost impossible to reproduce realistically using anything other than a physically based approach. This is essentially due to the fact that the exact details of the deformation are difficult to define, and they are perceived as more of a general effect rather than a specific type of deformation. However, subtle though they are, any attempt to reproduce them manually which does not capture their “essence” creates results that are unrealistic.

As well as displaying the specific deformation effects seen in real limbs, the model also successfully recreated the general behaviour of a real leg. For example, in the tests described in sections 6.3 and 6.4, flexing of the knee is achieved both due to the action of an external object which pushes the leg backwards, and also by activating the back muscle region. In the test described in section 6.6 both of the muscle regions are activated causing the forces generated on the leg to cancel one another out resulting in no motion about the knee joint. Finally, in the test described in section 6.7 the back muscle was activated but the presence of an external object prevented the leg from bending any further than ninety degrees. These are examples of behaviour that occurs in real legs, and is automatically recreated by the leg model due to the physically based approach used, but which is typically not included in geometric models. One surface deformation effect that wasn't recreated is wrinkling. This is due to the fact that this is predominantly a deformation effect of the skin rather than the flesh, and hence was not modelled. However, in the case of the leg this has little consequence since very little wrinkling occurs during normal situations.

7.5 Discussion of the Structural Model

Ideally, the meshing process used would have been one which automatically generated an unstructured tetrahedral mesh based on the surface descriptions of the leg model presented at the start of Chapter Three (or a simplified version of it). However, since

such a technique was not available, an alternative approach was developed. The resulting technique is based on approximating the surfaces using a regular cubic grid, and then subdividing the cubes into tetrahedrons. The technique was successful in that it created a mesh which fulfilled the requirements. Namely, that it produced a topologically accurate representation of the geometry provided. More specifically, the bone and muscle regions were all separate and surrounded by a region and non-muscular tissue. However, since the approach was based on a grid with a fixed overall resolution, then this naturally lead to the resulting model being a fairly crude approximation to the original geometry.

The crude nature of the model had a number of implications. Firstly, since the surface of the physical model was a “blocky” low resolution mesh, it was necessary to retain the original surface and embed it within the physical model. In the case of an unstructured mesh, this may not have been necessary since the mesh produced would have incorporated a suitably accurate approximation of the initial surface geometry within the structure of the mesh itself. A second implication was due to the fact that the path of the tendons were approximated by line segments based on the cubic grid points, and hence the resulting tendon lengths were longer than in the original model. This meant that the initial contraction of the muscle regions simply acted to “take up the slack” in the tendons and had no effect on the skeleton. As a result, the pre-tensioning approach was used in which the natural length of the tendons was shortened at the start of the simulation causing the tendons to become taut, and to adopt a more direct path through the non-muscular tissue. Another implication of the use of the structured meshing technique is that it probably resulted in a greater number of tetrahedrons than if an unstructured meshing approach was used. This is because the resolution of the mesh had to be made such that it could resolve the smallest feature in the model (which is the gap between the muscle and the bones). Unfortunately this then resulted in the same density of elements throughout the rest of the mesh, even in places where much larger elements could have been used, for example within the muscle regions, and in the non-muscular region at the back of the thigh. Hence the resulting mesh was inefficient.

An important question relating to the structural simplifications that arose in the meshing process is whether they had a direct effect on the results. In certain areas this was the case. For example, the deformation associated with the tendons was unrealistic in some areas, specifically behind the knee where we saw general deformation over a wide area rather than the distinct outline of the tendons below the skin surface as seen in real legs under the same conditions. The reason for this is primarily a structural one, in that even though the path of the tendons lies only a single “block” of elements away from the surface of the physical model, this is still a significantly larger distance than in a real case (a couple of centimetres as opposed to a couple of millimetres), and hence the resulting deformation becomes unrealistic.

Another problem was that the lack of bones in the knee caused the effect that was observed in the verification study where the region at the front of the knee became more rounded than in the real case when the leg was flexed. This lack of bone also implied an excess of flesh which lead to excessive bulging around the sides of the knee when it was flexed. Another area in which structural simplifications caused a problem was

during collisions due to differences between the physical and geometric skin surfaces. This resulted in a slight gap being visible between the contacting surfaces for both object-skin and skin-skin contact (i.e. behind the knee). However, in the latter case, this was reduced to acceptable levels by moving the grid points at the surface of the physical mesh closer towards the geometric surface during the meshing stage.

Apart from the cases described above, the structural simplifications seem to have had little effect on the overall deformation of the limb. For example, when looking at the images which show the deformation of the geometric model, there are no obvious artifacts to indicate that the physical simulation is taking place at a lower resolution than that of the visible surface. This is probably due to the fact that the mesh resolution is appropriate for the scale of the effects which occur. It also indicates that the “embedding” process which is used to deform the geometric model based on the deformation of the physical model is effective.

Although the meshing approach produces good results for a leg, this does not necessarily imply that it will produce equally good results for other parts of the body. For example, in the case of the arms and the face, the tissue regions tend to be thinner with relative to the thickness of the bones, which may lead to problems creating meshes which can resolve the individual tissue regions without containing excessively large numbers of elements. This is an area for future investigation.

Another area of possible future work is to investigate the possibility of using adaptive subdivision. For example, the non-muscular tissue within the calf region in the model experiences little deformation during most situations. Hence, it would be possible to use larger elements within this region than in other regions of high deformation. In certain situations however, deformation of this area might occur, for example if an external object came into contact with the skin around the calf, and hence small elements would then be required to model the deformation effectively. A scheme which could cater for both these situations would be one which subdivides and recombines elements based on the local degree of deformation. This would then increase efficiency by only using the minimum number of elements required in a given situation. However implementing a scheme of this kind is by no means trivial, and one of the most challenging problems would be to perform the subdivision/recombination without causing “popping” effects i.e. sudden changes in deformation due to differences in the new and old structures.

7.6 Discussion of the Material Model

The deformable material model described in Chapter Five is a new one based on the mass-spring-damper approach. The simulated uni-axial stress-strain tests that were carried out showed that it was suitably robust for use in the limb model, since it was capable of withstanding compressive and tensile strains beyond those that are typically encountered in limbs during normal use (i.e. nominal strains between -0.4 and +0.4). However, the behaviour of the model under other types of strain (i.e. shear and bending) is unknown, and hence this is another possible area of future work.

The new material model was used to represent the muscle and non-muscular tissue. An attempt was made to match the material properties of these tissue regions to those

of real tissue that were presented in Chapter Two. This was achieved by simulating the material in a uni-axial stress-strain test, and adjusting the material parameters until the stress-strain curve matched that of real tissue as closely as possible. This was done to a reasonable degree of accuracy. One approximation that was made, was to ignore the bi-phasic properties of real soft tissue by implementing only the first phase. This was deemed appropriate since the second phase is rarely encountered during normal functioning of a limb.

In the case of the muscle tissue, in order to enable it to generate the required contractile behaviour, force generating units were attached to those edges of the muscle elements which were aligned with the intended direction of contraction. An obvious limitation with this approach when combined with the structured meshing technique, is that the mesh as a whole has to be aligned with the contraction direction. However, in the case of modelling limbs this does not present problems, since the muscles typically contract along the length of the limb which is also the most appropriate orientation of the mesh. The particular muscle model used was a simple one based on a contractile force which is proportional to the muscle activation. This is a simplification since it fails to model the dependency of muscle strength on the overall extension, and the rate of change of length of the muscle. However, since the emphasis here was more on generating realistic deformation rather than realistic biomechanical behaviour, the simplification is considered appropriate. If the model was developed further for use in applications where muscle behaviour is an important aspect, then work could be carried out to implement a more complete muscle model, e.g. as presented by Zajac [107]. Another interesting area of study would be to then perform force-length and force-velocity experiments on the muscle model to see how well the results compare to those from real muscle samples (examples of which are shown in Chapter Two).

The tendons in the model successfully modelled the function of tendons in real limbs, which is to transmit contractile forces from muscles to attachment points on the skeleton. They also partially reproduced the effect that is seen in real life in which the shape of the tendon becomes prominent at the skin surface under certain conditions. However, a problem encountered was that it was not possible to model the true elastic modulus of tendon (typically 35MPa), since this would have meant using an extremely small step size in order to maintain stability. As a result, a much lower modulus was used (0.8 MPa). An alternative approach might be to model the tendon as a completely inelastic object rather than a stiff elastic one. This would involve maintaining the length of the tendon at its rest length when under tension by applying position constraints to the points along its length and at the end points.

7.6.1 Simulation issues

The model requires a large amount of processing power which leads to slow simulation. Typically this ranged from 5 minutes to 15 minutes per second of simulation time (depending on the state of the model) on a PC with a P200MMX processor. The high processor demands are due to a number of reasons. Firstly, the model incorporates a large number of deformable elements. Secondly, the soft tissue model used requires that the simulation step size has to be kept small when high degrees of compression are

occurring in the model, in order to maintain stability. Finally, other issues, for example collision detection, add significantly to the processing requirements. In order to try and improve simulation speed two optimisations were used. Firstly, collision detection was only applied in regions where collisions were likely to take place. Secondly, in order to reduce the computational demands of the soft tissue model, the techniques for collapse prevention and for limiting compression were only applied when the elements were compressed below a nominal strain of 0.6. Hence for areas of the model where little deformation took place e.g. in the calf, then only the basic material model was active. Another possible optimisation would be to simulate different parts of the model using different time steps based on the stiffness of the soft tissue present. Useful areas of future research could compare the effectiveness of the new material model to alternative approaches e.g. the finite element method. A possible candidate for an alternative deformable material model would be the technique presented by Smith [84], since here the particular approach used means that it doesn't suffer from the same problems when modelling stiff material. Indeed, the approach models large elastic deformations and quasi-rigid behaviour within the same theoretical framework.

Chapter 8

Conclusions

This thesis has presented a new physically based model of a human leg which incorporates several distinct tissue regions representing key anatomical features. These tissue regions interact with one another in a realistic manner with the result that they automatically recreate the behaviour that occurs in real human limbs. Namely, that muscles generate contractile forces which are transmitted to the skeleton via tendons, and in turn lead to motion of the limb, and deformation of the skin surface and the internal soft tissue regions. Tests have been carried out on the leg model to examine how it behaves under a number of different circumstances, and more specifically to observe the deformation effects that occur at the skin surface. The results from these tests show that the leg model correctly simulates the causes and effects of skin surface deformation, and as a result attains a level of realism that is not possible with the main existing approach of focusing solely on the shape of the skin surface, and the way in which this shape changes, rather than considering the physical aspects of the limb as a whole.

As a result of this work, the thesis has presented the argument that in order to achieve truly realistic computer generated images and animations of human limbs, a new methodology is required in which the focus is shifted away from *form* (the way in which it deforms) and *function* aspects (the way in which it operates), to *physical* aspects (the way in which it is assembled and the material properties of each region). The leg model clearly demonstrates the results of such an approach, however a number of other benefits are also apparent. For example, the process of constructing the model requires no explicit specification of either the way in which the model should work, or the deformation effects that it should exhibit. Rather, it is only necessary to define the structure of the tissue regions in the model, the way in which they are connected, and the material properties of each region. The result is then that the functional behaviour and the deformation effects **automatically** emerge from the simulation due the action of the tissue regions. Furthermore, as long as the structural and physical properties are accurate, then the model will behave realistically for the entire range of situations that might occur in the case of a real leg, rather than, as with geometric models, representing only the range of situations originally conceived during the construction of the model. These points support the main argument that focusing on physical aspects is an effective approach. However, the process of developing the leg model has also shown that it is

not possible to disregard form and function completely. This due to the fact that in order to make the model feasible (in terms of its implementation and its processing requirements) significant simplifications have to be made in terms of the number of tissue regions represented, their structure, and the material properties of the tissue regions. To do this effectively, it is necessary to have some understanding of function in order to know which tissue regions are relevant and which can be ignored, and also of form (i.e. deformation effects), in order to know which material properties are relevant and which can be ignored.

Although the leg model exhibits realistic behaviour, it does have a number of limitations. For example, it represents only a fraction of the complexity of real human legs. More specifically, it only models motion about the knee, and incorporates only two musculotendon systems rather than the twenty six that exist in real legs (for generating motion about the hip and knee). Also, the tissue regions in the (physical) model are crude approximations of those in the real leg, and this results in unrealistic deformation in some areas, for example around the front of knee due to the reduced amount of underlying bone. However, the fact that even this structurally simple model is able to show the correct range of causes and effects of skin deformation, indicates that inaccuracies in the results are due to the structural simplifications rather than any flaw in the overall approach.

Another major piece of work presented in this thesis, is a new model of soft tissue. This uses a mass-spring-damper approach to deformable material modelling which is attractive due to its simplicity and the fact that it is intuitive. However, many existing models which are based on this approach, suffer from the problem that they are extremely compressible, and are often not robust. This thesis has presented solutions to both of these problems, and as a result has created a model which can be used to effectively represent soft tissue. This is supported by the simulated uni-axial stress-strain tests that were carried out in Chapter Five, which showed the material behaving in a robust manner for the full range of strains that occur in the human body during everyday use. The deformation effects seen in the leg model e.g. bulging and creasing, also indicate that the model is effectively reproducing the relevant behaviour of real soft tissue. However, one problem with the model is that it requires large amounts of processing power. Also, the MSD approach in general, shows fundamental limitations when it is used to model stiff material since it necessitates the use of small time steps (thereby increasing the computational requirements needed to simulate at a particular rate) in order to prevent instability. Finally, since the model is based on enhancing basic MSD models, the approaches used to implement the collapse prevention and the incompressibility constraints could be considered ad hoc, rather than part of a low level unified approach. Further research could be done to compare this approach to other alternative soft tissue modelling approaches, for example FEM models based on elasticity theory. However, any alternative approach must be capable of modelling large strain deformation and limited compressibility in order to realistically model the skin deformation effects that were shown here.

An important part of the work presented here is the meshing process used. The approach involves initially approximating the structure of the tissue regions using a cubic grid and then subdividing this into tetrahedrons. The main problem with this

approach is that it produces a crude approximation to the real structure of the tissue regions, and is unable to represent any features on a scale lower than the mesh resolution. One implication of this was the decision to retain the original high resolution surface model that the mesh was generated from and use it to represent the visible surface of the model. Another is that the scale of the mesh has to be made such that it can resolve the smallest feature in the model, which proved to be inefficient since most regions in the leg do not require such small deformable elements. These problems could be alleviated if an unstructured tetrahedral mesh was used since it could contain elements of varying sizes. However, an investigation into existing techniques failed to find one that was capable of automatically meshing models of the relatively complex topology required here, without generating excessive numbers of elements.

Another mesh related issue is the ability to resolve *deformation* features during the simulation. Since the mesh resolution was fixed during the simulation, then as in the meshing process, only deformation features on a scale higher than the mesh resolution can be displayed. Fortunately, this did not seem to adversely affect the results presented. The main reason for this is that the scale of the deformation effects were typically large compared to the mesh resolution chosen. However, this may not be the case in other situations. One possible solution to the general problem of inappropriate mesh resolution, might be to consider the use of adaptive subdivision during simulation. Here, a scheme could be used in which regions of the mesh are altered based on the local strain, by subdividing or combining elements. However, a challenging problem that would have to be considered would be the issue of avoiding obvious “popping” (sudden motion) as the local mesh was changed.

Turning now to the suitability of the model in the various application areas, the ability of the technique used to automatically generate realistic motion and deformation of the leg would be extremely useful in animation applications. Unfortunately, its computational demands would preclude it from being used in an environment where the animator has to be able to interact with the model in realtime. However, an alternative scenario in which the model is used as a post processing stage to generate realistic flesh deformation based on predetermined skeleton motion could be feasible. One disadvantage with this approach however, is that the full capacity of the model to automatically generate flesh deformation and skeleton motion (and thereby simulate the general functioning of the limb) is not being utilised.

In order to fully utilise the model it would be necessary to control it via muscle activations. Attempting to do this manually, say, in order to get a character to walk, would be virtually impossible - the animator would be attempting a task that has required millions of years of evolution for our brains to be able to do. Hence some method of automatic control would be required to perform this, ideally based on high level commands (e.g. walk from A to B in a “lazy” style). Muscle based control is a current research topic [45] although, due to the substantial complexity of the problem, it is likely to be some time before it becomes a practical approach to character control in animation applications.

Another important issue relating to use in animation applications is how to achieve control over the deformation of the skin surface in the case where the animator wishes to create *unrealistic* results. One way in which this could be done is to alter the

physical properties of the model in order to create effects that are still physically valid, but which would not normally occur in the real world. For example, various kinds of Disney-style comic effects could be achieved by, say, increasing gravity to make flesh sag excessively, or reducing the damping in the material to make it “extra wobbly”. However, if the animator wished to generate results which were not within the scope of this approach, or were not physically valid, then in its current form the model would not be appropriate since it contains no means of directly affecting flesh deformation. A possible area of future research therefore would be to investigate the use of animatable constraints to enable this kind of deformation control.

Another application that was considered is the use of the model in biomechanical simulations of real limbs. Here, models of this type could be particularly useful since the information generated by the simulations contains not only the deformation of the tissue regions, but also the forces that are acting within the model at a particular time. This includes, forces in the tendons, forces which act on the bones in the joints, the stresses in the elements at the skin surface and so on. In the case of, say, a simulation involving a human sitting on a new design of car seat, these forces could be used to determine psychological issues like “comfort”. However, it is likely that an application of this kind would require a greater level of accuracy in terms of the model structure and the behaviour of the various tissue regions.

The third main application area discussed was use in medical applications. Here, as with realtime animation, the processing demands of the model would currently preclude it from being used in realtime medical applications e.g. simulated surgery. However, again it could be used “off line” to simulate surgical procedures, as with models of the face presented by Koch *et al.* [55] and in other work [75] [27]. Again however, the accuracy required would probably be significantly greater than that presented here.

A final part of the work was the verification study. The technique presented involved comparing images of the model to images of a real leg taken under similar circumstances. This is not a new idea by any means, and a similar form of visual analysis of humans and animals was carried out almost a hundred years ago by Edward Muybridge [70] who used a line of cameras connected to trip wires, to study the various gaits that are used by humans and animals.

In a few areas of the leg model, for example around the front of the knee, it was obvious before the study was made that the deformation exhibited there was significantly different to that in real limbs. In these cases, the study proved to be useful in defining exactly what the differences were, and even suggesting the reasons behind them. However, the study also proved to be useful in highlighting errors in other parts of the model that did not appear “wrong” when the images were viewed alone.

The verification study was shown to be useful. However, analysis of the behaviour of limb models in this way is something which is rarely carried out. This is primarily due to the fact that most of this research originates from the computer graphics and animation communities, which generally shows little regard for the verification of results (except in some radiosity studies and recent work involving modelling human gaits). However, this situation may not persist for much longer, since limb models of the type presented in this thesis have the potential to become invaluable tools in many safety critical medical and biomechanical applications. Hence if, or rather when, this happens,

effective testing and verification will not merely be a useful exercise, but will become an essential part of the modelling process.

Bibliography

- [1] J. Ruppert and R. Seidel. On the Difficulty of Triangulating Three-Dimensional Non-Convex Polyhedra. In *Discrete Computational Geometry (proceedings)*, 1992.
- [2] B. Chazelle. Triangulating a Simple Polygon in Linear Time. In *Discrete Computational Geometry (proceedings)*, 1991.
- [3] L. P. Chew. Constrained Delaunay Triangulations. In *Proceedings of the 3rd Annual ACM Symposium on Computational Geometry*, 1987.
- [4] J. Ruppert. A New and Simple Algorithm for Quality 2-Dimensional Mesh Generation. In *Proceedings of the 4th ACM-SIAM Symposium on Discrete Algorithms*, 1993.
- [5] B. Joe. GEOMPACK - A Software Package for the Generation of Meshes Using Geometric Algorithms. In *Advanced Engineering Software*, 1991.
- [6] *Siggraph '93 Course Notes: An Introduction to Physically based Modelling*, 1993.
- [7] William W. Armstrong and Mark W. Green. The dynamics of articulated rigid bodies for purposes of animation. In *The Visual Computer*. Springer Verlag, 1995.
- [8] Michael F. Ashby and David R. H. Jones. *Engineering Materials (Book 1)*. Pergamon Press, 1980.
- [9] Computer Graphics Lab (LIG) at the Swiss Federal Institute of Technology (EPFL). The humanoid project : A real-time and parallel system for the simulation of virtual humans, 1995.
- [10] R. J. Atkin. *An Introduction to the Theory of Elasticity*. Longman Group, London, 1980.
- [11] Francisco Azuola, Norman I. Badler, Pei-Hwa Ho, Ioannis Kakadiaris, Dimitri Metaxas, and Bond-Jay Ting. In *Proc. IMAGE VII Conf., Tuscon, AZ*.
- [12] N. Badler, D. Metaxas, B. Webber, and M. Steedman. The center for human modeling and simulation. *Presence*, 4(1):81–96, 1995.
- [13] N. I. Badler, J. O'Rourke, and H. Toltzis. A spherical representation of a human body for visualizing movement. *Proc. IEEE*, 67:1397–1403, October 1979.

- [14] Norman I. Badler, Cary B. Phillips, and Bonnie Lynn Webber. *Simulating Humans: Computer Graphics Animation and Control*. Oxford University Press, New York, 1993. ISBN 0-19-507359-2.
- [15] David Baraff and Andrew Witkin. Dynamic simulation of non-penetrating flexible bodies. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 303–308, July 1992.
- [16] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. *The Quickhull Algorithm. Technical Report GCG53*. The Geometry Center, University of Minnesota, Minneapolis.
- [17] A. Barr. Superquadrics and angle-preserving transformations. *IEEE Computer Graphics and Applications*, 1(1):11–23, 1981.
- [18] Alan H. Barr. Global and local deformations of solid primitives. In Hank Christiansen, editor, *Computer Graphics (SIGGRAPH '84 Proceedings)*, volume 18, pages 21–30, July 1984.
- [19] Andrew A. Biewener. Biomechanics of mammalian terrestrial locomotion. *Science*, pages 1097–1103, 1990.
- [20] Michel Carignan, Ying Yang, Nadia Magnenat Thalmann, and Daniel Thalmann. Dressing animated synthetic actors with complex deformable clothes. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 99–104, July 1992.
- [21] John E. Chadwick, David R. Haumann, and Richard E. Parent. Layered construction for deformable animated characters. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 243–252, July 1989.
- [22] David T. Chen and David Zeltzer. Pump it up: Computer animation of a biomechanically based model of muscle using the finite element method. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 89–98, July 1992.
- [23] Michael F. Cohen. Interactive spacetime control for animation. *Computer Graphics*, 26(2):293–302, July 1992.
- [24] Michael F. Cohen and Donald P. Greenberg. The hemi-cube, a radiosity solution for complex environments. In *Computer Graphics (SIGGRAPH '85 Proceedings)*, volume 19, pages 31–40, July 1985.
- [25] John D. C. Crisp. Properties of Tendon and Skin. In Y. C. Fung et al., editor, *Biomechanics: Its Foundations and Objectives*. Prentice Hall, 1972.
- [26] Agnes Daldegan, Nadia Magnenat Thalmann, Tsuneya Kurihara, and Daniel Thalmann. An integrated system for modeling, animating and rendering hair. In *Proceedings of Eurographics '93*. Eurographics, 1993.

- [27] Xiao Qi Deng. *A Finite Element Analysis of Surgery of the Human Facial Tissues*. PhD thesis, Columbia University, 1988.
- [28] James Cameron (Director). *Titanic*. Movie, 1997.
- [29] Joe Johnston (Director). *Jumanji*. Movie, 1995.
- [30] Francis A. Duck. *Physical Properties of Tissue : A Comprehensive Reference Book*. Academic Press, 1990.
- [31] G. Dumont, B. Arnaldi, and G. Hegron. Mechanics of solids for computer animation. In *PIXIM '89 (Conference proceedings) Paris*, pages 293–308, October 1998.
- [32] Dave Forsey. A kinematic model for collision response. Master's thesis, Department of Computer Science, University of British Columbia, Vancouver, 1994.
- [33] Y. C. Fung. *Biomechanics: Mechanical properties of Living Tissues*. Springer-Verlag, 1981.
- [34] Marie-Paule Gascuel and Claude Puech. Dynamic Animation of Deformable Bodies. In S. Coquillart, W. Straßer, and P. Stucki, editors, *From Object Modelling to Advanced Visual Communication*. Springer-Verlag, 1994.
- [35] Marie-Paule Gascuel, Anne Verroust, and Claude Puech. Animation with collisions of deformable articulated bodies. In *1st Eurographics workshop on Animation & Simulation*, 1990.
- [36] H. S. Gasser and A. V. Hill. The dynamics of muscular contraction. In *Royal Society of London Proceedings*, volume 96, pages 398–437, 1924.
- [37] T. Gibson, R. M. Kenedi, and J. E. Craik. The mobile micro-architecture of dermal collagen. *British Journal of Surgery*, 1965.
- [38] Cindy M. Goral, Kenneth E. Torrance, and Donald P. Greenberg. Modeling the interaction of light between diffuse surfaces. In *Computer Graphics (SIGGRAPH '84 Proceedings)*, volume 18, pages 213–222, July 1984.
- [39] Jean-Paul Gourret, Nadia Magnenat Thalmann, and Daniel Thalmann. Simulation of object and human skin deformations in a grasping task. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 21–30, July 1989.
- [40] Marc R. Grosso, Richard D. Quach, and Norman I. Badler. Anthropometry for computer animated human figures. In N. Magnenat-Thalmann and D. Thalmann, editors, *State-of-the Art in Computer Animation*, pages 83–96. Springer-Verlag, New York, NY, 1989.
- [41] R. D. Harkness. *Mechanical Properties of Skin in Relation to its Biological Function and its Chemical Components*. Wiley-Interscience, New York, 1977.

- [42] D. Herbison-Evans. Nudes 2: A numeric utility displaying ellipsoid solids, version 2. In *Computer Graphics (Proceedings of Siggraph '78)*, volume 12(3), pages 354–356, 1978.
- [43] D. Herbison-Evans. Real-time animation of human figure drawings with hidden lines omitted. In *IEEE Computer Graphics and Applications*, volume 2(9), pages 27–33, 1982.
- [44] A. V. Hill. The force-velocity relationship in shortening muscle. In *First and Last Experiments in Muscle Mechanics*, pages 23–41. Cambridge University Press, 1970.
- [45] Victor Ng 'Thow Hing. A biomechanical musculotendon model for animating articulated objects. Master's thesis, Department of Computer Science, University of Toronto, 1994.
- [46] J. K. Hodgins. Three-dimensional human running. In *Proceedings of the IEEE Conference on Robotics and Automation*, 1996.
- [47] J. K. Hodgins, W. L. Wooten, D. C. Brogan, and J. F. O'Brien. Animating human athletics. In *Computer Graphics (SIGGRAPH '95 Proceedings)*, 1984.
- [48] Jessica K. Hodgins and Nancy S. Pollard. Adapting simulated behaviour for new characters. In *Computer Graphics (SIGGRAPH '97 Proceedings)*, 1997.
- [49] R. J. Hollands and E. A. Trowbridge. A pc-based virtual reality arthroscopic surgical trainer. In *Proc. Simulation in Synthetic Environments*, pages 17–22, 1996.
- [50] Matthew Holton and Simon Alexander. Soft cellular modelling: A technique for the simulation of non-rigid materials. *Computer Graphics: Developments in Virtual Environments (Proceedings of Computer Graphics International '95)*, 1995.
- [51] Ken ichi Anjyo, Yoshiaki Usami, and Tsuneya Kurihara. A simple method for extracting the natural beauty of hair. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 111–120, July 1992.
- [52] Wayne F. Larrabee Jr. A finite element model of skin deformation - biomechanics of skin and soft tissue: A review. *Laryngoscope '86*, 1986.
- [53] H. Kardestuncer. *Finite Element Handbook*. McGraw-Hill, 1987.
- [54] R. M. Kenedi, T. Gibson, J. H. Evans, and J. C. Berbenel. Tissue mechanics. In *Physics in medicine and biology*, volume 20(5), pages 699–717, 1975.
- [55] R. M. Koch, M. H. Gross, F. R. Carls, D. F. von Büren, G. Fankhauser, and Y. Parish. Simulating facial surgery using finite element methods. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 421–428. ACM SIGGRAPH, Addison Wesley, 1996.

- [56] Ellen Kreighbaum and Katharine M. Barthels. *Biomechanics : A Qualitative Approach for studying human movement*. Burgess Publishing Company, 1995.
- [57] T. Kurihara, K. Anjyo, and D. Thalmann. Hair animation with collision detection. In *Models and Techniques in Computer Animation*, pages 128–138. Springer-Verlag, Tokyo.
- [58] Yuencheng Lee, Demetri Terzopoulos, and Keith Waters. Realistic modeling for facial animation. In Robert Cook, editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 55–62. ACM SIGGRAPH, Addison Wesley, August 1995.
- [59] Ming C. Lin and John F. Canny. Efficient collision detection for animation. In *3rd Eurographics workshop on Animation & Simulation*, 1992.
- [60] A. Liu and B. Joe. Relationship between tetrahedral shape measures. In *Proceedings of BIT*, volume 34, 1994.
- [61] Zicheng Liu, Steven J. Gortler, and Michael F. Cohen. Hierarchical spacetime control. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 35–42. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [62] Nadia Magnenat-Thalmann, Richard Laperriere, and Daniel Thalmann. Joint-dependent local deformations for hand animation and object grasping. In *Proceedings of Graphics Interface '88*, pages 26–33, June 1988.
- [63] Nadia Magnenat-Thalmann and Daniel Thalmann. The Direction of Synthetic Actors in the Film *rendez-vous à montréal*. *IEEE Computer Graphics & Applications*, pages 9–19, Dec 1987.
- [64] Delle Maxwell. Graphical marionette: a modern-day pinocchio. Master's thesis, MIT, Cambridge MA, 1983.
- [65] Thomas A. McMahon. Fundamental muscle mechanics. In *Muscles, Reflexes, and Locomotion*, pages 3–26. Princeton University Press, 1984.
- [66] R. McNeil and L. Alexander. *Animal Mechanics*. Sidgwick and Jackson, 1969.
- [67] Gavin S. P. Miller. The motion dynamics of snakes and worms. In John Dill, editor, *Computer Graphics (SIGGRAPH '88 Proceedings)*, volume 22, pages 169–178, August 1988.
- [68] S. Mitchell and S. A. Vavasis. Quality mesh generation in three dimensions. In *Proceedings of the ACM Computational Geometry Conference*, pages 212–221, 1992.
- [69] Matthew Moore and Jane Wilhems. Collision detection and response for computer animation. *Computer Graphics*, 22(4):289–298, August 1988.

- [70] Edward Muybridge. *Animals in locomotion*. London, 1899.
- [71] B. Nath. *Fundamentals of finite elements for engineers*. London, Athlone Press, 1974.
- [72] J. Thomas Ngo and Joe Marks. Spacetime Constraints Revisited. In *Proceedings of SIGGRAPH '93*, pages 343–350, 1993. In Computer Graphics proceedings, Annual Conference Series.
- [73] Margareta Nordin and Victor H. Frankel. Biomechanics of bone. In *Basic Biomechanics of the Musculoskeletal System (second edition)*, pages 3–29. Lea and Febiger, 1989.
- [74] National Library of Medicine. The visible human project. http://www.nlm.nih.gov/research/visible/visible_human.html, 1995.
- [75] Steven Pieper, Joseph Rosen, and David Zeltzer. Interactive graphics for plastic surgery: A task-level analysis and implementation. In David Zeltzer, editor, *Computer Graphics (1992 Symposium on Interactive 3D Graphics)*, volume 25:2, pages 127–134, March 1992.
- [76] Mark I. Pitman and Lars Peterson. Biomechanics of skeletal muscle. In *Basic Biomechanics of the Musculoskeletal System (second edition)*, pages 89–111. Lea and Febiger, 1989.
- [77] John C. Platt and Alan H. Barr. Constraint methods for flexible models. In John Dill, editor, *Computer Graphics (SIGGRAPH '88 Proceedings)*, volume 22, pages 279–288, August 1988.
- [78] S. M. Platt. Animating facial expressions. In *Computer Graphics (SIGGRAPH '81 Proceedings)*, volume 15, pages 245–252, August 1981.
- [79] T. E. Potter and K. D. Willmert. Three-dimensional human display model. In *Computer Graphics*, volume 9(1), pages 102–110, 1975.
- [80] E. H. Rose, L. M. Vistnes, and G. A. Ksander. A microarchitectural model of the regional variations in hypodermal mobility in porcine and human skin. In *Annals of Plastic Surgery*, volume 3, pages 252–266, 1978.
- [81] Mark A. Sagar, David Bullivant, Gordon D. Mallinson, Peter J. Hunter, and Ian W. Hunter. A virtual environment and model of the eye for surgical simulation. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 205–213. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [82] Ferdi Scheepers, Richard E. Paerent, Wayne E. Carlson, and Stephen F. May. Anatomy-based modeling of the human musculature. In *Computer Graphics (SIGGRAPH '97 Proceedings)*, pages 163–172, 1997.

- [83] Thomas W. Sederberg and Scott R. Parry. Free Form Deformation of Solid Geometric Models. *ACM SIGGRAPH Computer Graphics*, 20(4):151–160, Aug 1986.
- [84] Jonathan Smith. *Computer Modelling of the General Dynamics of Deformable Bodies*. PhD thesis, University of Bristol, 1992.
- [85] Demetri Terzopoulos and Kurt Fleischer. Deformable models. *The Visual Computer*, 4(6):306–331, December 1988.
- [86] Demetri Terzopoulos and Kurt Fleischer. Modeling inelastic deformation: Viscoelasticity, plasticity, fracture. *Computer Graphics*, 22(4):269–278, August 1988.
- [87] Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. Elastically deformable models. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 205–214, July 1987.
- [88] D. Thalmann, N. Magnenat-Thalmann, and P. Bergeron. Dream flight: A fictional film produced by 3d computer animation. In *Proceedings of Computer Graphics '82*, pages 353–368, 1982.
- [89] Nadia Magnenat Thalmann. Making 3d clothes for synthetic actors. In J. Vince and L. MacDonald, editors, *Interacting with Images*. John Wiley, 1994.
- [90] Nadia Magnenat Thalmann and Daniel Thalmann. *Computer Animation Theory and Practice*. Springer-Verlag, 1990.
- [91] F. Thomas and O. Johnston. *Disney Animation: The Illusion of Life*. Abberville Press, 1981.
- [92] R. T. Tregear. *Physical Functions of Skin*. Academic Press, London, 1966.
- [93] Russell Turner. Leman: A system for constructing and animating layered elastic characters. In *Computer Graphics: Development in virtual environments (Proceedings for CGI'95)*, 1995.
- [94] J. F. V. Vincent and J. D. Currey. *The Mechanical Properties of Biological Materials*. Cambridge University Press, 1980.
- [95] Pascal Volino, Martin Courchesne, and Nadia Magnenat-Thalmann. Versatile and efficient techniques for simulating cloth and other deformable objects. *Computer Graphics Proceedings, SIGGRAPH'95*, 1995.
- [96] Keith Waters. A muscle model for animating three-dimensional facial expression. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 17–24, July 1987.
- [97] Alan Watt and Mark Watt. *Advanced Animation and Rendering Techniques*. Addison-Wesley, 1993.

- [98] J. Wilhelms and B. A. Barsky. Using dynamic analysis to animate articulated bodies such as humans and robots. In M. Wein and E. M. Kidd, editors, *Graphics Interface '85 Proceedings*, pages 97–104. Canadian Inf. Process. Soc., 1985.
- [99] Jane Wilhelms. Using dynamic analysis for realistic animation of articulated bodies. *IEEE Computer Graphics and Applications*, pages 12–27, 1987.
- [100] Jane Wilhelms. Dynamic experiences. In Norman I. Badler, Brian A. Barsky, and David Zeltzer, editors, *Making them move: mechanics, control, and animation of articulated figures*, pages 265–279. Morgan Kaufmann, 1991.
- [101] Jane Wilhelms and Allen Van Gelder. Anatomically based modeling. In *Computer Graphics (SIGGRAPH '97 Proceedings)*, pages 173–180, 1997.
- [102] Peter Williams, Roger Warwick, Mary Dyson, and Lawrence Bannister. *Gray's anatomy*. Churchill Livingstone, thirty-seventh edition edition, 1989.
- [103] C. Withrow. A dynamic model for computer-aided choreography. Technical Report 70-103, Computer Science Department, University of Utah, 1970.
- [104] Andrew Witkin and Michael Kass. Spacetime constraints. *Computer Graphics*, 22(4):159–168, August 1988.
- [105] V. Wright. Elasticity and deformation of the skin. In H. R. Elden, editor, *Biophysical Properties of the Skin*. Wiley-Interscience, New York, 1977.
- [106] Shin ya Miyazaki, Masao Ishiguro, Takami Yasuda, Shigeki Yokoi, and Jun ichiro Toriwaki. A study of virtual manipulation of elastic objects. *Computer Graphics: Developments in Virtual Environments (Proceedings of Computer Graphics International '95)*, 1995.
- [107] Felix E. Zajac. Muscle and tendon: Properties, models, scaling and application to biomechanics and motor control. In *Critical Reviews in Biomedical Engineering*, volume 17(4), pages 359–411, 1989.
- [108] O.C. Zienkiewicz. *The Finite Element Method*. McGraw-Hill, 1982.

Appendix A

Guide to the CDROM

The files on the CDROM are in two main directories. The `\model` directory contains images and animations showing the structure of the various tissue regions within the model and is relevant to Chapter Four. The `\results` directory shows results from the material tests that are described in Chapter Five, the model test simulations described in Chapter Six, and the verification study in Chapter Seven. The correspondence between the figures in this thesis and the images on the CDROM is as follows:

Chapter 4 - Modelling the Structure of Limbs

Figure	Filename
4.2a	<code>\model\skin_1.tga</code>
4.2b	<code>\model\skin_2.tga</code>
4.2c	<code>\model\skin_3.tga</code>
4.2d	<code>\model\skin_4.tga</code>
4.3a	<code>\model\int_geom_1.tga</code>
4.3b	<code>\model\int_geom_2.tga</code>
4.3c	<code>\model\int_geom_3.tga</code>
4.3d	<code>\model\int_geom_4.tga</code>
4.10a	<code>\model\ext_phy_2.tga</code>
4.10b	<code>\model\ext_phy_1.tga</code>
4.10c	<code>\model\int_phy.tga</code>

Chapter 5 - Modelling the Physical Properties of Limbs

Figure	Filename
5.4	<code>\results\A_material_1\A_images.tif</code>
5.10	<code>\results\B_material_2\B_images.tif</code>
5.13	<code>\results\C_material_3\C_images_1.tif</code>
5.14	<code>\results\C_material_3\C_images_2.tif</code>
5.18	<code>\results\D_muscle\D_images_1.tif</code>

Chapter 6 - Results

6.1 Gravity Effects

Figure	Filename
6.1a	\results\E_gravity\E_before.tga
6.1b	\results\E_gravity\E_after.tga

6.2 Tendon Pre-tensioning

Figure	Filename
6.2a	\results\F_tendon_pre\F_int_phy_1.tga
6.2b	\results\F_tendon_pre\F_int_phy_2.tga
6.2c	\results\F_tendon_pre\F_int_geom_1.tga
6.2d	\results\F_tendon_pre\F_int_geom_2.tga

6.3 Bending of the Knee Caused by an External Object

Figure	Filename
6.3a	\results\G_flex_ext\G_phy_sk_1.tga
6.3b	\results\G_flex_ext\G_phy_sk_2.tga
6.3c	\results\G_flex_ext\G_phy_sk_3.tga
6.3d	\results\G_flex_ext\G_phy_sk_4.tga
6.4a	\results\G_flex_ext\G_geom_sk_1.tga
6.4b	\results\G_flex_ext\G_geom_sk_2.tga
6.4c	\results\G_flex_ext\G_geom_sk_3.tga
6.4d	\results\G_flex_ext\G_geom_sk_4.tga
6.5	\results\G_flex_ext\G_phy_crease.tga
6.6	\results\G_flex_ext\G_geom_crease_1.tga
6.7	\results\G_flex_ext\G_geom_crease_2.tga
6.8	\results\G_flex_ext\G_geom_crease_3.tga
6.9	\results\G_flex_ext\G_geom_skel.tga

6.4 Activating the “Hamstrings”

Figure	Filename
6.10a	\results\H_muscle_back\H_ext_phy_1.tga
6.10b	\results\H_muscle_back\H_ext_phy_2.tga
6.10c	\results\H_muscle_back\H_ext_phy_3.tga
6.10d	\results\H_muscle_back\H_ext_phy_4.tga
6.11a	\results\H_muscle_back\H_skin_1.tga
6.11b	\results\H_muscle_back\H_skin_2.tga
6.11c	\results\H_muscle_back\H_skin_3.tga
6.11d	\results\H_muscle_back\H_skin_4.tga
6.12a	\results\H_muscle_back\H_int_phy_1.tga
6.12b	\results\H_muscle_back\H_int_phy_2.tga
6.12c	\results\H_muscle_back\H_int_phy_3.tga
6.12d	\results\H_muscle_back\H_int_phy_4.tga
6.13a	\results\H_muscle_back\H_int_geom_1.tga
6.13b	\results\H_muscle_back\H_int_geom_2.tga
6.13c	\results\H_muscle_back\H_int_geom_3.tga
6.13d	\results\H_muscle_back\H_int_geom_4.tga
6.14a	\results\H_muscle_back\H_ext_phy_bk.tga
6.14b	\results\H_muscle_back\H_skin_back.tga
6.14c	\results\H_muscle_back\H_int_phy_bk.tga
6.14d	\results\H_muscle_back\H_int_geom_bk.tga

6.5 Activating the “Quadriceps”

Figure	Filename
6.15a	\results\I_muscle_front\I_skin_1.tga
6.15b	\results\I_muscle_front\I_ext_phy_1.tga
6.15c	\results\I_muscle_front\I_skin_2.tga
6.15d	\results\I_muscle_front\I_ext_phy_2.tga
6.15e	\results\I_muscle_front\I_skin_3.tga
6.15f	\results\I_muscle_front\I_ext_phy_3.tga
6.16a	\results\I_muscle_front\I_int_geom_1.tga
6.16b	\results\I_muscle_front\I_int_phy_1.tga
6.16c	\results\I_muscle_front\I_int_geom_2.tga
6.16d	\results\I_muscle_front\I_int_phy_2.tga
6.16e	\results\I_muscle_front\I_int_geom_3.tga
6.16f	\results\I_muscle_front\I_int_phy_3.tga
6.17	\results\I_muscle_front\I_musc.tga
6.18	\results\I_muscle_front\I_skin_front.tga

6.6 Isometric Contraction of Opposing Muscle Regions

Figure	Filename
6.19	\results\J_both_muscles\J_skin_off.tga
6.20	\results\J_both_muscles\J_skin_on.tga
6.21	\results\J_both_muscles\J_musc_off.tga
6.22	\results\J_both_muscles\J_musc_on.tga

6.7 Isometric Contraction of a Single Muscle Region

Figure	Filename
6.23	\results\K_isometric\K_skin_on.tga
6.24	\results\K_isometric\K_skin_off.tga
6.25	\results\K_isometric\K_close_off.tga
6.26	\results\K_isometric\K_close_on.tga
6.27	\results\K_isometric\K_int_off.tga
6.28	\results\K_isometric\K_int_on.tga
6.29	\results\K_isometric\K_back.tga

Chapter 7 - Verification and Discussion

Figure	Filename
7.1	\results\Verification\comp4.tga
7.2	\results\Verification\comp5.tga
7.3	\results\Verification\comp6.tga
7.4	\results\Verification\comp7.tga
7.5	\results\Verification\comp3.tga
7.6	\results\Verification\comp2.tga
7.7	\results\Verification\comp1.tga

Appendix B

Preventing Collapse Within Mass-Spring-Damper Models of Deformable Objects

This appendix contains a paper published in the proceedings of The Fifth International Conference in Central Europe on Computer Graphics and Visualization, held at the University of West Bohemia, Czech Republic, 10th-14th February 1997.

Preventing Collapse Within Mass-Spring-Damper Models of Deformable Objects

Lee Cooper (L.Cooper@dcs.shef.ac.uk)
Steve Maddock (S.Maddock@dcs.shef.ac.uk)
Department of Computer Science
University of Sheffield
211 Portobello St., Sheffield, S1 4DP, UK.

Mass-spring-damper (MSD) models have been used by a number of researchers in attempts to model deformable objects. One of the main advantages of this approach is its simplicity. This means that it can be implemented quickly and it is also relatively easy to integrate with other techniques e.g. for collision detection or to impose constraints. However, the simplicity of the MSD models commonly used means that they can exhibit unwanted behaviour under certain conditions. The worst aspect of this is the tendency of MSD models to collapse under relatively small compressive forces. This paper presents a new robust model which prevents collapse.

1. Introduction

Computer models of deformable objects have a wide range of potential applications in a diverse range of fields including animation for entertainment [3] and medicine[12][14][6][7][4]. In order to create such models, several methods have been proposed. One approach [8][2][4][9][15][17] has been to use Finite Element and Finite Difference Methods [11] (FEM and FDM) which have been used for many years in engineering applications to model mechanical components and structures. These methods produce accurate results but are computationally expensive and complex. In order to create models which are fast to simulate and easy to extend, we have been concentrating on an alternative approach using MSD models [18][10][3][5]. Our aim is to use these to represent deformable human tissue in models of human limbs. This is a particularly demanding application since the model has to cope with the extreme deformation that occurs around joints. This paper describes a new model that has been developed for this purpose.

2. MSD Models

Figure 1 shows a typical two dimensional MSD model of a deformable ball resting on a fixed surface. The model consists of a number of point masses (the nodes) connected by springs and dampers (the edges). Although several arrangements of springs and dampers have been used, the most common one consists of a spring and damper connected in parallel as seen in figure 2. Here F is the force exerted by the damped spring at the endpoints, x is the current length of the spring, and k and l are the spring and damping constants respectively. The force exerted by this spring/damper arrangement is given by:

$$F = k(x_o - x) - l\dot{x} \quad (1)$$

where x_o is the original (rest) length of the spring.

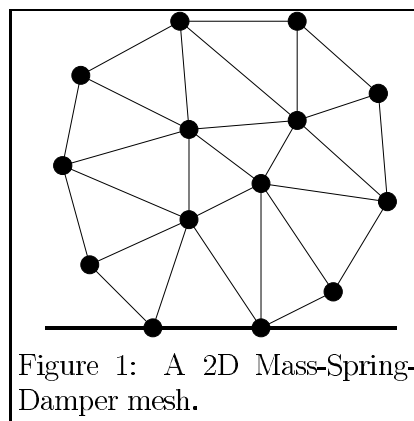


Figure 1: A 2D Mass-Spring-Damper mesh.

In order to simulate an MSD mesh the resultant force acting on each point mass by the springs and dampers is computed. The acceleration of each point mass is then determined using Newton's second law. Then the velocity and position of the point masses are found using a suitable numerical integration technique (See [1] for more details).




Figure 2: A damped spring.

3. Analysing Collapse in MSD Elements

In order to simplify this discussion we will deal with two dimensional meshes, although the ideas can be easily extended to three dimensions. We will also restrict the analysis to triangular meshes since meshes containing general polygons can be decomposed into triangular meshes. When modelling a deformable object using a triangular mesh each triangle can be thought of as an “element” of material. The behaviour of these basic elements (and the hence the model in general) can be altered by changing the way that the elements deform based on applied forces. In this paper, we will assume that the material being modelled is homogeneous and hence the behaviour of the elements throughout the material will be based on the same model even though their shape will vary.

Figure 3 shows a single element in a local coordinate system. Here the origin is made coincident with one of the vertices (A) and the x -axis is oriented such that it passes through a second vertex (B). The third vertex will then lie somewhere in the xy plane. In order to describe the current shape of a triangle, only three variables are required. In this case these are the x position of B and the x and y positions of C . If we consider the element shown in figure 3, then in general the element will exert a force on point C with the force vector \underline{F} being a function of the following variables:

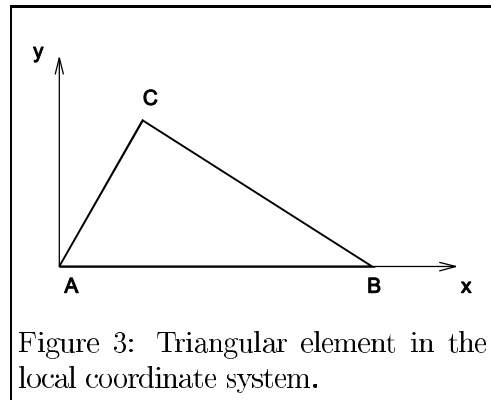


Figure 3: Triangular element in the local coordinate system.

$$\underline{F} = f(\underline{B}, \underline{C}, AB_o, BC_o, CA_o, k) \quad (2)$$

where \underline{B} and \underline{C} are the position vectors of points B and C , AB_o is the initial length of edge AB (and similarly for the other edges), and k is a variable relating to the elasticity of the material. In its initial state, the element is assumed to be in its rest position where all the forces on the vertices are zero. As it stands, equation 2 is dependent only on the position of the points and not their velocities. Velocity information is needed to model damping forces within the element, however this will be ignored for now since we are primarily interested in the final equilibrium state of the elements after any transient motion has died down. The force function (2) determines the force that is exerted by the elements on the point masses and hence the behaviour of the material. In order to compute the forces on a particular point mass within the mesh due to a particular element, we transform the element into the local coordinate system such that the point we are calculating the force for lies in the xy plane (corresponding to C) and the other vertices for the element are in positions corresponding to A and B (i.e. at the origin and aligned with the x -axis).

This is illustrated in figure 4. This method of applying a local coordinate system allows the problem of element design to be reduced to the problem of designing a suitable force function.

3.1 Standard Element

The simplest element design uses three linear springs along the edges of the element. This gives the following force function (note that damping forces are excluded):

$$\underline{F} = k(CA_o - |\underline{C}|) \frac{\underline{C}}{|\underline{C}|} + k(BC_o - |\underline{BC}|) \frac{\underline{BC}}{|\underline{BC}|} \quad (3)$$

where $\underline{BC} = \underline{C} - \underline{B}$ according to standard vector notation. In order to consider how the force on C changes as it moves with respect to A and B , the function was evaluated for the situation where $AB_o = 1$, $BC_o = 1$, $CA_o = \sqrt{2}$, $k = 1$, $\underline{B} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, and the point C is moved in the positive xy quadrant. These values of AB_o , BC_o and CA_o correspond to the rest state shown in figure 5.

The magnitude of the force on C is shown on the graph in figure 6. The graph shows two points at which the force is zero. One is at the initial position of $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and the other occurs on the x axis around $\begin{pmatrix} 0.6 \\ 0 \end{pmatrix}$. This second point occurs when the forces from the two springs (AC and BC) cancel each other out. The shape of this graph indicates why meshes composed of this type of element readily collapse. If we consider the element being compressed vertically, then we can see that if a suitably large force is applied it is possible to push C such that it lies on the x -axis. Consideration of the force function also shows that the graph below the x -axis (not shown) is a mirror image of that above. This means that there is a third point where the force acting on C becomes zero at $\underline{C} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ (due to the spring lengths becoming equal to their rest length) as illustrated in figure 7.

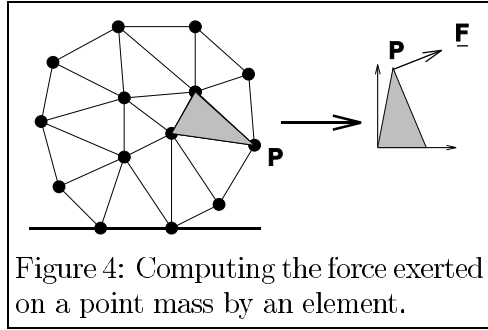


Figure 4: Computing the force exerted on a point mass by an element.

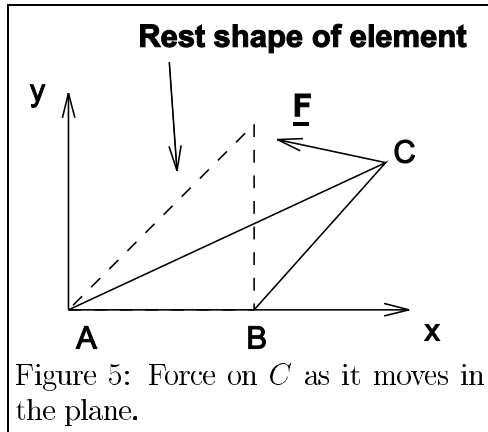


Figure 5: Force on C as it moves in the plane.




Figure 6: Magnitude of the force on C at different points within the plane.

This gives rise to two major problems. Firstly, it is possible to push the vertices past the opposite edge if a suitably large force is applied. In a realistic model this would not be possible since as the point approached the edge the material within the element would become more and more compressed producing reaction forces on C which tended to infinity as the point tended towards the edge. Secondly, the problem is compounded by the fact that if the point does cross the line through the opposite edge, then the element flips into an alternative state where the reaction forces act towards the rest point on the wrong side of the line. The term “flip” is appropriate since if these elements are simulated and this situation occurs the element suddenly changes shape. Another result of this “flipping” behaviour is the fact that even if the compression force is removed the element does not regain its original shape unless a stretching force is applied which is sufficient to flip the element back into the correct configuration. In order to try and prevent this collapse, we need to use a different force function. This is considered in the next section.




Figure 7: Alternative rest point.

4. Preventing Collapse

This section describes a number of methods that have been considered by ourselves and other researchers to try and prevent collapse occurring. It then describes a new method which we have developed.

4.1 Element Containing Springs to the Centroid

One approach that has been suggested [18] involves adding three more springs from the vertices to the centroid of the element. The reasoning behind this approach is that as the standard element is compressed its resistance to further compression reduces since the springs become less and less aligned with the crush direction. However, with springs to the centroid it is hoped that at least one of the springs should remain reasonably well aligned. The force magnitude plot for this element design is shown in figure 8. The graph shows that these extra springs make very little difference to the force function and the element will still compress fully if a suitably large force is applied. This is due to the fact that the linear springs used can be compressed fully. The next section considers the use of an alternative non-linear spring.




Figure 8: Force magnitude for springs along edges and to centroid.

4.2 Element Composed of Non-Linear Springs

In reality springs do not behave in a linear manner when highly compressed. Generally, as the spring length tends towards zero the spring will exert reaction forces which tend to infinity. In order to try and model this effect, we have devised the following function for a non-linear spring which resists collapse:

$$\begin{aligned} f &= k(x_o - x) + \frac{kx_o^2}{x} - kx_o \\ &= k\left(\frac{x_o^2}{x} - x\right) \end{aligned} \quad (4)$$

This combines the term for a linear spring with a non-linear term which tends to infinity as x tends to 0. The $-kx_o$ is also needed to ensure that $f = 0$ at $x = x_o$. A graph of the force exerted by this spring as a function of the displacement is shown in figure 9 (for $k = 1$, $x_o = 1$), together with the force from a linear spring. If we now replace the linear springs in the standard element with the new non-linear spring, we obtain the force magnitude plot shown in figure 10.

It can be seen that the force required to crush the element such that C is close to points A and B now becomes very large (tending to infinity as C tends towards the points themselves). However, there are still points along the x axis at which the force is finite, indicating that it is still possible to flip the element.

4.3 A New Element Design Based on Edge Repulsion Forces

In order to prevent this, a new type of spring is presented which exerts a force based on the perpendicular distance of the point from the opposite edge. The force is always away from the edge and its magnitude is given by the non-linear spring function in equation 4.

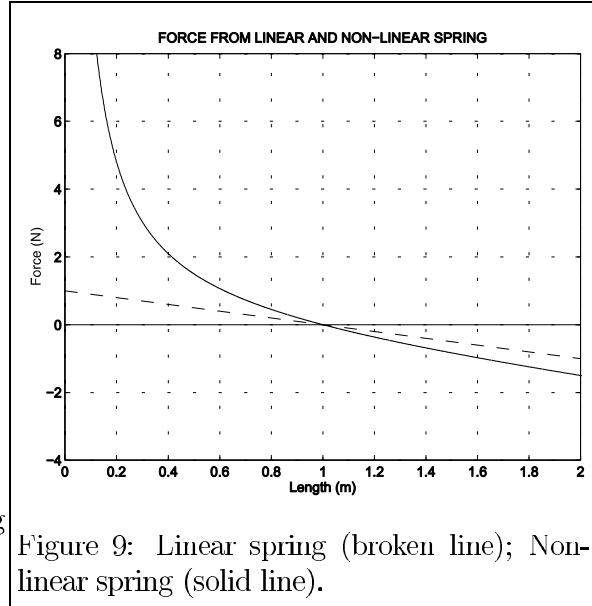


Figure 9: Linear spring (broken line); Non-linear spring (solid line).




Figure 10: Element using non-linear springs.

This method can essentially be thought of as the edges exerting a repulsion force on the points opposite which act directly away from the edges. This is illustrated in figure 11. A new element has been designed which consists of three damped linear springs along the edges and edge repulsion springs for each point. This will be referred to as the “ER” (edge repulsion) element. The new force function is given by,

$$\underline{F} = k(CA_o - |\underline{C}|) \frac{\underline{C}}{|\underline{C}|} + k(BC_o - |\underline{BC}|) \frac{\underline{BC}}{|\underline{BC}|} + k\left(\frac{EC_o^2}{|\underline{EC}|} - |\underline{EC}|\right) \frac{\underline{EC}}{|\underline{EC}|} \quad (5)$$

where E is the projection of C onto the line passing through AB , and has a position vector given by $\underline{E} = \begin{pmatrix} C_x \\ 0 \end{pmatrix}$.

The force magnitude graph for this new element is shown in figure 12. It can be seen that the reaction force now becomes infinite as we tend towards the x-axis, thus preventing the element from collapsing.

5. Results Using the New Model

This section describes a system that has been implemented to test the new element along with results that have been obtained. The simulation system uses a simple Euler method (see [13] for details) for computing the new position and velocity variables of the particles after each time step. We also implement a simple constraint method in

which point masses can be attached to rigid bodies (represented by closed polygons in 2D). In order to test the new elements, a highly simplified model of a limb was used (see figures 13 and 14). This consists of two rigid bodies (bones) surrounded by a layer of deformable material (tissue). The upper bone is held in a fixed position whilst the lower one is rotated around a point between the two bodies (representing a simple hinge joint). The first series of images shows the deformation which occurs when the standard elements are used (damped linear springs along each edge).

Figure 13a shows the limb in the initial straight position. At this stage, all the elements are in their rest state. The joint angle was then slowly increased at a rate of approximately 10 degrees per second of simulation time. At 45 degrees (figure 13b) some of the elements are already showing considerable compression, particularly around the inside corner of the joint. At 47 degrees (figure 13c) one of the elements collapses. At this stage the motion of the limb was stopped and the collapsed element eventually settled at the shape shown

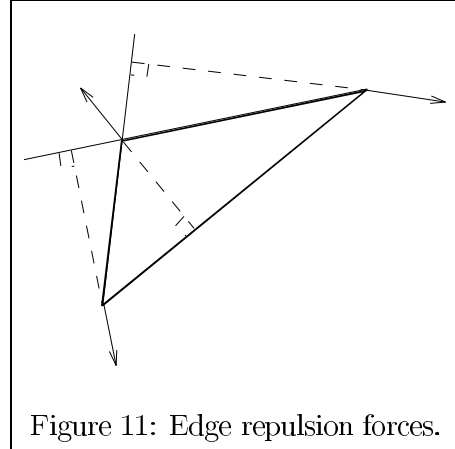


Figure 11: Edge repulsion forces.




Figure 12: Magnitude of force on C for ER-element.

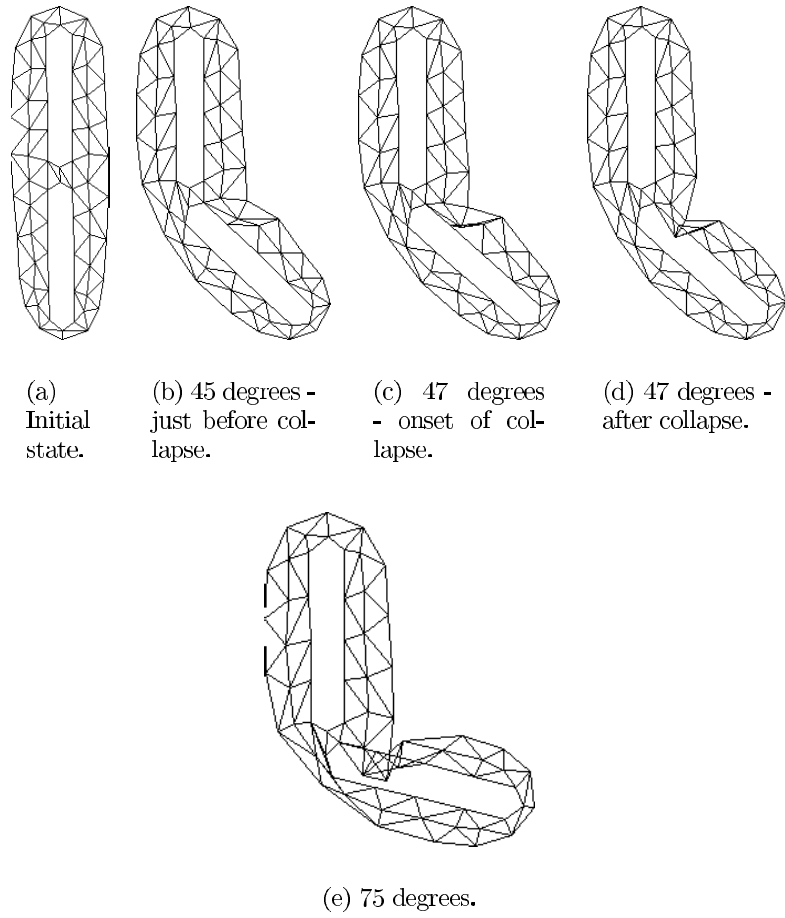


Figure 13: Deformation of standard elements at different joint angles.




Figure 14: Deformation of new elements at different joint angles.

in figure 13d. As the limb angle is increased further, other elements begin to collapse leading to significant divergence from the “expected” deformation (see figure 13e).

Figure 14 shows the deformation of a mesh composed of the ER-elements under exactly the same conditions. Figure 14a shows the deformation at a joint angle of 45 degrees, and it can be seen that there is significantly less compression compared to the model in 13b. As the joint angle is increased, some elements on the inside of the joint become highly compressed (see fig 14b), but they do not collapse. As well as remaining stable, the model also exhibits effects that occur with real deformable material. For example, slight bulging of the material can be seen and a natural crease appears on the inside of the joint. Finally, figure 14c shows the model at a joint angle of 110 degrees.

6. Conclusion

We have implemented a new element type for MSD models which solves the problem of collapse that occurs in existing models. Images are provided from a simulation (figures 13 and 14) which compare the performance of the new element with a standard element. Although this paper has concentrated on two dimensions, extending the model into three dimensions is straightforward. Rather than using a triangular element, tetrahedral elements are used having damped springs along each of the 6 edges. The edge repulsion now becomes face repulsion in which each of the vertices is repelled by the plane containing the opposite face in the tetrahedron.

Current work is concentrating on implementing the 3D model and we are also considering the development of new force functions which result in more realistic models. We are also interested in modelling other material properties, for example incompressibility, plasticity and fracture[16].

Acknowledgements

This research was funded by the EPSRC.

References

- [1] *Siggraph '93 Course Notes: An Introduction to Physically based Modelling*, 1993.
- [2] David Baraff and Andrew Witkin. Dynamic simulation of non-penetrating flexible bodies. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 303–308, July 1992.
- [3] John E. Chadwick, David R. Haumann, and Richard E. Parent. Layered construction for deformable animated characters. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 243–252, July 1989.
- [4] David T. Chen and David Zeltzer. Pump it up: Computer animation of a biomechanically based model of muscle using the finite element method. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 89–98, July 1992.
- [5] Lee Cooper. Modelling of human skin and tissue for realistic computer animation. Master’s thesis, Department of Computer Science, University of Manchester, 1995.

- [6] S. Delp, P. Loan, M. Hoy, F. Zajac, S. Fisher, and J. Rosen. An interactive graphics-based model of the lower extremity to study orthopaedic surgical procedures. *IEEE Transactions on Biomedical Engineering*, 37:8, August 1990.
- [7] Xiao Qi Deng. *A Finite Element Analysis of Surgery of the Human Facial Tissues*. PhD thesis, Columbia University, 1988.
- [8] Marie-Paule Gascuel and Claude Puech. Dynamic Animation of Deformable Bodies. In S. Coquillart, W. Straßer, and P. Stucki, editors, *From Object Modelling to Advance Visual Communication*. Springer-Verlag, 1994.
- [9] Jean-Paul Gourret, Nadia Magnenat Thalmann, and Daniel Thalmann. Simulation of object and human skin deformations in a grasping task. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 21–30, July 1989.
- [10] Matthew Holton and Simon Alexander. Soft cellular modelling: A technique for the simulation of non-rigid materials. *Computer Graphics: Developments in Virtual Environments (Proceedings of Computer Graphics International '95)*, 1995.
- [11] B. Nath. *Fundamentals of finite elements for engineers*. London, Athlone Press, 1974.
- [12] Steven Pieper, Joseph Rosen, and David Zeltzer. Interactive graphics for plastic surgery: A task-level analysis and implementation. In David Zeltzer, editor, *Computer Graphics (1992 Symposium on Interactive 3D Graphics)*, volume 25:2, pages 127–134, March 1992.
- [13] W.H. Press, B.P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, 1986.
- [14] Mark A. Sagar, David Bullivant, Gordon D. Mallinson, Peter J. Hunter, and Ian W. Hunter. A virtual environment and model of the eye for surgical simulation. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 205–213. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [15] Demetri Terzopoulos and Kurt Fleischer. Deformable models. *The Visual Computer*, 4(6):306–331, December 1988.
- [16] Demetri Terzopoulos and Kurt Fleischer. Modeling inelastic deformation: Viscoelasticity, plasticity, fracture. *Computer Graphics*, 22(4):269–278, August 1988.
- [17] Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. Elastically deformable models. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 205–214, July 1987.
- [18] Shin ya Miyazaki, Masao Ishiguro, Takami Yasuda, Shigeki Yokoi, and Jun ichiro Toriwaki. A study of virtual manipulation of elastic objects. *Computer Graphics: Developments in Virtual Environments (Proceedings of Computer Graphics International '95)*, 1995.